

WORKING PAPER · NO. 2019-42

Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data

John Grigsby, Erik Hurst, Ahu Yildirmaz

FEBRUARY 2019

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.

1 Introduction

Nominal rigidities are an important component of many models of aggregate fluctuations. A large literature has developed using micro data to establish key moments of output price adjustment, which have guided theories of the nature of nominal price rigidities.¹ However, the literature using micro data to document nominal wage adjustment is far less extensive. As a result, the nature of nominal wage stickiness remains a central question within both the macroeconomics and labor economics literatures. For example, at the 2014 Jackson Hole Symposium, Janet Yellen speculated that downward nominal wage rigidity was an important contributor both to why wages did not fall more during the Great Recession and why they did not increase at a faster rate during the subsequent recovery.

There are three reasons why the literature using micro data to measure nominal wage adjustment has remained underdeveloped. First, existing data sets are not well-suited to measure the extent of nominal wage rigidities. Household surveys often define the nominal wage by dividing self-reported earnings by self-reported hours. Any measurement error in either earnings, hours worked, or self-reported hourly wages can result in a substantial upward bias in the volatility of individual wage changes. Administrative datasets, on the other hand, have high quality panel data on quarterly or annual earnings but usually lack the measures of individual hours worked necessary to construct a wage.²

Second, the composition of compensation varies across workers and over time. For example, worker compensation includes their guaranteed contract earnings as well as commissions, tips, bonuses, performance pay, overtime premiums, and employer-provided fringe benefits. Existing household and administrative datasets do not decompose the different types of compensation into their components nor do they include measures of employer-provided fringe benefits. In many theories of employment dynamics, the *present value* of worker earnings determine labor market fluctuations. We document that the persistence and cyclicity of contract wage changes are orders of magnitudes higher than the persistence and cyclicity of changes in other forms of compensation such as bonuses. This suggests that the extent to which contract wages adjust may be a more informative moment for models of labor market fluctuations where the user cost of a worker is important. Adjustments in transient forms of

¹See, for example, Bils and Klenow (2004), Klenow and Kryvtsov (2008), and Nakamura and Steinsson (2008) for important contributions.

²There are some exceptions. Barattieri et al. (2014) attempts to correct for measurement error in self-reported hourly wages among SIPP respondents when examining nominal wage adjustments. Kurmann and McEntarfer (2018) and Jardim et al. (2019) use administrative data from Washington state which records measures of both earnings and hours to explore changes in earnings-per hour. Throughout the paper, we will discuss how having administrative data on actual worker wages contrasts with the results of these papers and other papers in the literature.

compensation, such as an annual bonus, will appear as flexibility in average hourly earnings but may not have allocative consequences in the aggregate, much in the way that sales have been considered largely irrelevant for aggregate fluctuations in the output price literature. Therefore, to the extent that different forms of compensation adjust differently, there are gains to understanding adjustment patterns of each form of compensation separately.

Finally, different models employ different notions of nominal wage rigidities. Many models of frictional labor markets require separate measures of nominal wage adjustments for those who remain with the same employer and for those who switch employers. In many of these models, it is the flexibility of new hire wages that is important for aggregate employment fluctuations. However, most New Keynesian models do not make a distinction between job-stayers and job-changers and instead require measures of aggregate nominal wage adjustments including both margins of adjustment. For example, in order to match aggregate employment and wage dynamics during and after the Great Recession, one would need measures of nominal wage changes inclusive of those who remain on the job and those who switch jobs. The variety of wage rigidity notions puts further requirements on data seeking to understand nominal adjustments.

In this paper, we use administrative data from ADP, LLC (henceforth known as ADP) – one of the world’s largest payroll processing companies – to produce a series of new facts about aggregate nominal wage adjustment in the U.S. over the last decade. Our data set is unique in that it: (1) includes administrative anonymized records of workers’ per-period nominal wage, (2) has a sample of about 20 million workers per month which are generally representative of the US population, (3) provides data for the universe of workers within a firm, (4) allows workers to be tracked both within and across firms over time, (5) includes administrative records on various other forms of compensation including bonuses and fringe benefits, and (6) spans multiple years so as to examine business cycle variation. The data allow us to compute measures of nominal wage adjustments separately for job-stayers and job-changers, as well as construct a composite aggregate measure including both stayers and changers. Additionally, we create measures of nominal wage adjustments both with and without bonuses and employer-provided fringe benefits. Finally, we explore the cyclicity of new hire wages. Collectively, our results paint a relatively complete picture of nominal wage adjustments for workers and firms in the U.S. during from 2008 to 2016.

We first describe the composition of compensation. We measure a worker’s “base wage” as either their hourly wage (for workers paid hourly) or their per pay-period contracted compensation (for salaried workers). A worker’s per pay-period contracted compensation is their contractually obligated annual salary divided by the annual number of pay-periods during the year: that is, their contracted weekly, bi-weekly or monthly earnings. For most workers,

“base earnings” comprise essentially all annual earnings (excluding employer-provided fringe benefits). However, for some workers, bonuses, commissions, performance pay and overtime premiums are also important. For about ten percent of workers, non-base pay accounts for 20% of annual earnings. We document that the share of earnings accruing from bonuses is monotonically increasing in the individual’s base wages. For individuals in the bottom forty percentiles of the base wage distribution, hardly any of their annual earnings comes from bonus pay. However, for the median household, the 80th percentile, the 95th percentile and the 99th percentile of the base wage distribution, bonuses account for about 3 percent, 5 percent, 10 percent and 16 percent of their annual earnings, respectively. We also measure the importance of employer-provided fringe benefits in workers’ total annual compensation.

The second part of the paper discusses the relevance of each form of compensation for aggregate fluctuations, borrowing heavily from the insights of both labor search and output price setting literatures. We show that bonuses are approximately i.i.d. at the individual level, while base wages are highly persistent. This, coupled with the observation that bonuses are relatively acyclical, suggests that base wage rigidity may be more relevant for determining aggregate fluctuations than rigidity in average hourly earnings, at least in standard models. On the other hand, models of incomplete markets, in which temporary shocks have allocative effects, or models understanding individual earnings dynamics, will require moments of the average hourly earnings adjustment distribution. The variety in purpose highlights the value of separately providing moments on base and non-base adjustments.

We then proceed by measuring the adjustment of base wages for a sample of workers who remain continuously employed with the same firm. We refer to this sample as our “job-stayer” sample. Our first main result is that nominal base wage cuts are exceedingly rare on-the-job. During our entire sample period, only 2.4% of all workers received a nominal base wage cut during a year. The number was slightly higher for salaried workers compared to workers paid hourly (3.6% vs. 1.8%). On average, about one-third of both hourly and salaried workers who remain on the job received no nominal wage adjustment during a given year. Therefore, about two-thirds of both hourly and salaried workers receive a positive nominal base wage increase during a given year. There is also a missing mass of small positive changes with many more workers receiving a nominal base wage increase of 2 to 4 percent than from 0.1 to 2 percent. The patterns are similar for both high and low wage workers and regardless of whether or not workers receives an annual bonus. Our results imply a duration of nominal base wages for the typical worker who remains continuously employed on the same job of about 6 quarters.

We next turn to the measurement of base wage adjustments for the aggregate economy including data on both job-stayers and job-changers. Almost all workers that transition

across jobs experience a nominal base wage change. Indeed, 38% of job-changers experience a *decline* in their nominal base wage. Given that job-switchers are a non-trivial share of the economy, we create a broader measure of nominal base wage flexibility pooling together both job-stayers and job-changers. Doing so, we find that roughly 24 percent of all workers experience a base wage change during a given quarter and 71 percent experience a base wage change during a given year. Including both the job-stayers and job-changers, 9 percent of workers experience a nominal base wage decline with most of the declines being driven by job-changers.

That the aggregate economy, including job-switchers, exhibits a substantially higher degree of base wage flexibility than the sample of job-stayers is another key insight of this paper. Models seeking to understand the muted fluctuations in mean nominal wages over the cycle must reckon with this finding that aggregate wages are made more flexible on the downside by the presence of job-changers. In addition, models without realistic job search components should be cautious about using wage rigidity estimates from job-stayer samples, as is standard in the literature, for doing so will overstate the degree of rigidity in the economy as a whole. Including both job-stayers and job-changers yields an average duration of nominal base wages (inclusive of base wages and bonuses) of about 5.5 quarters. However, there is still an asymmetry in adjustment with nominal wage increases being seven times more likely than nominal wage cuts.

The results on job-changers inform models of wage dynamics, but do not directly address the question of whether new hire wages, which determine employment fluctuations in many models, are flexible. It is difficult to measure the flexibility of new hire wages at business cycle frequencies given the importance of selection in who works over the business cycle.³ We exploit how wages evolve for job-changers relative to job-stayers at business cycle frequencies. To do so, we next benchmark job-changers who move from firm i to firm j between period $t - 1$ and t to a similar worker in firm j in period $t - 1$ based on their $t - 1$ wages and demographic characteristics. We then document that the evolution of wages for the job-changer is nearly identical to that of their matched counterpart, and their relative wages are nearly invariant to business cycle conditions. Collectively, these results suggest that new hire wages evolve similarly to incumbent workers within a firm at business cycle frequencies. These results complement the findings in Hazell and Taska (2018) which documents that posted wages on the near universe of online job boards display similar adjustment patterns as the base wage changes of existing workers.

We next turn to assessing how bonuses evolve for job-stayers. The transience of bonuses

³For the importance of selection in determining the wage cyclicalities, see Solon et al. (1994), Basu and House (2016), and Gertler et al. (2016).

means that they are unlikely to greatly affect the user cost of labor, and are thus akin to sales in the pricing literature. However, variation in bonuses are still of great interest for certain important applications, such as for disciplining models in which employee compensation is subject to financing constraints, or predicting the impact of transitory income shocks on consumption. We document that annual bonuses vary substantively from year to year. As a result, 16 percent of job-stayers, on average, receive a decline in nominal wages inclusive of bonuses during a given year, far more than implied from an examination of base wages. We further document that the additional flexibility in terms of nominal wage adjustments provided by bonuses differs markedly throughout the base wage distribution. For those in the bottom quartile of the base wage distribution bonuses are inconsequential. However, for those at the top of the base wage distribution, variation in annual bonuses is larger than that of annual base earnings, suggesting that there may be distributional consequences of shocks to financially-constrained firms. Finally, we show that fringe benefits also provide an additional but small margin of nominal wage adjustment for the average worker.

After documenting the difference between individual and aggregate wage rigidity, we examine the extent to which wages are able to adjust to shocks. We provide evidence that wage setting is state dependent. Even though nominal base wage cuts are very rare for job-stayers over our entire sample period, 6.6 percent of salaried workers received nominal base wage cuts during the Great Recession. Although the share of job-switchers, who are much more likely to see wage declines, fell during the recession, the aggregate propensity to receive a nominal base wage decline increased by 2.5 percentage points during the recession. However, we find that the propensity to receive a bonus and the size of the bonus is roughly acyclical. The fact that bonuses adjust much less to business cycle fluctuations further suggests that base wages adjustments may be a more relevant concept to discipline models of labor market fluctuations. We also document that industries hit hardest during the Great Recession (including both manufacturing and construction) were much more likely to cut nominal base wages during the recession relative to other industries. Finally, during the Great Recession, firms with declining employment were much more likely to reduce the nominal wages of their workers relative to firms with constant or increasing employment. These results suggests that any model with a constant fraction of wage adjustments will struggle to match the wage setting patterns over a business cycle.

There is an existing literature on measuring nominal wage adjustments using either household surveys or administrative datasets. Instead of reviewing that literature collectively at this point, we discuss the relevant literature in relationship to our results throughout the paper. This allows us to contrast our specific results with those from the literature. While some of our findings are qualitatively similar to results in the existing literature, they are often

quantitatively different in magnitudes. These quantitative differences may have large consequences for the behavior of canonical models with wage rigidities. As we discuss throughout, differences between the results in our paper using the ADP data and other results in the literature are consistent with both substantial measurement error in nominal wages in household surveys and the lack of high quality hours measures in administrative datasets.

Ultimately, this paper claims three principal contributions. First, we provide evidence that base wages are a better proxy of the user cost of labor than measures of compensation inclusive of nearly i.i.d. bonuses, which behave like sales. Our results emphasize that there is an important conceptual difference between *compensation* flexibility and *contract* flexibility. If contracts specify a base wage per unit of labor, as well as a schedule of bonuses as a function of performance, then the striking absence of base wage declines suggest that contracts may be subject to adjustment frictions, even if measured compensation per hour appears flexible. Second, we provide the highest quality measurement of wage adjustment for a large share of the U.S. workforce, and highlight exceptional asymmetry in base wage adjustment. What's more, the ability for workers to move across firms is a source of aggregate nominal wage flexibility, and represents a key conceptual difference between nominal wage adjustment and output price adjustments. Finally, we provide new evidence of state dependence in base wage adjustment at both the firm and aggregate level. The patterns presented here urge the development of theories in which wage adjustment is state dependent and asymmetric, as such asymmetries will engender different aggregate responses to shocks in downturns and booms. Lastly, our results suggest that researchers must think carefully about which notion of the wage - the base wage or a broader wage measure inclusive of bonuses - is the appropriate one to discipline their respective models.

The paper proceeds as follows. Section 2 describes the ADP data in detail. Section 3 describes the allocation of worker compensation across base pay, bonuses and fringe benefits. We discuss the conceptual difference between base wages and bonuses in Section 4. Section 5 presents key facts about nominal base wage adjustments for job-stayers. Section 6 present wage change statistics for job-changers while Section 7 presents our measures of aggregate nominal wage adjustments. Section 8 examines the cyclicity of new hire wages. Section 9 presents the adjustment patterns of bonuses and fringe benefits. Section 10 provides evidence of state dependence at the aggregate, industry, and firm levels. Section 11 compares our estimates with those of the existing literature while Section 12 concludes.

2 Data and Variable Definitions

2.1 Overview of ADP Data

We use anonymized administrative individual panel data provided by ADP. ADP is a large, international provider of human resources services including payroll processing, benefits management, tax services, and compliance. ADP has over 650,000 clients worldwide, and now covers payroll for over 20 million individual workers in the United States per month. The data to which we have access starts in May 2008 and extends through December 2016. During that period, ADP processed payroll for approximately one-eighth of the U.S. workforce.

The data contain monthly aggregates of anonymized individual paycheck information, as well as all relevant information needed for human resources management. Crucially, we observe, without measurement error, the statutory per-period payment rate for all employees. For hourly workers, this payment rate is simply the worker’s hourly wage. For salaried workers, it constitutes the pay that the worker is contractually obligated to receive each pay period (weekly, bi-weekly, or monthly). Given the data is aggregated to the monthly level, the per-period payment rate is measured as of the last pay period of the month.

In addition to the administrative wage information, the data contain all other information that would appear on the worker’s paycheck, such as the worker’s gross earnings per pay period, taxes paid, and any taxable benefits provided by the firm. Additionally, the data contain other payroll information including whether the worker is paid hourly, the frequency at which the worker is paid and the number of hours worked during the month. For hourly workers, the exact number of hours worked is reported. For salaried workers, hours information is provided by the firm’s HR administrator and often set to 40 hours. We also observe various additional geographic and demographic characteristics of a worker as well as details about the job, such as worker tenure, firm size, and industry. Selection into the ADP data is at the firm level. As a result, we can measure wage distributions within and across firms over time.⁴ Finally, the presence of consistently-defined anonymized worker identifiers permits the study of individual worker dynamics across ADP clients. However, given our sample size, movements from one ADP firm to another ADP firm are quite common.⁵

⁴Strictly speaking, our definition of a firm is an ADP-provided client code. This will usually be an autonomous firm, rather than any individual establishment. One possible exception to this rule arises if large conglomerates have multiple subsidiaries, all of which separately hire ADP to handle their payroll. In this case, each subsidiary would count as a separate ADP client. As a result, our ADP firms are a combination of both Census notions of firms and establishments.

⁵All worker and firm identifiers in the ADP data are unique and consistently defined over time. However, they are constructed in a way that all workers and firms remain anonymous for research purposes. Firm and worker names and detailed addresses as well as firm employer identification numbers and worker social security numbers are all purged from the data available to researchers. The ADP data use agreement

We make two major sample restrictions for our analysis. First, we restrict attention to prime age workers between 21 and 60 years old, inclusive. Second, ADP has two products separately targeted to firms with greater than or less than 50 employees. We only have access to data for firms with more than 50 employees throughout our sample period, and thus constrain our analysis to large employers.

The full dataset covers over 50 million unique individuals and over 141 thousand firms. To reduce computational burden, we create three random subsamples of the full data. The first chooses one million unique employees, and follows them through their entire tenure in the sample across all firms for which they work. This is the primary dataset for analysis. Second, we separately draw a sample of 1 million workers who change jobs during our sample period (“job-changers”). These are workers who show up in multiple firms during their time in the ADP database. This will allow us to explore fully the patterns of wage changes for workers who switch between ADP firms. However, these two datasets are ill-suited to study questions at the firm level; we therefore construct a third subsample of three thousand unique ADP clients, drawing all workers from those firms in the process. The random employee-level, job-changer and firm-level subsamples remain large, with roughly 25 million, 27 million, and 68 million unique employee-month observations, respectively.

2.2 Representativeness of ADP Data

There are two areas of concern regarding the representativeness of the ADP data. 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. However, given that 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.⁶

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.⁷ According to these same surveys, however, very large firms (firms with more than 10,000 employees) are less likely to outsource their payroll functions. As noted above, ADP processes payroll for about

prohibits using the data to explore wage patterns of any individual worker or firm.

⁶Furthermore, from 2013 onwards, we also have access to ADP’s data for firms with fewer than 50 employees. These data reinforce that any potential bias from excluding small firms from the main results in our paper is likely to be minor. We discuss these results in detail in the Appendix Appendix C.3.

⁷See, for example, https://www.feicanada.org/ajaxfilemanager/uploaded/RH_FERF_Benchmarking2014.pdf.

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.

Table 1 highlights the employment-weighted firm size distribution in our employee sample (column 1) and employees 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 91,500 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 Table 1 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.⁸ 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 has about 20 percent of employment (in our employee sample) in firms with more than 5,000 employees.⁹ As noted above, some of this difference also results from the fact that the ADP 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 subsequent analyses have been weighted so as 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.

Given this is the first paper using the ADP data, a deeper discussion of the representativeness of the ADP sample is warranted. We have relegated much of this discussion to the Online Appendix. In particular, we benchmark the demographic composition of the ADP

⁸According to BDS data, 72% of all U.S. employment during this time period is in firms with more than 50 employees.

⁹We 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 complementary underweight in retail trade, construction, and agriculture.

Table 1: 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	91,577	3,000	.
Number of Observations	24,831,316	68,267,166	.
% Firm Size: 50-499	37.8	31.3	29.5
% Firm Size: 500-999	13.6	13.9	7.3
% Firm Size: 1000-4999	25.1	22.2	17.5
% Firm Size: ≥ 5000	19.7	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.

sample to that of the CPS along a variety of dimensions. Additionally, we compare annual earnings dynamics in our ADP sample for people who remain continuously employed with the same firm for two years to the earnings dynamics in Guvenen et al. (2014) for a similarly defined sample. Further, we compare both levels and time series trends in the average hourly wage for workers paid hourly between our ADP sample and a similar CPS sample. We also examine patterns in nominal wage adjustments by firm size to explore potential biases from our data being under-representative of both really small and really large firms. Finally, we show the unweighted results for many of the paper's key findings. After performing all of these benchmarking exercises, we are confident that the ADP data provides a representative picture of nominal wage adjustments for U.S. workers over the 2008 to 2016 period.

