

NBER WORKING PAPER SERIES

ALLOCATIVE AND REMITTED WAGES: NEW FACTS AND CHALLENGES
FOR KEYNESIAN MODELS

Susanto Basu
Christopher L. House

Working Paper 22279
<http://www.nber.org/papers/w22279>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2016

We thank Tyler Atkinson and, especially, Rich Ryan for outstanding research assistance. We thank the editors, John Taylor and Harald Uhlig. We thank our formal discussants, Eric Sims and Andreas Westermark, for their insights. We received many useful comments and suggestions from seminar participants at the University of Michigan, the Handbook of Macroeconomics conference at the Becker-Friedman Institute, University of Chicago, Boston College, and the conference on New Developments in the Macroeconomics of Labor Markets at CREI, Pompeu Fabra University. This research was supported in part by funding from the Michigan Institute for Teaching and Research in Economics (MITRE). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Susanto Basu and Christopher L. House. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Allocative and Remitted Wages: New Facts and Challenges for Keynesian Models
Susanto Basu and Christopher L. House
NBER Working Paper No. 22279
May 2016
JEL No. E24,E3,E31,E32

ABSTRACT

Modern monetary business-cycle models rely heavily on price and wage rigidity. While there is substantial evidence that prices do not adjust frequently, there is much less evidence on whether wage rigidity is an important feature of real world labor markets. While real average hourly earnings are not particularly cyclical, and do not react significantly to monetary policy shocks, systematic changes in the composition of employed workers and implicit contracts within employment arrangements make it difficult to draw strong conclusions about the importance of wage rigidity. We augment a workhorse monetary DSGE model by allowing for endogenous changes in the composition of workers and also by explicitly allowing for a difference between allocative wages and remitted wages. Using both individual-level and aggregate data, we study and extend the available evidence on the cyclicalities of wages and we pay particular attention to the response of wages to identified monetary policy shocks. Our analysis suggests several broad conclusions: (i) in the data, composition bias plays a modest but noticeable role in cyclical compensation patterns; (ii) empirically, both the wages for newly hired workers and the "user cost of labor" respond strongly to identified monetary policy innovations; (iii) a model with implicit contracts between workers and firms and a flexible allocative wage replicates these patterns well. We conclude that price rigidity likely plays a substantially more important role than wage rigidity in governing economic fluctuations.

Susanto Basu
Department of Economics
Boston College
140 Commonwealth Avenue
Chestnut Hill, MA 02467
and NBER
susanto.basu@bc.edu

Christopher L. House
University of Michigan
Department of Economics
238 Lorch Hall
Ann Arbor, MI 48109-1220
and NBER
chouse@umich.edu

1 Introduction

Since at least Hume (1742), sluggish adjustment of wages and prices has been thought to be central for understanding the monetary transmission mechanism. This is certainly true in modern New Keynesian models, of either the textbook variety or in medium-scale models that attempt to match economic data.¹ Loosely speaking, models with nominal rigidities reproduce many of the patterns featured in partial-equilibrium settings, and can make demand-determined output fluctuations consistent with both basic business cycle facts and observed reactions to changes in monetary policy. Beyond monetary non-neutrality, it is now understood that models with nominal rigidities also behave differently in response to real shocks. For example, models with nominal rigidities can create business-cycle comovements in response to intertemporal shocks (such as news about future technology, uncertainty, or financial frictions that change expected capital returns), even when flexible-price models would not display such comovements.²

It is thus important to understand the extent and importance of nominal price and wage rigidities. While we discuss both wage and price rigidities in this survey, we focus mostly on wage rigidity, for several reasons. First, Christiano, Eichenbaum and Evans (2005; CEE henceforth) and many successors have found that wage rigidity is quantitatively more important than price rigidity for explaining the effects of monetary shocks and for explaining cyclical fluctuations more generally. Second, attempts to decompose the cyclical behavior of the “labor wedge”—the gap between the marginal product of labor and the marginal rate of substitution between consumption and leisure—typically find that sluggish wage adjustment accounts for a large fraction of the observed cyclical behavior of the total wedge. That is, the wage markup appears more cyclical than the price markup.³ Third, there is broad agreement among researchers on the basic empirical facts regarding price rigidity, but there is no such consensus regarding the nature of wage rigidity. Following the initial work of Bilts and Klenow (2004), a large number of recent papers have investigated the frequency and magnitude of price changes and the pass-through from costs to prices. By contrast, there are fewer studies of wage rigidity, and the ones that exist often do not relate their results to macroeconomic models. Part of the reason for the greater uncertainty regarding wage behavior is that wages

¹For a classic textbook treatment, see Woodford (2003). The canonical medium-scale models are due to Christiano, Eichenbaum and Evans (2005) and Smets and Wouters (2007).

²See, for example, Basu and Bundick (2017). The basic issues were pointed out by Barro and King (1984).

³See, for example, Galí, Gertler and López-Salido (2007). However, their conclusion depends sensitively on the wage measure used. See, for example, Bilts, Klenow and Malin (2014).

are harder to measure, and it is difficult to know whether observed wages are allocative.

