

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

John Grigsby Erik Hurst Ahu Yildirmaz

Preliminary and Evolving

July 2, 2018

Abstract

Using administrative payroll data from the largest U.S. payroll processing company, we document a series of new facts about the extent of nominal wage adjustments in the U.S.. First, we document that nominal wage cuts are exceedingly rare for job-stayers. Over the pooled 2008-2016 period, only 2% of workers who remain in a continuous employment relationships receive a nominal wage cut during a given year. Second, nominal wages are much more flexible (both up and down) for job changers. Aggregating job changers and job stayers shows that approximately 10% of workers receive a wage cut during a year. Third, the extent of wage rigidity is state dependent. Nominal wage adjustments are lower during recessions, in regions that suffered larger house price declines, and for firms that shed large amounts of workers. Moreover, nominal wage cuts are substantially higher during recessions. During the Great Recession, nearly 12 percent of workers received a nominal wage cut. Wage declines during the recession were particularly concentrated in areas with large house price declines and in shrinking firms. We end by discussing how measurement error in household level data and missing measures of hours in firm level data substantially bias existing measures of nominal wage rigidity.

*We thank Mark Aguiar, Steve Davis, John Haltiwanger, Alan Krueger, Marianna Kudlyak, John Shea, Rob Shimer, and Joe Vavra as well as seminar participants at the 2018 American Economic Association meetings, Chicago, the European Central Bank, Maryland, the San Francisco Federal Reserve, and USC for helpful comments. Authors contact information: jgrigsby@uchicago.edu, erik.hurst@chicagobooth.edu. and Ahu.Yildirmaz@ADP.com.

1 Introduction

The degree to which wages are flexible determines the extent to which economic shocks generate macroeconomic fluctuations in employment and earnings. In addition, modern New Keynesian models often show that the nominal wage rigidities dictate the aggregate response of real variables to monetary policy shocks. However, compared to the literature on price changes, there is only a relatively small literature which uses micro data to measure the extent to which nominal wages adjust. That different questions necessitate different measures of rigidity complicates this measurement exercise. Studies seeking to quantify the effect of downward nominal wage rigidity on employment flows will be concerned with within-job flexibility for job stayers, while studies of the aggregate response to monetary shocks require an aggregate measure of rigidity which accounts for the wage movements of both job-stayers and job-changers. Furthermore, most micro data sets are not well suited to measure nominal wage changes. 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 the self-reported hourly wages can result in a substantial upward bias in the volatility of individual level wage changes. Administrative datasets, on the other hand, have high quality panel data on quarterly or annual earnings but usually do not have measures of individual hours worked.¹ The lack of data on hours worked makes administrative datasets less than ideal when it comes to studying nominal wage fluctuations.

In this paper, we use administrative data from ADP – one of the world’s largest payroll processing companies – to document a series of facts about aggregate nominal wage adjustments for millions of U.S. workers during the last decade. Hundreds of thousands of firms per year contract with ADP to administer a variety human resource and payroll activities. Most of ADP’s clients use their payroll processing services, with ADP currently processing payroll checks for roughly 20 million workers each year, one-seventh of the U.S. workforce. We have access to data from 2008 through 2016. For workers paid hourly, the data records administrative measures of their hourly wage. For salaried workers, we observe administrative measures of the employees contracted earnings per pay period. Some salaried workers are paid weekly while others are paid bi-weekly or monthly. The administrative data records their contractually obligated per-period pay rate. We also have detailed administrative data on a workers gross monthly earnings (i.e., the sum of their monthly paychecks).

¹When individual hours worked are collected, they are often reported not by the individual but instead by a payroll administrator. For example, Washington is one state that uses hours worked in the prior year to determine UI eligibility. As we highlight later in the paper, such administrative reports of hours worked for salaried workers are likely measured with sufficient error to lead to overstated nominal wage adjustments if the measure of the nominal wage is earnings per hour.

Note, that monthly gross earnings will differ from a worker’s per-period contract rate. A worker’s monthly earnings will account for hours worked during the month (for hourly workers) and the number of pay periods per month (for salaried workers). Additionally, during the month the worker could receive overtime compensation, tips, commissions, meal and travel reimbursements, pay advances, signing bonuses, severance payments, cashed-out vacation time, quarterly performance pay, and annual bonuses which is accrued above and beyond a worker’s base pay.

This paper makes three principal contributions. First, the richness of the ADP data allows us to present new facts on the forms of compensation for American workers, and find that the overwhelming majority of earnings are in “base pay.” We refer to base pay as the part of earnings directly resulting from the worker’s per-period pay rate. We then define a worker’s nominal wage as their per-period contract rate. Second, we use our data to present several key facts about nominal wage adjustments for job-stayers, testing the common assumptions of Calvo, time dependent, and state dependent wage setting. We present substantial evidence for both state and time dependence in wage setting for job-stayers.² Finally, we measure the extent of aggregate nominal wage rigidity in the economy for the economy as a whole combining data on both job-stayers and job-changers. Given the fact that the nominal wages of job-changers are quite flexible, We find substantially more nominal rigidity for job-stayers than for the economy as a whole.

We begin the paper documenting how much worker compensation comes from different sources. For most workers, essentially all of their monthly earnings is attributed to their contractually obligated base pay. Focusing only on workers who are continuously employed with the same firm for a 12 month calendar period, the median hourly worker receives 97% of their annual earnings in base pay, while the equivalent number for salaried workers is 95%. Much of the additional income comes in the form of year end bonuses. We estimate that about one-out-of-five hourly workers and one-out-of-three salaried workers receive an annual bonus of at least 1 percent of their gross annual earnings. The average bonus size for hourly and salaried is, respectively about \$1,600 and \$5,400 per year. Given that the bulk of annual compensation for most workers is in base pay and annual bonuses, we focus our attention in the rest of the paper examining how base pay and bonus compensation evolves.

To measure aggregate nominal wage adjustments, we explore how nominal wages (as measured by per-period administrative contract rate) evolves for those workers who remain employed within the same firm. We refer to such workers as “job-stayers”. We then separately explore how nominal wages evolve for those workers who switch jobs. ADP’s large

²Bils and Klenow (1995) and Nakamura and Steinsson (2008) similarly present numerous facts about the nature of output price setting.

coverage of the US workforce allows us to observe workers migrating from one ADP firm to another ADP firm within a relatively short time period. Using data from Census's Job-to-Job flows which measures the amount of job-stayers relative to job-changers in the economy as a whole, we can aggregate our data to make a measure aggregate nominal wage adjustments. We complement our analysis by exploring how bonuses evolve for a given worker over time.

We begin our analysis by describing key facts of the wage setting process for job-stayers. Roughly 20 percent of job-stayers receive a nominal wage change over a quarter and about two-thirds of job-stayers receive a wage change during a year pooling over our full sample period. However, small wage changes of less than 2% are rare, consistent with menu cost models of wage setting. Unlike other studies in the literature, we find that nominal wage cuts for job-stayers are exceedingly rare. For our sample of job-stayers during the 2008-2016 period, only 0.9 percent of workers received a nominal wage cut during a given quarter and only 2.4 percent received a nominal wage cut during a given year. Essentially all wage changes within a given employment relationship are wage increases. This strong downward rigidity on the job may be an important factor for the large employment fluctuations observed during the Great Recession.

Wage adjustment patterns vary systematically by both firm size and industry. Smaller firms are much less likely to adjust wages than larger firms. Additionally, the frequency of wage adjustment is higher in the manufacturing, FIRE and services industries and lower in construction, retail trade and wholesale trade. These differences persist even conditional on a vector of individual and firm level controls.

We then provide evidence supporting the importance of time dependent wage setting, reflecting a model of Taylor (1980) staggered wage contracts at the individual level. In a sample of job-stayers, wage adjustments occur primarily at 12 month intervals. There is a small and constant probability of wage changes between 1 and 10 months and between 14 and 22 months suggesting that at the individual level, some wage changes occur at frequencies other than a year. The nature of these wage changes are different in that they are much larger than the wage changes that occur at annual frequencies. While, for the most part, workers receive a wage change every 12 months, the wage changes at the aggregate level are mostly smoothed out across individuals. Exploring the seasonality in wage changes, we find that wage changes are more common in January, April and July than in other months during the year. However, aggregating to quarters, the fraction of wage changes is roughly constant across quarters.

We next turn to the measurement of wage rigidity for the aggregate economy. Given the extensive nature of the ADP data, we can also measure wage changes for individuals who transition across firms. This is only possible for workers that transition from one firm who

uses ADP payroll services to another firm that uses ADP payroll services. Given our large sample sizes, we have many workers who transition across ADP firms. Almost all workers that transition across jobs experience a nominal wage change. Given that job switchers are a non-trivial share of the economy, we create a broader measure of nominal wage flexibility pooling together both job stayers and job changers. Doing so, we find that roughly 26 percent experience a wage change during a given quarter and 73 percent experience a wage change during a given year. Including both the job stayers and job changers, roughly 10 percent of workers experience a nominal wage decline with essentially all of the declines being driven by job changers, 38 percent of whom realize a wage decline during our sample. That the aggregate economy, including job switchers, exhibits a substantially higher degree of flexibility than the sample of job-stayers is a 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 relatively flexible on the downside given 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 result in overstating the degree of rigidity in the economy as a whole.

After documenting the difference between aggregate and individual wage rigidity, we examine the extent to which wages are able to adjust to shocks. We provide strong evidence that wage setting behavior is state dependent. Even though nominal wage cuts are very rare for job-stayers over our entire sample period, roughly 6.6 percent of salaried workers and 2.8 percent of hourly workers received nominal 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 wage decline year-over-year was 11.8% during the recession, compared with 9.7% during the recovery. Indeed, the mean wage growth for a worker rose from 2.7% during the recession to 5.2% during the recovery, though this change was mostly driven by the decline in the share of workers receiving cuts: the conditional mean size of wage increases and decreases did not vary much over the cycle. Moreover, we show that nominal wage cuts were much more likely in parts of the country that received large housing price declines. This suggests that any model with a constant fraction of wage adjustments (either in general or with respect to downward adjustment) will fail to match the wage setting patterns over a business cycle. The fact that wages adjust more on the downside during recessions serves to strengthen the puzzle of increased unemployment at business cycle frequencies. Furthermore, the presence of this state dependence in wage adjustment suggests that wages are indeed more downwardly flexible than much of the literature has assumed.

To complement our macro business cycle results, we explore cross-firm variation in wage setting in response to underlying firm level shocks. We document that firms with declining

employment were much more likely to reduce the nominal wages of their workers relative to either firms with constant or increasing employment during the recession. Following the recession, however, this pattern became much more muted, as growing firms and shrinking firms were equally unlikely to cut wages. Instead, shrinking firms during the recovery were much less likely than growing firms to increase workers wages. However, even firms with sharply declining employment raise the wages of many of their employees, both during and after the recession. The interaction between idiosyncratic and aggregate conditions for determining on-the-job wage adjustment patterns suggests that the value of workers' outside options are important for realized wage rigidity, and again urges considerations of models with state dependent wage adjustment.

We end the paper discussing how our results fit into the existing literature on nominal wage adjustments. In doing so, we highlight the importance of using administrative payroll data to measure the extent of nominal wage adjustments among workers. To help illustrate the benefits of administrative payroll data, we create two additional measures of nominal wage adjustments throwing away some of the strength of our data. First, we measure changes in quarterly earnings for a given worker who remains with a given firm. Then we measure quarterly earnings per hour. Both of these approaches use our administrative data on earnings instead of our administrative data on per-period contract rates. The latter of these new measures adjusts earnings by administrative measures of hours worked. We show that both of these measures lead to substantially higher amounts of perceived wage adjustment for workers who remain on the same job. Some of this is due to the fact that earnings includes lots of other forms of compensation like overtime payments, severance payments, commissions, tips, and annual bonuses. However, another important reason is that administrative measures of hours worked for salaried works are measured with substantial error.

To drive this point home, we compute changes in quarterly base earnings per hour for both hourly and salaried workers. By focusing on only base earnings per hour allows us to highlight explicitly the measurement error in hours within administrative datasets (including APD). For hourly workers, changes in quarterly base earnings per hour mimic nearly identically changes in per-period contract rates. However, for salaried workers, the standard deviation in changes in quarterly base earnings per hour far exceeds the standard deviation in changes in per-per contract rates. The only reason for this is the volatility in hours worked per quarter for salaried workers. Given that it is nearly impossible for firms to measure with any degree of accuracy hours worked for salaried workers, it is difficult to interpret with any confidence changes in quarterly earnings per hour as a measure of nominal wage adjustments for salaried workers. Collectively, these results highlight the importance of using administrative payroll

data to accurately measure nominal wage adjustments.

Overall, our results suggest important differences between aggregate and on-the-job wage rigidity. Removing the measurement error that plagues household surveys and administrative hours measures reduces measured nominal wage rigidity for job-stayers substantially, and reveals substantial downward on-the-job rigidity. However, since job changers have much more flexible wages, aggregate wages are much more flexible. Furthermore, we show complementary evidence supporting state dependence in wage setting, urging further consideration of menu cost models of wage adjustment. While there is no doubt that there is some short run stickiness, particularly on the downside, most workers do receive nominal wage adjustments within a year.

The paper proceeds as follows. Section 2 describes the ADP data in detail, and provides summary statistics to benchmark the data to existing data sources. Section 3 describes the allocation of worker compensation across base pay and bonuses. Section 4 presents key facts about wage adjustment for job-stayers, such as the distribution of wage changes, patterns by industry, and evidence of time dependence in wage setting. Section 5 present wage change statistics for job changers, and presents our measures of aggregate nominal wage adjustments. Sections 6 and 7 present evidence of state dependence at the aggregate and regional levels, respectively. Finally Section 8 compares our measures of rigidity to those found in the literature. Section 9 concludes.

2 Data

2.1 Overview of ADP Data

We use administrative individual panel data provided by the ADP Corporation. 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 currently 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, we observe payroll information for approximately 12% of the American workforce.

The data contain monthly aggregates of individual paycheck information, as well as all relevant pieces of 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, while for salaried workers, this constitutes the pay that the worker is contractually obligated to receive each

pay period (weekly, bi-weekly, or monthly). For much of our analysis, we consider hourly and salaried workers separately. 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, this is the exact number of hours worked. For salaried workers, these data are provided by the firm’s HR administrator and is often set to 40 hours. We also observe various additional worker characteristics including their zip code of residence, sex, and age, as well as details about the job, such as the start date of employment (and thus worker tenure), firm size, and industry. Selection into the ADP data is at the firm level. As a result, given unique firm identifiers, we can measure wage distributions within and across firms over time.³ Finally, the presence of consistently-defined worker identifiers permits the careful study of individual worker dynamics across firms. The one caveat is that we are only able to track workers if they move to another ADP-covered firm. However, given our sample size, movements from one ADP firm to another ADP firm are quite common.

We make two major sample restrictions for our analysis. First, we restrict ourselves to workers aged 21 through 60 years old. This restriction focuses our analysis on prime age workers. Second, we only make use of data from ADP’s “Autopay payment” product, which is marketed principally towards firms with over 50 employees. “Autopay” is ADP’s primary payroll processing product.⁴ Therefore, our data is restricted to include only firms with more than 50 employees.⁵

The full dataset that we have access to includes over 50 million unique individuals and over 141 thousand firms. To reduce computational burden, we create two random subsamples of the full data. The first chooses one million unique employees, and follows them through their entire tenure in the sample. This is the primary dataset for analysis. However, this dataset is ill-suited to study questions at the firm level; we therefore construct a second

³Strictly 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 particularly 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.