3 The Nature of US Worker Compensation

The ADP data includes many detailed administratively recorded measures of worker compensation. In this section, we describe the composition of worker compensation. We first exhibit the overwhelming importance of base earnings, before considering the size of bonuses and fringe benefits in workers' compensation.

3.1 Base Wages and Base Earnings

Employers participating in the ADP payroll services are required to report the contractually obligated per-period wage rate for each worker. For workers who are paid hourly, this is the workers' hourly wage. All salaried workers have an administrative field recording their

weekly, bi-weekly or monthly contracted salary rate depending on the frequency of their pay period. We refer to the contractually obligated per-period wage rate as a worker’s “base wage.”

A separate field in the data reports workers’ administrative “monthly gross earnings” (excluding employer-provided fringe benefits). A worker’s base pay is only one part of their monthly gross earnings. During a given month, a worker may also receive tips, commissions, overtime payments, performance pay, bonuses, cashed-out vacation days and meal and travel reimbursements.¹⁰ Monthly gross earnings is literally the sum of all paychecks (before taxes) earned by the worker during the month. To isolate the importance of base wages in worker earnings, we define the concept of a worker’s “monthly base earnings” using information on their base wage. If the worker is an hourly worker, their monthly base earnings is their base wage times the total number of hours worked during the month. If the worker is a salaried worker, their monthly base earnings is their base wage times the number of paychecks received during the month. Any difference between a worker’s monthly gross earnings and their monthly base earnings is the result of the worker earning some combination of bonuses, tips, overtime, reimbursements or other non-standard payments during the month.

Such non-standard payments are not likely to accrue every month for a given worker. To see how important these sources are for a typical worker, we aggregate our data to calendar years. When doing so, we restrict our analysis to workers who remain continuously employed with the same firm for all twelve calendar months of a given year. We refer to this sample as our “full-year” employee sample.

3.2 Measuring Overtime, Bonuses, and Commissions

Ideally, one would decompose workers’ “residual earnings” – gross earnings less base earnings – into various sub-components. However, the ADP data is not well-suited for such disaggregation. Firms are not required to separately report the different potential sub-components that comprise residual earnings. Despite the limitation, we make four refinements to our residual earnings measures. First, we impute the amount of monthly overtime premiums paid to hourly workers using an often-reported “overtime earnings” field in the data. For hourly workers, therefore, we can create a measure of monthly residual earnings net of overtime payments.¹¹

¹⁰The fact that meal and travel reimbursements can show up in workers’ paychecks implies that there is not a one-to-one mapping between monthly gross earnings in our dataset and monthly W-2 earnings. We do not have a worker’s W-2 earnings in the data provided to us by ADP.

¹¹We discuss this imputation procedure in greater detail in the Online Appendix. Firms are asked to report total hours worked (inclusive of overtime) and total earnings from hourly work (inclusive of overtime). Comparing these measures to our measures of base earnings allows for a crude imputation of overtime.

Second, we define large residual earnings to be any residual earnings net of overtime that accrue to a worker in a given month that exceeds 1% of their annual earnings. For example, if a worker earned \$50,000 during a given calendar year, we would classify that worker as having large residual earnings during a given month if residual earnings net of overtime pay exceeded \$500 during that month. By making this restriction, we exclude any small payments made to the worker during a given month such as small meal reimbursements or small measurement error in our overtime imputation.

Third, we compute the frequency of months a worker receives large residual payments during a given year. We define a worker to be a “commission worker” if they receive large residual earnings net of overtime in four or more calendar months during a given year. We are interpreting “commission workers” broadly in that these workers could have large residual payments in four or more months during a given year due to sufficiently frequent commissions, tips, performance pay, mis-measured overtime pay, or even large meal and travel reimbursements. We can then segment workers into “non-commission workers” and “commission workers”. According to this definition, roughly 10 percent of workers each year during our sample can be classified as commission workers.

Finally, for non-commission workers, we define a worker as having received a “bonus” if that worker received a large residual earnings payment net of overtime in at least one month but no more than three months during a given calendar year. Again, our definition of bonus is broad in that it applies to any large infrequent extra non-overtime payments received by workers during a year. During our sample period and given our definition, roughly 30 percent of non-commission workers receive a bonus during a given year.¹²

3.3 The Composition of Worker Compensation

Table 2 shows the importance of annual base pay and bonuses as a share of annual earnings for non-commission workers. The first two columns shows results pooling together hourly and salaried workers. Column 1 shows the share of total annual earnings that comes from base pay. About one-quarter of all workers receive essentially all of their annual compensation from base earnings. The median worker earns only 2.5 percent of their annual earnings from sources other than their base pay.¹³ This suggests that for most workers, base pay is their primary form of compensation. However, for some workers, other forms of compensation (e.g., bonuses, commissions, tips, overtime) comprise a more substantive portion of their

Consistent with anecdotal evidence, these overtime wage rates center on 1.5 and 2 times base wage rates.

¹²It should be noted that most of these extra payments occur in December, February and March suggesting that many of them are likely linked to annual bonuses.

¹³Including commission workers, the median worker earns about 3.7 percent of their annual earnings from sources other than their base pay.

Table 2: Share of Annual Base Earnings and Bonuses out of Annual Gross Earnings, 2009-2016, Non-Commission Workers

	All		Hourly		Salaried	
	Share		Share		Share	
	Share	Base+	Share	Base+	Share	Base+
	Base	Bonus	Base	Bonus	Base	Bonus
10 th percentile	89.7%	93.5%	91.6%	92.8%	86.6%	96.0%
25 th percentile	93.8%	96.8%	94.8%	95.7%	91.6%	99.0%
Median	97.5%	99.3%	97.8%	98.5%	96.5%	100%
75 th percentile	99.7%	100%	99.6%	99.8%	100%	100%
90 th percentile	100%	100%	100%	100%	100%	100%
Sample Size (thousands)	611	611	378	378	222	222

Notes: Table shows the distribution of the share of worker annual base earnings and annual base earnings plus annual bonuses out of annual worker gross earnings. We restrict our attention to our sample of non-commission workers who remain continuously employed with the same firm for all twelve months of a calendar year. The remaining columns show similar data separately for workers paid hourly (columns 3 and 4) and workers who are salaried (columns 5 and 6).

annual earnings. Ten percent of non-commission workers earn at least 10 percent of their annual earnings from sources other than their base pay.

Additionally, nearly all remaining compensation for non-commission workers is in what we classify as bonuses. Specifically, for the median non-commission worker, over 99% of all annual gross earnings are base earnings and what we classify as bonuses. This is important in that for the rest of the paper we are going to focus on nominal wage adjustments of base pay and bonuses. Doing so, captures essentially all of the compensation for most non-commission workers. The remaining gross earnings is in overtime pay and other small infrequent residual earnings payments. Throughout the paper, we present results on nominal wage adjustments separately for both non-commission and commission workers. The remaining columns of the table show patterns separately for hourly and salaried workers. Bonuses are more important for salaried workers relative to workers paid hourly. However, overtime earnings are more important for workers paid hourly.

3.4 Employer-Provided Fringe Benefits

Given our detailed data, we create a broader measure of worker compensation that includes employer-provided fringe benefits. The data contain all forms of fringe benefit that would appear on an employee's paycheck, including employer-provided health insurance and contributions to a retirement plan or pension - such as a 401(k) or Roth IRA - made by the

employer.¹⁴ Using these data, we create a measure of “annual fringe benefits” by summing the monthly employer-provided health benefits and retirement contributions over all months of a year. The fringe benefit measures were not-consistently reported prior to 2012. Starting in 2012, as part of the Affordable Care Act, employers were required to report their contributions to employee health benefits. Given this, when analyzing measures of broader compensation, our analysis is limited to workers who remain continuously employed with the same firm for a full year during the 2012-2016 period.

Using these measures, we construct a worker’s “total annual compensation” by summing their annual gross earnings with their annual employer-provided fringe benefits. Table 3 shows the distribution of the share of total compensation that is in fringe benefits for all workers in our full-year employee sample.¹⁵ Fringe benefits accounts for 8.2% of the median worker’s total compensation, but there is large variation around this number, with 10% of workers receiving more than 25% of their compensation through fringe benefits, and many workers receiving no fringe benefits at all. Hourly workers tend to receive fewer fringe benefits than do salaried workers: the median hourly worker has 6.5% of their total compensation in fringe benefits, compared with 9.5% for salaried workers. However, the right tail of fringe benefits for hourly workers is thicker for hourly workers than for salaried workers, with the 90th percentile of hourly workers receiving 27.1% of their compensation from special compensation, compared with 22.2% for salaried workers. Encouragingly, the numbers presented in this table match those found by the BLS in their Employer Cost for Employee Compensation (ECEC) reports. For example, the June 2016 report finds that 7.6% of workers’ total compensation is accounted for by the cost of health insurance, and 3.9% is accounted for by retirement and savings account contributions.¹⁶

3.5 Heterogeneity in Bonuses and Fringe Benefits Across Workers

The left panel of Figure 1 shows the share of annual bonus income out of annual total earnings sorted by workers’ base wage percentile.¹⁷ As with the results above, we exclude commission workers. Workers at the lower end of the base wage distribution receive, on

¹⁴We exclude all tax measures from our analysis including employer paid payroll taxes. Additionally, we note that our fringe benefit and bonus measures do not include stock options. ADP has data on when stock options are cashed in (for tax reasons), but not when the options were granted.

¹⁵The results between the full sample and the commission sample were nearly identical.

¹⁶See, for instance <https://www.bls.gov/news.release/ecec.nr0.htm>. The aggregate fringe benefit share does not match exactly, as the BLS includes paid leave, bonuses, and legally-required benefits such as social security payments, the first two of which will be included in our measure of gross earnings.

¹⁷To make the base wage percentile, we combine data on both hourly and salaried workers. For hourly workers, we use their base hourly wage. For salaried workers, to put things in the same hourly wage units, we divided their base weekly earnings by 40 hours.

Table 3: Share of Fringe Benefits out of Total Compensation, 2012-2016 Period

	All	Hourly	Salaried
10 th percentile	0.0%	0.0%	0.0%
25 th percentile	1.3%	0.0%	4.4%
Median	8.2%	6.5%	9.5%
75 th percentile	17.1%	18.3%	15.9%
90 th percentile	25.1%	27.1%	22.2%
Mean Total Fringe Share	10.7%	10.6%	10.8%
Mean Health Fringe Share	7.9%	8.7%	6.7%
Mean Pension Fringe Share	2.8%	2.0%	4.0%
Sample Size (Thousands)	630	370	260

Notes: Table shows the distribution of the share of worker fringe benefits out of total worker compensation during the 2012-2016 period for all workers in our full time employee sample (column 1) and then separately for hourly and salaried workers (columns 2 and 3).

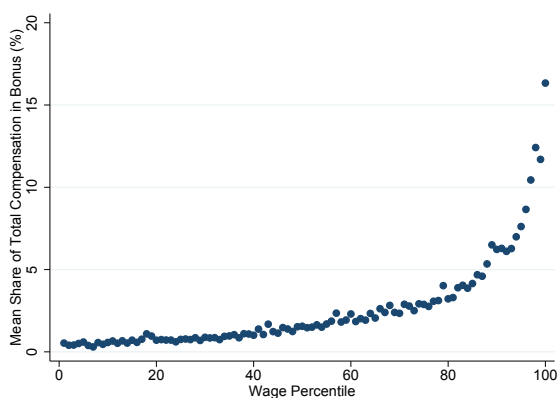
average, only less than 0.5% of their annual earnings in bonuses. The share of earnings in bonuses increases monotonically throughout the wage distribution. The median worker earns about 2% of their annual earnings in bonuses. Systematically, workers at the top of the wage distribution earn a substantial amount of their annual earnings in bonuses. Therefore, while annual bonuses are not an important form of compensation for most workers, they are substantial for high wage workers.

The right panel shows the share of annual fringe benefits provided by the employer out of a worker's annual total compensation (earnings plus fringe) as a function of the worker's base wage percentile. Fringe benefits are much less important for lower wage workers (below the 20th percentile). However, from the 20th percentile through the 95th percentile of the wage distribution, the share of total compensation in fringe is roughly constant in the 12 to 14 percent range. For top wage earners, fringe becomes a smaller fraction of total compensation. This is likely because fringe benefits are not usually provided on bonus income and because tax exempt employer-provided retirement contributions are capped.

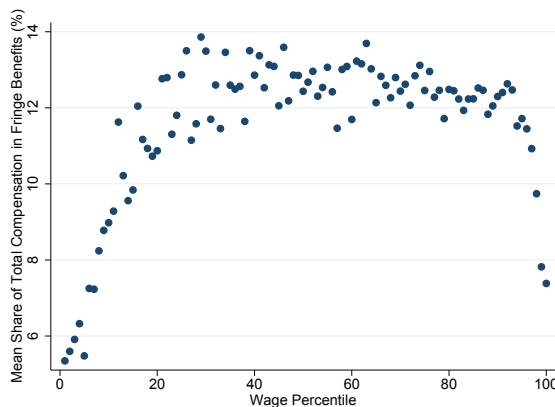
4 Setting the Stage: Base Wages vs. Bonuses

The remainder of the paper presents a host of new statistics concerning wage adjustment in the United States. We break aggregate earnings per hour fluctuations into three components: base wage adjustment for job-stayers, base wage adjustment for job-changers, and bonus adjustment for job-stayers. Throughout, we will be purposefully agnostic as to which

Figure 1: Bonus and Fringe Benefit Share by Employee Base Wage Percentile, Full-Year Job-Stayers, Excluding Commission Workers



PANEL A: SHARE IN BONUS
OUT OF TOTAL EARNINGS



PANEL B: SHARE IN FRINGE BENEFITS
OUT OF TOTAL COMPENSATION

Notes: Panel A shows the share of bonus earnings out of total earnings by worker percentile within the base earnings distribution. Panel B shows the share of fringe benefits out of total compensation by worker percentile within the base earnings distribution. Sample restricted to job-stayers who remain in the sample for a full calendar year. Commission workers are excluded.

moments should drive aggregate fluctuations, preferring instead to let modelers pick the relevant numbers for their calibration based on their specific model structure. In this section, however, we provide suggestions on the importance of each margin of adjustment, borrowing insights and intuition from both the wage and price setting literatures.

A large literature has developed arguing that the rigidity of new hires' wages determine fluctuations in employment (Pissarides, 2009). This is because firms and workers enter into long-term employment relationships, rendering the wage of new hires marginal for the decision of firms to post vacancies and the decision of unemployed workers to search for a job. As is highlighted in Kudlyak (2014), given the nature of such long-term employment relationships, it is not the *spot* wage of new hires that should matter for employment fluctuations but rather the *user cost* of labor. The user cost of labor is defined as the expected present value of costs to the firm associated with adding an additional worker in period t rather than waiting and adding the worker in period $t + 1$. In existing frameworks, there is no distinction between base wages and bonuses in the definition of the user cost. Indeed, it is not trivial to define the notion of a user cost when bonuses are used to incentivize unobserved effort as is assumed in many models of optimal contracting. For this reason and given our relatively short sample, we elect to not directly measure user cost rigidity in this paper.

Despite not directly measuring the user cost of labor, it is worth discussing which of our

Table 4: Annual Persistence of Bonuses vs Standard Wage, sample of full-year job-stayers, 2009-2016

	Log			
	December Base wage	Log Bonus	Linear Bonus	Share Bonus
	(1)	(2)	(3)	(4)
$y_{i,t-1}$	0.775*** (0.012)	0.016 (0.016)	0.018 (0.088)	-0.054** (0.023)
Job FE	Yes	Yes	Yes	Yes
Observations	586,023	145,904	586,023	586,023
R^2	0.995	0.926	0.764	0.780

Notes: Table reports OLS-estimated AR(1) coefficients from equation (1). White heteroskedasticity robust standard errors clustered at the employee-level reported in parentheses. ***, **, and * represent coefficients are statistically different from 0 at the 10%, 5%, and 1% levels, respectively. We use the employee sample restricting our analysis to workers who remain continuously employed with the same firm for two consecutive calendar years.

wage measures is conceptually better linked to the user cost notion. The persistence of our various wage concepts will be an important determinant of its impact on the present value of wages, a key input to the user cost. We therefore estimate autocorrelation coefficients for our various wage measures. Specifically, we estimate OLS regressions of the form:

$$y_{it} = \rho y_{it-1} + \alpha_i + \epsilon_{it} \quad (1)$$

where i indexes a worker-firm pair, t represents a year and α_i is a job fixed effect. y_{it} represents the value of a particular wage 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 adjustments in bonus pay.

Table 4 reports the estimated autocorrelation coefficient ρ of various wage measures at the individual level. We explore the persistence of log base wages (first column) and various specifications of bonuses (columns 2 through 4). As seen from the table, base wages have an annual autocorrelation of 0.775. By comparison, bonuses appear to be almost i.i.d., whether we consider log-linear or linear specifications. Indeed regressing the share of pay in bonuses (column 4) on its lagged value yields a negative coefficient, suggesting that high bonus years tend to be followed by low bonus years. These results suggests 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.

Bonuses could still be an important margin of adjustment if they fluctuated to match high-frequency productivity shocks. However, this is not the case. The strong seasonality of bonus payments reduces their usefulness as a source of flexibility. Conditional on receiving a bonus, most workers only receive a bonus once a year. Most bonuses are concentrated in only a few months (e.g., December, February and March). Indeed, at the firm-level, 28% of workers receive a bonus in their employer’s most common bonus-granting month, which is large since just 31% of workers receive a bonus in a given year. While bonuses may provide downward adjustments in workers’ total annual compensation (as we highlight below), they appear rigid over short time horizons.

These pieces of evidence lead us to conclude that base wage rigidity more closely resembles what we term “contract rigidity” - the ability of firms to adjust the terms governing their long-term relationship with their employees, which consist of a base wage and a bonus schedule as a function of performance - and thus is a more appropriate measure of rigidity to use for models seeking to understand aggregate employment fluctuations over the cycle. Measuring adjustments in base wages separately from other forms of compensation is thus a key contribution of this paper. Whether to use incumbent base wages (Section 5), job-changer base wages (Section 6), an aggregate of the two (Section 7), or new hire base wages (Section 8) is model dependent. For instance, Gertler et al. (2016) put forth a model in which the rigidity of incumbents’ wages is the key factor in determining cyclical fluctuations, while Pissarides (2009) makes a compelling case for the importance of new hire wages.

An analogue to the base wage versus bonus distinction can be drawn from the output price setting literature. Eichenbaum et al. (2011) show that nominal rigidities take the form of inertia in reference prices and costs. As a result, short-term sales of output prices may be ignored for assessing the impact of nominal rigidities on aggregate fluctuations. Indeed, much of the empirical literature has excluded sales from their measure of price adjustments (Nakamura and Steinsson, 2008) or directly incorporated a differential adjustment cost for sales relative to reference prices (Kehoe and Midrigan, 2008; Midrigan, 2011). Given the transient nature of bonuses, it is natural to liken them to a “sale” in the pricing literature.