While there is a tendency to discuss price and wage rigidity as independent phenomena, this is incorrect at the macroeconomic level. In some cases, one might be able to take the *microeconomic* rates of wage and price rigidity—for example, exogenous Poisson hazard rates for adjusting wages or prices—as independent parameters. But as macroeconomic models make clear, the inertia of the aggregate price level—the extent of *macroeconomic* price rigidity—depends heavily on the rigidity of wages. Since most models assume that target prices are set as a constant markup on nominal marginal costs, the inertia of the price level depends on sluggish adjustment of marginal cost. Wages are the largest component of the marginal cost of producing real value added, and thus wage stickiness naturally reinforces price stickiness. Indeed, in most medium-scale models, wage stickiness is essential for obtaining price level inertia and thus, for example, persistent real effects of nominal shocks. Similarly, wage setting, for example by monopoly unions, will also be influenced by expectations of future price inflation. In another chapter in this *Handbook*, Taylor (2016) discusses microeconomic evidence on staggered wage and price setting and its implications for macroeconomic models.

Our main purpose in this survey is to discuss several definitions and measures of wage stickiness and cyclical wage adjustment and then ask what they imply for wage and price rigidity in a prototypical medium-scale macroeconomic model. Although this survey concentrates on wage measurement, we believe that it is difficult to assess the implications of data without reference to theory. Thus, we construct a medium-scale DSGE model based on CEE, but with more extensive modeling of different concepts of wages. The model distinguishes between four concepts that we discuss below: average earnings, average earnings adjusted for labor force composition, wages of new hires, and the user cost of labor. These distinct wage concepts behave in different ways in response to the monetary shock we study, so we can use the model to predict the behavior of these different concepts of wages to a monetary policy shock, which is our measure of an archetypal nominal aggregate shock.

Our use of a model to motivate the measurement has the effect of focusing attention on the wage and price statistics that we believe are most relevant for macroeconomics. These are the responses of prices and wages to identified aggregate shocks—that is, conditional correlations—rather than average business-cycle correlations or the average frequency of wage or price change. When there are both idiosyncratic and aggregate shocks, micro wages and prices may change frequently for reasons unrelated to aggregate fluctuations, but they may change only slowly in response to aggregate shocks. To concentrate attention on the statistics

that matter most for macroeconomics, we focus on the responses of nominal wages and prices to a monetary shock, which is the standard example of a nominal aggregate shock. Our focus on monetary shocks does not reflect a judgment that these shocks cause a significant fraction of business-cycle fluctuations. On the contrary, most of the available evidence suggests that monetary shocks account for a relatively small fraction of output volatility. But because they are identified using a consensus set of restrictions, and because these shocks would be neutral absent some nominal rigidities, they provide a valuable opportunity for assessing the performance of macroeconomic models.

To preview our findings, we argue that recent research provides suggestive evidence that the conceptually correct measure of the allocative wage is strongly procyclical. This finding contrasts sharply with typical estimates in the macro literature, which often claim that the real wage is roughly acyclical. We then discuss the implications of the new facts about wages for models with nominal rigidities, and find that these models struggle to explain the empirical facts regarding the effects of monetary policy shocks. We show that standard DSGE models can be augmented with realistic features to reproduce many of the wage patterns found in the recent literature but that these features typically pose serious problems for the ability of DSGE models of the monetary business cycle to match estimated reactions of other variables to monetary shocks. We argue that additional evidence on measured adjustment for allocative wages and on propagation mechanisms for monetary models is needed to reconcile the micro data with our understanding of the monetary transmission mechanism. The search for this evidence should be a high priority for future research.

This survey is structured as follows. After an initial overview of different concepts of “the wage” and a survey of the relevant literature, we take the workhorse model of CEE and extend it along several dimensions. The model enables us to define several different concepts of wages in a precise manner, and to derive predictions for their behavior in response to a monetary policy shock. Importantly, the model produces output corresponding to each wage concept, even though only one is perceived by workers and firms as the allocative wage. Thus, we are able to use the model to predict the behavior of both allocative and non-allocative wages in response to a monetary shock. We then relate the model implications to existing micro data studies on wages and prices, and indicate points where further evidence is needed.

With the concepts from the model in mind, we discuss evidence on wage cyclicality, drawing mostly on research using U.S. micro data. We focus on three issues that are crucial for interpreting the data but are generally not incorporated into macroeconomic models: com-

position bias in aggregate wage measures, the distinction between wages of new hires and wages of workers in continuing employment, and the distinction between spot wages and the expected time path of wages. In order to draw our statistics from a common source and ensure that they are comparable to one another, we construct the composition-corrected wages, the wages of new hires and the user cost of labor using micro data from the National Longitudinal Survey of Youth 1979 (henceforth NLSY). We then show how the different measures of real wages respond to a monetary policy shock. Our main conclusion is that real wages, correctly defined and measured, are quite procyclical, in contrast to average hourly earnings, which are basically acyclical.

We conclude by confronting the model we have developed with the empirical evidence based on micro data. We find that the allocative wage needs to be quite flexible in order to match the behavior of the real wage we estimate from the micro data. However, in order to match the behavior of average hourly earnings, it is useful to combine the flexible allocative wage with a remitted (observed) wage that changes only infrequently, also in line with evidence from micro data. A model with sticky prices, flexible wages, and implicit labor contracts comes closest to matching the impulse responses of key variables to a monetary policy shock. However, a standard medium-scale DSGE model with flexible wages struggles to match the estimated high persistence of the output response to a monetary shock. Many recent models have used wage stickiness, justified in a variety of ways, as an important propagation mechanism for shocks. With the micro data indicating that wage flexibility is a better assumption than wage stickiness, macroeconomists need to search for new propagation mechanisms in order to match the observed persistence of output fluctuations, especially to monetary shocks.