⁴ADP does have a separate product called “Run” marketed to smaller firms. We have access to this dataset but only starting in July 2013. In the online appendix, we document that many of the main patterns documented for the smaller firms in our primary sample match the patterns for smaller firms in the “Run” sample for overlapping time periods.

⁵There are a few workers in the Autopay database that work at firms with less than 50 employees. We exclude these few workers from our analysis.

subsample of three thousand unique ADP clients, drawing all workers from those firms in the process. The random employee-level and firm-level subsamples remain large, with roughly 25 million and 68 million unique employee-month observations, respectively.

2.2 Representativeness of ADP Data

Table 1 highlights the firm size distribution for employees in our employee sample (column 1) and employees in our firm sample (column 2). For the results in this table, we pool our data together over the entire 2008-2016 period. The table also shows the number of employees and the number of firms in each of our samples. 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 data from the U.S. Census’s Business Dynamics Statistics (BDS) over the same time period. As discussed above, we have access over the entire 2008-2016 period for ADP’s product that is marketed to firms with more than 50 employees. Given this, we have no employees in our base samples that work at firms with less than 50 employees. According to BDS data, 28% of employment is in firms with less than 50 employees. For comparison with our data, column 3 shows the share of employment in firms of differing size relative to employment in all firms with at least 50 employees. As seen from the table, the ADP also slightly under-represents very large employees (those with at least 50 employees). The reason for this is that very large firms tend to have their own human resource department that processes their payroll. Additionally, large firms may be more likely to split into multiple distinct ADP clients. Despite this, ADP still has a large number of employees working in firms with over 5,000 employees. We also explore how the industry distribution of the ADP sample compares to the industry distribution in the BDS. 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.

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 in this paper have been weighted so as to match the BDS’s firm size by industry mix of employment shares for firms with greater

Table 1: Firm Size Distribution in ADP Samples and the BDS, Pooled 2008-2014 Data

	Pooled 2008-2016		
	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	24,831,244	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

Note:

than 50 employees. We compute our weights for each year between 2008 and 2016. By reweighting the data, we control for sample selection along these key observable dimensions. Although there may yet remain selection into the sample along unobservable dimensions (e.g., firms with high cash flow are more likely to hire ADP), we consider these potential selection issues to be small once controlling for firm size and industrial mix.⁶

Table 2 shows some additional summary statistics for our 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). As discussed in the Online Appendix, 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). About one-fifth of our sample is paid weekly while three-quarters is paid bi-weekly. Less than five percent are paid monthly.

For our sample, roughly 64 percent are paid hourly with the remaining 36 percent being classified as salaried workers. According to data from the CPS monthly supplements, only 57 percent of employed workers in the U.S. between the ages of 21 and 60 report being paid hourly during this time period. The difference between the CPS and ADP data may arise as the distinction between hourly and salaried workers is sometimes unclear within the ADP dataset. Some hourly 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

⁶In the Online Appendix that accompanies the paper we show a series of additional results. In particular, we show sample statistics for each individual year and report our sample weights. We also show our key results without imposing sample weights. Finally, we show that ADP is truly a national firm in that it has a very representative geographic coverage.

Table 2: Statistics for Employee Sample, Selected Years

	All	2008	2012	2016
Number of Workers	1,000,000	202,329	341,726	342,991
Number of Firms	89,350	89,350	89,350	89,350
Number of Observations	22,642,878	1,319,797	2,744,414	2,778,947
Age 21-30 (%)	24.9	25.2	23.9	26.4
Age 31-40 (%)	23.6	24.5	23.4	24.1
Age 41-50 (%)	23.3	24.4	23.7	21.7
Age 51-60 (%)	20.8	18.6	21.5	20.7
% Male	54.1	54.3	54.0	55.0
Average Tenure	66.8	73.2	67.8	61.4
% Paid Weekly	20.5	21.3	21.2	20.6
% Paid Bi-Weekly/Semi-Monthly	76.3	75.5	75.5	75.5
% Paid Monthly	3.3	3.1	3.3	3.8
% Hourly	65.2	64.2	65.4	65.0

wage workers. However, on many levels, these workers operate as if they were salaried: their actual hours are never recorded and 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 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 similar to the CPS averages.

Given that ADP is growing over time, so too is our sample. Of our 1 million workers, only 202,000 are in our sample in 2008 while 343,000 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.

2.3 Measuring Nominal Wage Adjustments

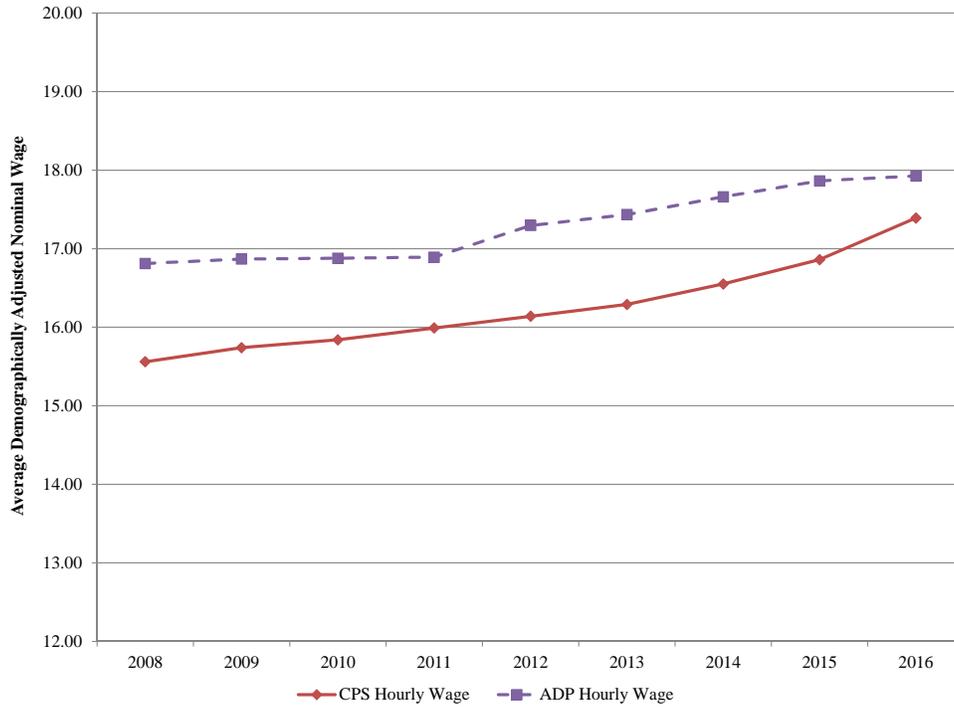
The focus of this paper is on the frequency and size of wage changes. We consider changes in the worker’s per-period payment rate as our measure of changes in the worker’s nominal wage. To reiterate, our nominal wage measure is the hourly wage for hourly workers and per-period earnings for salaried workers. We aggregate our unit of observation to the month. Specifically, our nominal wage measures are defined as the wage paid to the individual during the last pay period of the month. Throughout the paper, we explore one-month, three-month (one-quarter) and twelve-month (annual) wage changes.

Given that our nominal wage measures come from administrative HR records, there should be little, if any, measurement error in our wage measures. Despite this, there are some exceptionally small wage changes in our data resulting from salaried individuals earning annual amounts that do not easily divide into twelve months. As a result, we consider only wage changes of at least 0.1%. That is, if worker i earns wage w_{it} in period t , we consider the object $\Delta w_{it}^k = \log w_{it} - \log w_{it-k}$ for some k , and say that an individual has experienced a wage change in the previous k months if $|\Delta w_{it}^k| > 0.001$.

Figure 1 compares the average hourly 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.

Finally, for our analysis, we will separately analyze nominal wage adjustments for a sample of “job-stayers” as well as a separate sample of “job-changers”. Job-stayers are workers who remain employed at the same firm between the periods of t and $t + k$. Job-changers are workers who move from one ADP firm to a new ADP firm between the periods of t and $t + k$. For our job-changer sample, workers could have another job or be non-employed at some point between t and $t + k$. Given that we only measure job transitions from one ADP firm to another ADP firm, we cannot distinguish jobs at non-ADP firms separately from being non-employment. Finally, when measuring nominal wage changes for job-changers, we only focus on workers who transition from either hourly-to-hourly jobs or from salaried-to-salaried jobs.

Figure 1: 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. 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.

3 The Allocation of Worker Compensation

As discussed above, the primary focus of our paper is using a worker’s per-period contract rate to measure the flexibility of a worker’s wage over time. However, workers receive other forms of compensation including overtime payments, tips, commissions, and annual bonuses. In this section, we discuss how important these other margins are in terms of worker compensation. As we discuss below, the median worker receives 96.2 percent of her compensation from the per-period contract rates. However, for some workers annual bonuses are important. For another set of workers, a large part of their compensation comes from tips and commissions. The ADP dataset is one of the few datasets of which we are aware that allows researchers to measure the composition of worker compensation.

Specifically, in addition to measuring a worker’s per-period contract rate without error, the ADP data also include administrative measures of a worker’s monthly gross labor compensation (excluding benefits). The monthly gross labor compensation measure is the sum of the worker’s actual gross take home pay accrued during the month. We refer to all monthly earnings stemming from a workers per-period contract wage adjusted for either hours worked or the number of pay periods during a given month as being the worker’s monthly “base compensation”. For hourly workers, base compensation would be the product of their per-period hourly wage and the number of hours they worked during the month. Meanwhile, salaried workers who are paid weekly will have their monthly compensation be four or five times higher than their per-period weekly contract rate.

However, workers also receive compensation during a month above and beyond their base compensation as measured by their per-period contract rate. Some workers work overtime and receive additional overtime compensation (e.g., time-and-a-half). Other workers receive part of their compensation in the forms of tips and commissions, or after meeting contractually stipulated production or sales targets. Some workers receive meal and travel reimbursements within their paychecks accrued during a given month. Finally, workers on occasion receive bonuses from their firms which show up in their take-home pay.⁷

ADP firms are required to report a worker’s per-period contract wage, the worker’s hours worked (if they are hourly), and the worker’s gross wage earnings (which shows up as a worker’s gross (pre-tax) take home pay). However, the fact that firms are not required to provide detailed information on these other forms of compensation makes it difficult to ascertain with certainty the amount and composition of overtime, tips and commission, and bonus compensation. However, given the available information, we can make substantive progress on distinguishing both the size and composition of other forms of compensation.

⁷Meal reimbursements also show up in workers monthly take-home pay.

In particular, for both hourly and salaried workers, we can create a measure of monthly “residual earnings” by subtracting monthly base compensation from actual gross monthly earnings.⁸ These residual earnings will include meal and travel reimbursements, signing bonuses, severance pay, cash outs of unused vacation pay, overtime payments, commissions, tips, advance payments of paychecks, and bonuses accruing to the worker during a given month. One additional component of these residual earnings could be measurement error if the per-period contract rate changed during the month. Given we only have monthly aggregates and our monthly per-period contract rate is measured at the end of the month, any change in the wage during the month can cause a deviation between gross earnings and the per-period contract rate.

To help assess how important other forms of compensation above and beyond their base earnings is to the workers total annual compensation, we temporarily restrict our sample to only those workers who remain continuously with the same employer for a full calendar year. For this set of workers, we compute a measure of their gross annual earnings by summing together their monthly gross earnings across all calendar months. Additionally, we compute a measure of their gross annual base pay by summing together their monthly base pay earnings across all calendar months. By taking the ratio of the two, we can create a measure of the share of all annual gross earnings that comes from their base pay. If the ratio equals 1, the worker during the year only receives base pay. When the ratio is less than 1, the worker is compensated with some additional residual payments throughout the year.

Table 3 displays the distribution of the share of earnings accruing from base pay for our full sample (column 1) and our full-year sample (columns 2 and 3). Again, the full year sample is the sample of individuals who remain continuously employed with the same firm during a 12 month calendar year. We show patterns using monthly data (share of gross earnings during the month coming from base pay) and annual data (share of gross earnings during the year coming from base pay). The table shows that for most worker-month pairs, most earnings come from base pay. For example, even the 25th percentile of worker-month pairs has roughly 94 percent of earnings coming from base pay. However, for some worker-month pairs (around the 10 percentile), a substantial part of worker gross-monthly earnings comes from sources other than base pay. Again, this could stem from commissions, tips, overtime payments, signing bonuses, annual bonuses, meal and travel reimbursements, etc. The final column of the table shows the share of annual earnings coming from base pay. For the median worker during the 2008-2016 period, roughly 96 percent of all earnings comes from base pay. Even the 25th percentile of worker-year pairs has workers receiving over 90 percent of all gross earnings coming from base pay. It is this reason that we chose to focus

⁸See the Data Appendix for exact details on this procedure.

on base pay as our primary measure of worker compensation. Our nominal wage measure is directly related to base pay compensation.

While workers receive many additional types of compensation above and beyond base pay compensation, one explicit type of compensation we wish to explore more fully is bonus pay compensation. Again, bonuses are not consistently measured within the ADP data. However, we can create a proxy for bonuses by using our residual income measure. To do so, we define a worker bonus as occurring if residual income exceeds one-percent of annual earnings in either December, January, February or March. Most firms pay annual bonuses in December (as a “Christmas bonus”) or early in next calendar year.⁹ We exclude small residual income payments in these months (less than 1 percent of annual earnings) from our bonus measure so as to exclude small deviations between monthly earnings and monthly base pay due to things like small measurement due to within changes in the workers contracted pay amount.

Table 4 displays summary statistics for our broad bonus measure (top panel) summed over the 2008-2016 period. Again, we use our sample of workers who remain continuously employed with the same employer during a full 12 month calendar year. Give our broad bonus definition, roughly one-quarter of all workers (column 1), one-fifth of all hourly workers (column 2) and one-third of all salaried workers (column 3) receive an annual bonus. The mean and median size of the bonus, conditional on a bonus occurring, is roughly 5 percent and 3.4 percent, respectively, of annual gross earnings. Not surprisingly, the conditional bonus share is higher for salaried workers than for hourly workers. The fact that the mean bonus share is larger than the median suggest that some workers are receiving really large bonuses during the month. Bonuses explain one important reason why a worker’s annual gross pay exceeds their annual base pay.

Our broad bonus measure is likely an upper bound on true bonuses given that it our measure of residual earnings in these four months are still potentially contaminated with some overtime, commission an tip payments earned in those months. For example, workers who are paid with commissions will have potentially residual payments in many months during the year. To potentially account for this, we exclude from our broad bonus measure any worker who has residual monthly earnings greater than 1 percent of their annual earnings in 3 or more months during the year. These are people who persistently have large residual earnings throughout many months of the year. We refer to this as our narrow bonus measure. The bottom panel of Table 4 shows that roughly 16 percent of all workers receive large residual payments in December-March but receive no large residual payments in other months of the

⁹Over 85 percent of residual earnings (value weighted) accrued in these four months with the frequencies being by far the largest in December followed next by March.

Table 3: Share of Annual Compensation in Base Pay

	Sample		
	All	Full-Year	
	Monthly Data	Monthly Data	Annual Data
<u>All Workers</u>			
10 th Percentile Share in Base	78.6%	78.3%	80.3%
25 th Percentile Share in Base	93.7%	93.6%	90.1%
Median Share in Base	100%	100%	96.2%
75 th Percentile Share in Base	100%	100%	99.4%
90 th Percentile Share in Base	100%	100%	100%
<u>Hourly Workers</u>			
10 th Percentile Share in Base	83.4%	83.8%	84.7%
25 th Percentile Share in Base	93.2%	93.2%	92.0%
Median Share in Base	99.2%	98.9%	96.9%
75 th Percentile Share in Base	100%	100%	99.3%
90 th Percentile Share in Base	100%	100%	100%
<u>Salaried Workers</u>			
10 th Percentile Share in Base	65.2%	65.3%	74.1%
25 th Percentile Share in Base	95.6%	95.3%	86.6%
Median Share in Base	100%	100%	94.6%
75 th Percentile Share in Base	100%	100%	99.5%
90 th Percentile Share in Base	100%	100%	100%

Note: Table shows the distribution across households in the share of their total gross earnings that is base pay. Columns 1 and 2 focus on monthly shares while column 3 focuses on annual shares. Column 1 uses our full employee sample. Columns 2 and 3 restrict our sample to only those individuals who remain employed with the same employer for a full calendar year. The top panel includes data for all workers regardless of pay type while the bottom two panels include data for hourly and salaried workers separately.

