NEW EVIDENCE ON CYCLICAL VARIATION IN AVERAGE LABOR COSTS IN THE UNITED STATES

Grace Weishi Gu, Eswar Prasad, and Thomas Moehrle*

Abstract—We provide new evidence on the cyclicality of employers' real labor costs using BLS establishment job data for the 1982–2018 period. Average straight-time wages have become countercyclical since the financial crisis and the subsequent Great Recession. So have benefit expenditures and overall labor costs, as well as major benefit expenditures, including health insurance and Social Security. Consistent with prior literature, we find that total earnings–the sum of straight-time wages, bonuses, and overtime earnings–were procyclical before 2008; even earnings have become countercyclical since then. The increasing countercyclicality of labor costs is largely attributable to periods with below-trend GDP.

I. Introduction

THE cyclical behavior of real wages in the United States has received a great deal of attention in the macroeconomics and labor economics literature. The ability to match the observed correlation between real wages and the business cycle has long been a test for the empirical relevance of different classes of models. For instance, traditional equilibrium real business cycle models that emphasize technology shocks tend to generate strongly procyclical real wages. This implication seemed to run counter to a once widely held view that real wages are countercyclical, which was consistent with many Keynesian models and models of cyclical shifts in labor supply. The notion that wages are countercyclical was challenged by a subsequent wave of research arguing that this finding reflected changes in the skill composition of employment over the business cycle. But the debate remains far from settled, with alternative measures of wages often yielding conflicting results.

In this paper, we provide new evidence relevant to this debate using establishment-job-title data on average labor costs from the employer cost surveys conducted by the U.S. Bureau of Labor Statistics (BLS).[1] The evidence previously brought to bear on this subject has been based mainly on earnings data from individual or household surveys, which provide at best a partial picture of firms' labor costs. Over the past three decades, employer-provided nonwage benefit expenditures have accounted for a significant share of U.S. workers' overall compensation; this share stood at about 32% in 2018. The employer cost survey data allow us to explore the cyclical behavior of total average labor costs (i.e., total compensation) as well as straight-time wages, earnings (including overtime and bonuses), and nonwage benefit expenditures. The data set has a considerable amount of information about the characteristics of the sampled establishments and their jobs, but not about workers. We can therefore control for various firm and job characteristics but not for the observed or unobserved characteristics of employed workers. We have obtained annual data from the employer cost surveys for 1982 to 2018, which not only covers a long and continuous time horizon but also includes a substantial period after the global financial crisis.

Our main result is that real hourly straight-time wages, total benefit expenditures, and overall labor costs have all become significantly countercyclical in the aftermath of the global financial crisis and the ensuing Great Recession. The results are consistent across different indicators of the business cycle (based on GDP, the unemployment rate, NBER business cycle dating, and TFP shocks), the use of different deflators, and different approaches to calculating standard errors. In line with the prior literature, we find that earnings—the sum of straight-time wages, overtime, and bonuses—are procyclical over the period 1982 to 2007. But even earnings have become countercyclical since then. We find that the majority of the increasing countercyclicality of labor costs comes from periods with below-trend GDP, that is, the Great Recession, as well as some years during the subsequent recovery.

Moreover, consistently across our baseline specification and most robustness results, real total benefit expenditures have become countercyclical to a lesser degree than real wages. Using wage and earnings data alone, as is the case in the extensive prior literature on this topic, therefore provides an incomplete picture of the dynamics of firms labor costs and the adjustments of labor-cost structures. We also find that among the largest components of employer benefit expenditures, health insurance expenditures are countercyclical over the entire sample period, while Social Security contributions were strongly procyclical before 2008 and have become countercyclical since then. The BLS has also provided us access to quarterly data from the employer cost surveys, although starting only in 2004. Results using quarterly data for this shorter period confirm the increasing countercyclicality of average straight-time wages, total benefit expenditures, and overall labor costs.

Since we use establishment-job data, it is possible that our estimates of the cyclicality of labor costs are influenced by changes in the establishment-job composition of our sample and also by changes in the composition of workers within establishment-job units. The data set we use has very detailed information on the characteristics of establishments

Received for publication August 25, 2018. Revision accepted for publication June 12, 2019. Editor: Olivier Coibion.

*Gu: University of California Santa Cruz; Prasad: Cornell University, Brookings Institution, and NBER; Moehrle: U.S. Bureau of Labor Statistics.

We thank numerous seminar participants and colleagues, especially Eric Swanson and Maury Gittleman, for their helpful comments, and Kaiwen Wang for excellent research assistance. We are also grateful to the editor, Olivier Coibion, and four anonymous referees for their detailed and constructive comments.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/rest_a_00863.

[1]We use the terms *job* and *job title* and the terms *firm* and *establishment* interchangeably.

The Review of Economics and Statistics, December 2020, 102(5): 1–14

https://doi.org/10.1162/rest_a_00863

and jobs, which allows us to mitigate the first type of composition effects. The latter type of composition effects can also be attenuated in view of the data being based on highly disaggregated job titles. Still, we acknowledge that there could be additional composition bias related to systematic cyclical variation in the observed and unobserved characteristics of employed workers within establishment-job units. We cannot control for such composition effects since the data set does not have information on workers. However, since our focus is on the cyclicality of labor costs from the perspective of firms, not that of workers, this is less of a concern for our purposes. In other words, unlike the existing literature, our results provide an empirical characterization of the cyclicality of firms' labor costs for specific jobs/tasks.[2]

This paper makes three main contributions. First, we provide strong evidence of increasing countercyclicality in average real labor costs after the financial crisis. Second, unlike the previous literature that has focused on workers' earnings, we provide new results, from the perspective of firms, on labor cost cyclicality using data at the establishment-job level. These results have implications for models of job-level compensation-adjustment behavior by firms. Third, we characterize cyclical variation in overall labor costs as well as their components—including straight time wages and various nonwage benefit expenditures.

A. *Related Literature*

There is an extensive and rich literature on wage/earnings cyclicality. Abraham and Haltiwanger (1995) provide an excellent survey of papers based on U.S. data, and Brandolini (1995) surveys the evidence for a number of other countries. Nearly this entire prior literature has focused on labor income from the perspective of workers since previous researchers have not used detailed disaggregated data on labor costs from the perspective of firms. Moreover, even studies looking at worker incomes have not included nonwage benefits.

Studies using microlevel data, typically from the National Longitudinal Survey of Young Men (NLSY) and the Panel Study of Income Dynamics (PSID), have largely supported the view that real earnings are procyclical. For instance, using NLSY data for the period 1966 to 1980, Bils (1985) concludes that real earnings are more procyclical for individuals prone to moving between employers or in and out of the workforce. Kydland and Prescott (1988) use PSID data from 1969 to 1982 to construct a skill-weighted index of aggregate labor input and derive a measure of earnings per unit of labor input. They find this measure of the price of labor to be strongly procyclical. Shin (1994) shows that real wages are more procyclical in manufacturing than in the trade or services sectors. Solon, Barsky and Parker (1994) use PSID data over the period 1967 to 1987 and find that low-skill, low-wage workers have more procyclical employment than skilled, higher-wage workers, which induces a countercyclical composition bias in aggregate measures of earnings cyclicality.

In more recent literature, using CPS data for 1979 to 2011, Elsby, Shin and Solon (2016) find evidence of procyclical earnings (including incentive pay). Kudlyak (2014) constructs a measure of the user cost of labor based on NLSY data for 1978 to 2004 and finds that it is substantially more procyclical than average earnings.[3] However, Gertler et al. (2018) argue that this empirical pattern can be accounted for by changes in the composition of new hires (separated into those hired from unemployment versus job-to-job switchers). Using Survey of Income and Program Participation (SIPP) data from 1990 to 2012, they show that after controlling for composition effects involving procyclical job-to-job switching behavior, the wages of new hires are no more cyclical than those of existing workers.[4] In related work that builds on Daly, Hobijn and Wiles (2012), Daly and Hobijn (2016, 2017) use CPS data (over the period 1980 to 2015) to show that most of the wage procyclicality for those who remain fully employed over the business cycle also comes from job-to-job switchers.

Gertler, Huckfeldt, and Trigari (2018) suggest that when new hires and existing workers receive equal treatment, the current wage represents the user cost of labor and the behavior of existing workers' wages is a better guide to the cyclicality of the marginal cost of labor than the cyclicality of new hires' wages unadjusted for composition effects. Our use of labor cost data for existing workers is consistent with this argument. In addition, our data differ in another important respect from those of previous papers in that they are at the job-title level. On one hand, this reduces the effects of procyclical patterns of job-to-job switching that are present in worker-level data.[5] On the other hand, it constrains us to labor costs averaged across workers within a particular firm or job unit, which can induce other composition effects that we discuss in more detail later. At a minimum, even though marginal labor costs are more relevant in the context of most theoretical models, the cyclical behavior of average labor costs provides another useful moment for evaluating the empirical relevance of various models. For instance, Basu and House (2016) find similar results as Kudlyak (2014) did for user costs while also matching average earnings in the data.