2 Defining “the wage”

Macroeconomic models are typically populated by a large number of identical worker-consumers, who supply labor along the intensive margin in a spot market. In this setting, it is easy to define the wage: it is the current payment at time t for an extra unit of labor supplied in the same period. If the world were as simple as the model, “the wage” would be easy to measure. Unfortunately, nearly all of the assumptions about the labor market noted above are violated in important ways in the data, making the effort to measure wage behavior far more complicated.

First, workers are not homogeneous. This obvious fact would not necessarily create a

problem for measuring wage cyclicality if the hours of workers of different types increased and decreased in synchronized fashion. Then one could define a representative worker as a worker with human capital equal to the weighted average of the human capital of all workers in the population, and show that the average wage we observe in the data is also the wage commanded by the representative worker's fixed bundle of human capital characteristics. Unfortunately, the composition of the labor force changes over the cycle. Stockman (1983) conjectured and Solon, Barsky and Parker (1994) confirmed that the hours of low-paid workers are more cyclical than average. Hence, low-paid workers account for a larger share of labor payments in booms than in recessions. Thus, the cyclical behavior of the aggregate (average) real wage is not the cyclicality of the wage paid to a representative worker with fixed human capital characteristics, which is the implicit or explicit concept in almost all business-cycle models. As we shall see, correcting for this composition bias shows that the wage paid to a representative worker with fixed characteristics is considerably more procyclical than the average wage in the data. This is also the conclusion of an important early paper by Bilts (1985).

Second, since most workers are in long-term relationships with their employers, the labor market is not a spot market. Thus, their spot wages are not necessarily what the firm perceives as the marginal cost of labor, which is the key concept for most macroeconomic purposes. Barro (1977) used the idea of an implicit contract to criticize the "right to hire" model of wage stickiness, where workers propose a fixed, possibly nominal, wage, and firms choose employment (or hours) along their labor demand curves. He showed that other contracts would increase the payoff to both parties in the bargain, and suggested that an efficient contract would equate workers' marginal rates of substitution between consumption and leisure to firms' marginal products of labor in every period, with the total compensation for labor paid out to workers in "installment payments" over the lifetime of the worker-firm association. This reasoning follows the classic work of Becker (1962), who showed in a neoclassical setting that only the present discounted value of the wages paid by firms to workers over the length their association is allocative for employment. Holding the present value of wage payments constant, the time path of *remitted* wages can have any shape without affecting real outcomes.

Thus, one needs to know how the annuity value of the expected present value of wages, which one might conceptually regard as the "permanent wage," changes in response to changing economic conditions. By analogy to the permanent income hypothesis, the behavior of the permanent wage, not the current wage, is what matters to an optimizing worker or firm. Much of the search literature implicitly ties the behavior of the permanent wage to that of the wage

for new hires, but to the extent the two differ, the permanent wage matters more. If workers and firms are in long-term associations and the permanent wage is the correct measure of the cost of labor input, then it is possible for the observed average wage to appear insensitive to cyclical fluctuations (“sticky”) even if the correct allocative wage is flexible. This conclusion was anticipated by Barro (1977, p. 316), who wrote, “In fact, the principal contribution of the contracting approach to short-run macro-analysis may turn out to be its implication that some frequently discussed aspects of labor markets are a facade with respect to employment fluctuations. In this category one can list sticky wages”

3 Background and related literature

We survey the history of research on wage cyclicality, with an eye to distilling and interpreting the evidence on the cyclicality of the marginal cost of an efficiency unit of labor to the average firm. Given the central importance of this subject for macroeconomics and the vast number of papers written about it over several decades, we can only touch on the key ideas that are most closely related to our investigation of the topic. Fortunately, a number of fine surveys of wage cyclicality have been written over the years, and we refer the reader to those for more in-depth discussions of particular issues.⁴

Our survey ranges over estimates of both nominal wage rigidity and real wage cyclicality. Both are important for assessing modern “medium-scale” macroeconomic models, and especially the ability of these models to reproduce the real effects of monetary policy shocks as observed in the data.⁵ Ultimately what matters is the behavior of the “shadow” real wage facing firms. The shadow wage is the marginal cost of a unit of labor to the firm, which may or may not be what economists can readily observe in the data. To the extent that the shadow real wage is insensitive to changes in labor demand, it may be due to either real features of the economy (e.g., elastic labor supply) or to wage rigidity (even if notional labor supply is inelastic), or both. From the standpoint of a firm, both rationales can explain why total hours fluctuate nearly as much as GDP over the business cycle.⁶ However, one generally needs nom-

⁴For example, see Abraham and Haltiwanger (1995).

⁵The “narrative” approach to documenting monetary non-neutrality is exemplified by Friedman and Schwartz (1963) and is developed further by Romer and Romer (1989). The modern VAR literature on estimating the effects of monetary policy shocks originates with Bernanke and Blinder (1992). See Christiano, Eichenbaum and Evans (1999) for a survey of the VAR approach. For an alternative to both the narrative and VAR approaches to identification, see Romer and Romer (2004).