This does not imply that bonuses are unimportant for macroeconomics. Recent studies of average hourly earnings adjustments (Barattieri et al., 2014; Daly and Hobijn, 2014; Kurmann and McEntarfer, 2018; Jardim et al., 2019) have great value for different questions and models. Studies looking to understand earnings dynamics and the stagnation of average hourly earnings, or which care about temporary earnings shocks will naturally require earnings measures inclusive of bonuses. Furthermore, fuller models of nominal contract rigidity, which incorporate explicit optimal contracts inclusive of bonuses, or models of dynamic contracts, require knowledge of both base and bonus compensation. In essence, bonuses are

an important component of the spot wage for labor. To the extent that spot wages matter for macroeconomics – if, for instance, firms’ are subject to period-by-period financing constraints (Schoefer, 2016) or households are liquidity constrained (Kaplan and Violante, 2014) – then so too will aggregate wage adjustments *inclusive of bonuses*.

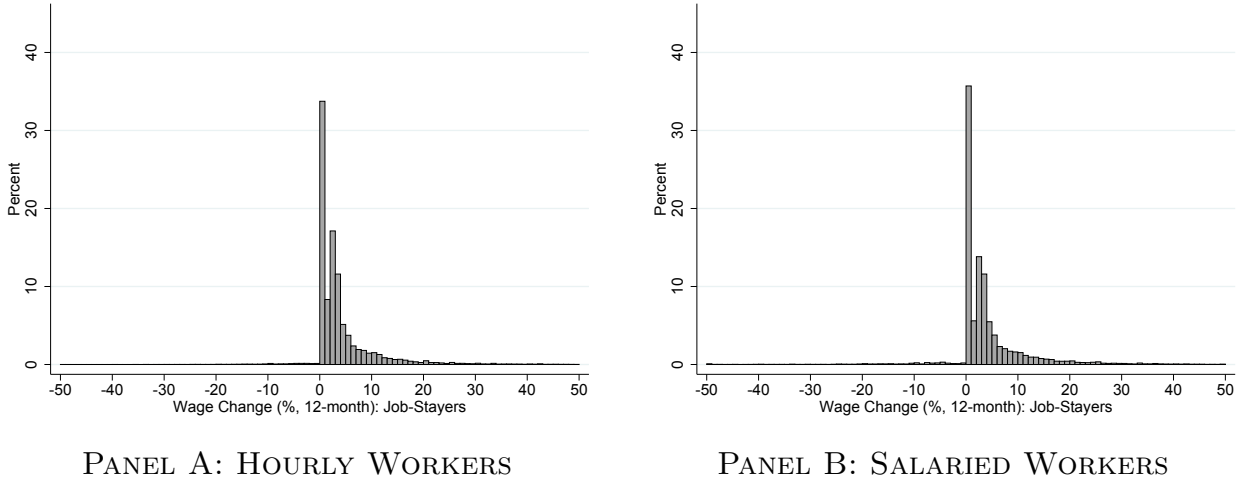
Given these issues, we begin the paper by documenting how base wages evolve for both job-stayers, job-changers, and new hires. We then discuss the evolution of wage measures inclusive of bonus variation, whose usefulness will be model-dependent.

5 Nominal Base Wage Adjustments for Job-Stayers

This section explores the nature of nominal base wage adjustments for workers who remain continuously employed in the same job. We use our employee sample and define “job-stayers” as those workers who remain continuously employed with the same firm between the two periods in which the nominal wage is measured. We present moments of nominal base wage adjustment at the monthly, quarterly and annual frequencies. For the monthly, quarterly and annual samples, we ensure that workers are continuously employed with the same firm for one, three and twelve consecutive months, respectively.

Figure 2 plots the distribution of 12-month nominal base wage changes for all job-stayers pooled over all years of our sample. As discussed above, base wages make up essentially all of annual compensation for most workers. Panel A plots the distribution for hourly workers, while Panel B plots the distribution for salaried workers. Four key observations are apparent from the figure. First, a large share of workers - 33% of hourly and 35% of salaried - do not receive a nominal base wage change in a given year. Second, the patterns of nominal base wage adjustments for hourly workers and salaried workers are nearly identical. Given this, we often pool the data for hourly and salaried workers together going forward when describing base wage adjustments. Third, there is a clear asymmetry in the base wage change distribution, with the overwhelming majority of changes being wage increases. Only 2.4 percent of workers (combining hourly and salaried) in the U.S. who remained continuously employed with the same firm for 12 months received a nominal base wage decline. Of the roughly 66% of all individuals who receive a nominal base wage change over a given 12-month period, only 3.6% received a nominal base wage cut (2.4/66). Finally, there are very few small nominal base wage changes for either hourly or salaried workers. Just 8.6% of all workers received a nominal base wage change of between 0.1 and 2 percentage points, compared with 27.1% receiving an increase between 2 and 4 percentage points. This missing mass of very small wage changes is consistent with the random menu cost models that are prevalent in the price setting literature.

Figure 2: 12-Month Nominal Base Wage Change Distribution, Job-Stayers



Notes: Figure shows the annual change in nominal base wages for workers in our employee sample (including commission workers) who remain employed on the same job for 12 consecutive months.

Table 5 provides a set of summary moments on the probability of base wage increases and base wage declines for three frequencies: monthly, quarterly and annual. The annual frequencies correspond to the data underlying Figure 2. While the asymmetry between nominal wage increases and nominal wage cuts is a feature of many existing empirical papers (see, e.g. Lebow et al. (2003); Kahn (1997); Card and Hyslop (1997)), the results in Figure 2 are quantitatively different from much of the existing literature. As we highlight below, measurement error in household data sets has resulted in estimates of nominal wage adjustments (both up and down) that are higher than our administrative payroll data suggest. In addition, the missing mass of small wage changes urges models of state dependent wage adjustment, which we explore in more depth in Section 10. Again, because of measurement error, this missing mass has been difficult to detect in prior work. Finally, most existing studies measure wages inclusive of both base wages and bonuses. As we also highlight below, the frequency of downward adjustment of base wages is quite different than the frequency of downward adjustment of wage measures inclusive of bonuses.

The patterns of nominal wage adjustments for job-stayers are fairly robust across workers who are compensated in different ways. In the Online Appendix, we show that the patterns in Figure 2 are nearly identical if we restrict our sample to only non-commission workers, only commission workers, only non-commission workers who receive a bonus and only non-commission workers who do not receive a bonus. These findings suggest that base wage adjustments do not differ across workers who receive other types of compensation. Additionally, we show that the patterns of base wage adjustment are nearly identical for those

Table 5: Probability of Base Wage Change, Pooled 2008-2016 Sample of Job-Stayers

	All Workers (1)	Hourly (2)	Salaried (3)
<u>Annual</u>			
Probability of Positive Base Wage Change (%)	63.9	65.3	61.6
Probability of Negative Base Wage Change (%)	2.4	1.8	3.6
<u>Quarterly</u>			
Probability of Positive Base Wage Change (%)	18.5	19.5	16.7
Probability of Negative Base Wage Change (%)	0.9	0.7	1.3
<u>Monthly</u>			
Probability of Positive Base Wage Change (%)	6.3	6.6	5.8
Probability of Negative Base Wage Change (%)	0.6	0.3	0.6

Notes: Table shows the frequency of base wage increases and base wage decreases at different horizons for our sample of job-stayers during the 2008-2016 period. The first column pools together hourly and salaried workers while the second and third columns, respectively, show the frequency of changes for hourly and salaried workers separately. The top panel shows results at the annual horizon while the middle and bottom panels show results at the quarterly and monthly horizons. We use our full employee sample for this analysis.

Table 6: Base Wage Change Statistics, Pooled 2008-2016 Sample of Job-Stayers

	Monthly	Quarterly	Annual
Mean Unconditional Change (%)	0.3	1.0	3.9
Median Unconditional Change (%)	0.0	0.0	2.4
Standard Deviation of Unconditional Change (%)	2.6	3.7	6.5
Mean Wage Change (%)	5.2	5.1	5.8
Median Wage Change (%)	3.1	3.1	3.4
Standard Deviation of Wage Change (%)	8.0	6.5	7.0
Mean Change, Conditional on Positive (%)	6.2	5.7	6.3
Median Change, Conditional on Positive (%)	3.3	3.3	3.5
Stan. Dev. Change, Conditional on Positive (%)	7.7	6.4	7.0
Mean Change, Conditional on Negative (%)	-10.7	-8.7	-7.3
Median Change, Conditional on Negative (%)	-8.3	-7.7	-6.6
Stan. Dev. Change, Conditional on Negative (%)	8.1	5.8	4.6

Notes: Table shows 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.

workers who are paid hourly and who have substantive movements in monthly hours worked throughout the year. Even for workers whose hours appear allocative, 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.

Table 6 shows additional moments of the base wage change distribution. For this table, we only report results pooling together both hourly and salaried workers given the frequency of adjustment distributions were similar between the two groups.¹⁸ During this period, mean and median nominal base wage growth for workers who remain on the same job equaled 3.9 percent and 2.4 percent, respectively.¹⁹ Conditional on a base wage change occurring, annual mean and median base nominal wage growth was 5.8 and 3.4 percent. A key statistic we will focus on throughout the paper is the standard deviation of nominal wage growth. Unconditionally and conditional on a base wage change occurring, the standard deviation of annual nominal base wage growth during the full 2008-2016 period was 6.5 percent and 7.0 percent, respectively.²⁰ Additionally, conditional on a positive base wage change occurring during a 12 month period, the mean and median size of the increase was 6.3 and 3.5 percent. The fact that the mean is much higher than the median reinforces the fact that some workers receive very large nominal base wage changes on the job, perhaps due to promotions. The mean and median size of a base wage cut, conditional on the worker experiencing a nominal base wage reduction, were both around 7 percent. While the frequency of base wage increases is much higher than wage cuts, the mean size of a base wage increase is very similar to the mean size of a base wage cut.

In Online Appendix, we show a set of additional results surrounding time dependence in base wage adjustments. Conditional on a job-stayer receiving a base wage change during a year, most receive only one change. Most firms change the base wages of their employees in the same month. Most job-stayers receive a base wage change one-year from their last wage change. Workers who receive a wage change off-cycle (in a different month from most workers within the firm) tend to get higher wage increases than those who receive a wage change on-cycle. Collectively, at the firm-worker level, these results show strong evidence of staggered contracts (see Taylor (1979)). However, different firms adjust the base wages of most of

¹⁸To limit the effect of extreme outliers when computing mean wage changes, we winsorize both the top and bottom 1% of nominal wages and the top and bottom 1% of wage changes. We only do this when computing the size of wage changes conditional on a wage change occurring. This does not affect our frequency of wage change results in any way.

¹⁹It should be noted that our wage growth for job-stayers includes a combination of cohort, time and age effects. The presence of age effects implies that wage growth for job-stayers are higher than the wage growth for the economy as a whole. See Beraja et al. (2016) who make a similar point when comparing time series and panel data wage growth patterns in the CPS during the Great Recession.

²⁰In the Appendix Table A2, we also report the skewness and kurtosis of the distribution.

their workers in different months. There is some monthly seasonality in the probability of a worker receiving a wage change with January, April, July and October being the months with most frequent adjustments. However, when aggregating the data to quarterly levels, there is little quarterly seasonality. Given these facts, we compute a measure of the average duration of a base wage increase for a worker assuming Calvo adjustment. This is a crude measure given the evidence of state dependence presented in Section 10. But, given that the literature focuses on a measure of average duration in Calvo models of wage adjustments, we feel it is a useful statistic to use when comparing the nature of nominal wage adjustment across our different wage concepts. The results in Table 5 suggest the average duration of a nominal base wage for job-stayers is about 6.0 quarters.

6 Nominal Base Wage Adjustments for Job-Changers

The prior section focuses on nominal base wage adjustments for individual job-stayers. However, models in which most base wage adjustment originates from movements across firms or due to the arrival of outside offers, such as many labor search models (Menzio and Shi (2011); Cahuc et al. (2006)), may be better calibrated to moments measuring nominal wage adjustments for job-changers. In this section, we use the ADP data to provide such moments.

There is a large literature documenting that wages of job-changers are more pro-cyclical than those of job-stayers (see, for instance, Bils (1985); Haefke et al. (2013); Pissarides (2009); Martins et al. (2012) and Gertler et al. (2016)). These studies show that the wages of employees entering new jobs tend to move almost one-for-one with labor productivity and that this high degree of pro-cyclicality persists even after controlling for detailed job characteristics; it does not appear to be completely due to pro-cyclical “job upgrading.” Our paper contributes to this literature by using high quality administrative data on wages to measure not only the *mean* wage adjustment of job-changers, but also provides a number of moments of the distribution of wage changes for job-changers as a whole.

The analysis in this section uses our job-changer sample. When measuring wage adjustment for job-changers, three 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, the notion of a “firm” within the ADP dataset is a unit that contracts with ADP. Sometimes, multiple establishments within a firm contract separately with ADP or 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.

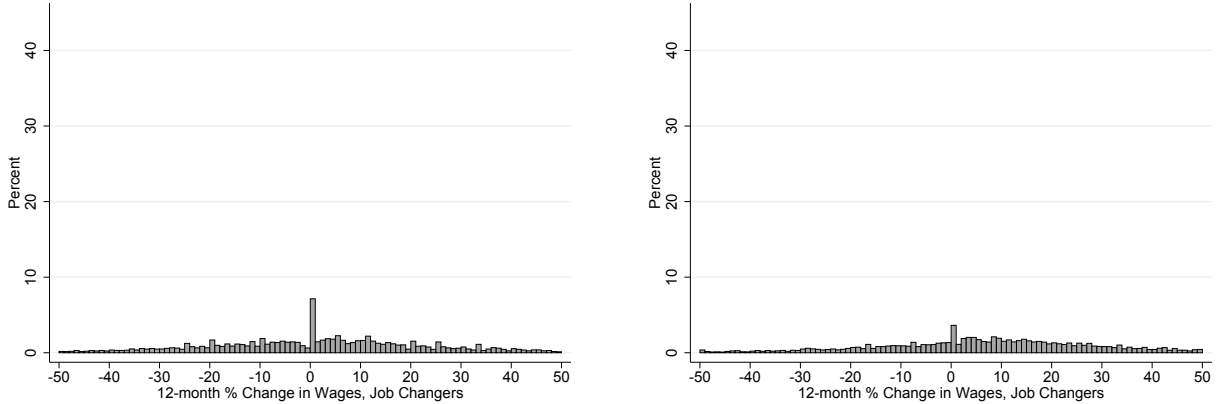
Finally, the choice of timing is more nuanced given the nature of our data. As with job-stayers, we can measure base wage changes for job-changers at one-month, one-quarter, and one-year frequencies. However, when we see a worker at firm i in month t and then see a worker at firm j in month $t + 12$, 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.²¹

Figure 3 plots the distribution of 12-month nominal base wage changes for a sample of job-changers. The patterns are strikingly different from the patterns in Figure 2. First, essentially all job-changers receive a base wage change over a given year. Only about 6.5 percent of hourly job-changers and 3.3 percent of salaried job-changers do not receive a year-over-year base wage change. Second, the propensity for a base wage cut is very high for job-changers with 39.2% of workers paid hourly and 31.2% of salaried workers receiving a base wage decline during a job-change. Finally, the distribution of base wage changes is more symmetric around zero. As seen from the figure, there are roughly as many small base wage increases (0-2 percent) as there are slightly larger base wage increases (2-4 percent). There is much more base wage adjustment for job-changers than there is for job-stayers.

Table 7 shows key statistics on the distribution of base wage changes for job-changers. Conditional on a job change and a base wage change, mean and median annual base wage growth was 8.0 and 4.6 percent accordingly. Base wage growth is much larger for job-changers than it is for job-stayers. As seen from Figure 3, there is a large amount of heterogeneity in base wage changes for job-changers. Job-changers whose nominal base wage increased over the year experienced, on average, a 26.1 percent increase. Job-changers whose nominal

²¹We 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, but generate large swings in base wages in almost all cases.

Figure 3: 12-Month Base Wage Change Distribution for Job-Changers



PANEL A: HOURLY-TO-HOURLY CHANGERS

PANEL B: SALARIED CHANGERS

Notes: Figure shows the 12 month change in nominal base wages for workers in our job-changer sample. We include all data from 2008-2016. See text for additional details.

base wage fell over the year experienced, on average, a 18.5 percent wage cut. Moreover, the standard deviation of annual nominal base wage changes for job-changers is 29.3 percent - almost five times larger than the standard deviation of annual nominal base wage changes for job-stayers.²²

7 Aggregate Nominal Wage Adjustments

We now combine our measures of base wage adjustments for job-stayers and job-changers into an aggregate measure of nominal wage adjustments.²³ This measure is appropriate for the study of movements of macro variables in models with no defined notion of a job-stayer or job-switcher, as is the case in canonical New Keynesian models such as Christiano et al. (2005) and Schmitt-Grohé and Uribe (2012). The large sample of both job-stayers and job-changers at a high frequency is a unique feature of the ADP data which allows us to construct such a measure for the first time. The inclusion of job-switchers vastly reduces the degree of realized nominal wage rigidity in the economy, particular on the downside, relative to the job-stayer benchmark which has been measured in the literature to-date.

²²The patterns of nominal wage changes of job-changers that we document using the ADP data are similar to the patterns found using French data in the 1990s as documented in Postel-Vinay and Robin (2002). Postel-Vinay and Robin (2002) document that about one-third of French workers experience a *real* wage decline as they move from job-to-job with no intervening unemployment spell during a period of relatively low inflation.

²³New entrants to the labor market also provide another margin of potential nominal wage adjustment. We are unable to measure job entrants within the ADP data so we abstract from them in our analysis.

Table 7: Nominal Base Wage Change Statistics, Pooled 2008-2016 Sample of Job-Changers

	Quarterly	Annual
Prob of Positive Change (%)	52.7	56.8
Prob of Negative Change (%)	37.4	38.0
Mean Unconditional Change (%)	6.3	8.0
Median Unconditional Change (%)	2.3	4.6
S.D. of Unconditional Change (%)	25.9	29.3
Mean Conditional Change (%)	7.0	8.5
Median Conditional Change (%)	5.0	6.0
S.D. of Conditional Change (%)	27.2	30.1
Mean Change, Conditional on Positive (%)	23.5	26.1
Median Change, Conditional on Positive (%)	16.7	18.5
Mean Change, Conditional on Negative (%)	-16.5	-18.5
Median Change, Conditional on Negative (%)	-13.6	-15.8

Notes: Table shows moments of the wage change distribution for job-changers for different horizons. For this table, we use our job-changer sample and pool together hourly and salaried workers.

To construct an aggregate measure of nominal wage flexibility, one must combine the patterns of wage adjustment for job-stayers with the patterns for job-changers. Were the universe of workers available, this would be a relatively easy task. However, as noted above, we can only measure job-changers who migrate between one ADP firm and another ADP firm. Given that ADP only has information on a subset of US workers, most job-to-job flows involve a non-ADP firm.