Table 4: Annual Bonus Distribution

	All	Hourly	Salaried
<u>Broad Bonus Definition</u>			
Fraction with Bonus in Year	26.8%	22.3%	33.5%
Mean Share of Pay in Bonus	1.2%	0.7%	2.0%
Mean Share of Pay in Bonus, Conditional > 0	4.8%	3.3%	6.3%
Median Share of Pay in Bonus, Conditional > 0	3.4%	2.2%	4.9%
Mean Size of Bonus, Conditional > 0	\$3,489	\$1,676	\$5,392
Median Size of Bonus, Conditional > 0	\$1,727	\$815	\$3,709
<u>Narrow Bonus Definition</u>			
Fraction with Bonus in Year	15.7%	13.3%	19.6%
Mean Share of Pay in Bonus	0.7%	0.4%	1.2%
Mean Share of Pay in Bonus, Conditional > 0	4.5%	3.2%	6.0%
Median Share of Pay in Bonus, Conditional > 0	3.1%	2.1%	4.7%
Mean Size of Bonus, Conditional > 0	\$3,149	\$1,577	\$4,877
Median Size of Bonus, Conditional > 0	\$1,507	\$751	\$3,226

Note: Table shows the distribution of “bonuses” as measured in our ADP employee sample. For this analysis, we restrict the sample to those employees who remain with the same firm for a full calendar year. We define “bonuses” as being positive when a worker has residual earnings in the months of December, January, February and March. See text for definition of residual earnings. Our broad measure of bonuses (top panel) defines bonuses as any positive residual earnings accruing in December, January, February or March. Our narrow measure of bonuses (bottom panel) only includes individuals who positive residual earnings in one calendar month during December, January, February or March.

year. Again, bonuses are more common for salaried workers relative to hourly workers.

Our narrow bonus measure is likely a lower bound on true annual bonuses received given that some of the commission workers that we may have excluded likely also receive a large annual end of the year bonus. Despite the potential measurement error concerns, we still think it valuable to explore the extent to which bonuses may vary over time for a given worker and how that variation evolves over the business cycle. So, in addition to just measuring nominal wage adjustments for a given worker (which comprises the overwhelming majority of their annual earnings for most workers), we are also going to explore bonus variation over time. Collectively, we think bonus pay and base pay comprise the two most important levers that firms can adjust with respect to the compensation of their workers.¹⁰

4 Nominal Wage Adjustments for Job-Stayers

In this section we present key facts about the nature of on-the-job wage adjustment. First, we measure the frequency and size of nominal wage adjustments for job-stayers, finding a strong asymmetry in wage changes on-the-job. We then explore whether nominal wage adjustments differ by firm size and industry. Finally, we highlight the importance of time dependence in wage setting at both the individual worker and firm level for job-stayers, and the extent to which bonuses provide a relevant margin of adjustment of workers' compensation.

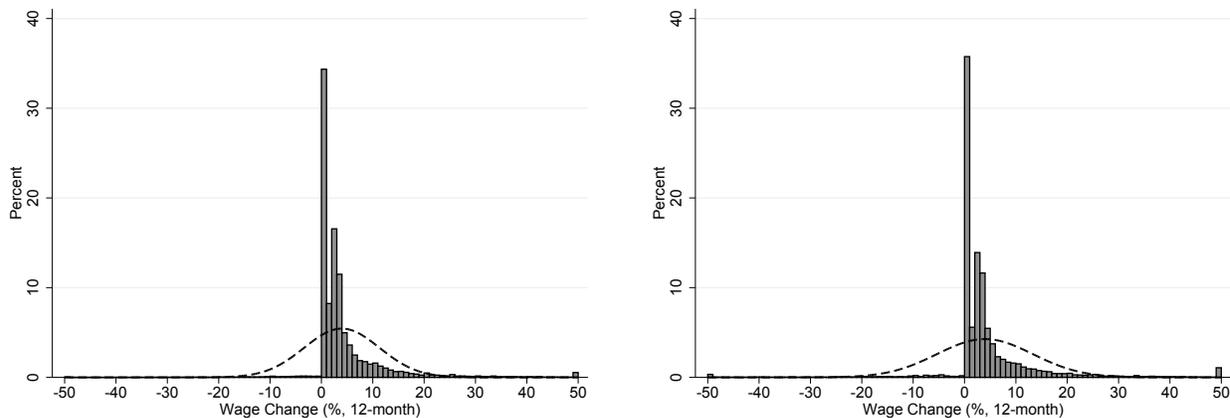
4.1 Main Results

Figure 2 highlights the first key set of facts of the paper. The figure plots the distribution of 12-month wage changes for "job stayers." As discussed above, our measure of the worker's nominal wage is their per-period contract rate. Panel A plots the distribution for hourly workers, while Panel B plots the distribution for salaried workers. Three 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 wage change in a given year. Second, there is a clear asymmetry in the wage change distribution, with the overwhelming majority of changes being wage increases. Of the roughly 66% of all individuals who receive a wage change over a given 12-month period, only 3.6% received a wage cut (2.4/66). Finally, there are very few small wage changes for either hourly or salaried workers. Just 8.6% of workers receive a wage change of between 0.1 and 2 percentage points, compared with 27.1% receiving between 2

¹⁰Ideally, we would like to also explore fringe benefits. The ADP data needs more processing before such a similar systematic analysis of fringe benefits can be conducted. We hope to explore such an analysis in future work.

and 4 percentage points. This missing mass of very small wage changes is consistent with the random menu cost models that are so prevalent in the price setting literature.

Figure 2: Nominal Wage Change Distribution for Job Stayers



PANEL A: HOURLY WORKERS

PANEL B: SALARIED WORKERS

Note: Figure shows the annual change in nominal wages for workers in our employee sample who remain employed on the same job. We use our employee sample for this analysis.

Table 6 provides a set of moments on the probability of wage increases and wage declines for three frequencies: monthly, quarterly and annual. The annual frequencies correspond to the underlying data shown in Figure 2. The first column pools together hourly and salaried workers while the second and third columns, respectively, show the frequency of wage changes for hourly and salaried workers separately. A few things are of note from the table. First, the frequency of wage changes is roughly similar between salaried and hourly workers. Roughly two-thirds of both receive annual wage changes over the entire sample period (summing over wage increases and wage declines). Second, while the average probability of a wage change is similar between the two groups in our sample of job stayers, salaried workers are more likely to receive a nominal wage cut. Over the entire sample, only 1.8% of hourly workers receive a nominal wage cut over a 12 month period while 3.6 percent of salaried workers receive a nominal wage cut. Third, one cannot simply extrapolate monthly nominal wage changes to quarterly or quarterly wage changes to annual. The probability of a quarterly nominal wage change is less than three times the monthly wage change and the probability of an annual nominal wage change is less than four times the quarterly change. This is likely due in part to well-known time aggregation issues arising from workers who receive multiple wage changes, as well as to the fact that the samples differ between the three horizons: for monthly wage changes, workers need to only remain with their employer for one month, while for annual wage changes workers need to remain with their employer for the full year.

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

	Monthly	Quarterly	Annual
<u>All Workers</u>			
Probability of Positive Wage Change (%)	6.3	18.5	63.9
Probability of Negative Wage Change (%)	0.4	0.9	2.4
<u>Hourly Workers</u>			
Probability of Positive Wage Change (%)	6.6	19.5	65.3
Probability of Negative Wage Change (%)	0.3	0.7	1.8
<u>Salaried Workers</u>			
Probability of Positive Wage Change (%)	5.8	16.7	61.6
Probability of Negative Wage Change (%)	0.6	1.3	3.6

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

Table 5 shows additional moments of the wage change distribution. For this table, we pool together both hourly and salaried workers.¹¹ During this period, mean and median nominal wage growth for workers who remain on the same job equaled 3.9 percent and 2.4 percent, respectively.¹² Conditional on a wage change occurring, annual mean and median nominal wage growth was 5.6 and 3.2 percent. A key statistic we will focus on throughout the paper is the standard deviation of nominal wage growth. Unconditionally and conditional on a wage change occurring, the standard deviation of annual nominal wage growth during the full 2008-2016 period was 6.5 percent and 6.9 percent, respectively. Consistent with the patterns in Figure 2, annual wage changes display very large amounts of both skewness and kurtosis. Conditional on a positive 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 wage changes, perhaps due to promotions. The mean and median size of a wage cut, conditional on the worker experiencing a nominal wage reduction were both around 7

¹¹The results were again roughly similar between hourly and salaried workers.

¹²To 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.

percent. While the frequency of wage increases is much higher than wage cuts, the mean size of a wage increase conditional on it happening is nearly identical to the mean size of a wage cut conditional on it happening.

The evidence presented in this section yield multiple lessons of importance for modelers and policymakers. First, we observe a large asymmetry in realized wage flexibility, with wages seemingly much more difficult to cut than raise. Although excessive downward nominal rigidities have been documented before in the literature (see, e.g. Lebow et al. (2003); Kahn (1997); Card and Hyslop (1997)), the presence of measurement error in hours and earnings in prior work has yielded quantitatively different magnitudes of this asymmetry. In addition, the missing mass of small wage changes urges consideration of models of state dependent wage adjustment, which we explore in more depth in Section 6. Again, this missing mass has been difficult to detect in prior work, for reasons explored in great detail in Section 8.

4.2 Nominal Wage Adjustments for Job-Stayers by Firm Size and Industry

Figure 3 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 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 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 wage adjustments. All of the variation across firm size groups is in the propensity to receive a nominal wage increase. While nominal wage cuts are rare for all workers, there is no systematic variation in the propensity of a nominal wage cut with firm size. Figure 3 also shows that there is a fair degree of heterogeneity across industries with respect to wage changes. For example, both hourly and salaried workers in the manufacturing industry are much more likely to receive a wage change than workers in construction during our sample period.

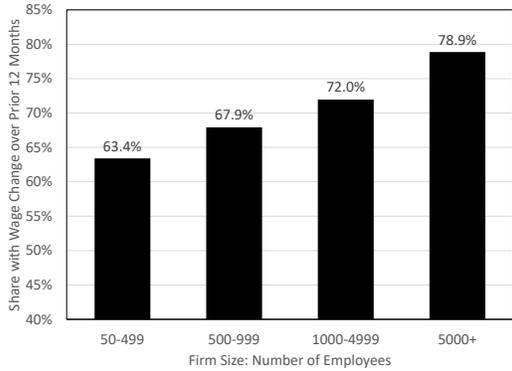
Given that firms within different industries also differ by size, a natural question is how much of the variation across industries is due to differences in firm size. To assess this, we regressed the probability of a nominal wage change during a given year on a vector of firm size dummies and a vector of industry dummies. We also included a vector of additional controls including a quadratic in worker age, a quadratic in worker tenure, an indicator of whether the

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

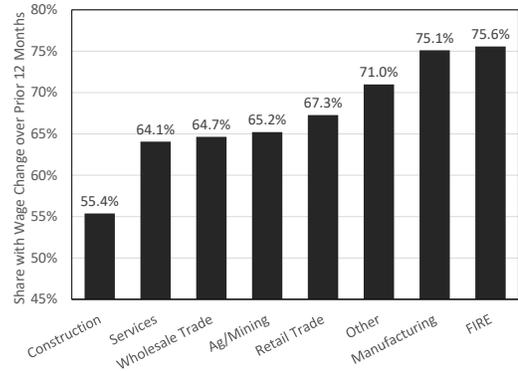
	Monthly	Quarterly	Annual
<u>Unconditional</u>			
Mean Wage Change (%)	0.3	1.0	3.9
Median Wage Change (%)	0.0	0.0	2.4
Standard Deviation of Wage Change (%)	2.6	3.7	6.5
Skewness of Wage Changes (%)	9.7	5.3	2.8
Kurtosis of Wage Changes (%)	175.4	49.5	14.4
<u>Conditional on Any Wage Change</u>			
Mean Wage Change (%)	5.0	4.9	5.6
Median Wage Change (%)	3.0	3.0	3.2
Standard Deviation of Wage Change (%)	8.1	6.5	6.9
Skewness of Wage Changes (%)	2.1	2.1	2.4
Kurtosis of Wage Changes (%)	15.9	13.6	12.1
<u>Conditional on Positive Wage Change</u>			
Mean Wage Change (%)	6.2	5.7	6.3
Median Wage Change (%)	3.3	3.3	3.5
Standard Deviation of Wage Change (%)	7.7	6.4	7.0
<u>Conditional on Negative Wage Change</u>			
Mean Wage Change (%)	-10.7	-8.7	-7.3
Median Wage Change (%)	-8.3	-7.7	-6.6
Standard Deviation of Wage Change (%)	8.1	5.8	4.6

Note: 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.

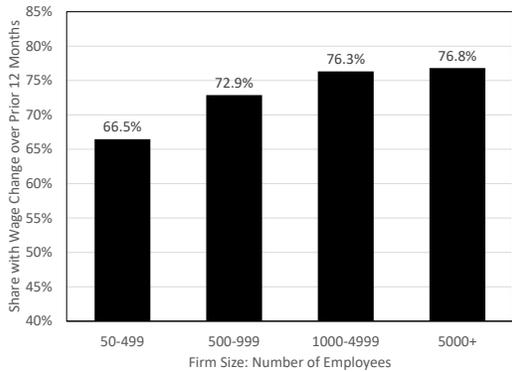
Figure 3: Share with Wage Change by Firm Size and Industry, All Years



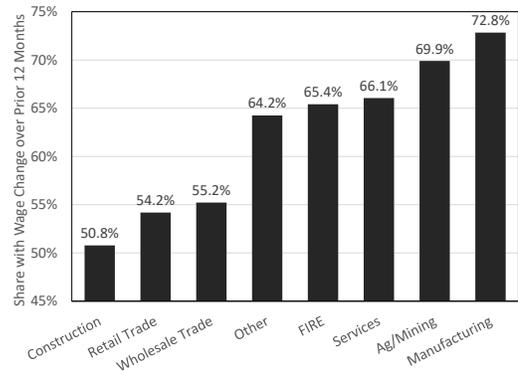
PANEL A: HOURLY WORKERS BY SIZE



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 wage change by firm size and industry for a 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.

worker is paid hourly, and a vector of state of residence \times month-year fixed effects. We also ran a version of the regression replacing the dependent variable with either the probability of a nominal wage increase during the year or the probability of a nominal wage cut during the year. The full results of these regressions are shown in the Online Appendix accompanying the paper. The results of the regression still show that there is a large and statistically significant gradient between firm size and the propensity of a nominal wage change. Workers in firms with over 5,000 employees are 10 percentage points more likely to experience a nominal wage change than workers in firms with 50-499 employees. Likewise, workers in the manufacturing sector were 7 percentage points more likely to receive a nominal wage change relative to workers in the construction or retail trade industries, conditional on observables.