⁶For a discussion of the basic statistical regularities of business cycles, see the survey by Stock and Watson

inal rigidity somewhere in the model to explain why a nominal shock has real effects. But even if nominal wages can adjust freely, acyclical real wages combined with nominal price rigidity can explain why monetary shocks are generally estimated to have sizable and persistent effects on output but little effect on nominal wages.⁷

In keeping with our focus on the cost of labor to a firm, we ignore a number of related and important topics on the behavior of wages. In particular, we do not survey the literature on the *reasons* for wage rigidity, such as efficiency-wage models or insider-outsider models. We also touch only briefly on search models of the labor market, although there is an important literature combining search models with New Keynesian macroeconomics.⁸

3.1 Wage rigidity in historical data

In the *General Theory*, Keynes (1936) made nominal wage rigidity the centerpiece of his theory of aggregate supply. His framework predicted that procyclical changes in prices, combined with money wage inertia, would result in countercyclical real wages. Since money wages and prices should move in the same direction, the *General Theory* predicted that nominal and real wage changes should be negatively correlated. This prediction was tested but not confirmed by Dunlop (1938) and Tarshis (1939), who took their findings as *prima facie* evidence against the hypothesis of nominal wage rigidity.⁹

On the other hand, a variety of papers have examined the historical data and find clear evidence of nominal wage rigidity in the late 19th and early 20th centuries. In a classic paper, Eichengreen and Sachs (1985) used cross-sectional data for 10 countries to show that over the Great Depression period there was a negative relationship between output and real wages. They also show that countries which remained on the gold standard had low output and high real wages, while countries that left gold early experienced high output and low real wages. Bernanke and Carey (1996) extended the Eichengreen-Sachs sample to 22 countries, examined dynamics by using panel data, and performed a number of other econometric and economic robustness tests, all of which supported the basic hypothesis of nominal wage rigidity. As Bernanke and Carey emphasized, they were studying the consequences of a purely nominal

(1999). For an interpretation in a neoclassical model, see King and Rebelo (1999).

⁷For a development of this argument, see Ball and Romer (1990), Kimball (1995), and Woodford (2003, chapter 3).

⁸See Walsh (2003), Ravenna and Walsh (2008) and Gertler and Trigari (2009).

⁹Pencavel (2015) discusses the early Keynesian literature on wages. Galí (2013) relates the controversies of the 1930s to modern New Keynesian analysis.

shock – the transition from the Gold Standard to a fiat money regime – which took place at different times in different countries. The fact that, when countries left the Gold Standard, real wages systematically fell while output rose suggests, first, that purely monetary shocks can have significant real effects, and second, that expansions in nominal aggregate demand raise output by lowering real wages, giving firms an incentive to employ more labor.

Evidence suggests, however, that nominal wage rigidity is not a universal feature of labor markets. Hanes (1993) argues that wages became inflexible around the time of widespread large-scale industrial production and episodes of labor unrest, which he dates to 1890. Hanes argues that nominal wages appear to stay rigid at least through World War I, although there is some weak evidence that they become somewhat more flexible starting in the 1970s. Basu and Taylor (1999) use both time-series and cross-country data to investigate wage cyclicality. They concentrate on real wages, and interpret their results in light of the prediction of countercyclical real wages in the *General Theory*. They find that there is no definite sign of the relation of real wage movements to the business cycle. For their sample of 13 countries, real wages were slightly procyclical in the period before World War I and somewhat countercyclical in the interwar period, before becoming decidedly more procyclical after World War II. Thus, their evidence supports the idea that real wages have become more procyclical over time.

Hanes (1996) and Huang, Liu and Phaneuf (2004) seek to explain the changing cyclicality of real wages over a long historical period of more than 100 years. Both papers propose an explanation that relies on *prices* becoming more sticky over time—due to a larger number of stages of processing for goods in Hanes’s case, and due to a larger output elasticity of intermediate inputs in the work of Huang *et al.* (Both papers also include mechanisms that deliver countercyclical price markups, which are important for the result.) Thus, in these works the change in real wage cyclicality over time emerges from changes that take place in the *product* market rather than the labor market. Whether or not this hypothesis is ultimately adjudged to be plausible, it is a sobering reminder that general-equilibrium effects complicate the interpretation of simple business-cycle correlations, especially in a macroeconomic setting.

3.2 Wage rigidity in modern data

In a benchmark survey of business-cycle facts, Stock and Watson (1999) find an almost zero correlation between detrended real average hourly earnings and detrended GDP in postwar

U.S. data.¹⁰ This and similar findings (for example, that labor productivity is also approximately acyclical in U.S. data), has led modelers to emphasize preferences or institutions leading to effective labor supply functions that are nearly infinitely elastic with respect to the wage.¹¹ Of course, a setting in which both nominal wages and prices are slow to adjust can also produce a real wage that is approximately acyclical regardless of preferences; this is the path taken by Christiano, Eichenbaum and Evans (2005) and Smets and Wouters (2007), among others.