To circumvent this problem, we use aggregate data on job-to-job flows published by the US Census Bureau using data from the Longitudinal Employer Household Dynamics (LEHD) database.²⁴ Using matched employee-employer records, Census creates measures of quarterly job flows. In particular, we use the Census’s Job-to-Job Flows Data (J2J) focusing on transitions between workers’ main jobs. For any given worker in quarter t whose main job is at firm i , the J2J data measure the probability that the worker’s main job in quarter $t + 1$ remained at firm i (job-stayers), was at a different firm j (job-changers), or that the worker was not employed in quarter $t + 1$ (become non-employed). These three probabilities sum to 1 within each quarter. Using data from 2008 through 2016, the quarterly job staying rate averaged 88.7 percent, the quarterly job switching rate averaged 4.6 percent, and the

²⁴See https://lehd.ces.census.gov/data/j2j_beta.html, accessed June 30, 2018. We focus on the job-to-job flows at the quarterly frequency allowing for at most short unemployment spells between the job transitions.

quarterly transition rate to non-employment was 6.9 percent. At the time of writing, the Census has not yet released annual job-to-job flows. As a rough approximation, we construct annual job changing rates by multiplying the quarterly rates by 4. Doing so implies that 18.5 percent of workers switch job annually.²⁵ In order to aggregate our job-staying and job-changing results, we weight our job-changing data by the fraction of job-changers in the LEHD data relative to one minus the fraction of job-changers. For quarterly data, we ensure that job-changers are weighted so that they represent 4.8 percent of workers ($0.046/(1-0.046)$) on average. For annual data, we ensure that job-changers are weighted so that they represent 22.7 percent of workers ($0.185/(1-0.185)$).

Table 8 shows statistics for the aggregate nominal base wage change distribution combining data from both job-stayers and job-changers. Column 1 of the table shows quarterly statistics on aggregate base wage changes while column 2 shows similar annual statistics. The table shows that there is much more aggregate nominal base wage flexibility than one would conclude from looking at job-stayers alone. Over the entire sample period, roughly 71.3 percent of workers receive a nominal wage change. Of those, nearly 9 percent received nominal base wage declines, compared with 2 percent of job-stayers. While nominal base wage declines are still rare in the aggregate relative to nominal wage increases, including data on job-changers quadruples the amount of nominal base wage cuts relative to looking at only job-stayers. Moreover, the standard deviation of base wage growth – both unconditionally and conditional on a wage change – is over twice as large in the aggregate as amongst job-stayers. For example, unconditionally, the standard deviation of nominal base wage growth in the aggregate is 12 percent while the standard deviation of base wage growth for job-stayers is about 6 percent.

Overall, the inclusion of job-changers in our measures of wage rigidity greatly increases realized flexibility in the economy. Although job-stayer wage rigidity may have a complicated equilibrium relationship with the decision to switch jobs, the evidence presented here has important consequences for the quantitative predictions of existing macro models. The excessive rigidity inferred by simply considering base wage adjustment for job-stayers will lead New Keynesians to overstate the pass-through of monetary policy to real quantities.

²⁵This approximation is consistent with aggregate data on job tenure. Hyatt and Spletzer (2016) use tenure supplements to the CPS and matched employer-employee data from the LEHD to document that roughly 20-25 percent of workers have tenure less than a year during the 2008-2014 period.

Table 8: Moments of Aggregate Wage Change Distribution Combining Job-Stayers and Job-Changers, Pooled 2008-2016

	Quarterly	Annual
Probability of Positive Wage Change (%)	20.6	62.7
Probability of Negative Wage Change (%)	3.2	8.7
Mean Unconditional Change (%)	1.2	4.4
Median Unconditional Change (%)	0.0	2.5
S.D. of Unconditional Change (%)	6.7	12.0
Mean Conditional Change (%)	5.3	6.2
Median Conditional Change (%)	3.1	3.5
S.D. of Conditional Change (%)	13.1	13.9

Notes: Table shows aggregate moments of base wage adjustment at different horizons combining data on both job-stayers and job-changers during the 2008-2016 period. For this table, we pool together both hourly and salaried workers. We use our employee sample for this analysis. See text for additional discussion of our job-changer sample.

8 The Cyclicalities of New Hire Base Wages

As Pissarides (2009) makes clear, the flexibility of new hire wages is the key determinant of aggregate employment fluctuations in many models. As a result, a large literature dating back to Bils (1985) has established that job-changers' wages are more cyclical than those of job-stayers, which we confirm with the ADP data above. However, the *cyclicalities* of new hire wages does not necessarily map directly into new hire wage *rigidity*. This is principally due to selection: if lower quality workers are more likely to be displaced during a recession, wages of new hires may look lower in recessions than in booms, even if the wage per efficiency unit of labor were perfectly rigid. Similarly, if firms hire higher quality workers for a given job in a recession, the cyclicalities of measured new hire wages may be different to that of new hire's wage per efficiency unit. This idea has been well established in the literature by, for instance, Solon et al. (1994, 1997), and more recently by Gertler et al. (2016) who show that much of the cyclicalities of new hire wages may be explained by pro-cyclical match quality.

A cleaner measure of new hire wage rigidity would compare the evolution of wages of new hires within a firm with that of similar incumbent workers within the same firm. If new hire wages are more flexible than incumbent workers, then job-changer wages should be lower than the wages of job-stayers who work in the same job during downturns, and higher (modulo tenure effects) in booms. By conditioning on worker type, this measure reduces the influence of selection on new hire wage cyclicalities, thereby better reflecting traditional notions of wage rigidity.

Controlling for selection in this way is challenging in practice. We proceed by constructing a matching estimator, which compares the year-over-year base wage adjustment of job-changers to a matched job-stayer, who is of a similar age and initial wage level (a proxy of the worker’s unobserved type) and works at the changer’s destination firm. Throughout this analysis, we focus on our job-changer sample. Define $w_{j,t-1}^i$ as the base wage (measured in nominal dollars per hour) of worker i who works at firm j in year $t - 1$. As above, we assume that salaried employees work 40 hours per week when making their hourly base wage measure. We treat this as a measure of a worker’s quality in $t - 1$. Define a_{t-1}^i as individual i ’s 5-year age bin (e.g. 21-25, 26-30, etc.) in year $t - 1$. Finally, define p_{t-1}^i to be individual i ’s percentile within the national wage distribution in year $t - 1$.

To construct our matching estimator, we begin by examining the base wage adjustment of job-stayers. Denote by $\mathcal{I}_t^j(p, a)$ the set of workers who work for firm j continuously between years $t - 1$ and t , and who lie in wage percentile p and age group a in $t - 1$:

$$\mathcal{I}_t^j(p, a) = \{i : (p_{t-1}^i = p) \cap (a_{t-1}^i = a) \cap (i \text{ works for } j \text{ in } t - 1 \text{ and } t)\}$$

Denote by $N_t^j(p, a)$ the size of this set: i.e. the number of job-stayers at firm j at percentile p and age a . For every combination of wage percentile and age bin within a firm, we construct the mean wage change for job-stayers:²⁶

$$\Delta_t^j(p, a) = \frac{\sum_{i \in \mathcal{I}_t^j} (\ln w_{j,t}^i - \ln w_{j,t-1}^i)}{N_t^j(p, a)}$$

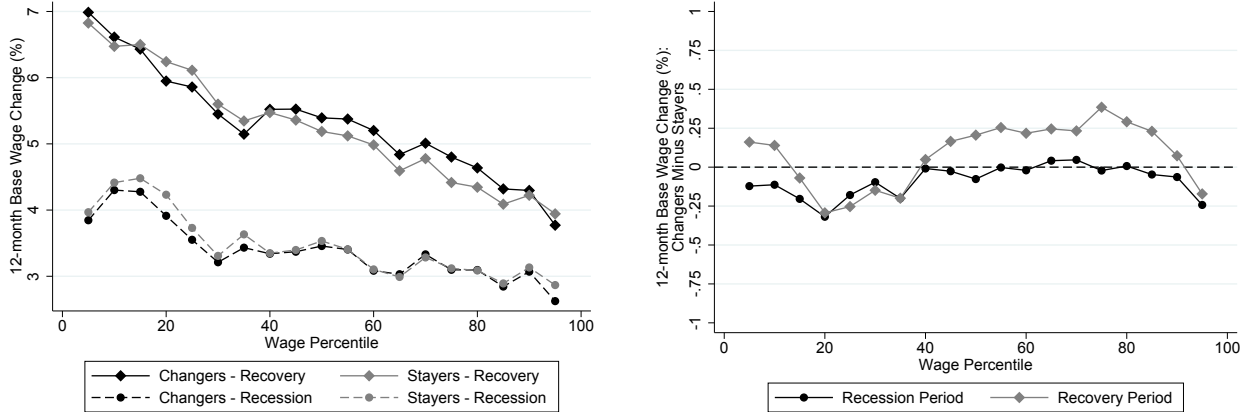
The goal is to compare this $\Delta_t^j(p, a)$ to the wage changes of similar job-changers into firm j . We therefore turn to our job-changer sample and construct the 12-month base wage change for every switcher in our sample as they move from firm j to j' , exactly as in section 6. Denote the change of switcher i as dw_i^c . Then we compare each job-changer i moving to firm j' with a matched job-stayer at that same firm with the same initial wage percentile and age group; that is we compare dw_i^c to $\Delta_t^j(p_{t-1}^i, a_{t-1}^i)$.

Figure 4 plots the results of this matching exercise. Panel A plots the mean wage changes of job-changers (black lines) and their matched job-stayers (gray lines) throughout the base wage distribution. The solid lines with diamond markers show the patterns for the recovery period from 2012-2016, while the dashed lines with circle markers present the mean wage changes in the recession period of 2008-2010.²⁷ The plot shows a negative relationship be-

²⁶To maximize power, we use the full ADP data for this exercise, rather than any of our subsamples.

²⁷We exclude the highest and lowest ventiles from the plot, because bottom earners (bottom five percent of the wage distribution) are constrained by institutional details such as the minimum wage, and the distribution

Figure 4: Job-Changers' Wage Changes Compared with Matched Job-Stayers, by Period



PANEL A: CHANGERS AND STAYERS SEPARATELY

PANEL B: CHANGERS MINUS STAYERS

Notes: Figure plots results of our matching exercise detailed in the text. Job-changers are matched to job-stayers in the same destination firm, at the same initial wage percentile, and same age. We then plot the mean 12 month wage changes of both changers and their matched stayer by wage ventile (Panel A), and the mean difference between changer and matched stayer (Panel B).

tween mean wage changes and initial wage percentile for both changers and stayers, in part reflecting life cycle effects.²⁸ However, this relationship is less pronounced during the recession. In the recovery period, an individual at the median of the national income distribution could expect a 5.5% increase in wages on average, but this fell to 3.5% during the recession, which mirrors the procyclicality of new hire wages documented in the literature.

One can also see from Panel A that job-changers' wages evolve similarly to those of their matched job-stayer counterparts: the black lines trace the gray lines very closely. Panel B plots the difference between changers' and matched stayers' year-over-year wage changes for the recession (dashed line) and recovery (solid line) periods. For the majority of the wage distribution, the recovery line is above the recession line, indicating that job-changers had lower wage growth relative to similar job-stayers in the recession than in the recovery. This suggests that new hire wages are more flexible than that of incumbent workers. However, the size of this difference is quite small in magnitude: on average, changers experienced a wage increase relative to their matched job-stayers which was 22 basis points lower in the recession than in the recovery. Indeed, if we regress, at the individual changer level, the gap between changers' wage growth and their matched stayer's wage growth ($\Delta_t^j(p_{t-1}^i, a_{t-1}^i) - dw_i^c$) on a

of top earners' income (top five percent of wage distribution) is quite disperse, reducing the reliability of our match based on wage percentiles.

²⁸Although the matching between job-changers and stayers is done with age, the wage percentiles are calculated using the *national* distribution. Thus Panel A includes life cycle effects in the mean wage change. In addition, we match based on percentiles, but collapse to mean changes by ventile for legibility.

recession indicator controlling for individual age, initial wage percentile, industry, firm size, sex, and an indicator for being hourly, we find that the gap between changers and stayers becomes a statistically significant 19 basis points more negative in the recession.²⁹ Given the individual standard deviation of this gap is 8 percentage points, and the mean gap is just 24 basis points over our whole sample, this constitutes a relatively small effect. For comparison, Gertler et al. (2016) find that a one percentage point increase in the aggregate unemployment rate is associated with a reduction of job-changer wage growth which is one percentage point larger than that of job-stayers.³⁰ We find that the roughly 3 percentage points higher unemployment rate during the crisis was associated with a disproportionate reduction in job-changer wage growth of about 20 basis points, implying a semi-elasticity of 0.06, roughly one-twentieth of Gertler et al. (2016)’s exercise which does not control for compositional shifts of changers over the cycle.

Collectively, our new hire results yield a few key insights. First, we confirm below that job-changers have more cyclical wages than do job-stayers. Furthermore, after controlling for selection as best we can, it remains true that job-changers’ wages respond more in downturns than do job-stayers, suggesting some higher degree of new hire flexibility. However, this residual differential flexibility for job-changers appears quite small, suggesting that the majority of the increased cyclical wage growth of job-changers can be explained through composition effects, as posited by Gertler et al. (2016). The results presented in this section complement those of Hazell and Taska (2018), who document that posted wages adjust about as frequently as do the wages of incumbent workers. New hire wages at the job level do not appear substantially more flexible than do those of job-stayers, suggesting that internal equity concerns may be important in wage setting.

9 Including Bonuses in Measures of Nominal Wage Adjustments for Job-Stayers

In this section, we gauge the importance of bonuses in providing an additional margin of flexibility for job-stayers. Bonuses have long been considered a potential source of additional wage cyclical flexibility and earnings flexibility (Shin and Solon, 2007).³¹ For this analysis, we restrict our employee sample to include only non-commission workers who remain con-

²⁹We cluster standard errors at the destination firm level for this exercise; $p < 0.001$.

³⁰A similar regression in our data shows that the response of 12-month base wage growth to a 1 percentage point change in aggregate unemployment is 1.4 percent larger for job-changers than for job-stayers.

³¹Makridis and Gittleman (2018) show that performance pay jobs have larger adjustments in compensation per worker over the business cycle than do fixed wage workers in the National Compensation Survey, but are unable to measure adjustments at the individual worker level.

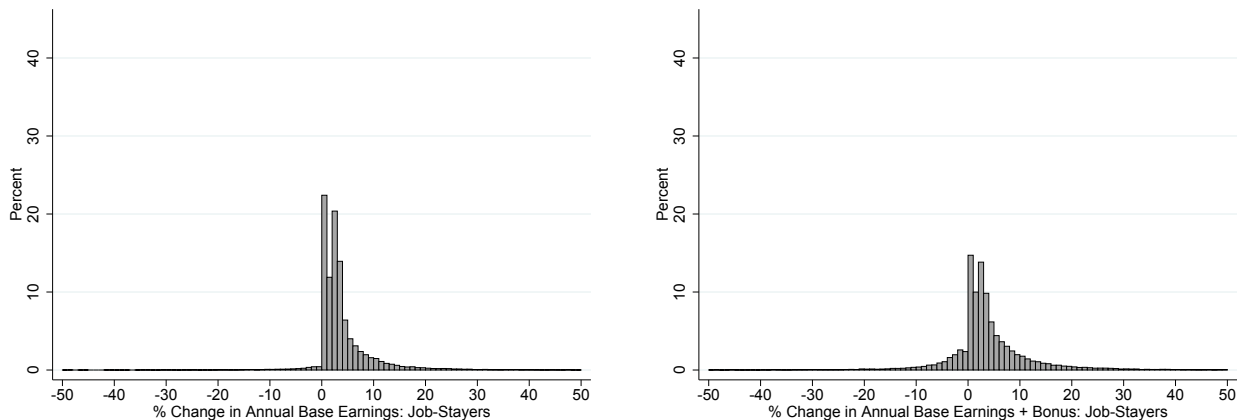
tinuously employed with the same firm for twenty four consecutive calendar months. The twenty-four consecutive calendar month restriction is necessitated by the fact that bonuses accrue annually. Given that our data starts mid-year through 2008, our bonus sample pools together workers for all two year periods between 2009 and 2016. During our sample period, 44.5 percent of non-commission job-stayers received no bonus both in year t and year $t + 1$ while 32.4 percent of non-commission job-stayers received a bonus in both year t and $t + 1$. The remaining 23.1 percent of workers received a bonus in only one of the two years.

To assess how bonus flexibility affects the frequency of nominal wage adjustment, our bonus and base wage measures must be in the same units. We therefore define a measure of “modified monthly base earnings,” which isolates fluctuations in a worker’s base pay which arise from changes in wages rather than hours. Since salaried workers have no meaningful hours adjustments, modified monthly base earnings for salaried workers is simply their monthly base earnings (their base wage times the number of pay periods per month). However, modified monthly base earnings for hourly workers is their monthly base wage times the *average* monthly hours worked over the relevant two year period. Specifically, when we explore changes in nominal wage adjustments inclusive of bonuses between year t and $t + 1$, average monthly hours worked for hourly workers is simply the total annual hours worked in both years t and $t + 1$ divided by 24. By fixing number of hours for hourly workers, we ensure that all the movements in modified base monthly earnings is coming from changes in the base wage; it is in this sense that the earnings are “modified”. Given that bonuses are measured annually, we then make a measure of “modified annual *base earnings*” by summing the monthly modified earnings for each worker over a 12 calendar month period.

Our key variable of interest is “modified annual *total earnings*” which is computed as the sum of modified annual base earnings plus annual bonuses. It is worth stressing that our *modified* annual total earnings variable differs from annualized total earnings in two ways. First, the measure includes only earnings from base wages and bonuses. It excludes any earnings from overtime premiums and other small infrequent payments. Second, it normalizes the hours for hourly workers to be the average annual hours worked over a given two year period. The benefit of this is that any movement in this earnings measure across the two years is attributed to a change in base wages or a change in bonuses.

Panel A of Figure 5 shows the 12 month change in modified annual base earnings for non-commission workers who remain continuously employed on the same job for two calendar years. Not surprisingly, this figure is very similar to those presented in 2. Both only measure variation in base wages over time. The only difference between the two figures results from selection on non-commission workers and time aggregation. Most of the difference stems from time aggregation. Given that Panel A of Figure 5 measures nominal wage adjustments

Figure 5: Annual Changes in Modified Base Earnings and Modified Annual Earnings: 24-month Job-Stayers, Excluding Commission Workers



PANEL A: MODIFIED BASE EARNINGS

PANEL B: MODIFIED ANNUAL EARNINGS

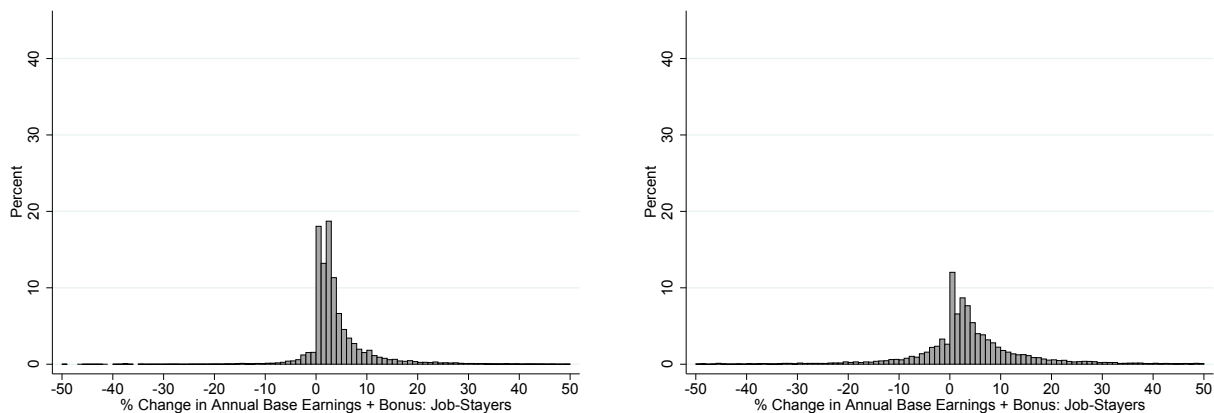
Notes: Figure plots the distribution of year-over-year changes in base wage (Panel A) and modified base earnings (Panel A) and modified annual earnings (Panel B). Modified base earnings are base earnings holding fixed hours for hourly workers. Modified annual earnings is modified base earnings plus annual bonus payments. Figure restricts attention to a sample of 24-month job-stayers excluding commission workers, between 2009 and 2016.

between all of $t - 1$ and all of t , wage changes that occur prior to December of $t - 1$ will show up as an additional nominal wage adjustment relative to Figure 2. As a result, the variance of changes is slightly higher for annual modified base earnings relative to 12-month changes in base wages.