The literature we have summarized tends not to explicitly distinguish between straight-time wages and earnings but tends to use these two concepts interchangeably. The

[2] Unless the kinds of tasks performed under a job title change over the business cycle. This can be partly controlled for by the job-task scores that rate each job title on four factors based on actual tasks performed: knowledge, job controls and complexity, professional contacts and communications the workers make, and physical environment of the job.

[3] Beaudry and DiNardo (1991) find that wages of new hires depend on labor market conditions at the time of hiring and are smoothed subsequently. Also see Pissarides (2009).

[4] Other papers on this include Martins, Solon, and Thomas (2012) and Hagedorn and Manovskii (2013).

[5] Real wage changes attributable to job switching are not a bias in terms of measuring a worker's real wage. However, our paper is about labor costs from the perspective of firms, in which case variations from job switching in individual earnings data would induce a procyclical bias in the labor cost accounting for a particular job in a given firm.

importance of variable elements of pay began to receive attention relatively recently. Using PSID data over the period 1970 to 1991, both Devereux (2001) and Swanson (2007) find that straight-time wage rates are acyclical but earnings, which include overtime payments and bonuses in addition to basic wages, are highly procyclical. They attribute this discrepancy to strong procyclical variation in variable pay margins such as bonuses, overtime, late shift premiums, and commissions.[6] Similarly, Shin and Solon (2007), using NLSY data, conclude that job stayers' real hourly earnings are substantially procyclical and that an important portion of that procyclicality is due to compensation beyond base wages.[7] Parker and Vissing-Jorgensen (2010) conjecture that the sharp increase in the cyclical variation of the earned income of high-income households in the United States over the past two decades may be attributable to nonwage compensation.

We are not the first to use the microlevel data set from the BLS employer cost surveys, although we are not aware of other papers that have used it to examine the issue of labor cost cyclicality. Other researchers have used it for different purposes. For instance, Gu (2018) and Makridis and Gittleman (2017) examine the relationship between firms' employment and labor costs.[8] Pierce (2001, 2010), and Monaco and Pierce (2015) find that inequality growth in broader measures of compensation exceeds wage inequality growth (also see Chung, 2003; Lettau, 2003; Gittleman & Pierce, 2013). Lebow, Saks, and Anne (2003) use the BLS data to reexamine the evidence on downward nominal wage rigidity. Hallock (2012) uses it to examine relative pay for major occupational groups by metropolitan area. In general, BLS data releases and reports, which characterize developments in aggregate trends in compensation, constitute the main source of publicly available analysis using the data.

II. Data Set

The employer cost survey (or National Compensation Survey, NCS) administered by the BLS covers a nationally representative random sample of establishments and jobs in all fifty states and the District of Columbia. The survey covers all sectors other than the federal government.[9] The survey collects detailed data on labor costs that are paid by establishments, including to wages, Social Security, unemployment benefits, pensions, health insurance premiums, and paid leave. It does not contain information about workers' received benefit coverages and their relevant preferences.

The sampling procedure has three stages. First, a probability sample of geographic areas is chosen based on employment size. Next, within each area, a probability sample of establishments is selected. For the private sector, an establishment usually refers to a single physical location that produces goods or provides services. Even if a sampled establishment is owned by a larger entity with many locations, only that specific establishment is included in the survey. For state and local governments, an establishment can include more than one physical location within a specific geographical area. In the third stage, a probability sample of jobs is chosen from each of the selected establishments. This is done by random sampling from a complete list of employees provided by the selected establishments. The job titles of the selected workers are then reported, and data are collected for all workers under each of those reported job titles. The number of jobs for which data are collected is based on total employment in the establishment.[10] In summary, the larger the employment size of an area or an establishment, or the number of workers holding a particular job within an establishment, the greater is its chance of being selected.[11]

The basic unit of observation in the data set is an establishment-job-title. The data set has a considerable amount of information about establishment characteristics, including employment size and state location, the establishment's industry according to the Standard Industrial Classification (SIC) system, and the North American Industry Classification System (NAICS), whether it is privately owned or operated by a state or local government, and whether it is for profit or nonprofit. Jobs are highly disaggregated at the level of job titles, which are differentiated by tasks performed and skills required. For instance, a senior accountant is a different job title from accounting manager. Job characteristics include the six-digit BLS Standard Occupational Classification (SOC) system, full-time versus part-time jobs, time-based versus incentive-based jobs, union versus nonunion jobs, and a job-task score based on knowledge required, controls and complexity, professional contacts (interactions with other workers), and physical environment. These detailed establishment and job categories tightly condition workers by their establishment-job attributes.

[6] Similar results of acyclical base wages but procyclical earnings (including overtime, commissions, bonuses, and other performance-related pay) have been reported for other countries such as the United Kingdom (Hart, 2008; Hart & Roberts, 2013) and Germany (Anger, 2011).

[7] There are other papers that also use the NLSY and discuss earnings beyond base wages. For instance, Bils (1985) contends that his wage measure from the NLSY data set includes overtime payments, but Keane, Moffitt, and Runkle (1988) refute this, noting that the NLSY does not contain data on overtime earnings and hours in every year. Hart, Malley, and Woitek (2009) impute average base and overtime earnings using BLS aggregate data and, using frequency domain analysis, report that basic wages and overtime pay comove with the business cycle.

[8] Gu (2018) uses labor market conditions to predict the level of benefit expenditures per worker, where she estimates a growth-to-level relation. Our paper examines the level-to-level or, equivalently, growth-to-growth, relation. Hence, these two papers' results are not directly comparable.

[9] More precisely, federal and quasi-federal agencies, the military, agricultural workers, private household workers, workers abroad, unpaid workers, self-employed, contractors, and long-term disability compensation recipients are excluded from the survey.

[10] The relationship between an establishment's total employment and number of jobs selected is as follows: 1 to 49 employees: up to four jobs; 50 to 249 employees: up to six jobs; and 250 or more employees: up to eight jobs. Exceptions include state and local government units, for which up to twenty jobs may be selected.

[11] For the period our analysis covers, there were actually a few slightly different sampling designs. But the majority of the period is covered by the procedure described here.

Data on labor costs represent the average across all workers employed in a particular job in a specific establishment, and there is no information about worker characteristics. In other words, observed and unobserved differences across workers in the same job in a given establishment, as well as measurement errors at the worker level, are averaged out. However, worker characteristics within an establishment-job unit could arguably change over the business cycle. For instance, the average experience and productivity level of workers at a given job may fluctuate over time. Since we do not have worker-level data, we cannot directly control for such within establishment-job composition effects.

The sample is divided into five panels that enter and exit the survey on a rotational basis. Each cohort is composed of establishment-job units that represent the population of all in-scope workers throughout the United States at the time of sampling and stays in the survey for about five years.[12] In other words, approximately one-fifth of the sample is refreshed each year in order to reflect changes in the distribution of employment across industries and geographic areas. Establishments that go out of business drop out of the sample, while those that decline to participate at some point in the survey period have their data imputed.

The cleaned-up version of the data set that we use, after dropping some observations with missing or BLS-identified unreliable data, covers an average of 6,600 establishments each year, with about four sample jobs for each establishment. The size of the sample rose between 1982 and 2018, yielding an average of 4,420 establishment-job observations in each of the first three years of the sample and an average of 38,853 in the last three years.

Some of the data collected through this survey are published in aggregated form in BLS publications and provided as part of data series such as the Employment Cost Index (ECI) and the Employer Cost for Employee Compensation (ECEC). External researchers need approval before gaining access to the database onsite at BLS headquarters in Washington, DC, with certain confidentiality restrictions.

For our purposes, this data set has many advantages. First, it has a long and continuous time series from 1982 to 2018 for annual data and 2004Q2 to 2018Q2 for quarterly data (thus including a substantial period after the global financial crisis), which makes it suitable for business cycle analysis. Second, it has data on detailed components of labor costs from the perspective of firms and distinguishes straight-time wages from other earnings. Third, the data set has a considerable amount of information about establishment-job characteristics.

A. *Measures of Labor Costs*

A key and unique feature of our data set is that it provides separate information on straight-time wages and various categories of nonwage benefit expenditures paid by firms. The hourly wage measure that we use is the straight-time hourly rate, which includes wages only during regular hours. Throughout this paper, we use the terms *wage* and *straight-time wage* interchangeably. Many previous papers have used weekly or annual earnings (including overtime pay) divided by straight-time hours, which we conjecture can induce a procyclical bias in the measured cyclicality of wages because they do not properly account for fluctuations in overtime hours.

Another virtue of our data set is that we can construct a comprehensive measure of hourly earnings that comprises straight-time wages, nonproduction bonuses, (overtime) premium pay, and shift differentials.[13] This helps make our results more comparable to many previous papers that use individual earnings data but often lack data on straight-time wages or total compensation. It is also worth noting one important detail: we measure hourly earnings as annual earnings divided by annual hours worked (i.e., annual straight-time hours plus overtime hours less all leave hours), not divided by straight-time hours alone.