A near-zero average correlation between output and real wages of course admits another interpretation. It might be the case that real wages fall in response to some shocks (perhaps expansionary monetary shocks) and rise with others (perhaps positive technology shocks). If the two types of shocks are roughly equally important in the data, then on average the real wage may be acyclical. Of course, this small average correlation could hide important conditional correlations that might be far from zero. The “multiple shocks” hypothesis could also explain the instability of the correlation between output and real wages in the historical data discussed above. The change in the correlation between the cyclical component of wages and the cyclical component of output might just reflect the changing contributions of the two types of shocks over different sub-periods.¹²

Sumner and Silver (1989) present evidence in favor of this hypothesis. They classify periods dominated by “demand shocks” as those in which output and the price level move in the same direction, while periods where the two variables move in opposite directions are classified as being dominated by “supply shocks.” They find that wages move countercyclically in response to demand shocks but procyclically in response to supply shocks, a finding that is consistent with an augmented version of the “Old Keynesian” model.¹³

Huang *et al.* (2004) argue against the “multiple shocks” interpretation of the changing correlation between output and real wages over the business cycle. Their main argument is that the observed change in cyclicalities applies to conditional correlations and not just

¹⁰Stock and Watson detrend both series using the band-pass filter, set to isolate fluctuations lasting between 6 and 32 quarters.

¹¹For models in which the wage is insensitive to output fluctuations, see Hansen (1985), Rogerson (1988), Greenwood, Hercowitz and Huffman (1988) and, in non-neoclassical settings, Solow (1979) and Hall (2005).

¹²Geary and Kennan (1982) present evidence from the manufacturing sectors of 12 OECD countries suggesting that wage cyclicalities change significantly depending on the time period studied. They also find that the choice of deflator (a consumer price or a product price index) can make a noticeable difference. Presumably a product price index is more appropriate for testing the hypothesis that employment and wages move along a stable labor demand curve.

¹³Fleischman (1999) comes to similar conclusions using a structural VAR with long-run restrictions to identify various categories of shocks.

simple correlations. For example, they cite the evidence of Eichengreen and Sachs (1985) and Bernanke and Carey (1996) discussed above to establish that real wages decline in response to expansionary monetary shocks during the interwar period, but then refer to evidence from structural VARs run on post-war data to show that real wages rise modestly in response to expansionary monetary shocks in the recent period. The empirical results in CEE, for example, show the real wage rising slightly several quarters after a monetary expansion and then declining slightly after 10 quarters, although at no horizon is the real wage response statistically significant in either direction. In their data, one can reject the hypothesis that real wages fall significantly in response to an expansionary monetary policy shock.

3.3 “The wage” in aggregate and micro data

Most papers in the historical and macro literatures examining the behavior of wages use aggregate wage data.¹⁴ Unfortunately, aggregate data are subject to a composition bias that makes aggregate (average) wage rates less procyclical than the wages of individual workers. Stockman (1983) conjectured that low-productivity (and hence low-wage) workers would have the most cyclical employment—they would be the most likely to be fired in recessions, but also the most likely to be hired in booms. If true, then the aggregate wage (either approximately or exactly the labor-income-weighted average of the individual wage rates) would be less procyclical than the wages of individual workers, since low-wage workers would earn a larger share of labor income in booms.

Bils (1985) used individual panel data from the National Longitudinal Survey of Young Men covering the period 1966-1980, and found that wages in micro data appear extremely procyclical: a one percentage point decline in the unemployment rate is associated with a rise in real wages of 1.5-2 percent. While Bils finds a countercyclical composition bias in aggregate wage measures, consistent with Stockman’s conjecture, he argues that this bias does not contribute significantly to his finding of a procyclical wage, since aggregate wage data also show a very procyclical real wage over this sample period. Other than the sample period, Bils attributes his finding of a procyclical wage to his inclusion of overtime earnings into his wage measure.

¹⁴Ironically, the historical literature is more likely to use disaggregated data, even though high-quality data are scarce for earlier periods. For example, Hanes (1993) and Hanes and James (2003) use fixed-weight indexes of wages in narrowly-defined occupations, a wage concept akin to the Employment Cost Index (ECI) produced by the Bureau of Labor Statistics.

Solon, Barsky and Parker (1994) also investigate Stockman’s hypothesis of composition bias in aggregate wage data using longitudinal microdata, in their case from the Panel Study of Income Dynamics (PSID) for the years 1967-87. Unlike Bils, they find that composition bias played a substantial part in reducing the apparent cyclical of the aggregate real wage over their sample period.¹⁵ Controlling for composition bias, they find that wages are about twice as procyclical as they appear in aggregate data. Solon, Barsky and Parker interpret their finding as consistent with movements of wages and employment along a stable aggregate labor supply curve with a labor supply elasticity between 1.0 and 1.4. They suggest that their finding is more consistent with models that predict procyclical real wages than is the usual stylized “fact” of acyclical wages, and note that both neoclassical and New Keynesian theories of the business cycle tend to predict that wages should be quite procyclical.