Panel B of Figure 5 presents the annual change in modified earnings (inclusive of bonuses) for all non-commission workers in our sample. Relative to Panel A, the change in annual modified earnings (inclusive of bonuses) is much more dispersed. While only 2.9 percent of workers received an annual modified base earnings decline, 15.7 percent of workers saw reduced annual nominal wages inclusive of bonuses. Additionally, accounting for bonuses increases the standard-deviation of nominal wage changes from 4.6 percent to 7.7 percent.

Bonuses provide firms with an additional margin of wage adjustment for their employees. As highlighted above, at the worker level, bonuses are not persistent. A worker's bonus in year t is not predictive of the amount of the worker's bonus received in year $t + 1$. Thus, while bonuses provide some flexibility to a worker's spot wage in a given year, they exert less influence on the user cost of a worker. However, bonuses are certainly important determinants of the volatility of a worker's earnings from year to year. Indeed, the results in Figure 5 are consistent with contemporaneous work by Kurmann and McEntarfer (2018) and Jardim et al. (2019) who use data from the US Longitudinal Employer Household Dynamics (LEHD) survey and Washington State Unemployment Insurance Records to examine nominal

Figure 6: Modified Annual Earnings Fluctuations by Base Wage Quartile, Two-Year Sample of Job-Stayers Excluding Commission Workers, 2009-2016



PANEL A: BASE + BONUS CHANGES:
BOTTOM QUARTILE WAGES

PANEL B: BASE + BONUS CHANGES:
TOP QUARTILE WAGES

Notes: Figure shows the distribution of modified annual earnings for our sample of workers who remain continuously employed with the same firm for two consecutive calendar years. Panels A and B show, respectively, patterns for those in the bottom and top quartiles of the base wage distribution, measured before the adjustment.

annual earnings-per-hour adjustments for a sample of job-stayers who reside in Washington state.³² Both Kurmann and McEntarfer (2018) and Jardim et al. (2019) find that roughly 20 percent of workers receive a decline in nominal earnings per hour during a given year. While their results are not directly comparable to ours because they are focused only on residents of Washington state and also include overtime premiums and commissions in their earnings measure along with base wages and bonuses, it is encouraging that our annual variation in annual modified earnings matches closely their measure of annual variation in earnings per hour. What distinguishes our broader results from theirs is our ability to highlight how much of the variation is coming from base earnings verses bonuses. As highlighted throughout, such distinctions are important for many potential applications.

If bonuses do provide firms with additional flexibility in adjusting the nominal wages of their workers, they only do so for higher income workers. Figure 6 presents the distribution annual changes in modified base earnings plus bonuses (modified annual earnings) for two groups of workers: those in the bottom quartile of the base wage distribution (left panel) and those in the top quartile of the base wage distribution (right panel). For most low-wage workers, bonuses are trivial share of income. As a result, accounting for bonuses has little effect on low-wage workers' nominal wage adjustments, leading them to hardly ever receive

³²They focus their sample on residents of Washington state because Washington requires employers to report the hours worked of their employees as part of their Unemployment Insurance program.

Table 9: Moments of the Annual Change Distribution Across Wage Notions, Two-year Job-Stayers, 2012-2016

	Modified Base Earnings	Modified Annual Earning	Aggregate Modified Annual Earnings
Prob of Positive Change (%)	80.4	75.0	68.5
Prob of Negative Change (%)	2.9	15.7	21.6
Mean Unconditional Change (%)	3.7	4.0	5.2
Median Unconditional Change (%)	2.7	2.7	2.8
SD of Unconditional Change (%)	4.6	7.7	16.9
Mean Conditional Change (%)	4.5	4.4	5.8
Median Conditional Change (%)	3.0	3.0	3.3
SD of Conditional Change (%)	4.7	7.9	18.6

Notes: Table plots key moments of the wage change distribution for the sample of two-year job-stayers. Each column presents the moments of the distribution of a separate notion of wage. Column 1 highlights modified base earnings for job stayers. Column 2 highlights modified annual earnings which combines base earnings and bonuses for job stayers. Column 3 highlights aggregated modified annual earnings which combines changes in modified annual earnings for job-stayers and changes in base wages for job-changers.

nominal wage cuts. However, for high-wage workers, bonuses provide a large amount of additional flexibility. Incorporating bonus adjustments results in 23.6 percent of workers in the top wage quartile experiencing a nominal wage cut during a year. This result is interesting given the fact that much of the displacement of workers at business cycle frequencies comes from the bottom of the wage distribution. It also contrasts with the results we highlight in the Online Appendix showing that the probability of base wage adjustments does not vary with a worker’s base wage. The results in Figure 6 suggest there are potential gains from modeling heterogeneity across workers in the extent to which bonuses are part of their compensation if bonuses are an important margin of adjustment within the model.

The first two columns of Table 9 summarizes results on the frequency of modified base earnings and modified annual earnings (inclusive of base and bonuses) for our sample of 2-year job-stayers. The third column of the table shows a measure of aggregate nominal wage adjustments inclusive of bonus income for job-stayers and accounting for the aggregate amount of job-changers. Specifically, we focus on annual modified earnings for job-stayers (column 2) and combine that with annual base wage changes for job-changers with the job-changers weighted accordingly using the aggregate job-to-job transition weights discussed

above.³³ In aggregate data, spot wages are more flexible both because of job-changers and because bonuses are moving around at annual frequencies. As seen from Column 3, in the aggregate, 21.6 percent of all workers receive a nominal wage decline during a given year when accounting for job-changers and bonuses received by job-stayers.

We conclude this section by briefly discussing how accounting for fringe benefits affects the frequency of nominal wage adjustments. When exploring the importance of fringe benefits, we created an additional wage measure which was defined as modified annual earnings (inclusive of bonuses) plus the annual value of employer-provided health insurance and contributions to deferred compensation plans. Incorporating fringe benefits into our analysis does not substantively alter the frequency of annual changes either up or down. The probability of an increase in base earnings plus bonuses and fringe was about 78 percent. The comparable probability for the increase in base earnings plus bonuses (column 2 of Table 9) is 75 percent. Both the size of the increase and the standard deviation of the change are larger after including fringe benefits with the standard deviation increasing from 7.7% to 10.9%. These patterns are consistent with benefits being a relatively constant fraction of a worker’s earnings. As a worker’s base earnings change, so too do their fringe benefits. Under this scenario, accounting for fringe does not alter the propensity for a worker to receive a wage change, but would scale up both the mean and the standard deviation of changes.

10 State Dependence in Nominal Wage Adjustment

In this section, we examine the extent to which aggregate wage adjustments move with aggregate and firm-specific conditions. Many macro models of wage adjustments assume a constant probability that a worker receives a wage adjustment. Using a variety of methodologies, we highlight that the probability of a base wage adjustment varies substantively with aggregate and firm level economic activity, but that bonuses are relatively acyclical.

10.1 Time Series Variation in Nominal Base Wage Adjustments

There are two principal reasons why one might observe state dependence in realized nominal wage changes for job-stayers. The first is if there is some explicit cost for firms to adjusting the wages of existing workers. Non-convex adjustment costs, or “menu costs,” are commonly employed in New Keynesian models of price setting in order to match moments of the price data. Fixed adjustment costs generate an inaction region whereby firms that are close

³³We cannot include bonuses for job-changers, as we do not observe a full year of data for job-changers that join their jobs in months other than January.

to their optimal price in a frictionless economy do not adjust their prices until they move sufficiently far away from their optimal price. Thus, with a menu cost of adjusting prices, the state of the firm - its distance from the optimal pricing rule - is central to price adjustment decisions. As a result, price changes are infrequent, and large when they occur. Although menu cost models of wage adjustments are rare, principally due to challenges arising from wage bargaining, the intuition gained from the output pricing literature helps guide analysis of state dependence in wage setting.

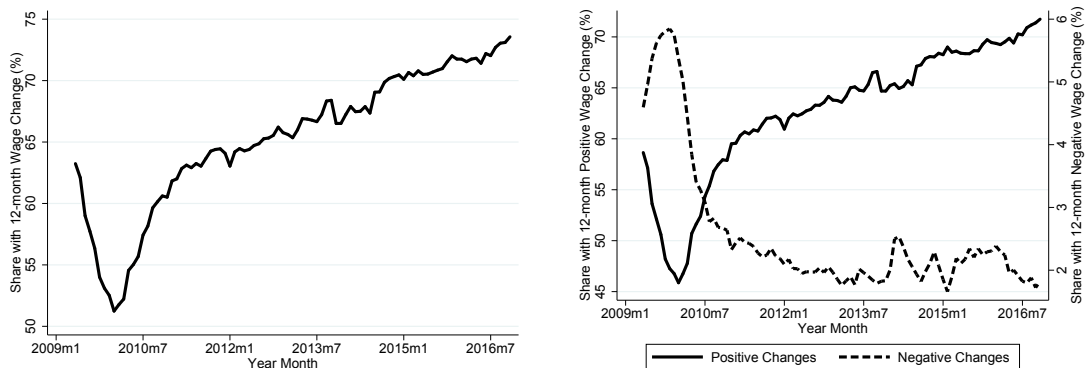
A second reason for state dependence in nominal wage adjustments for job-stayers might arise in a framework with asymmetric rigidity. For instance, suppose that it is harder for firms to cut wages than to raise them, possibly due to concerns over morale or because of union pressure. Under this scenario, firms receiving a negative productivity shock would have a lower probability of being able to adjust wages to the desired level than firms receiving a positive productivity shock. This would imply that wages would then appear less flexible in downturns than in booms.

Figure 7 plots the time series of base wage adjustments for job-stayers using our employee sample pooling together both hourly and salaried workers. The top panel plots the extensive margin of base wage changes: the percent of all employees in month t who have a different base wage from month $t - 12$. As a reminder, our data starts in May 2008. That means the first observation in each of the panels in Figure 7 is for May 2009 and measures the fraction of job-stayers who received a base wage change between May 2008 and May 2009. The fact that our data spans the Great Recession allows us to explore business cycle variation in the extent of base wage adjustments.

As seen from the left panel of Figure 7, wage adjustments of job-stayers exhibit striking pro-cyclicality. Only about 55 percent of continuing wage workers received a year-over-year wage change during the depths of the recession. However, after the recession ended, during the 2012 to 2014 period, between 65 and 70 percent of workers received a wage change. As of the end of 2016, nearly 75 percent of all workers received a nominal base wage change. While most of the time series variation was between the recession and non-recessionary periods, there is still a trend upwards in the share of workers receiving an annual base wage change between 2012 and 2016.

The right panel of Figure 7 separates the probability of a base wage change of job-stayers into the probability of a base wage increase (solid line - measured on the left axis) and the probability of a wage declines (dashed line - measured on the right axis). During the Great Recession, the propensity of base wage increases for job-stayers fell sharply while the propensity of base wage declines rose sharply. Although nominal base wage cuts are exceedingly rare for job-stayers during non-recessionary periods, upwards of 6 percent of all

Figure 7: Time Series of Nominal Base Wage Adjustments, Job-Stayers



PANEL A: HAS WAGE CHANGE

PANEL B: HAS WAGE CHANGE:
POS. VS NEG.

Notes: Figure plots the propensity to receive a 12-month base wage change over time for our employee sample of job-stayers between May 2009 and December 2016. The data are weighted to match the firm size \times industry mix found in the BDS.

continuing workers received a nominal base wage cut during late 2009 and early 2010.

Although not shown in the figure, the mean size of base wage changes, conditional on a wage change occurring, is also highly procyclical. During the Great Recession, the mean size of a nominal base wage change for those receiving a base wage change was about 5 percent. The corresponding number as of 2016 was 6.5 percent. Putting the above results together, within-firm nominal base wage growth for job-stayers is highly pro-cyclical. For employees who remain with the firm, both the probability and size of nominal base wage raises are increasing in business cycle conditions. There is a large literature on the cyclicity of aggregate wages.³⁴ Our results show that for a given worker on the job, nominal base wage *changes* are also strongly pro-cyclical.

The first two columns of Table 10 summarizes the business cycle differences in nominal base wage adjustments for job-stayers. We separate the sample into two periods: a period representing the depths of the Great Recession (May 2009–December 2010) and a period well into the recovery (January 2012 – December 2016). Consistent with the results in Figure 7, nominal wage cuts were more prevalent during the recession than during the recovery period. Base wages are less downwardly rigid during the Great Recession. 4 percent of job-stayers overall and 6 percent of salaried job-stayers received a base wage decline during the Great Recession. While base wages appeared more downwardly flexible during the Great Recession for job-stayers, the fraction of workers receiving a zero nominal base wage

³⁴See, for example, Solon et al. (1994), Gertler et al. (2016) and the recent survey by Basu and House (2016).

Table 10: Summary of 12-Month Wage Change Distribution During and After the Great Recession, Job-Stayers and Job-Changers

	Job-Stayers		Job-Changers		Aggregate	
	2009-10	2012-16	2009-10	2012-16	2009-10	2012-16
$Pr\{\Delta w < 0\}$	4.2	2.0	44.0	36.4	10.6	8.4
$Pr\{\Delta w = 0\}$	43.3	30.6	5.5	5.5	37.4	25.6
$Pr\{\Delta w > 0\}$	52.5	67.4	50.5	58.1	52.0	66.0
S.D. of Δw (%)	6.3	7.0	31.4	28.9	12.1	12.4

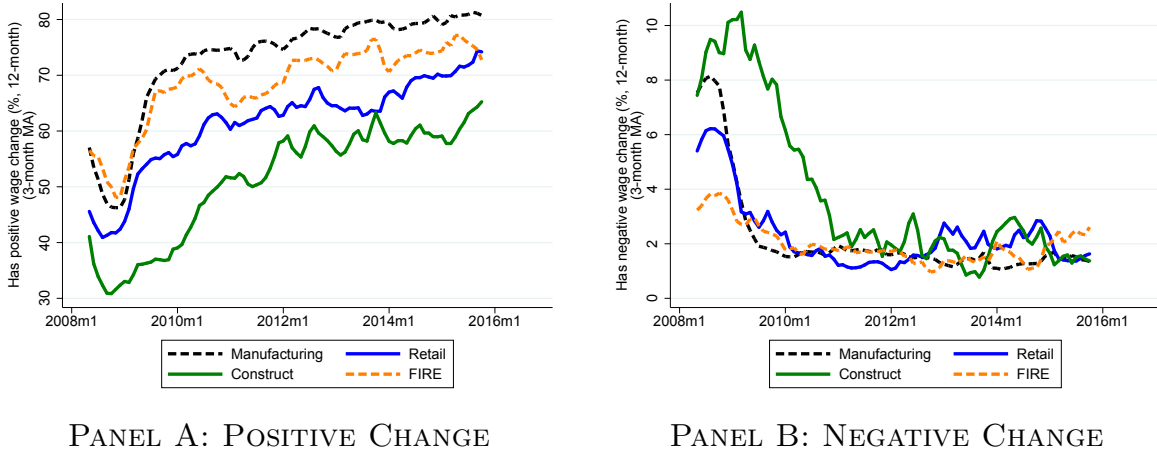
Notes: Table shows the distribution of 12-month wage adjustment for job-stayers (columns 1 and 2), job-changers (columns 3 and 4), and a combination of both (columns 5 and 6) over the cycle. The results are for a sample that pools together both hourly and salaried workers. We show statistics for the period between May 2009 and December 2010 and the period between January 2012 through December 2016.

change also increased for both salaried and hourly workers. Interestingly, the unconditional standard deviation of base wage changes fell slightly during the recession for all workers from 7.0 percent to 6.3 percent. In summary, overall nominal base wage adjustments fell during the Great Recession but downward adjustments increased.

Figure 8 shows time series trends in the probability of a nominal base wage increase (left panel) and a nominal base wage cut (right panel) for job-stayers by industry. During the Great Recession, manufacturing and construction were two of the hardest hit industries. Roughly 10 percent of construction workers and 8 percent of manufacturing workers who remained on their job received a year-over-year nominal base wage cut during 2009. The comparable numbers for retail and finance, insurance, and real estate (FIRE) were 6 and 3 percent, respectively. By 2012, continuing workers in all industries had a roughly 2 percent probability of receiving a nominal base wage cut. Note, the probability of a nominal base wage increase did not differ markedly across industries during the Great Recession. There are persistent level differences in the propensity of a nominal base wage increase across industries for job-stayers in all years. However, these differences remained relatively constant during the 2008-2016 period. These cross-industry patterns reinforce the time series patterns with respect to the state dependence of nominal base wage cuts of continuing workers. Not only were nominal base wage cuts more likely for job-stayers during the Great Recession, the propensity of nominal base wage cuts was highest in the industries hit hardest during the Great Recession. Firms in manufacturing and construction both were more likely to shed workers during the Great Recession and also were more likely to cut the base wages of the workers who remained with their firm.

As highlighted above, much of the flexibility in nominal base wage adjustments results from job-changers. The middle two columns of Table 10 show that the distribution of base

Figure 8: Time Series of Wage Changes by Industry, Job-Stayers



Notes: Figure shows the propensity to receive a 12-month wage increase (Panel A) and decrease (Panel B) for job-stayers over the cycle, broken out for select broad industry groups. This figure makes use of our employee sample, weighted to match the firm size distribution found in the BDS within each industry. “FIRE” refers to Finance, Insurance, and Real Estate.

wage adjustments for job-changers also varies over the business cycle. For the table, we report statistics of 12 month base wage changes for workers employed at firm j in t and then are subsequently employed at firm j' in $t + 12$. During the Great Recession, 44 percent of workers who changed jobs received a nominal base wage decline. The comparable number during the recovery period was 36.4 percent. Like job-stayers, there was more downward nominal adjustment during the Great Recession.

Aggregate nominal base wage flexibility is a function of both the base wage adjustments for job-stayers and job-changers. However, in order to measure the cyclical nature of aggregate wage adjustments, we need to know how the composition of job-stayers relative to job-switchers evolves at business cycle frequencies. We again use data from the Census’s Job-to-Job Flow Data (J2J) made from the underlying data of the LEHD. Appendix Figure A15 shows the quarterly share of job-stayers and job-switchers in the J2J data between 2000 and 2015. 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. When weighting the job-stayer and job-changer data, we ensure that 16 percent of workers were job-changers during the Great Recession

and roughly 20 percent were job-changers during the recovery. Since job-changers receive nominal wage changes and cuts at a substantially higher rate than job-stayers, this composition effect therefore pushes towards lower aggregate flexibility during the recession, even if both changers and stayers observe less downward rigidity in recession periods.