We also construct overall labor costs or, equivalently, compensation as the sum of wages and nonwage benefit expenditures. This measure of labor costs does not include firms' costs related to searching for, training, or firing workers.

B. *Stylized Facts*

U.S. firms' average benefit expenditures per worker rose from \$6,315 in 1982 to \$25,450 in 2018, a 53% increase in CPI inflation-adjusted terms. While there is a clear trend in the share of nonwage benefit expenditures in total labor costs, this share does not rise linearly over time and shows a significant amount of variation at business cycle frequencies (figure A.1 in the online appendix).

The employer cost surveys provide data on 23 categories of firms' benefit expenditures. Table A.1 (in the online appendix) provides basic information on the share of total compensation accounted for by each of the major benefit categories at the beginning and end of the sample. Certain benefits are mandated by law or negotiated, either partly or in full, as part of employment contracts. Payments for Social Security, Medicare, workers' compensation, and unemployment insurance are mandated by law at either or both the federal and state levels and not are subject to negotiation. Some of these legally required benefit expenditures can be cyclical due to policy mandates, as well as fluctuations in basic wages that they are based on. Other benefit expenditures such as retirement and insurance benefits, as well as vacations and other

[12] In 2012, the BLS began gradually introducing panels that have a span of only three years.

[13] Nonproduction bonuses are cash payments given to employees that might in some cases not be directly related to the productivity of the individual worker. They include Christmas or year-end, profit-sharing cash, referral, hiring, retention, and attendance bonuses. Nonproduction bonuses are included in the benefits portion of total compensation. A payment directly linked to sales or production is considered a production bonus and is included in wages. Below, we use the terms *bonuses* and *nonproduction bonuses* interchangeably.

paid leave, are typically open to negotiation between the employer and employee. These discretionary benefits account for most of the overall benefit expenditures and can also vary over the business cycle, perhaps as a consequence of bargaining outcomes.

Table A.1 shows that the importance of benefit expenditures in total labor costs has increased over time. The relative importance of different benefit expenditures in total labor costs has also changed considerably. The biggest shift is accounted for by a nearly 6 percentage point increase from 1982 to 2018 in the share of health insurance costs.

III. Empirical Framework

To examine the cyclicality of compensation, the most basic empirical model that one can use posits that the log level of labor costs (C_{ijt}) is related to a business cycle indicator such as log real GDP ($RGDP_t$). Here C_{ijt} is a generic notation used for various measures of labor costs such as straight-time wages, total benefit expenditures, and overall compensation at establishment i, job j, in time period t. A post-2007 dummy variable is added and interacted with the business cycle indicator to capture the post-Great Recession change in cyclicality:

$$\ln C_{ijt} = \beta_0 + \beta_1 \ln RGDP_t + \beta_2 \times post07 \times \ln RGDP_t + \beta_3 \times post07 + \mu_{ijt}. \quad (1)$$

The error term can be decomposed as

$$\mu_{ijt} = a_i + b_j + c_t + \upsilon_{ijt}, \quad (2)$$

where a_i is specific to establishment i; b_j is specific to job j; c_t is specific to time t; and υ_{ijt} is specific to establishment i, job j in time period t. We assume that a_i, b_j, c_t, and υ_{ijt} are independent of each other and that υ_{ijt} is independent of the regressors.

In our actual baseline specification, we take first differences of equation (2) to render the model stationary and then add establishment-job fixed effects, a linear trend term t, and establishment-job characteristics X_{ijt}:

$$\ln \frac{C_{ijt+1}}{C_{ijt}} = \beta_1 \ln \frac{RGDP_{t+1}}{RGDP_t} + \beta_2 \times post07 \times \ln \frac{RGDP_{t+1}}{RGDP_t} + \beta_3 \times post07 + \beta_{4ij} + \beta_5 t + \beta_6 X_{ijt} + \mu_{ijt+1} - \mu_{ijt} \quad (3)$$

$$\mu_{ijt+1} - \mu_{ijt} = (c_{t+1} - c_t) + (\upsilon_{ijt+1} - \upsilon_{ijt}). \quad (4)$$

In this specification, C_{ijt} is hourly or annual-per-worker real labor costs. The key coefficient of interest is β_2. A positive β_2 signifies more procyclical labor costs during the post-2007 period; a negative coefficient implies more countercyclical labor costs. In our analysis, we will also experiment with other measures of the business cycle using the same regression. The linear trend term t captures any trend that might exist in the growth of real labor costs over the sample period. Its unit is a year (or a quarter, as relevant). Other controls, represented by the term X_{ijt}, include establishment-job characteristics: major occupational group, NAICS one-digit industry, regional location, ownership, and union/non-union job. In later regressions, we also add controls for large/medium/small establishment based on employment size, profit/nonprofit firm, whether wages are paid based on time or incentives, job-task score, full-time/part-time job, wage percentile in the distribution, and lagged establishment employment growth. These additional controls are not included in the baseline regressions since their inclusion significantly reduces the number of observations. The list of control variables for each regression is specified under each result table. The error terms $\mu_{ijt+1} - \mu_{ijt}$ follow AR(1) processes.

The use of first-differenced measures of labor costs simplifies the error terms and helps avoid nonstationarity problems. It also helps mitigate composition effects that reflect shifts in the shares of different groups of firms and jobs in the economy over time. Because different groups of firms and jobs incur different levels of labor costs, level regression results can be biased if the shares of higher-labor-cost firms and jobs versus the shares of lower-paying counterparts change over the business cycle. For instance, suppose that the former groups of firms and jobs occupy a larger share in the economy during recessions but a smaller share during normal times; then this phenomenon can drive up average labor costs during a recession, not necessarily because most firms and jobs raise their labor costs but because of firm-job composition shifts. This effect would induce a countercyclical bias. Growth rates can reduce such bias, under the assumption that even though labor cost levels differ across firms and jobs, their growth rates are not as different. Hence, although the composition of establishments and jobs with different levels of labor costs can change over time, the impact of such changes is reduced, as endorsed by Abraham and Haltiwanger (1995).

All regressions are weighted with sampling weights, which corrects for heteroskedasticity and potential endogenous sampling issues, and estimated using yearly data unless otherwise mentioned. We report robust standard errors clustered at the two-digit NAICS industry level.[14]

IV. Empirical Results

We begin by presenting a set of baseline results characterizing the cyclical behavior of straight-time wages, nonwage benefit expenditures, and overall labor costs. We then conduct a number of experiments to evaluate the robustness of our results and relate our findings to the prior literature on this topic. Our primary objective is to provide a detailed

[14]Sampling weights are the product of the inverse of the sample selection probabilities at each stage of sampling, taking into account changes in the employment distribution.

TABLE 1.—LABOR COST CYCLICALITY: BASELINE RESULTS

Hourly Rate	(1) Wage	(2) Benefits	(3) Compensation
Δ Log RGDP	−0.0362	0.0165	−0.0225
	(0.0317)	(0.0922)	(0.0369)
Δ Log RGDP × Post07	−0.4460***	−0.3730***	−0.4260***
	(0.0467)	(0.0755)	(0.0491)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	741,949	741,949	741,949
R-squared	0.262	0.276	0.265
Annual Rate	**(4) Wage**	**(5) Benefits**	**(6) Compensation**
Δ Log RGDP	−0.0223	−0.0150	−0.0193
	(0.0465)	(0.0925)	(0.0481)
Δ Log RGDP × Post07	−0.4570***	−0.3370***	−0.4250***
	(0.0580)	(0.0746)	(0.0576)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	684,666	684,666	684,666
R-squared	0.287	0.290	0.290

Dependent variable is log-differenced real hourly (top panel) or annual (bottom panel) cost of labor (deflated by CPI). "Wage" refers to straight-time wages only, "Benefits" refers to total benefit expenditures paid by an employer, and "Compensation" refers to overall labor costs, the sum of the labor costs in the previous two columns. All regressions are weighted using sample weights and include establishment-job fixed effects, a linear trend, business cycle indicator, post-2007 dummy and its interaction, and other controls. Other controls include major occupational group, NAICS one-digit industry, regional location, ownership, and union or non-union job. In parentheses below coefficient estimates, we report robust standard errors that are clustered at the two-digit NAICS industry level. *, **, and *** denote statistical significance at the 5%, 1%, and 0.1% levels, respectively.

TABLE 2.—UNEMPLOYMENT RATE AS CYCLE INDICATOR

	(1) Wage	(2) Benefits	(3) Compensation
ΔUnemployment Rate (%)	0.0058***	0.0043**	0.0055***
	(0.0009)	(0.0012)	(0.0006)
ΔUR (%) × Post07	0.0019*	0.0014	0.0017*
	(0.0009)	(0.0013)	(0.0008)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	741,949	741,949	741,949
R-squared	0.263	0.276	0.266

Dependent variable is log-differenced real hourly cost of labor (deflated by CPI). Also see notes to table 1.

empirical characterization of cyclical variation in total labor costs and their major components rather than to test any particular theory.