It may appear that the finding of strongly procyclical real wages is at odds with the finding, discussed above, that U.S. labor productivity is roughly acyclical. In fact, there is no inconsistency. Once one admits that labor is heterogeneous, labor productivity needs to be measured in terms of output per efficiency unit of labor rather than output per raw labor hour. Since the lower-wage workers added in a boom contribute less in efficiency units of labor than their contribution of work hours would suggest, labor productivity correctly measured is also more procyclical than it appears in aggregate data. In fact, when it comes to measuring unit labor cost (the hourly wage divided by output per hour worked), the composition corrections for wages and labor productivity exactly offset. Thus, in the Cobb-Douglas case, the unadjusted unit labor cost measures used in the literature as a straightforward measure of the markup of price over marginal cost are not biased by cyclical changes in composition.¹⁶

Using data from the CPS, Daly and Hobijn (2016) come to similar conclusions regarding the “intensive” and “extensive” margins of wages. Along the intensive margin—wage changes of continuously employed individuals—wages are clearly procyclical. The extensive margin consists of cyclical changes in employment, which are concentrated among workers with lower-than-average earnings. The extensive margin makes the aggregate wage appear countercyclical. The two effects combine to make the aggregate real wage appear acyclical on average, although Daly and Hobijn note that the relative strength of the two margins varies

¹⁵Solon, Barsky and Parker (1994) argue that the estimates in Bils (1985) apply to composition bias within narrowly defined categories of workers but do not fully reflect compositional changes across groups, and thus understate the aggregate effects of compositional changes.

¹⁶For measures of unit labor costs interpreted as the markup, see Rotemberg and Woodford (1991, 1999), Galí and Gertler (1999), Sbordone (2002) and Nekarda and Ramey (2013).

over time, and so does the cyclicality of the aggregate wage.

Elsby, Shin and Solon (2016) revisit the issues of wage cyclicality and composition bias, focusing on the experience of the U.S. and the U.K. during the Great Recession of the 2000s. They use longitudinal microdata for both countries, but note that in many respects the U.K. data are preferable, first because of the larger sample size, and second because the data on earnings and work hours come from the payroll data of employers, which are generally thought to be significantly more accurate than workers' recollections.

Elsby *et al.* report somewhat nuanced findings. They confirm the earlier microdata-based result for the U.S., that men's real wages are significantly procyclical, but find that their wages were less cyclical in the Great Recession than the experience of previous large recessions would suggest. Women's real wages, which had been rising sharply in the period since 1979, stagnated during the Great Recession. However, Elsby *et al.* find some hints that women's wage growth was declining prior to the last recession, and thus conclude that it is too early to tell whether the lack of wage growth in the most recent recession is due to women's wages being highly procyclical or whether it is due to a shift to a new trend of slow wage growth. In one respect, the findings for the U.K. for both men and women are similar to those of the U.S. men: real wages fell significantly in the Great Recession. But the variation in wage cyclicality across recessions is more or less the opposite in the two countries: wages in the U.K. were much more procyclical in the Great Recession than in previous recessions, while the opposite was true in the U.S., at least for men. Another major difference is that composition bias appears to matter much less for measuring wage cyclicality in the U.K. than it does in the U.S. These differences are important to bear in mind when drawing lessons from the empirical results we report in this chapter, which are based exclusively on U.S. data.

3.4 Downward nominal wage rigidity

A long strand of Keynesian analysis is based on the hypothesis that wages are more rigid downward than upward. For example, Tobin's (1972) Presidential Address suggested that workers care about relative wages, implying that they would tend to resist asynchronized wage cuts but might tolerate a neutral mechanism like inflation that cuts all real wages proportionally. This hypothesis of an asymmetry between wage increases and decreases has been the focus of a substantial literature in labor economics. One of the first researchers to address this question using micro data is McLaughlin (1994), who failed to find much evidence

of asymmetry. Later work by Card and Hyslop (1997), Kahn (1997), and Lebow, Saks and Wilson (1999) found evidence of downward nominal wage rigidity (DNWR), including a large spike in the observed wage change distribution at zero (unchanged nominal wages), and a smaller number of small wage declines than small wage increases. Gottschalk (2005) performed an analysis of micro data on wage changes using an econometric procedure to correct for measurement error in self-reported wages, and found substantial downward nominal rigidity. The more recent papers thus suggest significant downward rigidity of nominal wages and, given the low-inflation environment that has prevailed since the mid-1990s, of real wages as well.

Hanes and James (2003) examine historical data on individual wage changes in another low-inflation period, the years 1841-1891. Applying the tests for asymmetry in wage changes developed in the literature analyzing modern wage data, they find no evidence of DNWR. They interpret their results as suggesting that an aversion to nominal wage cuts is not a fundamental feature of worker preferences. They note, however, that their results do not contradict the hypothesis that institutions may have changed in such a way as to make DNWR desirable in the modern era, perhaps as a boost to worker morale and hence productivity, as suggested by Bewley (1999). Another cautionary note in interpreting the consequences of DNWR comes from Elsby (2009). Elsby begins by assuming that an aversion to nominal wage cuts is indeed a feature of workers' preferences, but then shows that preferences of this unusual form often have only a small effect on equilibrium outcomes. The reason is that dynamically optimizing firms, when confronted with a workforce that exhibits DNWR preferences, will delay nominal wage increases, thus keeping a cushion that allows real wages to rise without causing substantial employment declines if the constraint on nominal wage declines comes into play. Elsby's model suggests that it is possible to find substantial evidence of DNWR in micro data while observing few macro consequences of such asymmetric behavior. (Indeed, the evidence supporting the macroeconomic implications of DNWR does not seem to be overwhelming: see, for example, Akerlof, Dickens and Perry (1996).)