The final two columns of Table 10 shows the cyclical patterns of aggregate nominal base wage adjustments combining data on both job-stayers and job-changers. As with Table 10, we break our sample into two periods: May 2009-December 2010 and January 2012-December 2016. Focusing on the annual aggregate nominal base wage adjustments, 10.6 percent of workers in the aggregate economy received a nominal base wage decline during the Great Recession. The comparable number during non-recession times was 8.4 percent. While downward adjustments were slightly more common, upward adjustments were less common with only 52 percent of workers receiving a base wage increase during the 2009-2010 period. The unconditional standard deviation of base wage changes was slightly lower during the Great Recession.

Downward nominal wage rigidity has received a substantial attention as an explanation for why aggregate wages did not fall more during the Great Recession. The results above show that 10.6 percent of workers received nominal base wage cuts and another 37.4 percent received no nominal base wage increase. Much of the downward adjustments occur through job-changers. Moreover, both job-stayers and job-changers experienced more nominal base wage declines during the Great Recession than during the 2012-2016 recovery. However, the job-changing propensity also fell during the Great Recession reducing some of the aggregate flexibility in nominal base wage adjustments. The results in Table 10 provide a set of moments for researchers to calibrate models to assess whether the moments of the base wage adjustment distribution can lead to sufficient rigidities to explain why aggregate wage growth did not fall more during the 2008-2012 period. The results in this subsection show that the composition of nominal base wage adjustment for varies over the business cycle. Any model that assumes a constant hazard of base wage adjustments for job-stayers over the business cycle is at odds with the underlying wage setting data, and may lead to incorrect conclusions regarding the responsiveness of the economy to countercyclical monetary expansions.

10.2 The Cyclicity of Compensation Components

Thus far, we have focused on base wages under the presumption that bonuses do not provide substantial flexibility in the allocative user cost of labor. In this section, we study the cyclicity of bonus and base wage adjustment separately. Unlike the probability of base wage changes, both the probability of bonus receipt and the average size of bonus receipt is

Table 11: Cyclicity of Various Forms of Compensation - State level regressions, 2009-2016

	% With Pos. Wage Change (1)	% With Neg. Wage Change (2)	% Receiving Bonus (3)	Share in Bonus (4)	Log Mean Bonus (5)
Unemp Rate (%)	-2.22*** (0.45)	0.49*** (0.16)	0.20 (0.54)	-0.04 (0.06)	-0.03 (0.03)
Observations	350	350	350	350	350
Mean of Dep. Var.	75.04	1.93	38.88	2.98	8.10
S.D. of Dep. Var.	9.60	1.63	5.30	0.67	0.43

Notes: Table reports coefficients estimated from equation (2). An observation is a state-year. Columns 1 and 2 examine the probability that an individual in that state-year receives a positive and negative December base wage change year-over-year, respectively. Columns 3-5 show the cyclicity of the probability of receiving a bonus, the mean share of income in a bonus, and the log mean bonus payment (inclusive of 0s). All regressions include state and year fixed effects. Standard errors clustered at the state level reported in parentheses. ***, **, and * represent that the coefficients are statistically different from 0 at the 1%, 5%, and 10% level, respectively.

roughly acyclical during the 2009-2016 period. Since our sample period contains only one aggregate recession, it is hard to tease out cyclical movements from our short time series. To get around this problem, we exploit cross-region variation in business cycle conditions during the 2009 to 2016 period. We focus on our sample of 12-month job-stayers and collapse the ADP data down to the state-year level. We then regress various measures of adjustment - the share of employees receiving a positive or negative annual base wage change (measured in December), the share of workers receiving a bonus, the mean share of workers' income accounted for by bonuses, and the log of the mean bonus payment in a given year - on the average unemployment rates prevailing in a given state-year.³⁵ Specifically we estimate OLS regressions of the form

$$y_{st} = \beta u_{st} + \delta_s + \gamma_t + \epsilon_{st} \quad (2)$$

where y_{st} is a measure of wage or bonus adjustment in state s in year t , u_{st} is the average unemployment rate in state s in year t , and δ_s and γ_t are state and year fixed effects.

Estimates of β are presented in Table 11. The table shows that states with a one percentage point higher unemployment rate is associated with approximately 2.2 percentage points fewer nominal base wage increases, and 0.5 percentage points more nominal base wage decreases, conditional on state and year fixed effects. These changes are both statistically and

³⁵One may be concerned that bonuses are not determined by the state-level performance of the firm, but rather the firm's national performance. To test this, in unreported results, we estimate aggregate time series regressions using the aggregate unemployment rate as the measure of cyclicity. The results have larger standard errors as a result of the short time series, but are qualitatively similar.

economically meaningful, accounting for 0.23 and 0.3 of a standard deviation in state-level base wage increase and decrease probabilities, respectively. In contrast, however, bonuses exhibit no statistically-significant cyclical pattern, whether “bonuses” are measured as the percent of a state’s population receiving a bonus, the mean share of compensation being granted in bonuses, or the natural logarithm of the mean bonus earnings.³⁶

To some degree, this result is unsurprising. Many optimal contracting models suggest that incentive pay, such as bonuses, should filter out fluctuations in performance driven by factors outside of the control of the agent, such as the aggregate state of the economy. Our result suggests that, while bonuses may provide a natural buffer for employers to cut compensation for unexpected idiosyncratic underperformance, they are unlikely to systematically “grease the wheels” of the macroeconomy by alleviating downward nominal rigidity in the allocative wage. In contrast, although base wages appear relatively rigid and are rarely adjusted downwards, the fact that their adjustment patterns respond to aggregate conditions is consistent with base pay being a decent proxy of the allocative wage.

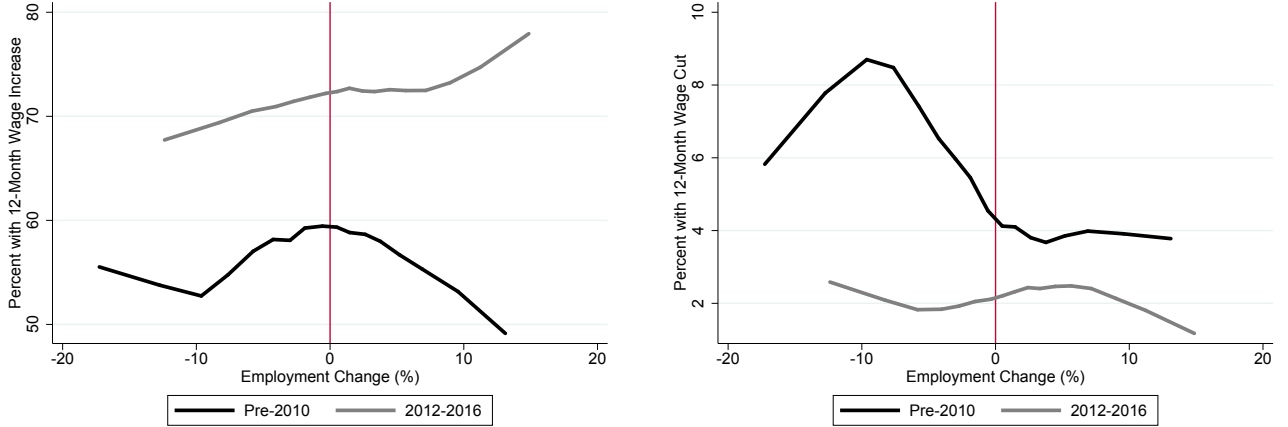
10.3 Within Firm Variation in Nominal Wage Adjustments

To further shed light on the extent to which nominal wage adjustments respond to firm-level shocks, we explore patterns in nominal wage adjustments by growing and contracting firms. To do so, we consider state dependence at the firm level by comparing moments of the wage change distributions for firms with differing levels of employment growth. We view this analysis as being a reduced form combination two potential effects. First, firms who experience negative employment growth are more likely to have received a negative firm-level shock. Firms receiving negative shocks are likely to adjust their labor inputs on multiple margins. As a result, firms experiencing negative employment growth may also be more likely to cut the wages of their existing workers and be less likely to give their existing workers a wage increase. Second, there may be a trade-off between adjusting wages and adjusting employment. For example, a financially-constrained firm receiving a negative cash flow shock may choose to cut wages or layoff workers in order to reduce their labor costs. Such a trade-off could in theory result in a negative relationship between firm employment growth and firm propensity to cut nominal wages.

Figure 9 plots the propensity for a firm to increase their workers nominal base wages (left panel) and the propensity for a firm to cut their workers’ nominal base wages (right panel) as a function of observed firm size growth. To make this figure, we use our firm-level sample.

³⁶In unreported results, we also find no systematic cyclical pattern in the *dollar-weighted* share of pay in bonuses, or in the mean year-over-year bonus changes. As a result, the cyclicity of base plus bonus pay is less than that of base earnings.

Figure 9: Probability of Wage Changes by Firm Growth Status



PANEL A: $\text{PR}\{\text{WAGE INCREASE}\}$

PANEL B: $\text{PR}\{\text{WAGE DECREASE}\}$

Note: Figure plots Locally-Weighted Scatterplot Smoothing (LOWESS) estimates of the relationship between the probability that a worker receives a wage increase (Panel A) or decrease (Panel B) and the firm’s backward-looking year-over-year growth rate. The black line plots the patterns for the recession period of May 2009 through December 2010, while the gray line shows the conditional expectation function for the recover period, defined as January 2012 through December 2016. Our firm sample is used to construct this figure, and weighted to match the BDS’ firm size \times industry mix. Firm size is calculated *after* firm growth is taken into account.

All firm-level changes are at the annual level. As a result, the x-axis is the observed change in firm employment between months $t - 12$ and t , while the y-axis is the share of the firm’s workers which had their wages adjusted (up or down) over the same window. Each panel has two lines, reflecting LOWESS smoothed regression lines. The darker line restricts our firm samples only to firm behavior during the Great Recession (2009-2010) while the lighter line restricts the sample to firm behavior during the recovery (2012-2016).

While not overly surprising, firms that grew were slightly more likely than firms that shrank to increase their workers wages during the 2012-2016. Specifically, firms that grew by 10 percent between 2012 and 2016 increased about 73 percent of their workers wages while firms that contracted by 10 percent during that period increased only about 68 percent of their workers wages. What is more surprising, however, is that during the Great Recession there was no systematic relationship between firm size growth and the propensity to rise workers wages. Across all firms, the propensity to raise wages was systematically lower during the Great Recession regardless of firm size growth.

During the Great Recession, however, contracting firms were much more likely to cut the nominal wages of their workers compared to firms that were stagnant or growing. Firms, whose employment fell by 10 percent over the year during the Great Recession, cut roughly 8 percent of their worker’s wages. However, growing firms cut only 4 percent of their worker’s

wages. The propensity to cut wages again appears highly state dependent. The right panel of Figure 9 also shows another interesting fact. During non-recession times, there is no systematic relationship between firm size growth and the propensity to cut a worker’s wage. Contracting firms in 2012-2016 were no more likely to cut nominal wages than growing firms. Moreover, growing firms during the Great Recession were much more likely to cut the nominal wages of their workers than growing firms during the recovery. This pattern would be consistent with aggregate conditions during the Great Recession being such that it is easier for any firm to cut the nominal wage of their workers when many other firms are also doing so.

The evidence presented in this section shows an important interaction between idiosyncratic and aggregate conditions for determining on-the-job wage adjustment patterns. This suggests that the value of workers’ outside options are important for realized wage rigidity, a point which has been raised in, for instance, Christiano et al. (2015). The evidence here supports the hypothesis of state dependence in wage setting, which yields procyclical downward adjustment, and countercyclical upward adjustment. Further research is required to assess the impact of idiosyncratic firm shocks on aggregate wage movements over the cycle.

11 Discussion

Collectively, the results in the paper provide a set of high quality statistics on the wage change distribution for job-stayers, job-changers, and for the aggregate economy. We end the paper by highlighting how our estimates contrast with many of the existing estimates of nominal wage rigidity within the literature. Our results differ in two important ways. First, our administrative data minimizes the impact of measurement error. Second, we separately compute measures of a worker’s base wage and a broader wage measure inclusive of bonuses.³⁷

Table 12 summarizes many of our key findings and the compares them to other prominent estimates of nominal wage adjustments in the literature. We separate the table into research estimating nominal wage adjustments for job-stayers and research estimating nominal wage adjustments for the economy as a whole (inclusive of both job-stayers and job-changers). Much of the literature focuses on measures of earnings per hour that incorporates base wage compensation, bonuses, and other earnings such as overtime premiums.

Using the panel component of the CPS, Daly and Hobijn (2014) report that roughly 85

³⁷There is a long literature, surveyed by Bewley (2004) and Howitt (2002), examining the root causes of nominal wage rigidity. In a series of interviews with business managers responsible for compensation policy, studies have documented that the primary resistance to wage cuts arises from concerns over damaging worker morale. See, for instance, Kaufman (1984), Blinder and Choi (1990), Agell and Lundborg (1995, 1999), Campbell III and Kamalani (1997), and Bewley (1999).

percent of wages of job-stayers (measured by earnings per hour) change annually during a sample period that overlaps with ours.³⁸ As noted above, we find that only about two-thirds of job-stayers receive an annual nominal base wage change during the 2008-2016 period. It is difficult to pinpoint why the CPS results are different than ours. First, it is well documented that there is a large amount of measurement error in both earnings and hours in household survey.³⁹ The fact that measurement error in earnings and hours is high in household surveys can explain the higher variance of wage changes in the CPS. Second, it is unclear in household surveys whether workers are reporting only their base earnings or a broader measure of earnings inclusive of bonuses. This ambiguity makes it harder to make comparisons between household surveys and administrative data sets like ADP.

Using data from the Survey of Income and Program Participation, Barattieri et al. (2014) try to account for the measurement error in wages and hours in household data by looking for structural breaks in their individual hourly wage series. Their primary focus is on workers who are paid hourly. Given that they focus on an individual's self-reported hourly wage (reported in dollars per hour), their wage measure is similar in concept to our measure of a base wage for hourly workers. When they make their correction for measurement error, they find that the frequency of quarterly wage changes for job-stayers falls from over 50 percent to between 15 and 20 percent - depending on their adjustment procedure - during their period of study. Our quarterly frequency of base wage changes for job-stayers who are paid hourly is 20 percent which is at the upper range of their estimates. More importantly, however, they estimate a much larger fraction of downward wage adjustments. Specifically, they find that 12 percent of all quarterly wage changes for job-stayers are downward changes. That is three and half times larger than our administrative data reports. As Barattieri et al. (2014) highlight, there is substantial measurement error in household surveys with respect to measuring how nominal wages adjust. The fact that their patterns still differ relative to the ADP results is consistent with some residual measurement error remaining even after implementing their structural break procedure.⁴⁰

³⁸One of the earliest papers to estimate the extent of nominal wage rigidity using household level data was Kahn (1997), who used data from the Panel Study of Income Dynamics (PSID) to find that about 92% of workers receive a nominal wage change during a given year.

³⁹There is a literature documenting sizable measurement error in both income and hours in household surveys using audit studies. Bound and Krueger (1991) finds that just about forty percent of cross-sectional variance of the change in income for men in household surveys can be explained by measurement error. Bound et al. (1989) documents that the measurement error in reported hours in household surveys is even larger than the measurement error in income.

⁴⁰A more recent literature has emerged using firm-level data to measure wage stickiness. Both Lebow et al. (2003) and Fallick et al. (2016) use data from the BLS's Employment Cost Index (ECI) to measure nominal wage rigidity. Unlike the household surveys or other administrative payroll data, the unit of analysis in the ECI is a job not a worker. To the extent that workers who populate a specific job are heterogeneous with respect to underlying skills, nominal wage variation could occur due to shifting sampling of different quality

Table 12: Comparison with Literature

Paper	Data Source	Concept	Period	$\Pr\{\Delta w \neq 0\}$	$\Pr\{\Delta w < 0\}$
<u>Job-Stayer Estimates</u> This Paper This Paper	ADP ADP	Base Wage (Stayers) Base + Bonus (Stayers)	2008-2016 2008-2016	66.3% 90.7%	2.4% 15.7%
Daly and Hobijn (2014) Barattieri et al. (2014)	CPS SIPP (Corrected)	Average Hourly Earnings Base Hourly Earnings	2010-2011 1996-1999	84% 62.2%	– 7.7%
Kurmann and McEntarfer (2018) Jardim et al. (2019)	Washington State Washington State	Annual Earnings/Hour Annual Earnings/Hour	2005-2006 2005-2015	92.5% 94.0%	20% 20.6%
Altonji and Devereux (2000) Altonji and Devereux (2000)	Payroll, One Firm Payroll, One Firm	Base Wage, Salaried Base + Bonus, Salaried	1996-1997 1996-1997	88.6% –	0.5% 8.69%
<u>Aggregate Estimates</u> This Paper This Paper	ADP ADP	Base Wage (Aggregate) Base + Bonus (Aggregate)	2008-2016 2008-2016	71.3% 90.1%	8.5% 21.4%
Christiano et al. (2005) Smets and Wouters (2007) Linde et al. (2016) Christiano et al. (2014)	Model-Implied Model-Implied Model-Implied Model-Implied	Average Hourly Earnings Average Hourly Earnings Average Hourly Earnings Average Hourly Earnings	1965-1995 1966-2004 1966-2007 1985-2010	83.2% 71.6% 61.1% 57.0%	– – – –
Beraja et al. (2016)	Regional ACS	Average Hourly Earnings	2007-2011	76.0%	–

Notes: Table reports estimates of the frequency of price changes from our ADP data as well as from existing papers in the literature. All probabilities are expressed in annual terms. When papers report the quarterly probability of wage change, q , we compute the annual probability by assuming a constant hazard of wage change, so that the annual probability of change is given by $1 - (1 - q)^4$. Barattieri et al. (2014) report the quarterly probability of wage declines, which we aggregate to the annual level by assuming a constant share of wage changes are negative.

Administrative data sets containing measures of both earnings and hours are exceedingly rare in the United States. As discussed above, contemporaneously, Kurmann and McEntarfer (2018) and Jardim et al. (2019) use data from the US Longitudinal Employer Household Dynamics (LEHD) and Washington State Unemployment Insurance Records to examine nominal earnings-per-hour adjustments for a sample of job-stayers who reside in Washington state over a two year period. They focus their sample on residents of Washington state because Washington requires employers to report the hours worked of their employees as part of their Unemployment Insurance program. Their findings with high quality administrative data from Washington are consistent with our patterns for modified annual earnings (which includes both base and bonus compensation). While such measures are very informative about earnings per hour volatility, they are not well-suited to consider whether fluctuations are driven by base wage changes or other sources like bonuses and overtime premiums. As discussed above, understanding the movement of base wages versus bonuses may be important for disciplining models of aggregate fluctuations in the same way that understanding the distinction between reference prices and sales is important in the pricing literature.

Our paper is closest in spirit to the results in Altonji and Devereux (2000) and Fehr and Goette (2005). Altonji and Devereux (2000) uses administrative payroll data similar to ours for one large financial service company during 1996 and 1997 while Fehr and Goette (2005) uses administrative payroll data from two Swiss firms in the early 1990s. The patterns of base (administrative) wage adjustment for job-stayers these authors document within their selected companies closely match the patterns we document for the whole U.S. economy during the 2008-2016 period. They also find that bonuses allow for more downward flexibility in earnings per hour. That firm-level payroll data from an earlier period within the U.S. and Switzerland broadly matches the job-stayer results from the recent ADP data highlights the importance of using administrative payroll data to measure base wage stickiness.