A. *Baseline Results*

Table 1 shows the full-sample results for regressions of growth in straight-time wages, benefit expenditures, and total compensation on real GDP growth and other controls. The top panel shows results based on hourly measures of labor costs, while the lower panel shows results for annual per worker measures. All of the labor cost measures are deflated by the seasonally adjusted CPI for all urban consumers.

The first column of the upper panel shows that hourly real straight-time wages are acyclical over the period 1982 to 2007 but then turn significantly countercyclical after 2007. For the post-2007 period, the coefficient estimate indicates that a 1 percentage point increase in real GDP growth reduces real hourly wage growth by 0.45 percentage points more in the latter period relative to the earlier period.[15] The pre-2008 result is consistent with the findings of the previous literature that distinguishes straight-time wages from total earnings (e.g., Devereux, 2001; Swanson, 2007). The striking new finding is the pattern of countercyclical wage variation after 2007, the focus of this paper.

In the next two columns, we show estimates of similar regressions for real nonwage benefit expenditures and overall real compensation (i.e., total labor costs), respectively. These measures have also turned countercyclical, with coefficient estimates slightly smaller than the one for wages. A 1 percentage point increase in real GDP growth reduces growth in real hourly total labor costs by 0.43 percentage points more during the post-2007 period than in the earlier period. This suggests that nonwage benefits do not vary systematically in a way that counteracts the cyclicality of the variation in wages. Indeed, when we run a similar regression for average real nonwage benefit expenditures, the coefficient on the interaction between real GDP growth and the post-2007 dummy is again significantly negative (column 2), indicating that benefit expenditures have also become countercyclical. A 1 percentage point increase in GDP growth reduces growth in benefit expenditures by 0.37 percentage points more during the post-2007 period than in the earlier period.

The lower panel of table 1 presents a similar set of results using annual per worker wages, benefit expenditures, and compensation data. It is possible that employers adjust labor input at the intensive margin and thereby influence the cyclicality of their labor input costs. Weekly hours or weeks worked in a year for existing employees might be easier to adjust in response to cyclical variation in demand than employment levels or hourly wages. In fact, the results using annual measures are similar to those in the top panel, confirming increasingly countercyclical variation in wages, nonwage benefit expenditures, and compensation, independent of the unit of the data. Moreover, in all of the results above, we find both total benefit expenditures and overall labor costs turning slightly less countercyclical than wages.

We have so far used real GDP as an indicator of the business cycle. In table 2, we report regression results using changes in the unemployment rate as the cyclical indicator. This is the business cycle indicator a number of previous papers have focused on. But we use this as a secondary indicator since unusual changes in labor market dynamics in recent years appear to have altered the traditional relationship between unemployment and GDP. The coefficient estimates indicate that a 1 percentage point increase in the unemployment rate is associated with about a 0.2 percentage point greater increase in wage and compensation growth during the post-2007 period than in the earlier period. Thus, our baseline results about increasing countercyclical variation in wages and compensation after 2007 are preserved when we use a labor market measure of the aggregate business cycle, although the effect is weaker.

[15]The regression specification implies a symmetric effect when real GDP growth declines.

There are two differences relative to the results in table 1. One is that in the regression for benefit expenditures (second column), the coefficient on the interaction term is not statistically significant, although it has a positive sign. The second difference is that all three measures of labor costs appear countercyclical even before 2008. This is, however, not a robust result as shown in our later results. Abraham and Haltiwanger (1995) also note that in some periods, their results regarding earnings cyclicality are affected by whether they use employment or the industrial production index as the business cycle indicator. The bottom line is that although the cyclicality of labor costs before 2008 appears sensitive to whether we use GDP or the unemployment rate to characterize the business cycle (and to our later robustness checks), our key result focuses on the *change* in the cyclicality of labor costs—the increasing countercyclicality of labor costs from the pre-2008 period to the post-2007 period in our sample. This is an important theme that we investigate through a variety of sensitivity tests that we turn to next. We do not intend to draw any conclusion about the absolute level of cyclicality for the pre-2008 period.

B. *Robustness*

Abraham and Haltiwanger (1995) note that the following factors can influence the measured cyclicality of real wages/earnings: (a) time period; (b) method of detrending; (c) the choice of cyclical indicator; (d) use of consumption versus production deflators; and (e) composition effects. We now evaluate the robustness of the main result from our baseline regressions—that the cyclicality of labor costs has changed markedly in the post-2007 period—in a variety of dimensions, including the factors just mentioned. We briefly summarize the results of these robustness tests below, relegating detailed results to the online appendix.

Time period. First, to confirm the change in labor cost cyclicality over our sample period, we ran the baseline regression separately for subsamples before 2008 and after 2007. The coefficients on real GDP growth are consistent with our key result that real labor costs have turned countercyclical since 2008. In addition, in order to control for possible changes in labor cost trends after 2007, we included a post-2007 interaction with the time trend in the baseline specification. This made little difference to our key result.

Alternative detrending methods and cyclical indicators. We then experimented with different approaches to detrending GDP and using alternative business cycle indicators. We first replaced GDP growth with the cyclical component of real GDP derived using the Hodrick-Prescott (HP) filter. Consistent with our main result, results using this business cycle indicator showed that both real wages and total compensation have become more countercyclical since 2008, although the coefficient magnitudes are different. It is not surprising that different detrending methods can yield different coefficient magnitudes across time periods. As Abraham and Haltiwanger (1995) note, wage cyclicality measured using HP-filtered GDP is substantially more procyclical over the period 1970 to 1993 than over the period 1949 to 1969, while the subperiod results using first differences are less clear cut.

Next, we experimented with additional business cycle indicators—NBER business cycle chronologies and TFP growth. These results confirmed the increasing countercyclicality of labor costs in the postcrisis period. For instance, using NBER business cycle dates, we find that an economic contraction (from peak to trough) is associated with a 0.2 to 0.3 percentage points greater increase in real wage, benefits, and compensation growth during the post-2007 period than in the earlier period. Similarly, a 1 percentage point increase in the growth of TFP, as estimated by Fernald (2014), is associated with a 0.3 to 0.4 percentage point smaller increase in real wage, benefits, and compensation growth during the post-2007 period than in the earlier period.

To examine if state-specific business cycle conditions affect our results, we used state-level business cycle indicators—changes in state unemployment rates and employment growth rates—in place of national-level indicators. The results confirmed the pattern of increasing countercyclicality of labor costs since 2008. The magnitudes of the post-2007 interaction coefficients using changes in the state unemployment rates were similar to those based on changes in the aggregate unemployment rate.

Alternative deflators. For the next set of experiments, we used alternative price deflators in place of the aggregate CPI to deflate labor costs. First, we used the PPI, which may be more relevant for considering the real cost of labor from the perspective of firms. Wages, benefits, and compensation deflated by the PPI have become even more strongly countercyclical post-2008, relative to the baseline results using the CPI as the price deflator. This echoes the findings of Abraham and Haltiwanger (1995) that using the PPI rather than the CPI to deflate real wages tends to make wage variation look less procyclical. This result is also consistent with that of Swanson (2004), who deflates wages with sectoral average product prices and concludes that most industries pay real wages that vary countercyclically with the state of their industry.

We also attempted to account for the fact that inflation may vary across regions, affecting real wage dynamics in ways different from that when we use aggregate CPI. In order to capture regional variations in price developments, we divided the United States into four regions—Northeast, Midwest, West, and South—and then deflated observations on wages, benefit expenditures, and compensation using region-specific CPI indexes. Real labor costs remained more countercyclical during the post-2007 period.

Alternative unit and measurement errors. In the next experiment, we aggregate up sample-weighted real hourly labor costs to the occupation-industry level for each year and then estimate regressions with the aggregated panel data. Assuming measurement errors at the establishment-job level are random, the aggregation can mitigate such errors. We use eight broad occupation categories (executive and managerial, professional and technical, production and craft, machine operators, clerical, transportation, helpers and laborers, sales, and service); industries are broken down to the three-digit NAICS level. Consistent with the baseline results, real wages and total compensation are countercyclical in the post-2007 period and acyclical before that.[16]

Alternative standard errors. In the baseline regressions, we clustered the standard errors at the NAICS two-digit industry level. Our next robustness experiment relates to alternative approaches to adjusting standard errors to account for possible cross-sectional correlations. We widen the cluster to the NAICS one-digit industry level to allow for possible correlation across two-digit industries. In separate regressions, we also use Driscoll-Kraay standard errors, which allow for more general patterns of correlations across observations.[17] Whatever the procedure used for calculating standard errors, the strongly significant coefficients for the post-2007 interaction terms are preserved.

Finally, in order to help control for within-state correlations, in the last robustness check, we add state unemployment rates or state employment growth rates as an additional control in the baseline specification. Note that these additional state-level variables are not intended here as cyclical indicators. Rather, they help reduce potential correlation in regression residuals induced by common state-specific factors that may cause labor costs of firms in a given state to be correlated. Again, our result about the increasing countercyclicality of real labor costs is preserved.