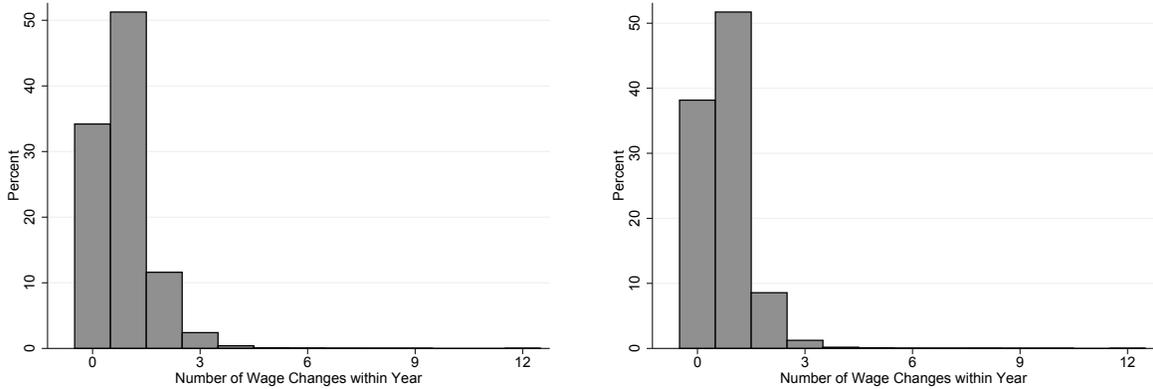
4.3 Time Dependence in Nominal Wage Adjustments, Job-Stayers

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 sub-section, we use our detailed micro data to explore evidence of time dependence in wage adjustment for our sample of job-stayers. The purpose of doing so is two-fold. First, this section provides some background summary statistics on the frequency and nature of wage adjustment for job-stayers. Second, the presence of time dependence in wage setting informs the models of wage setting that should be considered in labor and macroeconomics going forwards.

Figure 4 plots the average number of wage changes during a given year for workers in our employee sample. As seen from Table 5, roughly 35 percent of job-stayers receive no wage change during a 12 month period. Over 50 percent of both hourly and salaried workers receive exactly one wage change during a 12 month period where they remained continuously on the job. Between 10 and 15 percent of job-stayers receive multiple wage changes during a given year. The take away from Figure 4 is that roughly 90 percent of job-stayers receive either zero or one nominal wage change during a given year. Multiple nominal wage changes within a year are rare for continuing employees who remain on the same job.

To begin formally studying time dependence in wage setting, we exploit the individual level micro data and estimate an individual duration model of wage changes. Figure 5 plots the resulting hazard functions of wage adjustment for the subset of job-staying employees who experience at least two wage changes over our sample period. Specifically, the figure shows the probability of a one month wage change between $t - 1$ and t conditional on the worker surviving to month t without a wage change at the same firm.

Figure 4: Number of Nominal Wage Changes over 12 month period, Job-Stayer Sample

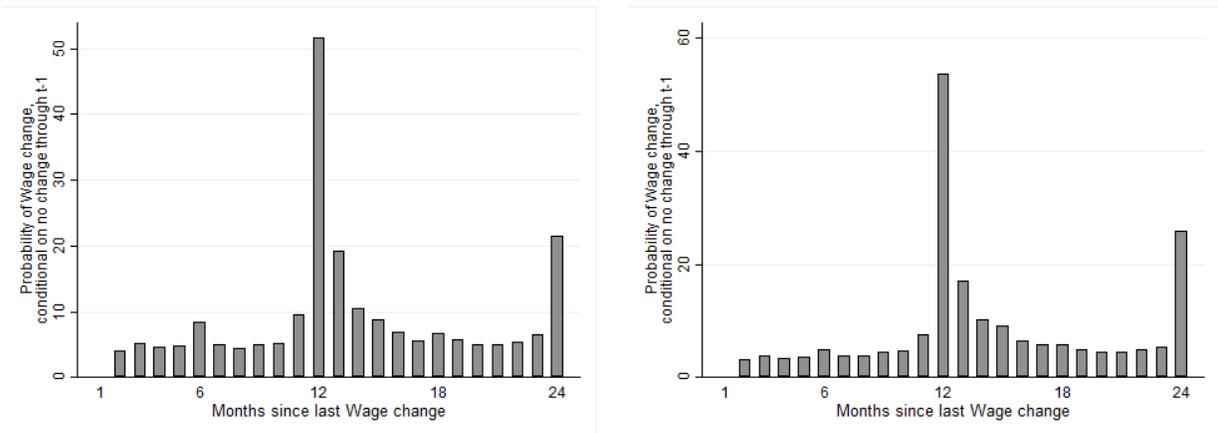


PANEL A: HOURLY WORKERS

PANEL B: SALARIED WORKERS

Note: Table shows the average number of nominal 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.

Figure 5: Hazard Function of Wage Change, Job-Stayers



PANEL A: HOURLY WORKERS

PANEL B: SALARIED WORKERS

Note: Figure shows the hazard rate of a wage change between $t - 1$ and t conditional on surviving to t without a wage change at the same firm. Sample only includes individuals with at least two 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.

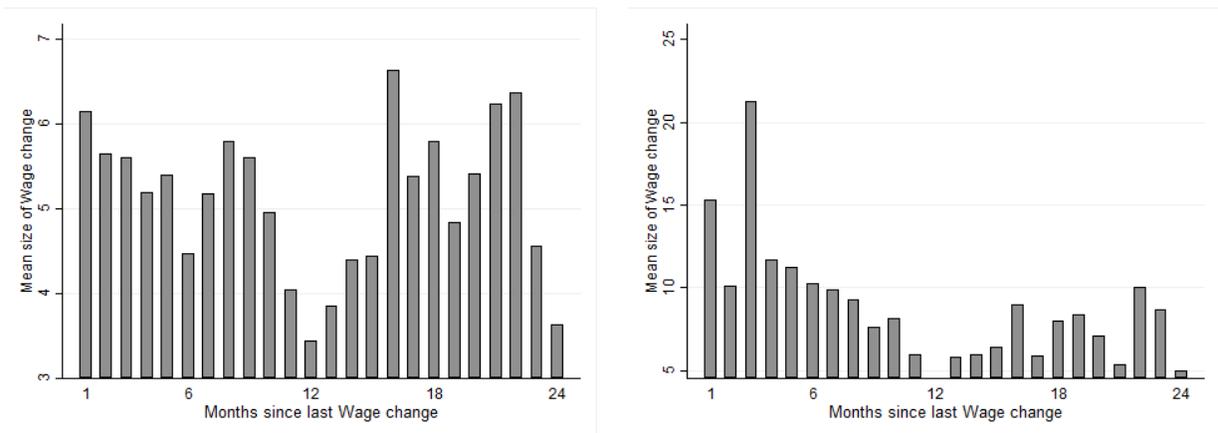
The figure rejects the Calvo prediction that the probability of wage change is constant over time at the individual level for job-stayers. In most months, the probability of a wage change is roughly constant at about 3-4%. However, roughly 12 months after the last wage increase, individuals are much more likely to get another wage increase. Conditional on making it to month 11 with no wage change, there is over a 50% probability than an individual gets a 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 and a more modest spike at 36 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 wage change. We can then ask how many of these workers got their next wage change 11-13 months later. Of consistently employed workers, 30% receive their next wage change exactly one year after their prior wage change.

Figure 5 provides some evidence of time dependence in wage adjustment. The majority of 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 6 shows the average size of the wage change for job stayers by the time since last wage change. Since the vast majority of wage changes for job stayers are positive, this figure only includes workers who received a positive wage change. While most wage changes occur at 12 month frequencies, Figure 6 shows that the size of the 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 wage changes that occur off-cycle are positively selected in some way, this could explain why they receive higher wage increases. For example, if the worker receives an outside offer, the firm may have to raise the worker's wage earlier than their annual cycle in order to retain the worker. Or, if a worker is promoted internally and the promotions are distributed throughout the year, it is not surprising that workers who receive a wage change off cycle also get larger wage changes.

Figure 7 shows the time dependence in wage setting at the firm level. For this analysis, we use our sample of 3,000 unique firms. 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 wage change

Figure 6: Mean Size of Wage Changes by Time Since Last Change, Job-Stayers



PANEL A: HOURLY WORKERS

PANEL B: SALARIED WORKERS

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

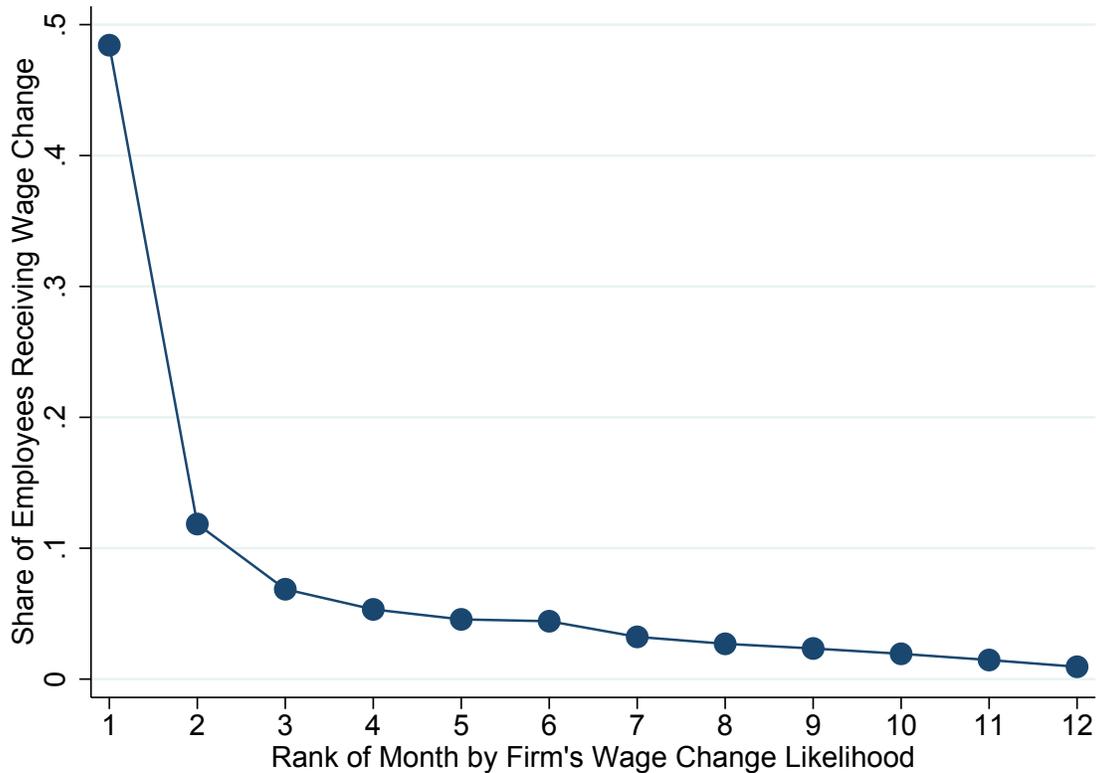
during each calendar month. We then rank the months within a given firm-year pair from the month with the highest fraction of nominal wage changes to the month with the lowest fraction of nominal wage changes. For example, for some firms the highest month may be September while for other firms the highest month may be January. We then take the simple average probability of a worker receiving a wage change across firm-year pairs for each ranked month.¹³

The figure shows that when a firm tends to adjust wages, it makes all their wage adjustments during one particular month of a given year. For example, a typical firm adjusts 50 percent of their workers wages in the month where they make the most wage changes. Given that only about 65 percent of workers get a wage change (in the population as a whole) and the fact that we are averaging over firms and not workers, the figure suggest that firms do most of their wage changes in one month out of the year.¹⁴ As a point of contrast, firms only adjust roughly 10 percent of their workers wages in the second highest ranked month. The fact that the share of wages adjusted are roughly flat between the second highest ranked month and the 12 highest ranked month is consistent with the worker data where some adjustments are occurring off-cycle at a roughly constant hazard. These changes are likely

¹³We also restrict our sample to only firm-year pairs where the firm adjusted at least 25 percent of their workers wages at some point during the calendar year. We do this to focus on firms who are adjusting wages to avoid a problem with firms no wages during the year. This restriction is not too binding as 91% of firm-year pairs in our sample adjusted at least 25 percent of their workers wages during the year.

¹⁴This 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.

Figure 7: Share Receiving 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 wage change during that month. We then rank the months within a given firm-year pair from highest month of wage changes to lowest month of 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 wages at some point during the calendar year.

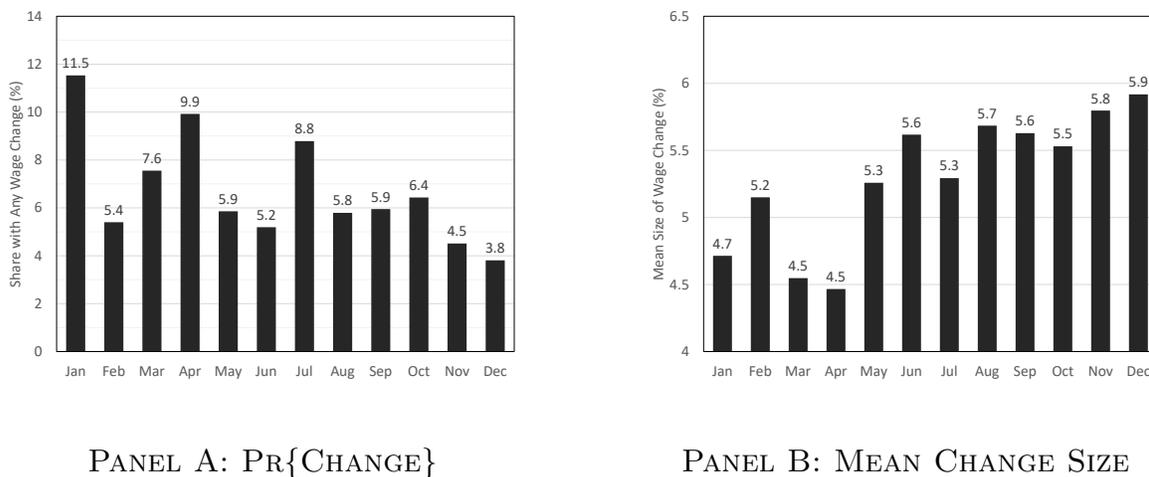
due to promotions and/or the response to outside offers.

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 changed 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 wage changes are coordinated within a given calendar month.

Figure 8 shows the probability of wage changes by calendar month pooling together

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 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 wage change in the first quarter of the year while 21.1 and 21.5 percent of workers receive a wage change in the second and third quarters. Only 16.6 percent of workers receive a wage change in the last quarter of the year.

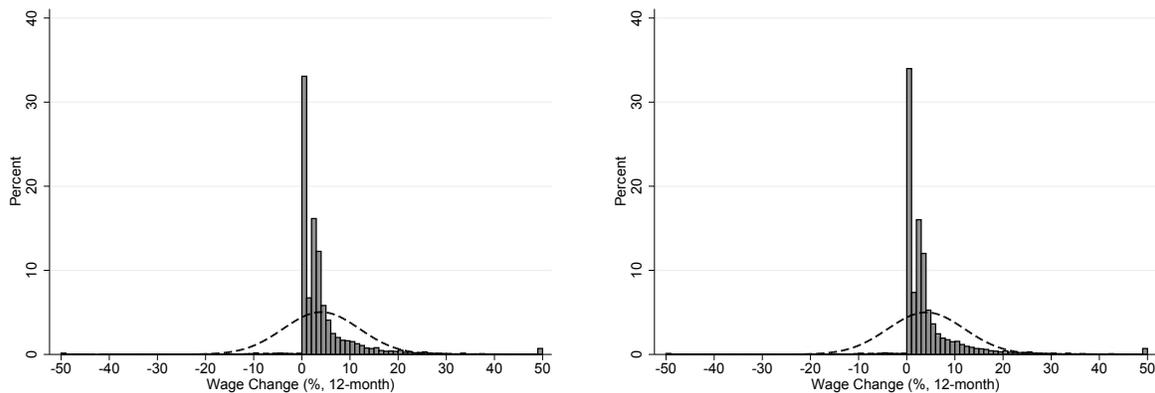
Figure 8: Seasonality in Wage Changes, Job Stayers, All Years



Note: Figure plots moments of the 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 wage adjustment. The majority of 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 wage adjustment across the four quarters of the year, suggesting that models of Calvo adjustment may be a reasonable approximation of the wage adjustment process. However, as we will document in Section 6, wage 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 wage cuts and increases (see, e.g. Schmitt-Grohé and Uribe (2012)).