Most of the above results focus on the nominal wage adjustments for job-stayers. There is a separate literature that combines aggregate time series data and medium scale DSGE models to infer the extent of aggregate nominal wage stickiness. In this literature, nominal wage adjustments are generally modeled as being Calvo. We highlight four such papers in Table 12. As can be seen from the table, estimates of the frequency of annual nominal wage adjustments vary markedly across papers across papers. Christiano et al. (2005) uses time series data and their model to infer that 83% of wages adjust annually. However, Christiano et al. (2014) infers that only 57% of wages adjust annually. Like with the pricing literature, estimates of nominal wage stickiness are not well disciplined in time series data. This is a

workers over time. Consistent with this fact, the nominal wage variation in the ECI for a given job is much larger than what we document in the payroll data for job-stayers.

point that is highlighted in Beraja et al. (2016). Using regional variation in annual earnings per hour using data from the American Community Survey (ACS), Beraja et al. (2016) find that nominal wages are more flexible than implied by recent time series estimates from medium scale DSGE models. They conclude that nominal wage stickiness cannot explain why aggregate nominal wages did not fall more sharply during the Great Recession. The estimates of annual nominal wage adjustments using cross-region variation in Beraja et al. (2016) are similar to our aggregate annual base wage adjustment found using ADP data.

12 Conclusion

This paper measures nominal wage adjustments for millions of US workers from 2008 to 2016 using administrative payroll records. Measurement error in household surveys and missing hours in administrative earnings data have prevented the development of a reliable set of moments for the distribution of worker compensation adjustment. Exploiting payroll data from the largest payroll processing company in the US allows us to measure various forms of worker compensation, including the per period contract wage, without error.

Conceptually, our paper highlights the importance of thinking about the nature of compensation when using various moments to discipline models of nominal wage adjustments. For most workers, essentially all of their compensation comes from their base (contract) wage. However, for some workers, bonuses are also quantitatively important. We provide evidence suggesting that the extent to which base wages adjust is the appropriate moment to discipline models where the present value of worker earnings determine labor market fluctuations. Base wage changes are highly persistent and co-vary strongly with both the aggregate and regional business cycles. Bonuses, on the other hand, display very little persistence and are much less cyclical. This is consistent with the large literature in personnel economics where bonuses are used to incentivize worker effort. We conclude that bonuses are more akin to sales in the output price setting literature and should be ignored when calibrating many models of business cycle dynamics.

However, bonuses are important for providing flexibility in workers spot wages. Over our whole sample period, approximately 2 percent of workers who remain consistently employed with the same firm receive a base wage reduction during a given year. For a broader measure of the wage inclusive of bonuses, we document that nearly 16 percent of job-stayers receive a nominal wage decline during a given year. Such measures of wage adjustment inclusive of bonuses will be useful for models of incomplete markets, in which temporary shocks have allocative effects, or for models seeking to understand individual earnings dynamics. The differential adjustment patterns of various compensation forms urges careful theoretical

consideration over the correct notion of “wage rigidity”. The evidence presented in this paper suggests that workers’ *contracts*, which specify base wages and a bonus schedule, may be quite rigid, particularly on the downside. However, workers’ *realized compensation*, which includes both base pay and bonuses, may be quite flexible, as stipulated by the (explicit or implicit) contract governing their employment relationship.

We also provide a series of statistics on the wage adjustments of job-changers. Approximately 40 percent of job-changers receive a nominal base wage decline within a year. Since wage adjustment in the aggregate is a combination of the wage movements of both job-stayers and job-changers, this margin is key to determining the role that nominal wage rigidity plays in explaining aggregate wage dynamics. Additionally, we show that the wages of new hires into the firm are about as cyclical as similar job-stayers. Depending on the model, the appropriate empirical moment may be the wage adjustments of job-stayers, job-changers, new hires, or a combination of all the different groups.

Finally, we show that worker wage adjustments are state dependent. During some months of the Great Recession, 6 percent of job-stayers on average received a base wage cut. In some industries, upwards of 10 percent of job-stayers received a base wage cut. Worker wage growth is procyclical for job-stayers. The state dependence is also shown using firm-level employment growth variation with shrinking firms being much more likely to cut nominal wages relative to growing firms, particularly during the Great Recession. Moreover, during the Great Recession, growing firms were much more likely to cut the wages of their workers than similarly growing firms during the 2012-2016 recovery.

The goal of the paper is to provide a set of moments that can be used to both guide and discipline models of nominal wage adjustments. The nominal wage flexibility we estimate using detailed micro data is broadly consistent with the flexibility of wages estimated in Beraja et al. (2016) using cross-region variation. Beraja et al. (2016) show that wage adjustments of the sort we identified cannot explain why aggregate wage growth remained robust during the Great Recession. However, one drawback of the model put forth by Beraja et al. (2016) is that it has no role for the type of asymmetry in nominal wage adjustments that this paper highlights. Going forward, research should explore whether the asymmetry in nominal wage adjustments substantially alters the conclusions of models with wage rigidity. Our state dependence results also suggest reason to move away from Calvo models of wage adjustment, which are prevalent in many modern New Keynesian models, and instead develop menu cost models of wage adjustment. Finally, our results highlight that different types of compensation have different degrees of flexibility and the compensation composition differs markedly across workers. Understanding what drives such differences and whether or not such differences have macro implications is a fruitful area for future work.

References

- Agell, Jonas and Per Lundborg**, “Theories of Pay and Unemployment: Survey Evidence from Swedish Manufacturing Firms,” *The Scandinavian Journal of Economics*, 1995, 97 (2), 295–307.
- and –, “Survey evidence on wage rigidity and unemployment: Sweden in the 1990s,” 1999.
- Altonji, Joseph G. and Paul J. Devereux**, “The Extent and Consequences of Downward Nominal Wage Rigidity,” in Solomon Polachek and Konstantinos Tatsiramos, eds., *Research in Labor Economics Volume 19*, Emerald Group Publishing Limited, 2000, pp. 383–431.
- Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk**, “Some Evidence on the Importance of Sticky Wages,” *American Economic Journal: Macroeconomics*, 2014, 6 (1), 70–101.
- Basu, Sustanto and Chris House**, “Allocative and Remitted Wages: New Facts and Challenges for Keynesian Models,” in John B. Taylor and Harald Uhlig, eds., *John B. Taylor and Harald Uhlig, eds.*, Vol. 2 of *Handbook of Macroeconomics*, Elsevier, 2016, pp. 297 – 354.
- Beraja, Martin, Erik Hurst, and Juan Ospina**, “The Aggregate Implications of Regional Business Cycles,” Working Paper 21956, National Bureau of Economic Research February 2016.
- Bewley, Truman**, “Fairness, Reciprocity, and Wage Rigidity,” 2004.
- Bewley, Truman F.**, *Why Wages Don't Fall During a Recession*, Cambridge, MA: Harvard University Press, 1999.
- Bils, Mark and Pete Klenow**, “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 2004, 112 (5), 947–985.
- Bils, Mark J.**, “Real Wages over the Business Cycle: Evidence from Panel Data,” *Journal of Political Economy*, 1985, 93 (41), 666–689.
- Blinder, Alan S. and Don H. Choi**, “A Shred of Evidence on Theories of Wage Stickiness Author,” *The Quarterly Journal of Economics*, 1990, 105 (4), 1003–1015.
- Bound, John and Alan B. Krueger**, “The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?,” *Journal of Labor Economics*, 1991, 9 (1), 1–24.
- , **Charles Brown, Greg J. Duncan, and Willard L. Rodgers**, “Measurement Error In Cross-Sectional and Longitudinal Labor Market Surveys: Results From Two Validation Studies,” NBER Working Papers 2884, National Bureau of Economic Research, Inc March 1989.

- Brown, Charles and James Medoff**, “The Employer Size Wage Effect,” 1989.
- Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin**, “Wage Bargaining with On-the-Job Search: Theory and Evidence Author,” *Econometrica*, 2006, *74* (2), 323–364.
- Calvo, Guillermo**, “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 1983, *12*, 383–398.
- Campbell III, Carl M. and Kunal S. Kamalani**, “The Reasons for Wage Rigidity: Evidence from a Survey of Firms,” *The Quarterly Journal of Economics*, 1997, *112* (3), 759–789.
- Card, David and Dean Hyslop**, “Does Inflation ”Grease the Wheels of the Labor Market?,” 1997.
- Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans**, “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 2005, *113* (1), 1–45.
- , **Martin S Eichenbaum, and Mathias Trabandt**, “Unemployment and Business Cycles,” 2015.
- Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno**, “Risk Shocks,” *American Economic Review*, January 2014, *104* (1), 27–65.
- Daly, Mary C and Bart Hobijn**, “Downward Nominal Wage Rigidities Bend the Phillips Curve,” 2014.
- Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo**, “Reference Prices, Costs, and Nominal Rigidities,” *American Economic Review*, February 2011, *101* (1), 234–62.
- Fallick, Bruce C, Michael Lettau, and William L Wascher**, “Downward Nominal Wage Rigidity in the United States During and After the Great Recession,” 2016.
- Fehr, Ernst and Lorenz Goette**, “Robustness and real consequences of nominal wage rigidity,” *Journal of Monetary Economics*, 2005, *52* (4), 779 – 804.
- Gertler, Mark, Christopher Huckfeldt, and Antonella Trigari**, “Unemployment Fluctuations, Match Quality, and the Wage Cyclicity of New Hires,” 2016.
- Güvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song**, “What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?,” 2015.
- , **Serdar Ozkan, and Jae Song**, “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 2014, *122* (3), 621–660.
- Haefke, Christian, Marcus Sonntag, and Thijs van Rens**, “Wage Rigidity and Job Creation,” *Journal of Monetary Economics*, 2013, *60*, 887–899.
- Hazell, Jonathon and Bledi Taska**, “Posted Wage Rigidity,” *MIT Working Paper*, 2018.

- Howitt, Peter**, “Looking Inside the Labor Market: A Review Article,” *Journal of Economic Literature*, 2002, 40 (March), 125–138.
- Hyatt, Henry R. and James R. Spletzer**, “The Shifting Job Tenure Distribution,” 2016.
- Jardim, Ekaterina S., Gary Solon, and Jacob L. Vigdor**, “How Prevalent Is Downward Rigidity in Nominal Wages? Evidence from Payroll Records in Washington State,” 2019.
- Kahn, Shulamit**, “Evidence of Nominal Wage Stickiness from Microdata,” *The American Economic Review*, 1997, 87 (5), 993–1008.
- Kaplan, Greg and Giovanni L. Violante**, “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 2014, 82 (4), 1199–1239.
- Kaufman, Roger T.**, *British Journal of Industrial Relations*, 1984, 22, 101 – 112.
- Kehoe, Patrick J. and Virgiliu Midrigan**, “Temporary Price Changes and the Real Effects of Monetary Policy,” 2008.
- Klenow, Peter J. and Oleksiy Kryvtsov**, “State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?,” *The Quarterly Journal of Economics*, 2008, 123 (3), 863–904.
- Kudlyak, Marianna**, “The cyclical of the user cost of labor,” *Journal of Monetary Economics*, 2014, 68, 53–67.
- Kurmann, Andre and Erika McEntarfer**, “Downward Wage Rigidity in the United States: New Evidence from Administrative Data,” 2018.
- Lebow, David E, Raven E Saks, and Beth Anne Wilson**, “Downward Nominal Wage Rigidity: Evidence from the Employment Cost Index,” *Advances in Macroeconomics*, 2003, 3 (1).
- Linde, Jesper, Frank Smets, and Rafael Wouters**, “Challenges for Macro Models Used at Central Banks,” Working Paper, Sveriges Riksbank Research Paper Series 2016.
- Makridis, Christos and Maury Gittleman**, “On the Cyclical of Real Wages and Employment: New Evidence and Stylized Facts from Performance Pay and Fixed Wage Jobs,” *Working Paper*, 2018.
- Martins, Pedro S., Gary Solon, and Jonathan P. Thomas**, “Measuring What Employers Do about Entry Wages over the Business Cycle: A New Approach,” *American Economic Journal: Macroeconomics*, 2012, 4 (4), 36–55.
- Menzio, Guido and Shouyong Shi**, “Efficient Search on the Job and the Business Cycle,” *Journal of Political Economy*, 2011, 119 (3), 468–510.
- Midrigan, Virgiliu**, “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations,” *Econometrica*, 2011, 79 (4), 1139–1180.

- Nakamura, Emi and Jón Steinsson**, “Five facts about prices: a reevaluation of menu cost models,” *The Quarterly Journal of Economics*, 2008, *123* (4), 1415–1464.
- Pissarides, Christopher A.**, “The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?,” *Econometrica*, 2009, *77* (5), 1339–1369.
- Postel-Vinay, Fabien and Jean-Marc Robin**, “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, 2002, *70* (6), 2295–2350.
- Schmitt-Grohé, Stephanie and Martin Uribe**, “The Making of a Great Contraction with a Liquidity Trap and a Jobless Recovery,” 2012.
- Schoefer, Benjamin**, “The Financial Channel of Wage Rigidity,” *Working Paper*, 2016.
- Shin, Donggyun and Gary Solon**, “New Evidence on Real Wage Cyclicity Within Employer-Employee Matches,” *Scottish Journal of Political Economy*, 2007, *54* (5), 648–660.
- Smets, Frank and Rafael Wouters**, “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *The American Economic Review*, 2007, *97* (3), 586–606.
- Solon, Gary, Robert Barsky, and Jonathan Parker**, “Measuring the Cyclicity of Real Wages: How Important is Composition Bias,” *Quarterly Journal of Economics*, 1994, *109* (1), 1–25.
- , **Warren Whatley, and Ann Huff Stevens**, “Wage Changes and Intrafirm Job Mobility over the Business Cycle: Two Case Studies,” *ILR Review*, 1997, *50* (3), 402–415.
- Taylor, John B.**, “Staggered Wage Setting in a Macro Model,” *American Economic Review*, 1979, *69* (2), 108–113.
- , “Aggregate Dynamics and Staggered Contracts,” *Journal of Political Economy*, 1980, *88* (1), 1–23.

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

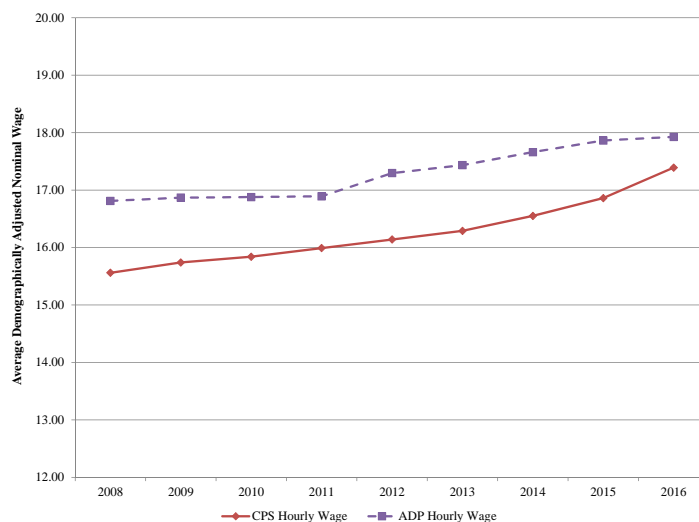
Table A1: Statistics for Employee Sample, Selected Years

	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.

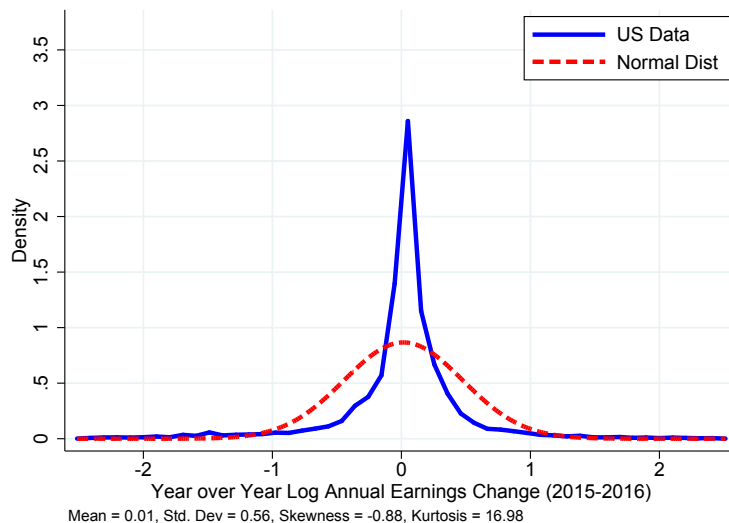
The central challenge in benchmarking our data to annual earnings records is that the ADP data do not follow a worker if they move to employers who are not clients of ADP. This can generate large swings, both positive and negative, in annual earnings, which are not observed in datasets with the universe 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).⁴¹ We interpret these small differences to be the result of

⁴¹Varying 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

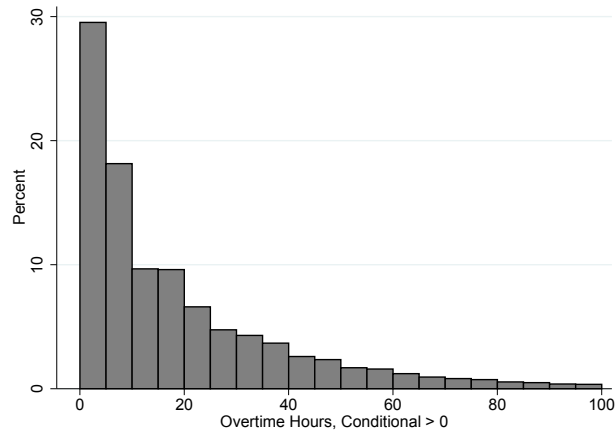
winsorize wage rates below the federal minimum wage for service workers who receive tips. Some of these individuals may be unpaid interns who receive, for instance, transportation benefits from their employer. Third, we drop employees whose status codes indicate that their employment has been terminated. 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.⁴² 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.

⁴²When 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

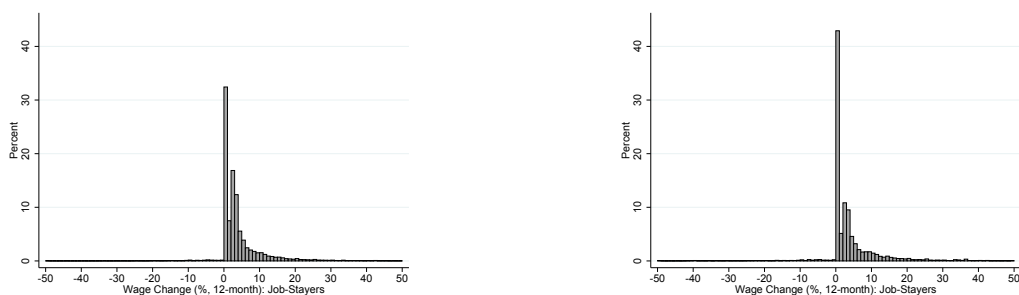
Appendix C.1 Similarity in Patterns across Compensation Arrangements

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.