C. *Composition Effects*

One concern in analyzing the cyclicality of labor costs relates to possible composition effects. Levels of wages and broader measures of labor costs differ across workers, jobs, and firms, and the composition of each of these may vary systematically over the business cycle. Prior evidence on the size and direction of these composition effects is mixed. For instance, Bils (1985) finds that aggregate data on average real wages have a countercyclical bias. This is induced by the countercyclical average skill level of the employed workforce—employers tend to disproportionately lay off unskilled or lower-paid workers during recessions. Similarly, Daly and Hobijn (2017) use CPS data and find that systematic cyclical variations in net exits out of full-time employment among workers with lower earnings impart a countercyclical bias to aggregate real earnings. By contrast, Keane et al. (1988) use NLSY data and find evidence of a procyclical bias in aggregate data that arises if workers with high transitory earnings are more likely to lose their jobs during recessions. In their survey, Abraham and Haltiwanger (1995) argue that although composition effects may be important over specific time periods, it is an open question whether such composition effects have become more or less important over time and in which direction they affect the cyclicality of labor costs.

We separate potential composition effects that could affect our analysis into two categories. One has to do with changes in the composition of jobs and firms, the other with changes in the composition of worker characteristics within a job and firm. In the following sections, we use the richness of our data set to mitigate such composition effects to the extent possible but do not claim to completely eliminate them.

Between establishment-job composition effects. Before we address the first type of composition effects, it is worth noting at the outset that with the probability sampling survey procedure that the BLS uses, the larger an establishment is and the more workers in a particular job, the more likely it is for that establishment and job to get sampled. This attenuates the potential composition effects in our data from cyclical fluctuations in the entry and exit of small firms and of jobs that are not core functions or employ few workers in an establishment. For instance, several papers have shown that small firms tend to have more procyclical entry and more countercyclical exit probabilities than large firms (Lee & Mukoyama, 2012; Tian, 2018). As a simple test of whether establishment sizes have implications for our results, we interacted the business cycle indicator with dummies for different establishment sizes. The coefficients on the interaction terms were small and insignificant, suggesting that establishment size does not affect the cyclicality of labor costs (table C.10 in the online appendix). Hence, establishment-size related composition effects are unlikely to be important for our main results.

Nevertheless, within the constraints of our data set, we adopt several methods to address potential firm-job composition effects that could influence our results. Previous papers using individual-level data have addressed similar issues by using wage growth rates instead of wage levels (endorsed by Abraham & Haltiwanger, 1995) and by controlling for the observed characteristics of workers. In all of our regressions, we too use growth rates, rather than levels, of labor costs. This helps us reduce the impact of between-establishment-job composition changes over the business cycle since these growth rates differ less than labor cost levels across establishments and jobs.

In addition, we add more establishment-job controls to mitigate between-establishment-job composition effects. The regression results we have already discussed contain controls

[16]The occupation-industry level results were similar when we used Driscoll-Kraay standard errors to control for cross-sectional correlations.

[17]Using Driscoll-Kraay errors does not allow us to use sample weights for the weighted regressions; hence, we did not use them for the baseline regressions and others. The regressions in the lower panels of table B.8 in the online appendix are unweighted.

TABLE 3.—COMPOSITION: MORE ESTABLISHMENT-JOB CONTROLS

Annual Data (1982–2018)	(1) Wage	(2) Benefits	(3) Compensation
Δ Log RGDP	0.5880***	0.2530	0.5000***
	(0.0982)	(0.1570)	(0.1010)
Δ Log RGDP × Post07	−1.0100***	−0.5530**	−0.8990***
	(0.1400)	(0.1590)	(0.1350)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	492,276	492,276	492,276
R-squared	0.529	0.373	0.488
Quarterly (2004Q2–2018Q2)	(4) Wage	(5) Benefits	(6) Compensation
Δ Log RGDP	−0.0328	0.0357	0.0095
	(0.0820)	(0.1800)	(0.1050)
Δ Log RGDP × Post07	−0.7150***	−0.7530***	−0.7510***
	(0.0769)	(0.1820)	(0.0970)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	766,234	766,234	766,234
R-squared	0.147	0.068	0.135
Four-Quarter Growth	(7) Wage	(8) Benefits	(9) Compensation
Δ Log RGDP	−0.2870***	−0.3100***	−0.2570***
	(0.0405)	(0.0691)	(0.0413)
Δ Log RGDP × Post07	−0.1670***	−0.0397	−0.1710***
	(0.0409)	(0.0699)	(0.0417)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	500,576	500,576	500,576
R-squared	0.423	0.248	0.383

Dependent variable is log-differenced real hourly cost of labor (deflated by CPI). Other controls include major occupational group, NAICS one-digit industry, regional location, ownership, union/non-union job, large/medium/small establishment employment size, profit/nonprofit firm, whether wages are paid based on time or incentive, job-task score, full-time/part-time job, wage percentile, and lagged establishment employment growth. Also see notes to table 1.

for establishment-job characteristics, which help reduce the impact of firm-job composition effects. We now expand the list of such controls to include a job's wage percentile (and its square), full-time or part-time status, establishment employment size (large/medium/small), and lagged establishment employment growth. In particular, the change in the wage percentile variable, which is based on an establishment-job unit's position in the cross-sectional wage distribution in a given period, potentially absorbs some of the bias caused by shifts in the wage distribution when the composition of establishments and jobs changes. These additional controls cut our sample size by about one-third, which is why we did not include them in the baseline regressions.

The results are reported in table 3. The regressions in the top panel use annual data from 1982 to 2018 as in the baseline. For the pre-2008 period, real labor costs are now procyclical/acyclical. Comparing this to the baseline results in table 1 suggests that composition effects could be relevant for that period. However, our key result of the countercyclicality of labor costs in the post-2007 period is not affected.

For the remaining panels, we use quarterly data from 2004Q2 to 2018Q2 and add more controls that become available starting in 2004Q2: profit or nonprofit status of the firm, whether a job is based on time or incentive pay, and job-task scores. The job-task score rates each job title in an establishment based on four factors that reflect actual tasks performed: knowledge, job controls and complexity, professional contacts and communications the workers make, and physical environment of the job. This helps control for any changes over the business cycle in tasks performed under a job title in an establishment. The middle panel reports results from regressions for quarterly growth in real labor costs, while the bottom panel reports results using four-quarter growth rates. No matter the frequency of data and how many establishment-job controls we add, this table shows that real wages and compensation have consistently become more countercyclical during the post-2007 period.

Within establishment-job composition effects. The second type of composition effects that could potentially influence our results is related to worker composition changes within a job-firm unit. Differences in labor cost variation over the cycle for different workers within the same firm-job unit could affect our results as the within-firm-job-unit composition of workers changes. However, it is not obvious in which direction this factor might affect the cyclicality of labor costs. On one hand, the least experienced workers are more likely to become unemployed during recessions. Xu and Couch (2017) find that younger workers are more likely to move from employment to unemployment than other workers as the economy worsens. Forsythe (2017) also shows that during recessions, the hiring rate falls faster for young workers than for more experienced workers. This "seniority rule" might induce a countercyclical labor cost pattern within an establishment-job (title) since less experienced workers may have lower wages.

Haltiwanger et al. (2017), however, show that both mobility across jobs and the wage increases associated with voluntary job switching tend to fall during recessions. This cyclical change in the share of workers within a firm-job unit with external offers may cause a procyclical labor cost pattern. In addition, Gertler et al. (2018) and Daly and Hobijn (2016, 2017) also show that the cyclical patterns of the composition of new hires and the turnover of job-to-job switchers have significant procyclical impacts on wages.

Previous papers using individual-level data have addressed such issues by controlling for the observed characteristics of the employed. We cannot directly control for cyclical variation in worker characteristics within firm-job units. However, job-firm units in the data set are highly detailed; in particular, they are sampled by job titles within an establishment. The jobs are differentiated by full-time or part-time status; time based versus incentive based; union on nonunion; supervisory or nonsupervisory; and further differentiated by tasks performed and skills required. Given the highly detailed categorization of jobs, any systematic cyclical fluctuations of worker characteristics, observed or unobserved, are tightly controlled within the scope of the tasks and skills associated with specific establishment-job units. It is much less likely that average worker characteristics within the same

TABLE 4.—COMPOSITION: RECESSION SEVERITY

	(1) Wage	(2) Benefits	(3) Compensation
Δ Log RGDP	−0.0058	−0.0059	−0.0154
	(0.0464)	(0.0869)	(0.0451)
Δ Log RGDP × Post-07	-0.4430^{***}	-0.3130^{**}	-0.3980^{***}
	(0.0616)	(0.0876)	(0.0630)
Δ Log RGDP × High-ΔUR States	−0.0516 (0.0525)	0.0432 (0.0923)	−0.0082 (0.0481)
Δ Log RGDP × High-ΔUR States × Post-07	−0.0153 (0.0547)	−0.1170 (0.0911)	−0.0614 (0.0561)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	741,949	741,949	741,949
R-squared	0.262	0.276	0.265

Dependent variable is log-differenced real hourly cost of labor (deflated by CPI). Additional controls include high-unemployment-rate-increase state dummy and their interactions. High-unemployment-rate-increase states are identified as those that experienced higher (lower) unemployment rate increases (decreases) relative to the national unemployment rate increases (decreases) over our sample period. The set of such states can vary over time. Also see notes to table 1.