Figure 9: Wage Change Histograms by Whether Individual Receives Bonus in Year, Job-Stayers



PANEL A: WITH BONUS

PANEL B: NO BONUS

Note: Figure shows the distribution of annual nominal wage changes for a sample of households who receive bonuses (left panel) and a sample of households who do not receive bonuses (right panel). We use our broad bonus definition when separating the sample. Our base sample includes data on job-stayers from our employee sample. We include all data between 2008 and 2016 and pool together both hourly and salaried workers.

4.4 Bonus Payments and Nominal Wage Adjustments

Before concluding this section, we explore the extent to which Bonus payments provide flexibility for some workers in their annual earnings. As discussed above, most workers do not receive a bonus payment during a given year. However, for those that do, bonuses may provide a way for firms to adjust the compensation of their workers.

We begin our analysis by exploring whether the annual nominal wage adjustments for those job-stayers who received a bonus during the year differed from the annual nominal wage adjustments for those job-stayers who did not receive bonus during the year. One may conjecture that if a firm offers a bonus and that the bonus is a margin of adjustment for worker wages, the patterns of wage adjustments for firms that give bonuses could systematically differ from those firms that do not pay bonuses. That is simply not the case. As seen from Figure 9, the annual wage change distribution of those job-stayers who do receive a bonus during the year (left panel) is essentially identical to the annual wage change distribution of those job-stayers who do not receive a bonus during the year (right panel). We explored additional moments of the wage change distribution between the two groups and there were no meaningful differences between the nominal wage adjustments of job-stayers with and without bonuses.

4.5 Summary of Nominal Wage Adjustments for Job-Stayers

Collectively, the results in Tables 5 and 6 provide a set of high quality statistics on the wage change distribution for job-stayers that can be used to calibrate many different types of macroeconomic models. At this stage, however, it is worth reflecting on the appropriate use of these moments. These moments will be appropriate for any model for which workers only observe earnings growth on-the-job, or models in which there is a clear distinction between on-the-job growth and earnings growth through job search or ladders. One example would be Hall (2005)'s model of employment fluctuations with partially smoothed wages. However, models in which the majority of wage growth arises from moving up a job ladder or transitioning to new employers, or macro models where households may frictionlessly supply units of labor to the market may be better served by considering more aggregate measures of nominal rigidity. We now turn to the construction of our aggregate nominal wage flexibility measures.¹⁵

5 Aggregate Nominal Wage Adjustments

In the prior section, we focus on nominal wage adjustments for individual job-stayers. Most of the existing literature looking at nominal wage changes focuses only on how wages evolve during a year for a sample of job-stayers (Altonji and Devereux, 2000; Lebow et al., 2003; Sigurdsson and Sigurdardottir, 2016). However, workers need not realize wage changes within a job in order to achieve earnings growth. In this section, we explore how wages change for workers who switch jobs. We then create an aggregate wage adjustment measure combining both job-stayers and job-switchers.

5.1 Nominal Wage Adjustments, Job-Changers

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 the patterns of wage adjustments for workers who migrate across ADP firms are similar to the patterns of wage adjustment for workers who migrate to and from non-ADP firms.

¹⁵Kudlyak (2014) and Basu and House (2017) have both argued for caution in considering rigidity solely in remitted wages. Kudlyak (2014) argues that a firm deciding whether to hire a new worker will consider the *user cost* of that worker, while Basu and House (2017) highlight that there might be differences between allocative and remitted wages, and that these differences matter in the context of a medium-scale DSGE model.

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 a firm will look like a job-change. To account for such flows, we measure the percent of job changers leaving a given firm i in month t and showing up at another ADP firm in month $t + 1$ or month $t + 2$. If more than 20 percent of job changers leaving firm i and subsequently show up in firm j with no intervening employment spell elsewhere, we treat those as within firm movements and do not include them in our job-changer sample.¹⁶ In addition, if a worker’s reported tenure does not reset after switching firms, we exclude that worker from the job changer sample.

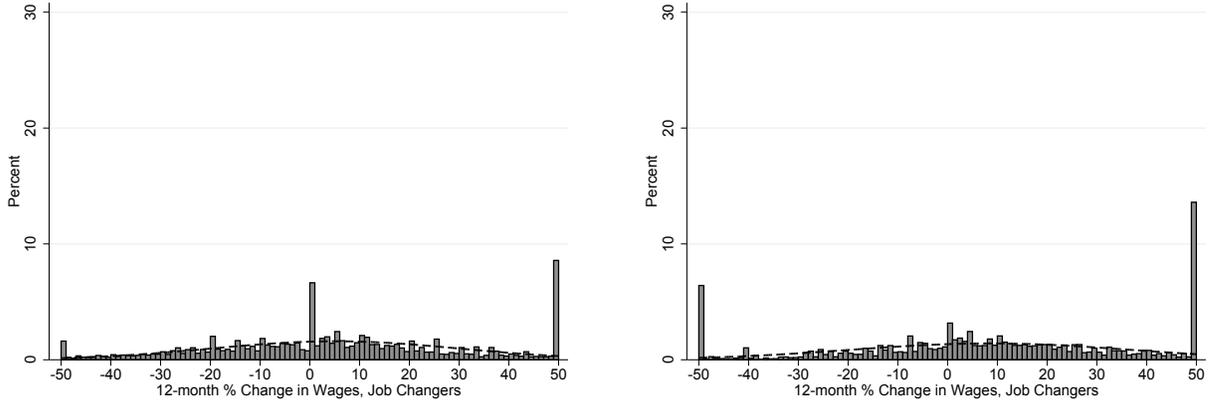
Finally, the choice of timing is more nuanced given the nature of our data. As with job-stayers, we can measure wage changes for job-changers at one-month, one-quarter, and one-year frequencies. However, when we see a worker at firm i in month t and then see a worker at firm j in month $t + 12$, the 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 important to keep such timing issues in mind when interpreting our wage adjustment measures for job-changers.¹⁷

Figure 10 plots the distribution of annual nominal wage changes for a sample of job-changers. The patterns are strikingly different from the patterns in Figure 2. First, essentially all workers receive a nominal wage change over a given year if they change jobs. Only about 6 percent of hourly job-changers and 3 percent of salaried job-changers do not receive a year-over-year nominal wage change. Second, the propensity for a nominal wage cut is very high for job-changers. Specifically, one-third of workers receive a nominal wage decline during a job-change. Finally, the distribution of nominal wage changes is much more symmetric around zero. As seen from the figure, there are roughly as many small nominal wage increases (0-2 percent) as there are slightly larger nominal wage increases (2-4 percent).

¹⁶We have experimented with a 10 percent cutoff and a 30 percent cutoff instead of a 20 percent cutoff. Our results are robust across these three cutoff levels.

¹⁷When measuring “quarterly” nominal wage changes for job-changers, we include workers who show up in another ADP firm between 1 and 5 months after leaving their original ADP firm. When measuring “annual” nominal wage changes for job-changers, we include workers who show up in another ADP firm between 10 and 14 months after leaving their original ADP firm. We create wider bins to include more job-changers in our analysis. When we examine patterns by year later in the paper, having larger overall samples is helpful for power reasons. Additionally, we restrict our analysis 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.

Figure 10: Annual Wage Change Distribution for Job Changers



PANEL A: HOURLY-TO-HOURLY CHANGERS

PANEL B: SALARIED CHANGERS

Note: Figure shows the annual change in nominal wages for workers in our employee sample who switched jobs. We include all data from 2008-2016. See text for additional details.

There is much more nominal wage flexibility for job-changers than there is for job-stayers. Table 7 summarizes the propensity for job-changers to receive a nominal wage increase and a nominal wage decline at different time horizons.

Table 8 shows key statistics on the distribution of wage changes for job changers. Conditional on a job change and a nominal wage change, mean and median annual nominal wage growth was 9.1 and 6.3 percent accordingly. Nominal wage growth is much larger for job-changers than it is for job stayers. As seen from Figure 10, there is a tremendous amount of heterogeneity in nominal wage changes for job changers. Job-changers whose nominal wage increased experienced over the year, on average, about a 27 percent increase. Job-changers whose nominal wage fell over the year experienced, on average, a 20 percent wage cut. Moreover, the standard deviation of annual nominal wage changes for job-changers is 30 percent - almost five times larger than the standard deviation of annual nominal wage changes for job stayers. Job changing provides much more nominal wage flexibility than remaining on the same job.¹⁸

The moments presented in this section illustrate the importance of not simply measuring the rigidity of job-stayers. Indeed, for many applications, the relevant notion of rigidity may not include job-stayers at all. Models in which most wage increases arise from movements across firms or due to the arrival of outside offer, such as many labor search models (Menzio and Shi (2011); Cahuc et al. (2006)), may be better calibrated to the moments presented

¹⁸Faberman and Justiniano (2015) documents a tight correlation between aggregate nominal wage growth and aggregate job-switching rates.

Table 7: Probability of Wage Change, Pooled 2008-2016 Sample of Job Changers

	Quarterly	Annual
<u>All Workers</u>		
Probability of Positive Wage Change (%)	52.7	56.8
Probability of Negative Wage Change (%)	37.6	38.0
<u>Hourly Workers</u>		
Probability of Positive Wage Change (%)	51.1	55.1
Probability of Negative Wage Change (%)	38.5	39.1
<u>Salaried Workers</u>		
Probability of Positive Wage Change (%)	60.0	63.0
Probability of Negative Wage Change (%)	33.1	34.2

Note: Table shows the frequency of wage increases and wage decreases at different horizons for our sample of job-changers during the 2008-2016 period. The top panel pools together hourly and salaried workers while the middle and bottom panels, respectively, show the frequency of changes for hourly and salaried workers separately. The first and second columns show results at the quarterly and annual horizons, respectively. We use our employee sample for this analysis. See text for additional discussion of our job-changer sample.

in Tables 7 and 8 than those presented in Section 4 or in much of the existing literature on wage stickiness.

5.2 Combining Job-Stayers and Job-Changers

We now combine our measures of wage rigidity into an aggregate measure of nominal rigidity. Aggregate nominal wage flexibility is a function of the wage flexibility of job-changers and the flexibility of job-stayers.¹⁹ 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 models such as Christiano et al. (2015), 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. As we show below, the inclusion of job-switchers vastly reduces the degree of realized nominal wage rigidity in the economy, particularly on the downside, relative to the job-stayer benchmark which has been measured in the literature to-date.

To construct an aggregate measure of nominal wage flexibility, one must combine the

¹⁹New 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 8: Wage Change Statistics, Pooled 2008-2016 Sample of Job-Changers

	Quarterly	Annual
<u>Unconditional</u>		
Mean Wage Change (%)	6.5	7.9
Median Wage Change (%)	2.2	4.5
Standard Deviation of Wage Change (%)	27.0	30.4
Skewness of Wage Changes (%)	1.0	0.9
Kurtosis of Wage Changes (%)	5.3	4.9
<u>Conditional on Any Wage Change</u>		
Mean Wage Change (%)	7.6	9.1
Median Wage Change (%)	5.4	6.3
Standard Deviation of Wage Change (%)	28.2	30.8
Skewness of Wage Changes (%)	0.9	1.0
Kurtosis of Wage Changes (%)	4.7	5.1
<u>Conditional on Positive Wage Change</u>		
Mean Wage Change (%)	24.6	27.0
Median Wage Change (%)	17.7	20.0
Standard Deviation of Wage Change (%)	23.0	24.8
<u>Conditional on Negative Wage Change</u>		
Mean Wage Change (%)	-17.2	-19.7
Median Wage Change (%)	-14.3	-16.7
Standard Deviation of Wage Change (%)	12.8	14.4

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

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.²⁰ 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). In particular, we use J2J’s job flow measures allowing for a short non-employment spell between jobs. Additionally, the J2J data focuses on a given worker’s main job. For any given worker in quarter t whose main job is at firm i , the J2J data measures whether the worker’s main job in quarter $t + 1$ remained at firm i (job-stayers), whether the worker’s main job in quarter $t + 1$ was at a different firm j (job-changers), and whether the worker was not employed in quarter $t + 1$ (become non-employed). The sum of these three measures sum to 1 within each quarter. Using data from 2008 through 2016, the quarterly job staying rate average 88.7 percent, the quarterly job switching rate averaged 4.6 percent, and the quarterly become non-employed rate was 6.9 percent. As of now, 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.²¹ 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 ($0.046/(1-0.046)$). 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 9 shows statistics for the aggregate nominal wage change distribution combining data from both job-stayers and job-changers. there is much more aggregate nominal wage flexibility than one would conclude from looking at job-stayers alone. Over the entire sample period, roughly 73 percent of workers receive a nominal wage change. Of those, nearly 10 percent received nominal wage declines, compared with 2 percent of job-stayers. While nominal wage declines are still rare in the aggregate relative to nominal wage increases, including data on job-changers quadruples the amount of nominal wage cuts relative to looking at only job-stayers. Moreover, the standard deviation of wage growth – both unconditionally and

²⁰See https://lehd.ces.census.gov/data/j2j_beta.html, accessed June 30, 2018.

²¹This 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.

conditional on a wage change – is over twice as large in the aggregate compared with only looking at data on wage changes. For example, unconditionally, the standard deviation of nominal wage growth in the aggregate is 13.6 percent while the standard deviation of wage growth for job stayers was about 6 percent.

Overall, the inclusion of job-changers in our measures of wage rigidity greatly increases realized flexibility in the economy. This has important consequences for the quantitative predictions of existing macro models. The excessive price rigidity inferred by simply considering wage adjustment for job-stayers will lead New Keynesians to overstate the passthrough of monetary policy to real quantities. Similarly, those studying the extent to which downward nominal rigidities could contribute to sluggish wage growth should be aware that roughly 10% of workers received an annual wage cut from the period 2008-2016.²² Even if wages for job-stayers appear exceptionally downwardly rigid, aggregate wage levels have not been as inflexible over the past ten years, owing to the high churn of employment in the US economy. We now turn to a discussion of wage rigidity over the cycle, and the extent to which wages may respond to real economic shocks.

6 State Dependence in Aggregate 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 parameter for the probability that a worker receives a wage adjustment. Using a variety of methodologies, we highlight that the probability of a wage adjustment varies substantively with business cycle conditions. We decompose the cyclical variation in aggregate wage adjustments into a component due to wage adjustments for job-stayers, a component due to wage adjustments for job-changers as well as accounting for the cyclical variation in the ratio of job-stayers to job-changers. We also explore the cyclical variation in bonus receipt.

²²This explanation for slow wage growth is at the front of policymakers minds. For instance, in her 2014 Jackson Hole Symposium, then-Chairwoman of the Federal Reserve Janet Yellen stated that The sluggish pace of nominal and real wage growth in recent years may reflect the phenomenon of ‘pent-up wage deflation.’ The evidence suggests that many firms face significant constraints in lowering compensation during the recession and the earlier part of the recover because of downward nominal wage rigidity – namely, an inability or unwillingness on the part of firms to cut nominal wages.