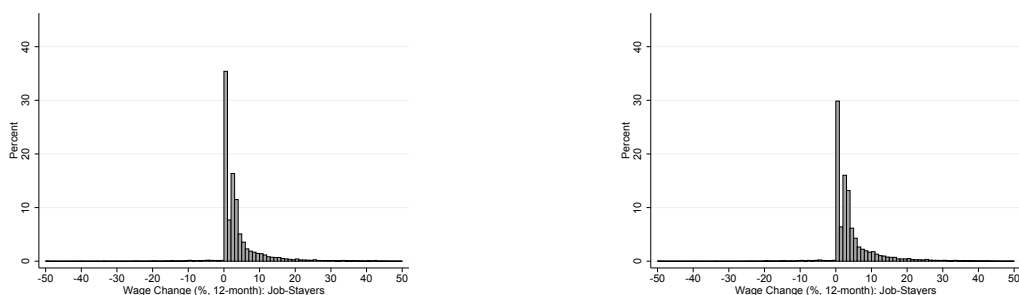
Figure A5 explores the extent to which nominal base wages are allocative. Specifically, we focus on our sample of hourly workers whose monthly hours worked fluctuates over the year. The number of pay weeks in the month varies over time, so we adjust our monthly hours for the number of pay periods making 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



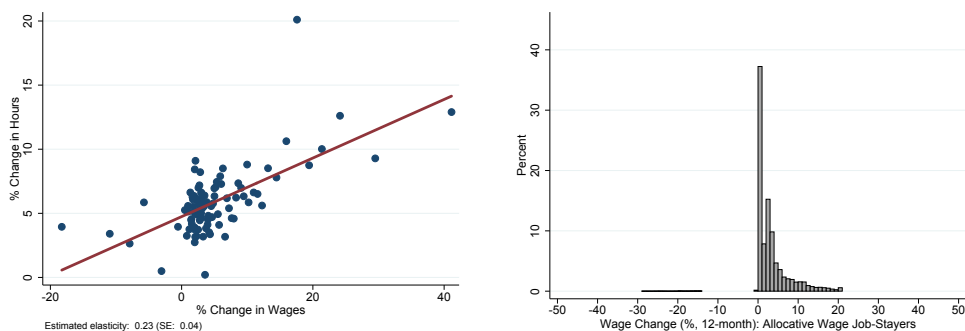
PANEL A: NON-COMMISSION IN YEAR $t - 1$ PANEL B: WITH COMMISSION IN YEAR $t - 1$



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.

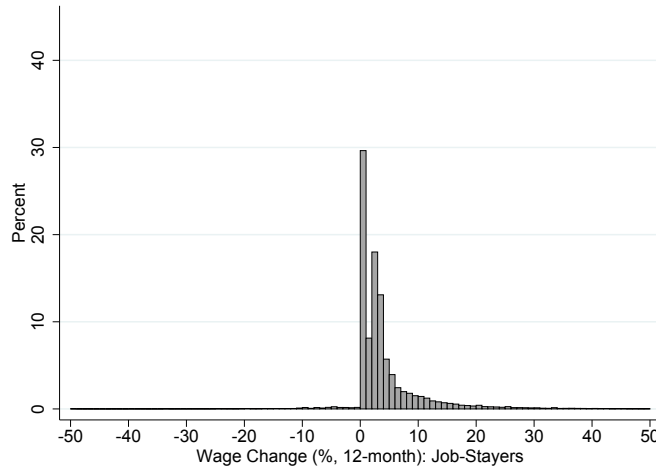
Appendix C.3 Robustness to weighting and sampling

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



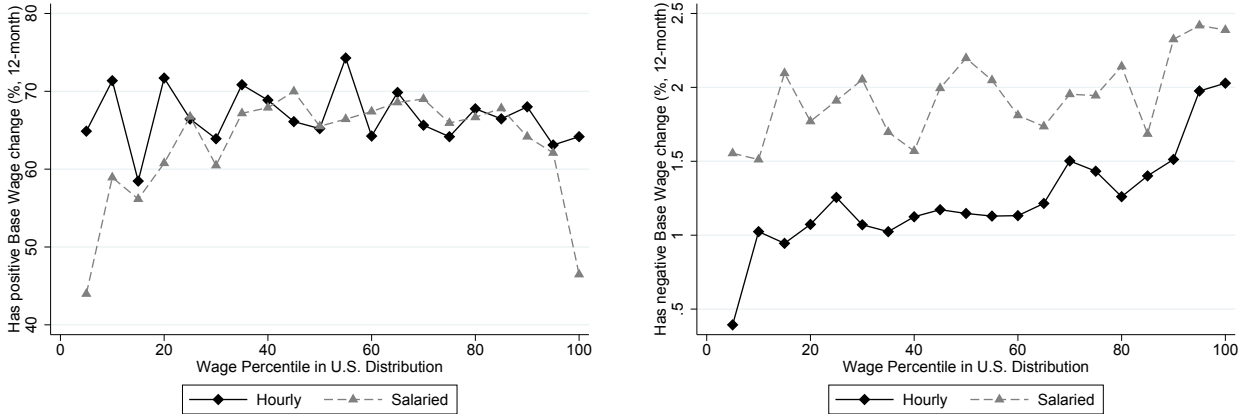
Note: Figure plots the unweighted distribution of 12-month base wage changes for job-stayers.

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.

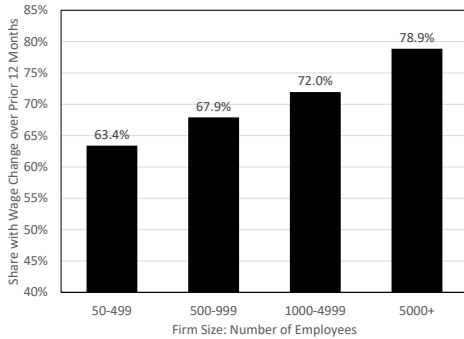
Figure A8 also shows that there is some degree of heterogeneity across industries with respect to base wage changes. For example, both hourly and salaried workers in the manufacturing industry are much more likely to receive a base wage change than workers in construction during our sample period. This 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 2014-2016.⁴³ 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.

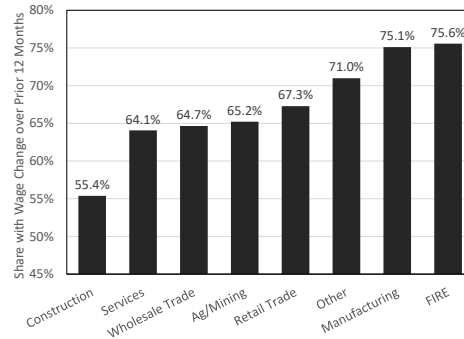
How much can the exclusion of small firms from our main analysis bias our results? The BDS shows that 27.1% of workers were employed in small firms in 2016. We can use the results in Appendix Figure A9 to compute a new measure of the probability of nominal base wage adjustments inclusive of workers in 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

⁴³We 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.

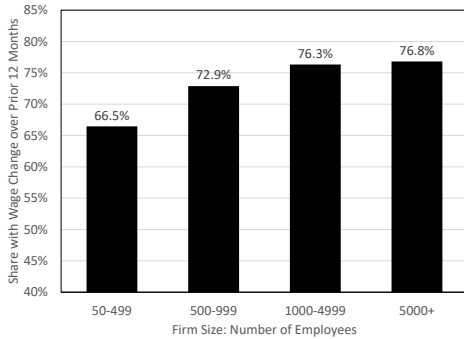
Figure A8: Share with Base Wage Change by Firm Size and Industry, All Years



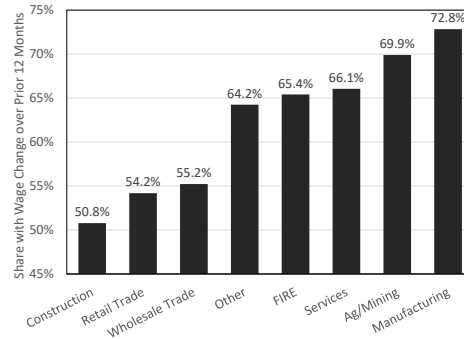
PANEL A: HOURLY WORKERS BY SIZE



PANEL B: HOURLY WORKERS BY INDUSTRY



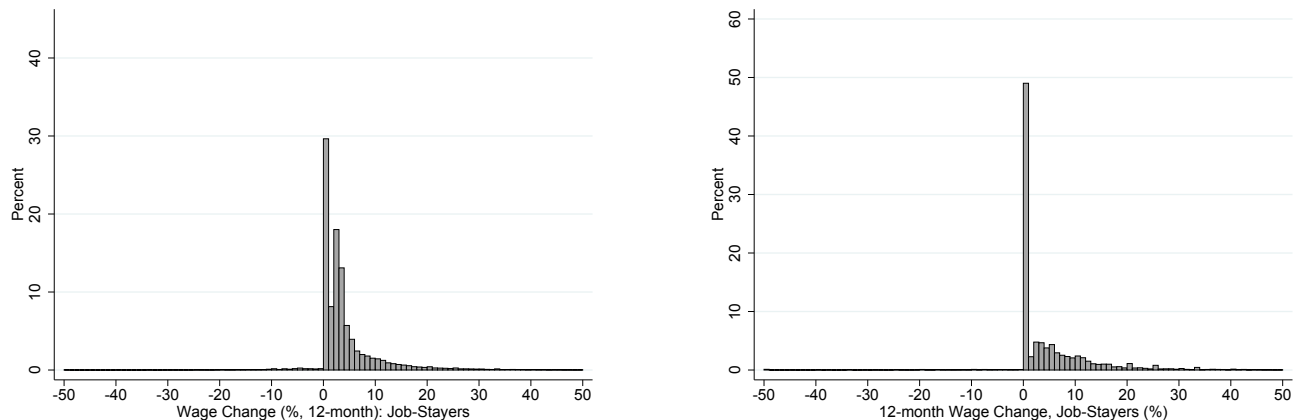
PANEL C: SALARIED WORKERS BY SIZE



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: ≥ 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.

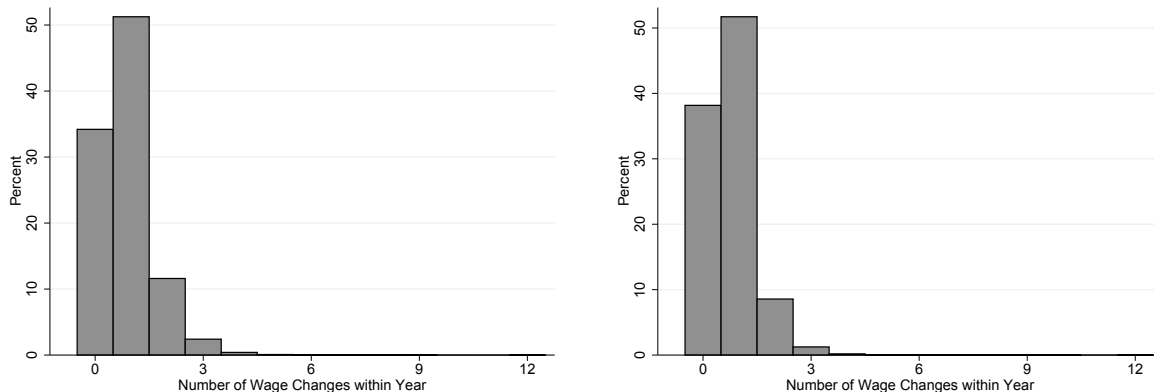
very close to those reported in the main text including only workers in firms with more than 50 employees (66.3% vs. 62.2% and 2.4% vs. 2.3%). To summarize, we conclude that the omission of workers in firms with less than 50 employees is not biasing our results substantively.

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

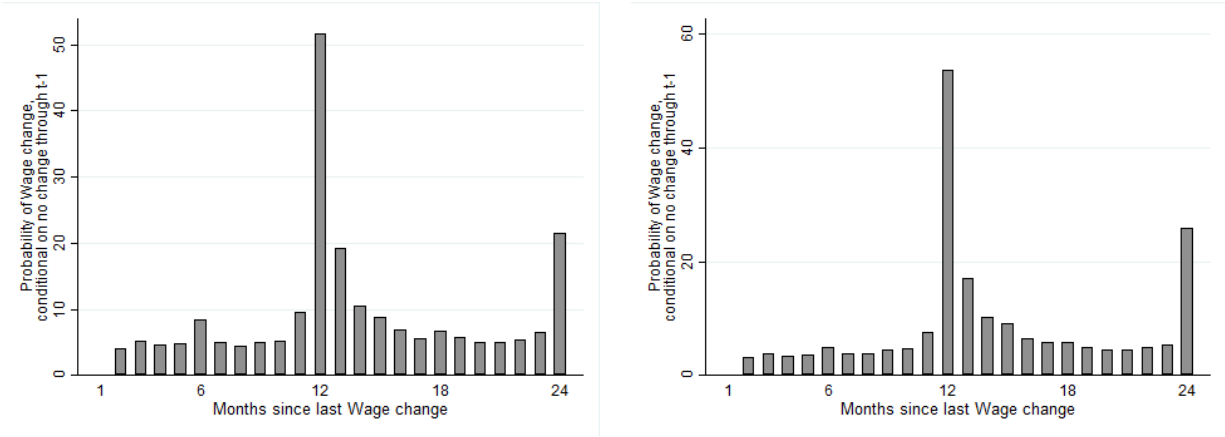
Note: Table shows the average number of nominal base wage changes for hourly workers (left panel) and salaried workers (right panel). We use our employee sample for this analysis and restrict our sample to those workers who remain continuously employed with the same firm during a 12 month calendar year.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

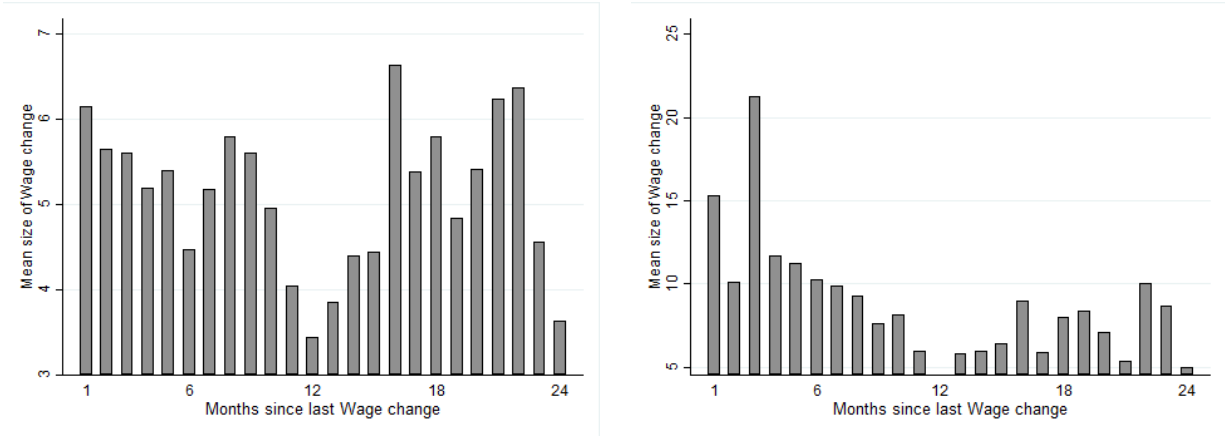
Note: Figure shows the hazard rate of a base wage change between $t - 1$ and t conditional on surviving to t without a base wage change at the same firm. Sample only includes individuals with at least two base wage changes. We use all data between 2008 and 2016 for this analysis, and weight the data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

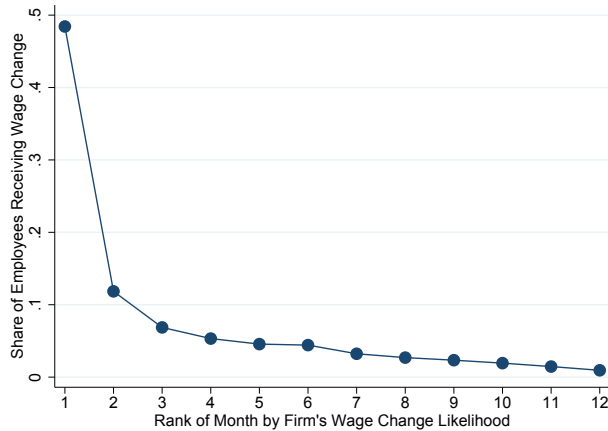
if a worker is promoted internally and the promotions are distributed throughout the year, it is not surprising that workers who receive a base wage change off cycle also get larger base wage changes.

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.⁴⁴

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,

⁴⁴We 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.⁴⁵ 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.

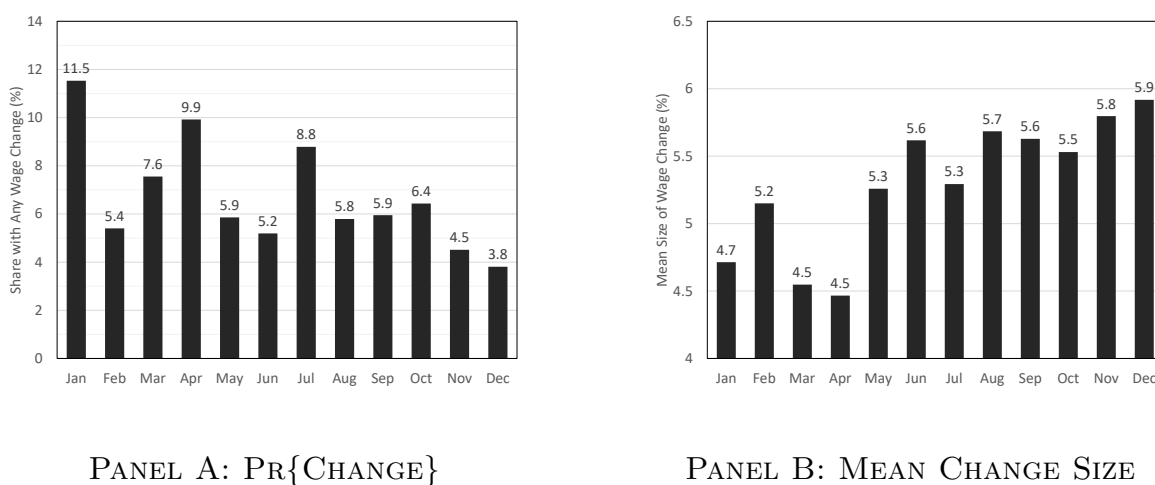
While the Calvo predictions may be rejected at the individual and firm level, Calvo may still be a good approximation for the aggregate macro economy if firms stagger the months in which they adjust wages. Indeed, this is the underlying intuition behind the staggered wage contract model. Instead of each 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-

⁴⁵This 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.

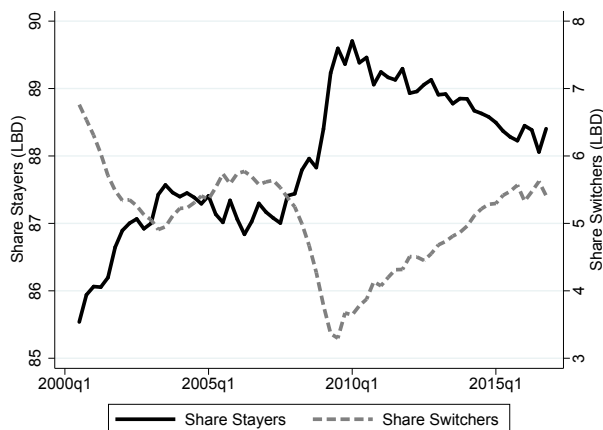
Figure A14: Seasonality in Base Wage Changes, Job-Stayers, All Years



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.