TABLE 5.—EXTENSION: PRIVATE SECTOR ONLY

Private Sector	(1) Wage	(2) Benefits	(3) Compensation
Δ Log RGDP	0.0061	0.2520^{**}	0.0788
	(0.0509)	(0.0778)	(0.0515)
Δ Log RGDP × Post07	-0.4630^{***}	-0.4960^{***}	-0.4750^{***}
	(0.0689)	(0.0947)	(0.0721)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	599,569	599,569	599,569
R-squared	0.273	0.296	0.278

Dependent variable is log-differenced real hourly cost of labor (deflated by CPI). Also see notes to table 1.

firm-job-title unit vary as much over the short term as average worker characteristics across the economy.

Moreover, our emphasis throughout this paper is on the cyclical behavior of labor costs from the perspective of firms. In other words, what we document in this paper, which distinguishes it from the previous literature, is the cyclicality of real (hourly) labor costs for a firm to accomplish the same kind of tasks or, equivalently, the real labor costs associated with a given job. We acknowledge that we cannot control for possible changes in the composition and productivity of workers within specific job titles that are systematically related to business cycle conditions, although there is no evidence we are aware of that this phenomenon is quantitatively important.

Recession severity. It is possible that severe recessions, such as the Great Recession, that entail pronounced job destruction can induce larger composition effects in general on the measured cyclicality of labor costs. To examine the importance of such effects, we exploit state-level differences in business cycle conditions.

First, we create a dummy variable that takes the value of 1 if the change in a state's unemployment rate is greater than the change in the national unemployment rate. This dummy variable captures states that experienced higher (lower) unemployment rate increases (decreases) relative to the national unemployment rate increases (decreases). Note that the set of states for which the dummy takes the value of 1 can vary over time. We then interact this dummy variable with the national RGDP growth and post-2007 variables in the baseline specification (table 4). We find that states with more severe recessions did not experience significantly different changes in average labor cost cyclicality relative to those with less severe recessions; this is true both before and after 2007.

Then we zoom in on the post-2007 period and construct similar dummy variables for just that period by comparing states' unemployment rate changes to the national unemployment rate changes. Now the unemployment rate changes used for this dummy are from 2008 to 2010 only. The set of states with high unemployment rate changes does not vary over the post-2007 period and includes Indiana, Michigan, the Carolinas, Florida, Alabama, Idaho, Arizona, Nevada, Oregon, and California. The results are in the top panel of table C.11 in the online appendix, and they again suggest that different recession severity levels across states do not matter for the result about post-2007 countercyclicality.

In addition, we use two other formulations of state dummies (in the bottom two panels of table C.11). One is a dummy for states with higher unemployment rate changes than the cross-sectional median of those changes; the other is for states whose unemployment rate changes are among the top quantile of all states' unemployment rate changes. Neither modification alters our results about the insignificant impact of recession severity on average labor cost cyclicality. Furthermore, we conduct robustness checks with national unemployment rate changes and state-level unemployment rate changes as the business cycle indicator, respectively, and report the results in the online appendix (tables C.12 and C.13). There, we continue to find that increasing countercyclicality in the post-2007 period is not driven by states with especially severe recessions.

D. *Extensions*

Up to this point, we have established the robustness of our baseline results and accounted for composition effects to the extent possible. We now extend our baseline results in a variety of dimensions, exploiting the disaggregated data from the BLS employer cost surveys as well as the detailed data that we have on firms' benefit expenditures.

Private sector only. First, we limit our analysis to the private sector, which accounts for about 81% of our baseline observations. As shown in table 5, the result of increasing countercyclicality of real labor costs during the post-2007 period is preserved. More specifically, benefit expenditures are the component of labor costs whose cyclicality is most affected by confining the sample to the private sector. Comparing the results in table 5 with those in the baseline, we find that benefit expenditures are more procyclical in the private sector than in the public sector during the pre-2008 period, and they become increasingly countercyclical after 2007 in the private sector just as in the public sector.

TABLE 6.—EXTENSION: POST-2007 EVERY YEAR

Quarterly Data	(1) Wage	(2) Benefits	(3) Compensation
Δ Log RGDP × 2008	−1.9740***	−2.0520***	−1.9960***
	(0.0540)	(0.0475)	(0.0429)
Δ Log RGDP × 2009	0.1800***	0.2870*	0.2020**
	(0.0438)	(0.1170)	(0.0554)
Δ Log RGDP × 2010	0.7840**	−0.1020	0.5230**
	(0.2230)	(0.3420)	(0.1800)
Δ Log RGDP × 2011	0.2920***	0.4290	0.3550**
	(0.0741)	(0.2390)	(0.0964)
Δ Log RGDP × 2012	−0.6970*	−0.0658	−0.7250*
	(0.2960)	(0.7590)	(0.3190)
Δ Log RGDP × 2013	−0.4100*	0.1430	−0.2890*
	(0.1750)	(0.2790)	(0.1350)
Δ Log RGDP × 2015	−0.9030***	−1.1560***	−1.0150***
	(0.0896)	(0.2640)	(0.1160)
Δ Log RGDP × 2016	−2.8710***	−4.7330**	−3.4520***
	(0.5880)	(1.2760)	(0.3550)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	1,004,197	1,004,197	1,004,197
R-squared	0.083	0.097	0.088

Dependent variable is log-differenced real hourly cost of labor (deflated by CPI). Additional controls include year dummies and their interactions with the quarterly business cycle indicator. Also see notes to table 1.

Post-2007 drivers of labor cost cyclicality. The results in section 4.3.3 suggest that the post-2007 countercyclicality of labor costs is not driven purely by the Great Recession. To further investigate which years drive our key finding, we now exploit the quarterly labor cost data to explore within-year cyclicality by estimating regressions using year dummies and their interactions with quarterly real GDP growth. The results in table 6 show that the recession year 2008 indeed contributes significantly to the post-2007 countercyclicality of labor costs, but so do some of the subsequent years in our sample period.

Booms versus busts. Another interesting question is whether our results about countercyclical wages are driven more by labor cost dynamics during business cycle booms versus busts. To explore this question, we create a dummy variable for periods in which GDP is below trend or not, using a measure of trend output constructed using the Hodrick-Prescott filter. We then estimate regressions for the pre-2008 and post-2007 periods, respectively, that include interactions of this dummy variable with GDP growth (table 7). Before 2008, periods of below-trend GDP tend to have slightly more countercyclical wages and benefit expenditures than periods of above-trend GDP (top panel).

After 2007, this cyclical asymmetry becomes sharper (middle panel), with labor costs countercyclical during periods of below-trend GDP and procyclical when GDP is above trend. In other words, real labor costs increase as GDP rises during periods with above-trend GDP; however, they do not decrease as much or even increase as GDP declines during periods with below-trend GDP. This asymmetry of labor cost cyclicality between above-trend GDP periods and below-trend GDP periods has become stronger since 2008.

TABLE 7.—EXTENSION: ASYMMETRY IN CYCLICAL VARIATION

Pre-2008	(1) Wage	(2) Benefits	(3) Compensation
Δ Log RGDP	0.1720	0.4270**	0.2320§
	(0.1340)	(0.1280)	(0.1260)
Δ Log RGDP × Below-trend	−0.2420§	−0.3620§	−0.2540
	(0.1400)	(0.2060)	(0.1510)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	360,012	360,012	360,012
R-squared	0.238	0.260	0.243
Post-2007	**(4) Wage**	**(5) Benefits**	**(6) Compensation**
Δ Log RGDP	0.2780***	0.3790	0.3170**
	(0.0425)	(0.2290)	(0.0975)
Δ Log RGDP × Below-trend	−1.0410***	−0.9490***	−1.0300***
	(0.0580)	(0.2330)	(0.0995)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	381,937	381,937	381,937
R-squared	0.309	0.298	0.307
Post-2007: Excluding 2008–2009	**(7) Wage**	**(8) Benefits**	**(9) Compensation**
Δ Log RGDP	0.3300***	0.3500	0.3500***
	(0.0467)	(0.2120)	(0.0895)
Δ Log RGDP × Below-trend	−0.7410**	−0.1200	−0.5980*
	(0.2160)	(0.6140)	(0.2290)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	303,003	303,003	303,003
R-squared	0.335	0.331	0.335

Dependent variable is log-differenced real hourly cost of labor (deflated by CPI). Additional controls include RGDP-below-trend dummy and its interaction. The symbols §, *, **, and *** denote statistical significance at the 10%, 5%, 1%, and 0.1% levels, respectively. Also see notes to table 1.