A number of observers have suggested that DNWR is a good explanation for the recent observation that inflation has been slow to decline during protracted and severe recessions (for example, in Japan starting in the 1990s and the Great Recession in the U.S. and other countries in the 2000s). To our knowledge, no formal evidence of this connection has been offered. However, Schmitt-Grohé and Uribe (2013) suggest that if DNWR exists, then there is a strong case for higher inflation in the Eurozone to lower real wages and stimulate employment.

3.5 Wage change frequency in micro data

Canonical New Keynesian models, such as CEE and Smets and Wouters (2007), follow Blanchard and Kiyotaki (1987) in assuming that wages for each type of worker are set by a monopoly union. Like the monopolists in the product market, the monopoly unions are subject to the Calvo friction when changing nominal wages. Thus, just as the frequency of price adjustment is important for quantifying the significance of nominal price rigidities, the frequency of wage changes in micro data is important for assessing the plausible degree of inertia in nominal wage rates. However, unlike the large literature on the rigidity of micro-level prices, there are few studies of the frequency of change of individual wages.

Barattieri, Basu and Gottschalk (2014) provide one such study using micro data from the U.S. Survey of Income and Program Participation (SIPP). One advantage of the SIPP is that participants are surveyed three times a year, unlike participants in the PSID, who are surveyed annually. SIPP data are thus more suitable for high-frequency analysis of individual wages.¹⁷ Because all large surveys of U.S. micro data on individual wages use self-reported wages, a substantial fraction of the paper of Barattieri *et al.* (2014) is devoted to proposing a method to correct for measurement error in such a way that one can recover a consistent estimate of the frequency of individual wage adjustment. Such studies can be carried out more easily in countries where one can obtain access to administrative data, which presumably have less measurement error. Individual wage change probabilities have been analyzed for France by Le Bihan, Mortornès and Heckel (2012), for Luxembourg by Lünemann and Wintr (2009), and for Iceland by Sigurdsson and Sigurdardottir (2011). (Researchers can access confidential administrative data sets for the U.S. as well, but these generally provide information on total earnings rather than hourly wage rates, which were the focus of Barattieri *et al.*)

As we shall see in the model of the next section, the frequency of changes in the observed wage at the individual level is an important parameter for calibrating implicit-contracting models of the labor market. (This is true even if, as in the model we present below, the observed wage need not be the allocative wage.) The estimate reported by Barattieri *et al.* is not directly applicable to the full U.S. labor market, since these authors restricted their sample to hourly paid workers. Here, we present new estimates for the frequency of wage

¹⁷Other well-known sources of micro wage data, the Current Population Survey (CPS) and the Employment Cost Index (ECI), do not provide sufficiently long time-series data on the wages of individual workers to be useful for this purpose.

changes for salaried workers using the methodology of Barattieri *et al.*¹⁸ The results are in Table 1.

The results for hourly paid workers, the first column, reproduce the first three lines of results for the “Overall” sample in Barattieri *et al.* (2014, Table 6). The new results for salaried workers are in the second column. Earnings per hour change even more frequently for salaried workers in the raw, self-reported data than they do for hourly paid workers. Nearly three-quarters of hourly earnings for salaried workers change each quarter. However, applying the iterative procedure of Gottschalk (2005) to correct for measurement error in wages reduces the estimate of the quarterly probability of actual wage changes for salaried workers to 6.1 percent. Unfortunately, this is not a consistent estimate of the desired probability due to the presence of Type I and Type II errors. Using the adjustment for the signal-to-noise ratio based on the work of Gottschalk and Huynh (2010), as presented in Barattieri *et al.* (2014), the final estimate of the quarterly probability of a change in earnings per hour of salaried workers is 20.9 percent. This figure is remarkably close to the probability of 21.1 percent for hourly paid workers in Table 1. In our model calibrations below, we generally set the quarterly frequency of an observed change in the remitted wage to 21 percent.

3.6 Implicit contracts, adjustment costs and real wage cyclicality

In a classic paper, Becker (1962) showed in a neoclassical setting that only the present discounted value of the wages paid by firms to workers over the length their association is allocative for employment. Holding the present value of wage payments constant, the time path of *remitted* (observed) wages could have any shape without affecting real outcomes. For example, firms and workers might agree to an implicit contract in which remitted wage payments are smoothed relative to changes in the allocative present value of wages, but the fact that the observed wage is smooth would not affect real outcomes. Barro (1977) used the idea of an implicit contract to criticize the “right to hire” model of wage stickiness, where workers propose a fixed wage, and firms choose employment (or hours) along their labor demand curves. He showed that other contracts would increase the payoff to both parties in the bargain, and suggested that an efficient contract would equate workers’ marginal rates of substitution between consumption and leisure to firms’ marginal products of labor in every period, with the total compensation for labor paid out to workers in “installment payments”

¹⁸We are greatly indebted to Alessandro Barattieri for these estimates.

over the lifetime of the worker-firm association.