Table 9: Probability of Aggregate Wage Change Combining Job Stayers and Job Changers, Pooled 2008-2016

	Quarterly	Annual
<u>Probability of Wage Change</u>		
Probability of Wage Change (%)	26.0	72.7
Probability of Positive Wage Change (%)	21.8	62.8
Probability of Negative Wage Change (%)	4.1	9.9
<u>Unconditional</u>		
Mean Wage Change (%)	1.4	4.6
Median Wage Change (%)	0.00	2.5
Standard Deviation of Wage Change (%)	8.1	13.6
Skewness of Wage Changes (%)	3.6	1.5
Kurtosis of Wage Changes (%)	36.9	11.4
<u>Conditional on Any Wage Change</u>		
Mean Wage Change (%)	5.4	6.4
Median Wage Change (%)	3.1	3.5
Standard Deviation of Wage Change (%)	15.3	15.7
Skewness of Wage Changes (%)	1.3	1.1
Kurtosis of Wage Changes (%)	9.5	8.5
<u>Conditional on Positive Wage Change</u>		
Mean Wage Change (%)	9.1	9.8
Median Wage Change (%)	4.0	4.1
Standard Deviation of Wage Change (%)	12.8	13.1
<u>Conditional on Negative Wage Change</u>		
Mean Wage Change (%)	-15.4	-16.6
Median Wage Change (%)	-12.7	-13.9
Standard Deviation of Wage Change (%)	11.0	12.0

Note: Table shows aggregate moments of wage adjustment 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. The first column shows results at the quarterly horizon while the second column shows the results at the annual horizons. We use our employee sample for this analysis. See text for additional discussion of our job-changer sample. We use data from the LBD to compute the weights for job-changers and job-stayers. See the text for additional details.

6.1 Time Series Variation in the Nominal Wage Adjustments, Job Stayers

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 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. The presence of fixed adjustment costs generates an inaction region in which firms that are close to their optimal price in a frictionless economy will not adjust their prices until they move sufficiently far away from their optimal price in a frictionless economy. Thus with a menu cost of adjusting prices the state of the firm - its distance from the optimal pricing rule - is crucial in determining price adjustment decisions. As a result, price changes are infrequent, and relatively 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 11 plots the time series of wage adjustments for job-stayers pooling together both hourly and salaried workers. The top panel plots the extensive margin of wage changes: the percent of all employees in month t who have a different 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 11 is for May 2009 and measures the fraction of job stayers who received a 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 wage adjustments.

As seen from the left panel of Figure 11A, wage adjustments of job-stayers exhibits 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, between 65 and 70 percent of workers received a wage change. While most of the time series variation was between the recession and non-recessionary periods, there is still a slight trend upwards in the share of workers receiving an annual wage change between

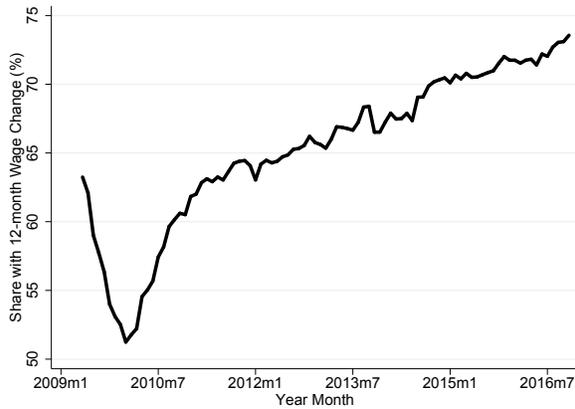
2012 and 2016. The top right panel of Figure 11 separates the probability of a wage change of job-stayers into the probability of a wage increase (solid line) and the probability of a wage declines(dashed line). During the Great Recession, the propensity of wage increases for job-stayers fall sharply and the propensity of wage declines increases sharply. One of the key findings of the paper is that while nominal wage cuts are exceedingly rare for job-stayers during non-recessionary periods, nearly 6 percent of all continuing workers received a nominal wage cut during late 2009 and early 2010.

The bottom panel of Figure 11 plots the mean size of wage changes, restricting attention to those who have indeed received a wage change in the prior 12 months. The bottom left panel pools together all wage changes while the bottom right panel separately looks at wage increases and wage declines. The overall mean wage change size is highly pro-cyclical. Strikingly, however, the mean size of wage *cuts* does not move much over the cycle. The large fall in conditional mean nominal wage change for job-stayers during the Great Recession is mostly driven by an increased share of wage changes being negative, rather than changes in the size of cuts themselves.

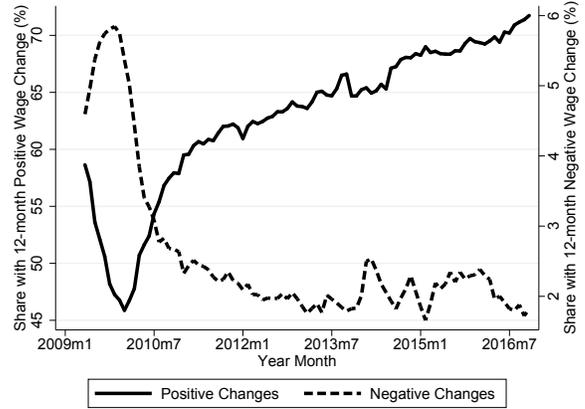
Table 10 summarizes the business cycle differences in nominal wage adjustments for job-stayers. We separate the sample into two periods: a period representing the depths of the Great Recession (May 2009-December 2010) and a period well into the recovery (January 2012 - December 2016). A few things are of note from the summary table. First, 6.6 percent of salaried workers who remained on their job received a nominal wage cut during the Great Recession. During the post-recession period, only 2.8 percent of salaried workers received a nominal wage cut. Again, for salaried workers, wages are much less downwardly rigid during the Great Recession. The propensity of a nominal wage cut during the Great Recession for hourly workers was twice as high as it was during the post-2012 period. While wages were more downwardly flexible during the Great Recession for job stayers, the fraction of workers receiving a zero nominal wage change increased for both salaried and hourly workers. Interestingly, the unconditional standard deviation of wage changes fell slightly during the recession for all workers from 7.0 percent to 6.3 percent. In summary, overall nominal wage adjustments fell during the Great Recession but downward adjustments increased.

Figure 12 shows time series trends in the probability of a nominal wage increase (left panel) and a nominal 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 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

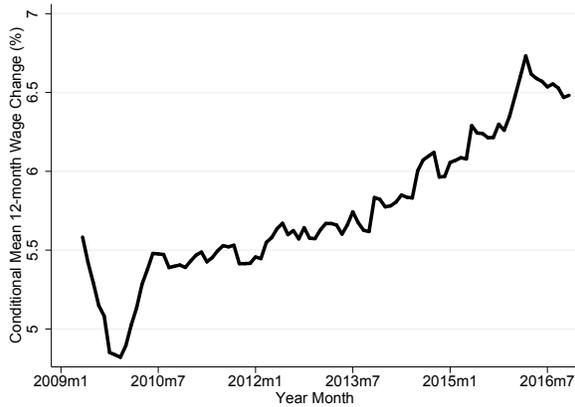
Figure 11: Time Series of Nominal Wage Adjustments, Job Stayers



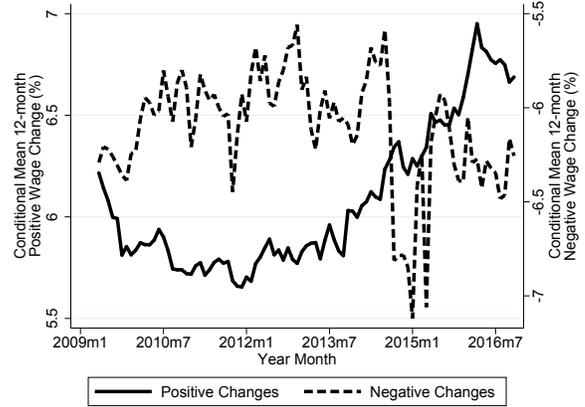
PANEL A: HAS WAGE CHANGE



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



PANEL C: MEAN WAGE CHANGE SIZE



PANEL D: MEAN WAGE CHANGE SIZE - POS. VS NEG.

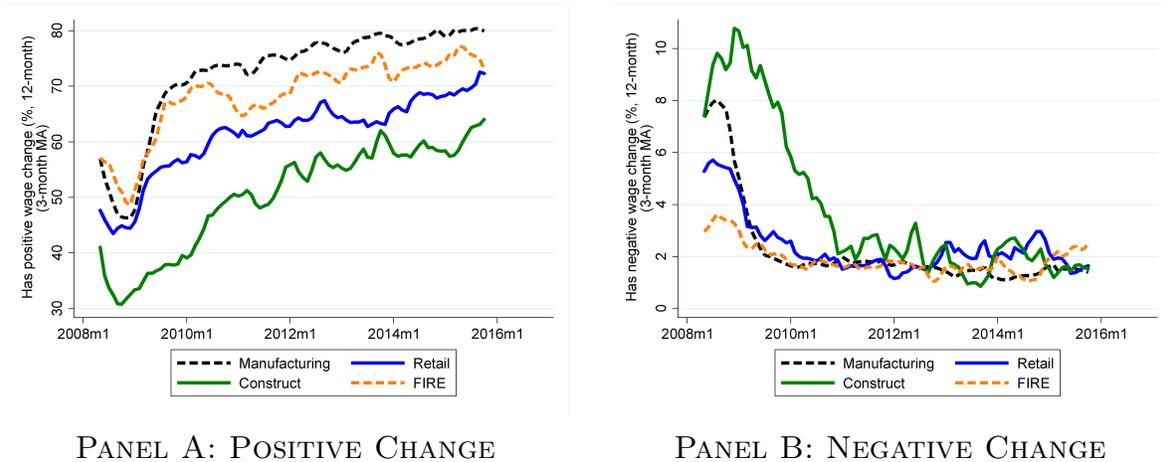
Note: Figure plots the propensity to receive a 12-month wage change (Panels A and B) and the mean size of wage changes (Panels C and D) over time for our employee sample of job-stayers between May 2009 and December 2016. The data are weighted to match the firm size \times industry mix found in the BDS. Since the first month of our ADP data is May 2008, we may only observe 12-month wage changes beginning in May of 2009.

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

	% Decline	% Zero	% 0-5	% 5-10	% 10+	Mean Cut	Mean Raise	S.D. of Change
May 2009 - December 2010								
Hourly	2.8	42.1	37.0	9.8	8.3	-5.7	5.6	6.1
Salaried	6.6	45.2	29.1	9.6	9.6	-6.5	6.5	7.4
All	4.2	43.3	34.0	9.7	8.8	-6.2	5.9	6.3
January 2012 - December 2016								
Hourly	1.5	29.8	44.1	11.7	12.9	-5.8	6.2	6.8
Salaried	2.8	31.9	40.0	11.7	13.6	-6.2	6.7	7.2
All	2.0	30.6	42.58	11.7	13.2	-6.0	6.4	7.0

Note: Table shows the distribution of 12-month wage adjustment for job-stayers over the cycle. The top panel shows the distribution for workers during the recession - from May 2009 through December 2010 - while the bottom panel plots the distribution of wage changes during the recover period, defined here as January 2012 through December 2016. This table makes use of our employee sample, weighted to match the firm size \times industry mix found in the BDS. Since the first month of our ADP data is May 2008, we may only observe 12-month wage changes beginning in May of 2009.

Figure 12: Time Series of Wage Changes by Industry, Job Stayers



Note: Table 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. Since the first month of our ADP data is May 2008, we may only observe 12-month wage changes beginning in May of 2009. "FIRE" refers to Finance, Insurance, and Real Estate.

a nominal wage cut. Note, the probability of a nominal wage increase did not differ markedly across industries during the Great Recession. As highlighted above, there are level differences in the propensity of a nominal wage increase across industries for job-stayers. 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 wage cuts of continuing workers. Not only were nominal wage cuts more likely for job-stayers during the Great Recession, the propensity of nominal 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 wages of the workers who remained with their firm. Below we show that the individual firms that reduced their employment the most during the Great Recession were also much more likely to cut the wages of their workers.

The results in this subsection show that the composition of nominal wage adjustment for job-stayers varies over the business cycle. Any model that assumes a constant hazard of 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.

6.2 Time Series Variation in Aggregate Nominal Wage Adjustments

As highlighted above, much of the flexibility in nominal wage adjustments results from job-changers. Table 11 shows that the distribution of wage adjustments for job-changers also varies over the business cycle. For the table, we report statistics of 12 month wage changes for workers employed at firm i in t and then are subsequently employed at firm j in $t + 12$.²³

During the Great Recession, 45 percent of hourly workers and 56 percent of salaried workers who changed jobs received a nominal wage cut. The comparable numbers during the recovery period were 38 percent and 32 percent, respectively. There was much more downward wage adjustment for job changers during the Great Recession. During the Great Recession, the distribution of wage changes for job-changers essentially shifted left with the unconditional standard deviation of nominal wage changes remaining roughly constant. The job-switchers are providing even more downward flexibility during the Great Recession.

Table 11: Summary of 12-month Wage Change Distribution During and After the Great Recession, Job-Changers

	% Decline	% Zero	% 0-5	% 5-10	% 10+	Mean Cut	Mean Raise	S.D. of Change
May 2009 - December 2010								
Hourly	44.8	7.3	5.2	7.4	35.3	-19.6	28.9	31.7
Salaried	56.3	3.1	4.5	6.8	29.3	-20.0	28.6	32.5
All	47.2	6.4	5.0	7.3	34.1	-19.7	28.8	31.9
Jan 2012 - Dec 2016								
Hourly	38.4	5.6	7.3	7.8	40.9	-18.8	26.5	30.8
Salaried	31.7	2.8	8.7	6.3	50.4	-24.0	29.3	33.9
All	37.0	5.0	8.6	7.5	43.0	-19.7	27.2	31.5

Notes:

Aggregate nominal wage flexibility is a function of both the wage adjustments for job-stayers and job-changers. However, in order to measure the cyclical nature of aggregate wage adjustments, we also need to know how the composition of job-stayers relative to job-switchers evolves at business cycle frequencies. As discussed above, we cannot measure an

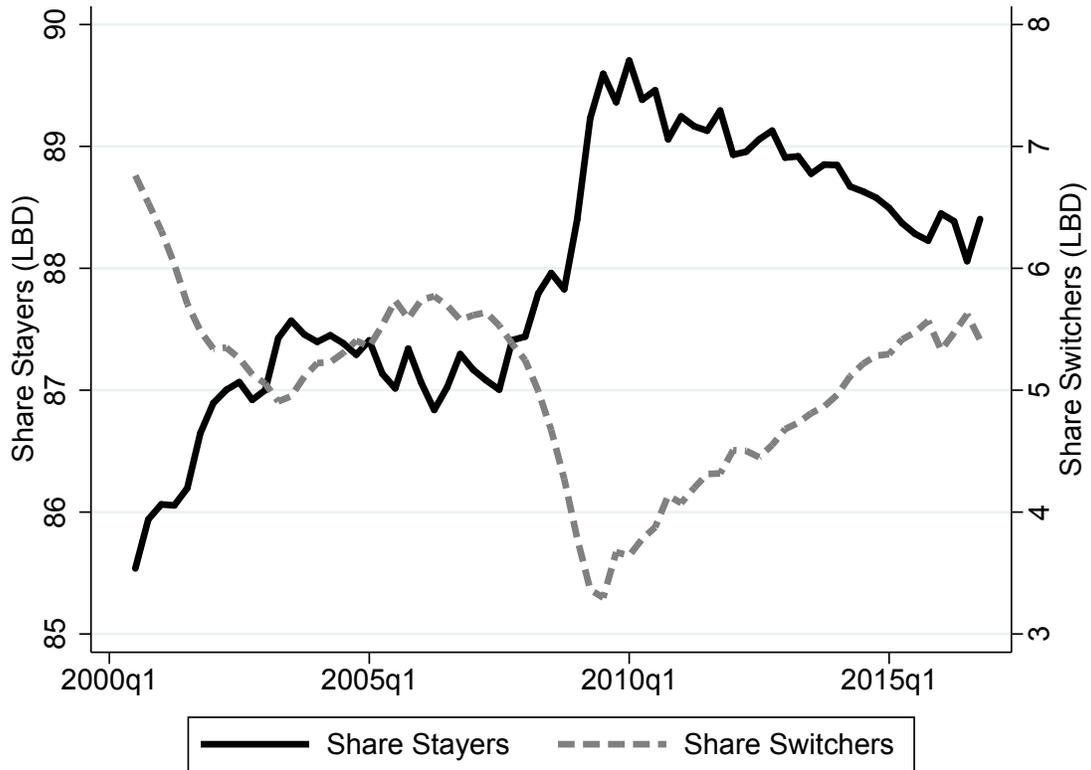
²³In the Online Appendix, we also show quarterly changes in the wages of job changers. The quarterly wage changes display very similar cyclical patterns. In the work below, we compute aggregate wage adjustments - combining job-stayers and job-changers - at both the quarterly and aggregate frequencies.

aggregate amount of job-switching within the ADP data because we can only track worker movements from one ADP firm to another ADP firm. We again use data from the Census's Job-to-Job Flow Data (J2J) made from the underlying data of the LEHD. Figure 13 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 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.