To examine whether this post-2007 shift in the degree of asymmetry is solely driven by the Great Recession, in the bottom panel of table 7, we exclude the years 2008–2009.[18] The asymmetry in the cyclicality of labor costs becomes weaker when we exclude the Great Recession, but the increase in the asymmetry during the post-2007 period is still preserved for wages and total compensation.

In the online appendix, we report results from a variety of robustness checks for the greater asymmetry of labor cost cyclicality in the post-2007 period (tables D.14–D.16). For one, we use unemployment rate changes instead of GDP growth as the business cycle indicator and interact the new indicator with the previous dummy variable for GDP-below-trend periods or with a new dummy variable for periods with unemployment rate increases, respectively. Both exercises show that the asymmetry of labor cost cyclicality between above-trend GDP (or falling unemployment rate) periods and below-trend GDP (or rising unemployment rate) periods has become stronger since 2008, and this result is not just driven by the Great Recession. For another robustness check, with everything else being the same as in table 7, we replace the dummy for the below-trend-GDP periods with a dummy variable for periods with negative GDP growth. We again find similar asymmetry results as before, but we could not

[18]Below-trend periods include troughs as well as the period leading to them and the period right after them. Hence, not all below-trend years have negative GDP growth.

TABLE 8.—EXTENSION: PERIODS WITH ABOVE- VERSUS BELOW-TREND GDP

Below-Trend GDP Periods	(1) Wage	(2) Benefits	(3) Compensation
Δ Log RGDP	−0.0332	0.1360	0.0236
	(0.0530)	(0.1650)	(0.0806)
Δ Log RGDP × Post-2007	−0.7160***	−0.7080***	−0.7320***
	(0.0364)	(0.1060)	(0.0535)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	359,390	359,390	359,390
R-squared	0.488	0.492	0.493
Above-Trend GDP Periods	**(4) Wage**	**(5) Benefits**	**(6) Compensation**
Δ Log RGDP	0.3000	0.4860*	0.3440§
	(0.2110)	(0.1750)	(0.1950)
Δ Log RGDP × Post-2007	0.0000	−0.1840	−0.0363
	(0.2230)	(0.3050)	(0.2240)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	382,559	382,559	382,559
R-squared	0.525	0.535	0.523

Dependent variable is log-differenced real hourly cost of labor (deflated by CPI). The symbols §, *, **, and *** denote statistical significance at the 10%, 5%, 1%, and 0.1% levels, respectively. Also see notes to table 1.

TABLE 9.—EXTENSION: EARNINGS AND MAJOR BENEFITS

	(1) Earnings	(2) Health	(3) Social
Δ Log RGDP	0.1450§	−0.0049**	0.0065***
	(0.0762)	(0.0015)	(0.0007)
Δ Log RGDP × Post07	−0.6060***	−0.0023§	−0.0103***
	(0.1270)	(0.0013)	(0.0007)
Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	741,949	741,949	741,949
R-squared	0.344	0.242	0.261

Dependent variable is log-differenced real hourly cost of labor (deflated by CPI). Total earnings include basic straight-time wages, premium pay, shift differentials, and nonproduction bonuses. "Health" and "Social" refer to health insurance premiums and Social Security paid by the employer, respectively. The symbols §, *, **, and *** denote statistical significance at the 10%, 5%, 1%, and 0.1% levels, respectively. Also see notes to table 1.

conduct the test excluding 2008–2009 because they overlap with the periods of negative growth rates.[19]

These results suggest that average labor costs are relatively more procyclical during booms and relatively more countercyclical during busts, and that this asymmetry has increased since 2008. To examine which of these two types of periods contributes more to the increased asymmetry, we run separate regressions for them. The results are shown in table 8. Labor compensation has become more countercyclical since 2008 for periods of both booms and busts, but the change is greater for periods when GDP is below trend. Hence, the increased asymmetry in cyclicality of labor costs between boom-bust periods as well as the increased overall countercyclicality of labor costs since 2008 are mostly attributable to the periods of below-trend GDP that are composed of both the Great Recession and some years during the subsequent recovery, which is consistent with our results from table 6.

Earnings and major benefits. Lastly, we run the baseline regressions for earnings and major benefit expenditures (table 9). Barring a few exceptions that we have already noted, the previous literature on wage cyclicality has largely focused on total earnings that are the sum of overtime earnings and bonuses (this is the measure available, for instance, in the PSID). Our results for earnings are more comparable to those of the previous literature than our prior results using straight-time wages. One key aspect in which our paper differs from the prior literature is that our data set allows us to break down the different components of earnings and examine the cyclicality of straight-time wages and total earnings side by side.

To compare our results with the previous literature, in table 9 (the first column), we report results using a measure of earnings, defined as the sum of hourly straight-time wages, bonuses, and overtime pay. Before 2008, earnings are only weakly procyclical but still more procyclical than straight-time wages (compared to the baseline results). This shows that by using a definition of "wages" similar to that in previous papers, we get a result using our data set that matches what other papers have documented using individual-level data sets and controlling for composition effects. As noted earlier, our measure of hourly earnings is based on total hours worked, including straight-time hours plus overtime hours less all leave hours. This is a key difference relative to many previous papers that, due to data constraints, use only straight-time hours to compute hourly earnings. Not taking into account the increase in overtime hours during booms can render hourly earnings, calculated based on just straight-time hours, artificially high in booms. This difference contributes to the weaker procyclicality in our earnings result compared to those in previous papers.

However, since 2008, earnings have also become countercyclical, similar to our results for straight-time wages. While the recent literature has emphasized the procyclicality of earnings, in fact there is evidence of changes in earnings cyclicality across different periods. For instance, Abraham and Haltiwanger (1995) note that real earnings were procyclical in the 1970s and 1980s and acyclical or countercyclical in the 1950s and 1960s. Sumner and Silver (1989) find that real earnings were countercyclical from 1948 to 1971 but procyclical from 1966 to 1980. The former period is also studied by Neftci (1978) and Geary and Kennan (1982), who report countercyclical earnings variation, and the latter period is covered by Bils (1985) and many others who report procyclical earnings.

Table 9 also shows the patterns of cyclicality for two major benefit expenditures, health insurance and Social Security. Health insurance benefits are countercyclical over the entire sample period and have become even more countercyclical after 2007. Social Security expenditures were strongly procyclical before 2008 but have become increasingly countercyclical since then. Not all benefit expenditures have become

[19]We do not use negative-GDP-growth periods in our main extension results because there are fewer years with negative GDP growth than years with below-trend GDP.

more countercyclical; table E.17 in the online appendix contains results showing which ones have experienced changes in cyclicality.

V. Conclusion

In this paper, we analyze the cyclical variation of total labor costs, including expenditures on employer-provided benefits, from the perspective of firms. Using only wage and earnings data, as is the case in the extensive prior literature on this topic, provides an incomplete picture of the cyclical dynamics of firms' labor costs. Drawing on BLS employer survey data over the period 1982 to 2018, we find that real straight-time wages have become countercyclical since the financial crisis. This is also the case for average total nonwage benefit expenditures and average total labor costs, although to a lesser extent, which indicates that firms adjust wages and benefit expenditures to slightly different degrees over the business cycle. We find that the increasing countercyclicality of labor costs not only prevailed during the Great Recession but also continued into the subsequent recovery and is largely attributable to periods with below-trend GDP rather than those with above-trend GDP.

To ensure our results are not specific to our data set, we also checked if using alternative sources of aggregate data makes a difference to our key conclusion. Results from these experiments are available in the working paper version of this paper (Gu & Prasad, 2018). We used two alternative data sources used in previous papers on this topic: CPS data on the median usual weekly earnings of full-time wage and salary workers and Current Employment Survey (CES) data on the average hourly earnings of production and nonsupervisory employees. The latter data set is the one that Abraham and Haltiwanger (1995) used. We used data from 1980 to the middle of 2017 (2017Q2 for the CPS and 2017M7 for the CES) and a similar variety of detrending procedures and deflators as in our analysis using the BLS NCS data. All of these results point to the countercyclicality of earnings after the financial crisis. These results are, of course, only suggestive since we do not control for composition effects using the microdata. Nevertheless, they reinforce the point made by Abraham and Haltiwanger (1995) that the time period appears to be the crucial factor in determining the cyclicality of wages and earnings.

Our finding of stronger countercyclical variation in average labor costs does not imply that their cyclicality has changed permanently after 2007, as the literature review shows that patterns of cyclicality can vary over time periods. However, our finding may reflect a combination of stronger downward nominal rigidities and procyclical inflation, potentially reflecting the rising relative importance of aggregate demand shocks. In particular, due to the decline in inflation, the Great Recession led to an upward spike in real wages despite rising unemployment, intensifying the countercyclical variation in firms' average labor costs. More research is needed on how nominal rigidities react to booms and busts based on the method in Barattieri, Basu and Gottschalk (2014) and Fallick, Lettau and Wascher (2016).