Table 12 shows the cyclical patterns of aggregate 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 wage adjustments, 12 percent of workers received a nominal wage decline during the Great Recession. While downward adjustments were more common, upward adjustments were less common with only half of workers receiving a wage increase during the 2009-2010 period. This is due to the composition effect noted above. The unconditional standard deviation of 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 12 percent of workers did receive nominal wage cuts and another 40 percent received no nominal wage increase. Much of the downward adjustments occur through job changers. Moreover, both job-stayers and job-changers experienced more nominal 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 wage adjustments. The results in Table 12 provide a set of moments for researchers to calibrate models to assess whether the moments of the wage adjustment distribution can lead to sufficient rigidities to explain why aggregate wage growth did not fall more during the 2008-2012 period. Again, we stress the importance of considering measures

Figure 13: Time Series of the Share of Workers who are Stayers versus Switchers, LBD 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 LBD's Job-to-Job (J2J) flows data over the period 2000Q1 through 2016Q4.

Table 12: Aggregate Nominal Wage Adjustments Combining Job Stayers and Job Changers, 2009-2010 vs. 2012-2016

	Quarterly		Annual	
	March 09- Dec. 10	Jan. 12- Dec. 16	March 09- Dec. 10	Jan. 12 - Dec. 16
<u>Probability of Wage Change</u>				
Share Positive Wage Change (%)	17.7	23.5	51.2	66.3
Share Negative Wage Change (%)	5.1	3.9	11.8	9.7
<u>Unconditional Size of Wage Change</u>				
Mean Wage Change(%)	0.7	1.6	2.7	5.2
Median Wage Change(%)	0.0	0.0	1.0	2.8
Stan. Deviation of Wage Change (%)	8.1	8.2	12.7	14.2
<u>Conditional Size of Any Wage Change</u>				
Mean Wage Change(%)	3.3	6.0	4.4	6.9
Median Wage Change(%)	3.0	3.3	3.2	3.5
Stan. Deviation of Wage Change (%)	16.8	15.0	15.8	16.0

Note: Table shows aggregate moments of wage adjustment combining data on both job-stayers and job-changers during the March 2009-December 2010 period and from the January 2012-December 2016 period. For this table, we pool together both hourly and salaried workers. The first two columns shows results at the quarterly horizon while the second two columns shows the results at the annual horizons. We use our employee sample for this analysis. See text for additional discussion of our job-changer sample. We use data from the LBD to compute the weights for job-changers and job-stayers. In addition, we weight to match the firm size \times industry mix of employment found in the BDS. Weights for job-changers are assigned to match the destination firm's size and industry. See the text for additional details.

of *aggregate* wage flexibility when assessing such claims: both the level and trend of rigidity implied by the job-stayer sample is substantially different than those found in the aggregate economy.

6.3 Time Series Variation in Bonuses

We have also explored the time series patterns in bonuses. The propensity to receive a bonus and the conditional size of a bonus, for the most part, did not vary over the 2009-2016 period within the ADP data. For example, there was only a very slight decline in the propensity for workers to receive an annual bonus during the Great Recession relative to the 2012-2016 period (26 vs 27 percent). Mean and median conditional bonus size were only about one hundred dollars less in 2009-2010 relative to other years. These findings are consistent with news reports that documented that many workers in finance still received large bonuses during the Great Recession.²⁴ Given that the bonus patterns did not vary substantively in the ADP data, we have relegated some of these patterns to the Online Appendix. However, we are continuing to explore these patterns as the paper is evolving.

7 Regional Variation and Within Firm Variation in Nominal Wage Adjustments

To further shed light on the extent to which nominal wage adjustments respond to local shocks, we exploit cross-region variation and explore patterns by growing and contracting firms. In both cases, we see that there is some state dependence in nominal wage adjustments.

7.1 Regional Variation in Nominal Wage Adjustments During the Great Recession

Although the time series variation is suggestive of state dependence in wage adjustment, it is by no means conclusive. One may easily conjure models without state dependence in wage adjustment that could generate some of the patterns observed above. Suppose, for instance, that wages are set according to a Calvo adjustment process whose aggregate Calvo parameter evolves according to some stochastic process. One could rationalize the patterns observed in Figure 11 and Table 10 in such a model if the Calvo probability of receiving a wage change randomly draws a low value at the beginning of 2009. Wages would adjust

²⁴See, for example, “Bankers Reaped Lavish Bonuses During Bailouts” in the Wall Street Journal, July 30th 2009.

less on average, not because wage adjustment is state dependent *per se*, but because the stochastic aggregate state was such that adjustment is randomly more difficult.

To address this potential identification concern, we use regional variation to difference out aggregate shocks to wage stickiness. Specifically, we use the decline in local house prices between 2007 and 2010 as our measure of the severity of the recession.²⁵ We are not interested in a causal channel of house prices on the nature of worker wage adjustment. Instead, we are using house price declines as a potential shock to local labor demand. We are, however, interested in how wage setting varies across regions that experienced more and less severe recessions.²⁶ Our goal is to determine whether realized wage adjustment responds to local economic conditions.

Figure 14 shows a simple scatter plot relationship between local house price declines at the CBSA level and the probability of a job-stayer receiving a wage increase (left panel) and the probability of a job-stayer receiving a wage decline (right panel). CBSA's with the largest declines in house prices were much less likely to raise nominal wage of job-stayers and were much more likely to cut nominal wages of job-stayers during the Great Recession. Unconditionally, a 30 percent decline in house prices was associated with a 6.8 percentage point decline in the probability of increasing wages and a 1.4 percentage point increase in the probability of a nominal wage cut. The base probability of increasing and decreasing wages during this time period was 51.6 percent and 5.6 percent, respectively. These cross region patterns match well the aggregate time series patterns discussed above.

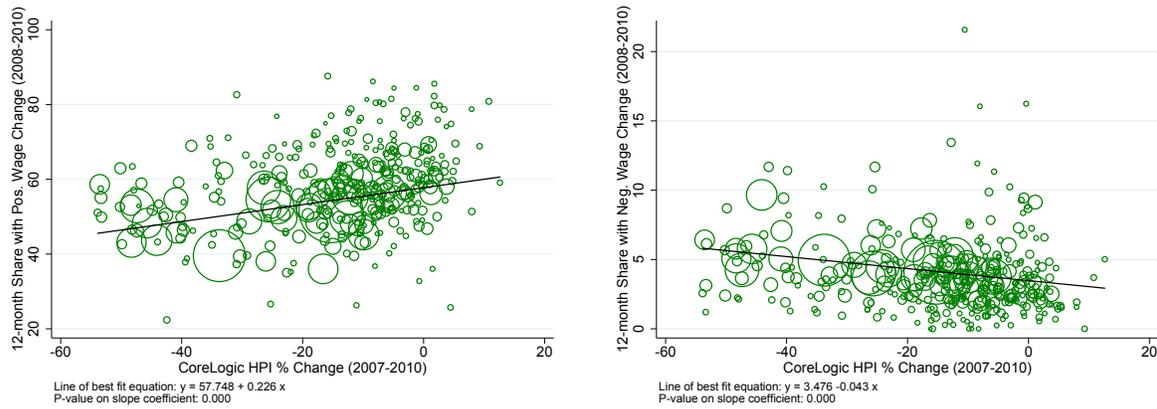
7.2 Firm Level Employment Growth and Wage Adjustment for Job-Stayers

Although suggestive of state dependence in wage setting, the regional evidence provided above does not precisely emulate the decision problem faced by an individual firm. We therefore 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 reduced form and combining two potential effects. First, firms who expe-

²⁵There is a large literature showing that housing price declines explain much of the regional variation in economic activity during the Great Recession. See, for example, Mian et al. (2013).

²⁶Our results are similar if we use changes in the local unemployment rate as our measure of the severity of the local recession. We prefer housing price declines as a potentially more exogenous measure of local economic activity with respect to wage changes given that the potential for local wage flexibility to affect local employment changes. Using employment changes as a measure of recession severity, for instance, might pose a reverse causality problem if firms which are most unable to cut wages are most likely to lay off employees. While this may be a valid concern *ex-ante*, the fact that our results are so similar when we use house price declines and changes in the unemployment rate as our measure of recession suggests that this concern is not warranted during this period.

Figure 14: Probability of 12 Month Wage Changes by Local House Price Decline, Job-Stayers



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

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

Note: Figure plots the relationship between a CBSA's CoreLogic Home Price Index (HPI) change from 2007 to 2010, and the probability of job-stayers receiving a year-over-year wage increase (Panel A) or decrease (Panel B). When collapsing to the CBSA-level propensity to receive a wage change, we weight to match the aggregate firm size \times industry mix found in the BDS. Regression equation for the line of best fit reported beneath each plot. Scatters and regressions weighted by 2007 CBSA population estimates from the Census Bureau.

rience negative employment growth should have been more likely to receive a negative firm level shock. Firms receiving negative shocks are likely to adjust their labor inputs on multiple margins. As a result, firms receiving 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 is some trade-off between wage adjustments and employment adjustments within a firm conditional receiving a shock of a given size. For example, a firm who receives a given size shock may choose to cut more of the wages of their workers thereby reducing the number of workers they have to layoff. Such differences in the extent to which firms are willing to cut the wages of their workers conditional on a given sized shock could result in a negative relationship between firm employment growth and firm propensity to cut nominal wages.

Figure 15 plots the propensity for a firm to increase their workers wages (left panel) and the propensity for a firm to cut their workers wages (right panel) as a function of observed firm size growth. To make this figure, we use our firm level sample. All firm level changes are at the annual level. As a result, the x-axis is the observed change in firm employment during a 12 month period while the y-axis is the propensity to adjust wages (up or down) during a 12 month period. Each panel has two lines, reflecting LOWESS smoothed regression lines. The LOWESS estimator provides a non-parametric estimate of the conditional expectation function of the propensity to receive a wage change conditional on the firm's growth rate.

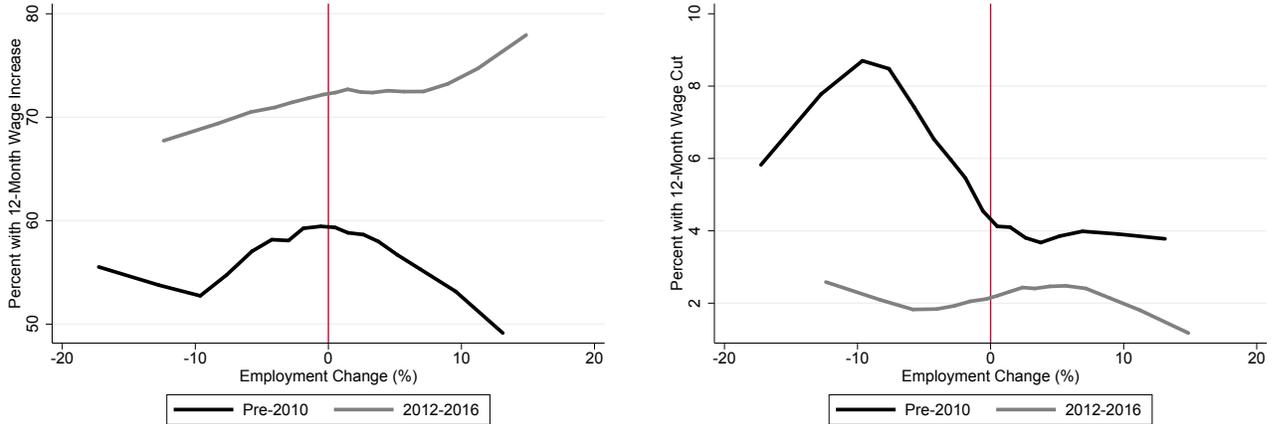
The darker line restricts our firm samples only to firm behavior during the Great Recession (2009-2010) while the lighter line restricts the sample to firm behavior during the recovery (2012-2016).

While not overly surprising, firms that grew were slightly more likely than firms that shrunk 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. Additionally, contracting firms were again slightly less likely to raise their workers wages relative to firms whose size remained constant. However, growing firms not that different in their propensity relative to contracting firms.

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 15 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 rigidity, and countercyclical upward rigidity. Further research is required to assess the impact of idiosyncratic firm shocks on aggregate wage movements over the cycle.

Figure 15: Probability of Wage Changes by Firm Growth Status



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

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

Note: Figure 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.

8 Discussion

Our work contributes to a growing literature examining nominal wage stickiness.²⁷ Across most papers, a consensus has emerged that most job-stayers experience a nominal wage change during a given year and that wages are more downward rigid than upward rigid for job-stayers. However, the differences across papers surround magnitudes. These differences stem from the fact that measuring nominal wage changes in household and administrative datasets is notoriously difficult given the presence of measurement error (household surveys) or missing hours (administrative datasets). That same measurement error also makes it difficult to track the wage changes of job-changers. In this section, we discuss how our work relates the existing literature measuring nominal wage adjustments using either household data or administrative data. We end with a discussion of how we can use our data to compare variation in quarterly earnings-per-hour to the variation we have highlighted through the paper in per-period contract rates. We show that variation in earnings per hour can over

²⁷There 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 Kamlani (1997), and Bewley (1999).

inflate variation in nominal wages (or compensation more broadly) given the coarseness of the data.

It is worth emphasizing that the findings in our paper for job-stayers are dramatically different than essentially all of the papers that use household data to measure nominal wage adjustments. These papers all define nominal wages as self-reported earnings divided by self-reported hours worked. For example, using the panel component of the CPS, Daly and Hobijn (2014) report that roughly 85 percent of wages of job-stayers change annually during this time.²⁸ As noted above, we find that only about two-thirds of job-stayers receive an annual nominal wage change during the 2008-2016 period. Using the methodology of Daly and Hobijn (2014), the San Francisco Federal Reserve has created a “Wage Rigidity Meter”. They report higher wage flexibility for salaried workers relative to hourly workers. Again, this finding is inconsistent with the findings in our paper. But, 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. Additionally, the fact that hours are likely measured with more error for salaried workers would generate more measured wage volatility for salaried workers relative to hourly workers in household surveys.