Our paper constitutes a first step in understanding firms' job-level average labor cost dynamics, which has potentially important implications for models of job-level compensation-adjustment behavior by firms. It is possible that recessions induce labor market adjustments in worker composition at the within-job level, a phenomenon that the data set used in this paper cannot capture. More work is needed to understand within-job worker-productivity and worker-composition changes over the business cycle. We leave a more detailed investigation of explanations for our findings for future research. Finally, we note that our findings are consistent with and provide a different perspective on the puzzle of the vanishing procyclicality of labor productivity (McGrattan & Prescott, 2012; Gali & van Rens, 2017) and could have implications for the effectiveness of monetary policy in counteracting increases in unemployment during business cycle downturns.

REFERENCES

Abraham, Katharine G., and John C. Haltiwanger, "Real Wages and the Business Cycle," *Journal of Economic Literature* 33 (1995), 1215–1264.

Anger, Silke, "The Cyclicality of Effective Wages within Employer–Employee Matches in a Rigid Labor Market," *Labour Economics* 18 (2011), 786–797.

Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk, "Some Evidence on the Importance of Sticky Wages," *American Economic Journal: Macroeconomics* 6 (2014), 71–101.

Basu, Susanto, and Christopher L. House, "Allocative and Remitted Wages: New Facts and Challenges for Keynesian Models" (pp. 297–354), in John B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics* (Amsterdam: Elsevier, 2016).

Beaudry, Paul, and John DiNardo, "The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Micro Data," *Journal of Political Economy* 99 (1991), 665–688.

Bils, Mark J., "Real Wages over the Business Cycle: Evidence from Panel Data," *Journal of Political Economy* 93 (1985), 666–689.

Brandolini, Andrea, "In Search of a Stylised Fact: Do Real Wages Exhibit a Consistent Pattern of Cyclical Variability?" *Journal of Economic Surveys* 9 (1995), 103–163.

Chung, Wankyo, "Fringe Benefits and Inequality in the Labor Market," *Economic Inquiry* 41 (2003), 517–529.

Daly, Mary C., and Bart Hobijn, "The Intensive and Extensive Margins of Real Wage Adjustment," Federal Reserve Bank of San Francisco working paper (2016).

——— "Composition and Aggregate Real Wage Growth," *American Economic Review* 107 (2017), 349–52.

Daly, Mary C., Bart Hobijn, and Theodore S. Wiles, "Dissecting Aggregate Real Wage Fluctuations: Individual Wage Growth and the Composition Effect," Federal Reserve Bank of San Francisco working paper (2012).

Devereux, Paul J., "The Cyclicality of Real Wages within Employer-Employee Matches," *Industrial and Labor Relations Review* 54 (2001), 835–850.

Elsby, Michael W. L., Donggyun Shin, and Gary Solon, "Wage Adjustment in the Great Recession and Other Downturns: Evidence from the United States and Great Britain," *Journal of Labor Economics* 34 (2016), S249–S291.

Fallick, Bruce C., Michael Lettau, and William L. Wascher, "Downward Nominal Wage Rigidity in the United States during and after the Great Recession," Federal Reserve Board of Governors finance and economics discussion series 1 (January 2016).

Fernald, John, "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Federal Reserve Bank of San Francisco working paper 2012-19 (2014).

Forsythe, Eliza, "Why Don't Firms Hire Young Workers during Recessions?" UIUC working paper (2017).
Gali, Jordi, and Thijs van Rens, "The Vanishing Procyclicality of Labor Productivity," University of Warwick economics research paper series 1062 (2017).
Geary, Patrick T., and John Kennan, "The Employment-Real Wage Relationship: An International Study," *Journal of Political Economy* 90 (1982), 854–871.
Gertler, Mark, Christopher Huckfeldt, and Antonella Trigari, "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires," New York University and Cornell University working paper (July 2018).
Gittleman, Maury, and Brooks Pierce, "An Improved Measure of Inter-Industry Pay Differentials," *Journal of Economic and Social Measurement* 38 (2013) 229–242.
Gu, Grace W., "Employment and Cyclical Cost of Worker Benefits," *Review of Economics Dynamics* 28 (2018), 96–120.
Gu, Grace Weishi, and Eswar Prasad, "New Evidence on Cyclical Variation in Labor Costs in the U.S.," NBER working paper 24266 (January 2018).
Hagedorn, Marcus, and Iourii Manovskii, "Job Selection and Wages over the Business Cycle," *American Economic Review* 103 (2013), 771–803.
Hallock, K. F., *Pay: Why People Earn What They Earn and What You Can Do Now to Make More* (Cambridge: Cambridge University Press, 2012).
Haltiwanger, John, Henry Hyatt, Lisa B. Kahn, and Erika McEntarfer, "Cyclical Job Ladders by Firm Size and Firm Wage," NBER working paper 23485 (June 2017).
Hart, Robert A., "Piece Work Pay and Hourly Pay over the Cycle," *Labour Economics* 15 (2008), 1006–1022.
Hart, Robert A., and J. Elizabeth Roberts, "Industrial Composition, Methods of Compensation and Real Earnings in the Great Depression," *National Institute Economic Review* 226 (2013), R17–R29.
Hart, Robert A., James R. Malley, and Ulrich Woitek, "Real Earnings and Business Cycles: New Evidence," *Empirical Economics* 37 (2009), 51–71.
Keane, Michael, Robert Moffitt, and David Runkle, "Real Wages over the Business Cycle: Estimating the Impact of Heterogeneity with Micro Data," *Journal of Political Economy* 96 (1988), 1232–1266.
Kudlyak, Marianna, "The Cyclicality of the User Cost of Labor," *Journal of Monetary Economics* 68 (2014) 53–67.
Kydland, Finn E., and Edward C Prescott, "The Workweek of Capital and Its Cyclical Implications," *Journal of Monetary Economics* 21 (1988), 343–360.
Lebow, David E., Raven E. Saks, and Beth Anne Wilson, "Downward Nominal Wage Rigidity: Evidence from the Employment Cost Index," *B.E. Journal of Macroeconomics* 3 (2003), 1–30.
Lee, Yoonsoo, and Toshihiko Mukoyama, "Entry, Exit, and Plant-Level Dynamics over the Business Cycle," University of Virginia working paper (2012).
Lettau, Michael K, "New Estimates for Wage Rate Inequality Using the Employment Cost Index," *Journal of Human Resources* 38 (2003), 792–805.
Makridis, Christos, and Maury Gittleman, "Does 'Performance Pay' Pay? Wage Flexibility over the Great Recession," Stanford University and Bureau of Labor Statistics working paper (2017).
Martins, Pedro S., Gary Solon, and Jonathan P. Thomas, "Measuring What Employers Do about Entry Wages over the Business Cycle: A New Approach," *American Economic Journal: Macroeconomics* 4 (2012), 36–55.
McGrattan, Ellen R., and Edward C. Prescott, "The Labor Productivity Puzzle," Federal Reserve Bank of Minneapolis working paper 694 (2012).
Monaco, Kristen, and Brooks Pierce, "Compensation Inequality: Evidence from the National Compensation Survey," *Monthly Labor Review* (July 2015), 1–11.
Neftci, Salih N., "A Time-Series Analysis of the Real Wages-Employment Relationship," *Journal of Political Economy* 86 (1978), 281–291.
Parker, Jonathan A., and Annette Vissing-Jorgensen, "The Increase in Income Cyclicality of High-Income Households and Its Relation to the Rise in Top Income Shares," *Brookings Papers on Economic Activity* (2010), 1.
Pierce, Brooks, "Compensation Inequality," *Quarterly Journal of Economics* 116 (2001), 1493–1525.
——— "Recent Trends in Compensation Inequality" (pp. 63–98) in Katharine G. Abraham, James R. Spletzer, and Michael J. Harper, eds., *Labor in the New Economy* (Chicago: University of Chicago Press, 2010).
Pissarides, Christopher A., "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?" *Econometrica* 77 (2009), 1339–1369.
Shin, Donggyun, "Cyclicality of Real Wages among Young Men," *Economics Letters* 46 (1994), 137–142.
Shin, Donggyun, and Gary Solon, "New Evidence on Real Wage Cyclicality within Employer-Employee Matches," *Scottish Journal of Political Economy* 54 (2007), 648–660.
Solon, Gary, Robert Barsky, and Jonathan A. Parker, "Measuring the Cyclicality of Real Wages: How Important Is Composition Bias?" *Quarterly Journal of Economics* 109 (1994), 1–25.
Sumner, Scott, and Stephen Silver, "Real Wages, Employment, and the Phillips Curve," *Journal of Political Economy* 97 (1989), 706–720.
Swanson, Eric T., "Measuring the Cyclicality of Real Wages: How Important Is the Firm's Point of View?" this REVIEW 86 (2004), 362–377.
——— "Real Wage Cyclicality in the Panel Study of Income Dynamics," *Scottish Journal of Political Economy* 54 (2007), 617–647.
Tian, Can, "Firm-Level Entry and Exit Dynamics over the Business Cycles," *European Economic Review* 102 (2018), 298–326.
Xu, Huanan, and Kenneth A. Couch, "The Business Cycle, Labor Market Transitions by Age, and the Great Recession," *Applied Economics* 49 (2017), 5370–5396.