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 wage series. When they make their correction, the frequency of quarterly wage changes for job stayers falls from over 50 percent to about 15 percent. Our quarterly frequency of wage changes for job stayers is about 20 percent. Their correction seems to over-correct for the frequency of wage changes. Despite lower flexibility overall, they also 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. As discussed in Section 4, we estimate only 4.6 percent of all quarterly wage changes for job stayers are downward changes (0.9/19.4). Finally, unlike our results, they find no differences in wage change probabilities across occupations and industries. The differences between our results and Barattieri et al. (2014) are consistent with a non-trivial amount of residual measurement error remaining even after their procedure to remove measurement error from their wage adjustment measures.

A more recent literature has emerged using firm level data to measure wage stickiness.²⁹

²⁸One 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.

²⁹There are many additional studies examining the nature of earnings dynamics using high-quality administrative data. See, for example, Postel-Vinay and Robin (2002), Bonhomme and Robin (2010), and Guvenen et al. (2014).

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 the 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 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.

Kurmann and McEntarfer (2017) use data from the US Longitudinal Employer Household Dynamics (LEHD) 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. The hours measures are reported at the quarterly level and are self reported by the firm administrator filling in the unemployment insurance records. For salaried workers, the reported hours worked are often but not always set to 40 hours per week. Le Bihan et al. (2012) and Sigurdsson and Sigurdardottir (2016) also use administrative establishment level data from France and Iceland, respectively, to measure nominal wage rigidities. Both datasets have administrative measures of earnings for a given employee. Also, both datasets have some measure of firm reported hours for their employees. To make hourly wages, they divide administrative earnings by an administrative report of hours. The extent to which reported hours are measured with error for salaried workers will artificially inflate the variance of nominal wage changes using administrative earnings data. As a result, most of these papers find larger amounts of nominal wage adjustments for job stayers relative to what we find with the ADP data.³⁰

Our paper is closest in spirit to Altonji and Devereux (2000), who use administrative payroll data similar to ours for one large financial service company during 1996 and 1997. The patterns of wage adjustment for job-stayers they document within this large financial company closely match the patterns we document for the whole U.S. economy during the 2008-2016 period. In particular, they document that only 0.5% of workers receive a nominal wage cut during a given year. The fact that payroll data from an earlier period broadly matches the job-stayer results from the recent ADP data highlights the importance of using administrative payroll data to measure wage stickiness.

To illustrate the point that measurement error in hours worked can significantly alter patterns of wage flexibility, we treat our data similarly to that of other administrative data

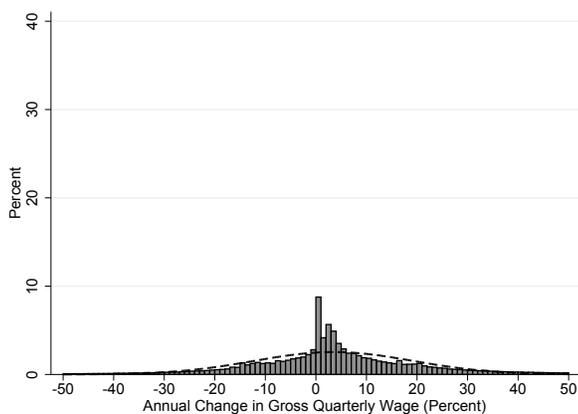
³⁰Sigurdsson and Sigurdardottir (2016) also find evidence of state-dependence in nominal wage adjustments of job-stayers in Iceland.

sources. In particular, we forgo the pay-rate level variation in our data and aggregate total earnings of a worker to the quarterly level for our sample of job-stayers. Panel A of Figure 16 plots the distribution of wage changes that would be inferred if we used our administrative payroll data to carry out a similar procedure to that done in administrative records that track worker quarterly earnings without adjusting for hours. For our quarterly earnings measures, we compare changes for a worker who work at firm i in a given quarter in year t to the same worker who continues to work at firm i during the same quarter in year $t + 1$. So, these are annual changes in quarterly earnings measures for a sample of job-stayers. Panel A plots the distribution of quarterly earnings changes, not correcting for variation in hours. Unsurprisingly, failing to account for variation in hours leads to a far more dispersed earnings change distribution than is shown by the wage change distribution in Figure 2. Using this measure, 32.2 percent of job-stayers receive a nominal wage cut during a given period. Moreover, the unconditional standard deviation of nominal earning adjustments is 20.0 percent.

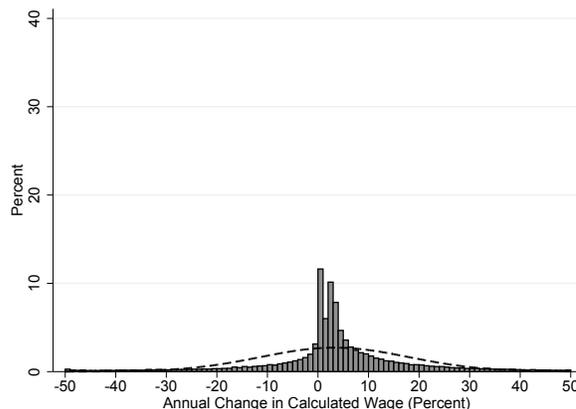
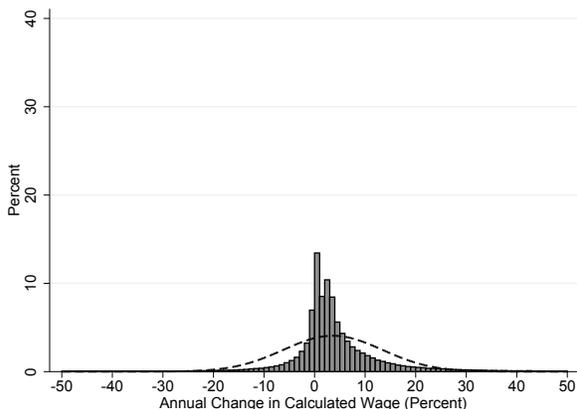
To partially address the problem of not accounting for hours, Panels B and C plots the distribution of changes in a worker's imputed quarterly earnings per hour separately for hourly workers and salaried workers, which are calculated by dividing a worker's quarterly earnings by the hours worked as reported in their payroll records. For hourly workers, this is the actual hours they worked. For salaried workers, this is often set at 40 hours. Even with an hours measure which is as high quality as can be reasonably expected - the number of hours reported on a worker's paycheck - the figure shows a substantially more dispersed distribution of annual wage changes than is accurate. Although we observe that 10-15% of workers receive no imputed earnings per hour change over a given year, this is much smaller than the 35% of workers who receive no actual wage change during a quarter as discussed in Table 5. Moreover, with this earnings per hour measure, 21.2 percent of wage workers and 25.3 percent of salaried workers receive a nominal wage cut during the year over our sample period. Recall, the comparable number using our detailed nominal wage measures is about 2.5 percent per year. However, we wish to stress that controlling for hours reduces the unconditional standard deviation of nominal wage adjustments to 19.2 percent, and 15.9 percent for hourly workers.

Why do nominal earnings-per-hour vary so much relative to what we document using our more accurate nominal wage measures described above? There are three reasons for the difference. First, similar to the household data sets, hours are not measured accurately for salaried workers in administrative data sets. The noise in the hours measures for salaried workers results in spurious variation in measures of earnings per hour. This is the benefit of using nominal earnings per pay period - which does not rely on hours - to measure nominal

Figure 16: Imputed Quarterly Earnings Per Hour Changes



PANEL A: QUARTERLY GROSS EARNINGS



PANEL B: IMPUTED EARNINGS PER HOUR
HOURLY WORKERS

PANEL C: IMPUTED EARNINGS PER HOUR
SALARIED WORKERS

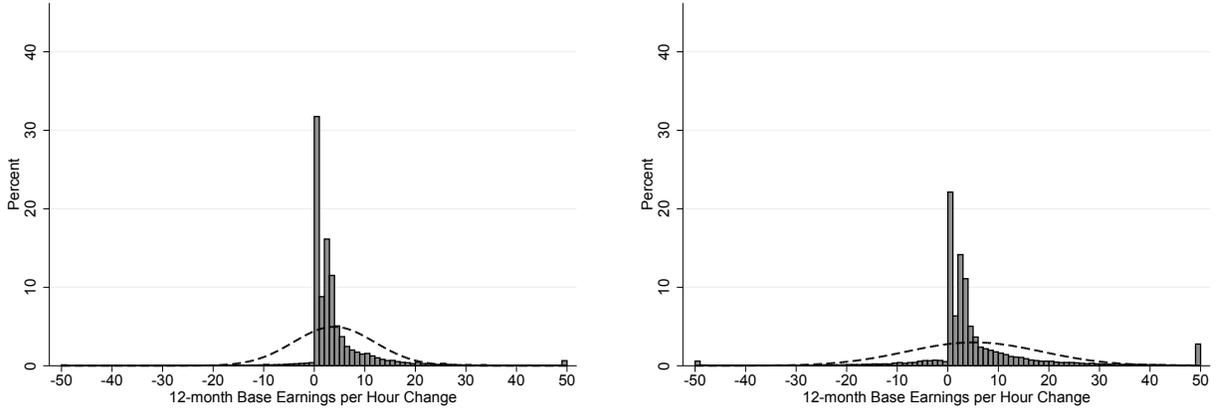
Note: Panel A plots the distribution of four-quarter earnings changes for our sample of job-stayers pooled between 2008 and 2016. Panels B and C report the distribution of quarterly earnings per hour changes, where earnings per hour are imputed by dividing total quarterly earnings by total hours worked in a quarter, for hourly and salaried workers, respectively. Data are weighted to match the firm size \times industry mix found in the BDS.

wage stickiness. We illustrate this point more forcefully below. Second, firms often provide both hourly and salaried workers with payments for sick days, holidays, and vacations. Under this scenario, individuals continue to have positive earnings with no measures of hours. A given individual with the same quarterly earnings across two quarters may have measured changes in quarterly earnings per hour if the individual took differing amounts of sick days, vacation and holidays during the two quarters. The existence of variation in hours of this sort which are compensated as part of an employer's benefits program will make it appear that firm wage setting behavior is more volatile than it is in practice. Lastly, as we highlight above, individual earnings accrued during a quarter includes many forms of compensation. For example, the primary form of compensation for most workers is their base earnings. However, some workers also accrue overtime compensation. Other individuals receive tips, commissions and variable performance compensation which can vary with both the individual's effort and other business cycle conditions. Finally, as we have documented throughout, many employers pay their workers annual bonuses which may be included in their quarterly earnings.

To help shed light on why the variation in quarterly earnings per hour differs so dramatically from the variation in our nominal wage measures, we perform a few additional analyses. Specifically, we exploit our detailed data on base pay compensation. For each individual, we derive total quarterly base pay compensation divided by total hours worked during the quarter. We further exclude any worker regularly receive tips or commission payments. This measure should exactly match the patterns in Figure 2 for hourly workers. For hourly workers quarterly base pay divided by quarterly hours should match the quarterly per-period contract wage. However, this measure could differ substantially for salaried workers if hours of salaried workers are measured with noise.

Figure 17 shows quarterly base-pay per hour changes over a 12 month period for both hourly workers (panel A) and salaried workers (panel B) during our sample period restriction our sample to include only those workers who receive standard base pay compensation during the month. Notice for hourly workers, the patterns of quarterly earnings per hour changes using this measure of earnings is essentially identical to the quarterly nominal wage change described above. Again, this has to be the case given our measurement. The only differences that accrue are due to time aggregation given that some of the wage change will occur in the middle of the quarter. However, for salaried workers, base earnings per hour is far more volatile than our per-pay period nominal wage measures. This is not the results of bonuses, commissions, tips, or any other residual earnings. The only reason for the difference in patterns for salaried workers between Figures 2 and 17 is the adjustment for hours worked. As seen from the figure, there are lots of really large differences in base wage per hours for

Figure 17: 12-month Changes in Base Earnings per Hour, Job-Stayers



PANEL A: HOURLY WORKERS

PANEL B: SALARIED WORKERS

Note: Figure plots the distribution of four-quarter base earnings per hour changes for our sample of job-stayers pooled between 2008 and 2016. Panels A reports the patterns for hourly workers, while Panel B plots the distribution for salaried workers. Base earnings defined to be the product of contract wage rate and total hours worked for hourly workers, and weekly pair multiplied by four or five for salaried workers. Data are weighted to match the firm size \times industry mix found in the BDS.

salaried workers. The unconditional standard deviation of nominal wage changes in this figure is 19.7 percent (compared with 8.5 percent for hourly workers). Part of the reason for this is that hours for salaried workers do not have much meaning. Because the number of hours worked has no bearing on take-home pay for these workers, neither the worker, nor the firm has strong incentives to accurately report true hours worked. This manifests in large swings in hours worked for these workers, which drastically changes the imputed wage. Although the studies which use administrative earnings data to measure earnings rigidity are very useful contributions, one should therefore be cautious about using their estimates to discipline traditional models with wage stickiness, in which the nominal rigidity is on the *per hour* payment rate for a worker's time.

Lastly, we want to continue to emphasize that our work differs from all the existing papers because our goal is to measure a broad measure of aggregate nominal wage adjustments. In doing so, we measure nominal wage adjustments along a variety of dimension including job-stayers, job-switchers, and worker bonuses. Such a decomposition is not possible in many other datasets.

9 Conclusion

Coming soon

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, Susanto and Christopher L House**, “Allocative and Remitted Wages: New Facts and Challenges for Keynesian Models,” 2017.
- 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 Peter J Klenow**, “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 1995, 112 (5), 947–985.
- 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.
- Bonhomme, Stephane and Jean-Marc Robin**, “Generalized Non-Parametric Deconvolution with an Application to Earnings Dynamics,” *The Review of Economic Studies*, 2010, 77 (2), 491–533.
- 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. Kamlani**, “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.
- Daly, Mary C and Bart Hobijn**, “Downward Nominal Wage Rigidities Bend the Phillips Curve,” 2014.
- Faberman, Jason and Alejandro Justiniano**, “Job Switching and Wage Growth,” 2015.
- Fallick, Bruce C, Michael Lettau, and William L Wascher**, “Downward Nominal Wage Rigidity in the United States During and After the Great Recession,” 2016.
- Guvenen, Fatih, Serdar Ozkan, and Jae Song**, “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 2014, *122* (3), 621–660.
- Hall, Robert E.**, “Employment Fluctuations with Equilibrium Wage Stickiness,” *The American Economic Review*, 2005, *95* (1), 50–65.
- 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.
- Kahn, Shulamit**, “Evidence of Nominal Wage Stickiness from Microdata,” *The American Economic Review*, 1997, *87* (5), 993–1008.
- Kaufman, Roger T.**, *British Journal of Industrial Relations*, 1984, *22*, 101 – 112.
- 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,” 2017.
- Le Bihan, Herve, Jeremi Montornes, and Thomas Heckel**, “Sticky Wages: Evidence from Quarterly Microeconomic Data,” *American Economic Journal: Macroeconomics*, 2012, *4* (3), 1–32.
- 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).
- Menzio, Guido and Shouyong Shi**, “Efficient Search on the Job and the Business Cycle,” *Journal of Political Economy*, 2011, *119* (3), 468–510.
- Mian, Atif, Kamalesh Rao, and Amir Sufi**, “Household Balance Sheets, Consumption, and the Economic Slump,” *Quarterly Journal of Economics*, 2013, *128* (4), 1687–1726.

- 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.
- 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.
- Sigurdsson, Josef and Rannveig Sigurdardottir**, “Evidence of Nominal Wage Rigidity and Wage Setting from Icelandic Microdata,” 2016.
- 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.