Models where workers and firms have an implicit contract over the present value of wages clearly require some assumptions about the ability of the parties to commit. In some models, such as the one we present later, one simply assumes that commitment is feasible. An alternative is to assume adjustment costs to dissolving the match for one or both parties. Absent such costs, the party that is “ahead” in the installment payments would dissolve the match. The most popular current model of labor adjustment costs is based on search. Hall (2005) addressed Barro’s (1977) critique of allocative wage stickiness by showing that the allocative wage could be history-dependent and hence sticky within the Diamond-Mortensen-Pissarides model of search in the labor market, as long as the preset wage remains within the Nash bargaining set generated by that model. (This argument addressed the critique of the DMP model due to Shimer (2005), who identified the sharp procyclicality of the wage as the central reason why this canonical model fails to match the volatility of the unemployment and vacancy rates.) Hall and Milgrom (2008) showed that some wage stickiness could emerge from alternative-offer bargaining. Pissarides (2009) and Gertler and Trigari (2009) showed that in the search setting, the key allocative wage is that of new hires. Haefke, Sonntag and van Rens (2013) examine data from the Current Population Survey and conclude that the wages of newly hired workers are in fact much more procyclical than average hourly earnings of all employed workers.

Relative to the literature on composition bias, the main contribution of the search-based papers is to concentrate attention on a subset of wages, namely the wages of new hires. Thus, for example, Gertler and Trigari (2009) argue that the key statistic is whether new hires receive the same wages as workers currently employed by the firm they are joining, or whether they can be hired at different wages that better reflect current economic conditions.

Assuming that newly-hired workers expect to stay with their current employer for a significant length of time, it is intuitive that their expected labor compensation consists of the expected present value of the wages they will receive over the length of the association. In this case, what matters is actually not even the cyclicity of the spot wage of new hires *per se*, but the cyclicity of the expected present value of wage payments to new hires.

In an important recent paper, Kudlyak (2014) uses such a framework to observe that one way to measure the opportunity cost of hiring a worker this period is, apart from discounting, the cost of hiring the same worker in the next period. If the labor market is a spot market, then this difference is just the current-period wage. But if there are implicit contracts, the

difference of present values can differ significantly from the wage. Kudlyak observes that the object of interest, which she terms the “user cost of labor” can be constructed by using panel data on workers to estimate the present discounted value of wages at time t and $t+1$, correcting for both observed differences in human capital characteristics and for unobserved differences by estimating a worker fixed effect. Using data from the NLSY, she presents such estimates for the period 1978-1997. Kudlyak finds that the user cost is significantly more procyclical than average hourly earnings, and more procyclical than even the wage for new hires. In the empirical component of the paper, we also construct the user cost of labor using NLSY data and a procedure much like Kudlyak’s, and find very similar results.

Kudlyak shows that her user cost of labor is the right measure of the allocative wage in a large range of search models of the labor market. Thus, she calls into question search models based on sticky allocative wages, as in many of the papers discussed above. We embed Kudlyak’s insight into a standard New Keynesian model, and find that the user cost is also the allocative wage in that framework.

Kudlyak’s empirical finding was foreshadowed in two important earlier papers by Beaudry and DiNardo (1991, 1995), who found that the “permanent” wage might be significantly more procyclical than the wage at a point in time. They found that workers hired when the unemployment rate was high received persistently lower wages, even after the economy recovered. Thus, while the spot wage was cyclical, the present value of the wage fluctuated even more. Beaudry and DiNardo interpreted their finding as support for the Becker-Barro hypothesis of implicit contracts with costly worker mobility. In a sense, Beaudry and DiNardo approached the problem from the workers’ side, asking why a worker would take a job in a recession, since the effective (permanent) wage that he or she receives is so low. Our approach (like Kudlyak’s) examines the same facts from the firms’ side, asking why firms do not hire more in recessions, since the effective cost of hiring a worker in a downturn appears to be low. (Beaudry and DiNardo also argue that the data favor a model where workers cannot commit fully to a time path of future wages since, in addition to the unemployment rate that prevailed when the worker was hired, wages seem to depend positively on the minimum unemployment rate observed since the hiring date.)

Hagedorn and Manovskii (2013) argue that much of the observed history dependence of current wages can be understood by appealing to labor search when workers face a job ladder. In Hagedorn and Manovskii’s search model, wages are completely determined by current labor market conditions but because workers gradually “climb” the job ladder, wages

appear to be history dependent. To see this, define an employment cycle as the length of time between spells of involuntary unemployment (Wolpin, 1992). *Ceteris paribus*, the longer an employment cycle, the more job offers the worker has received. Consequently, the current wage must be relatively high to outbid the other competing offers. An individual who enters a period of involuntary employment (thus breaking an employment cycle) falls off the job ladder, and thus his reservation wage falls. Moreover, when a worker begins a new employment cycle, his initial wage offer is determined by current labor market conditions. Workers who start an employment cycle during an expansion, start relatively high up on the job ladder because they receive relatively more offers initially. Workers who start an employment cycle during a recession receive relatively fewer offers and thus accept a lower wage initially.

In Hagedorn and Manovskii’s model, the match quality of a job can be proxied by including the cumulative labor market “tightness” during the employment cycle in the wage regression. Labor market tightness is the ratio of vacancies to unemployment. Intuitively, during an employment cycle, a worker gradually climbs up the match-quality ladder. How fast he or she climbs is determined by current aggregate labor-market tightness. Ultimately, how high the person gets is given by *cumulative* labor-market tightness over the employment cycle. Hagedorn and Manovskii use empirical work based on their model to criticize the conclusions of Beaudry and DiNardo (1991, 1995). They find that when they augment wage regressions with empirical proxies for match quality based on the job ladders model, they no longer find a significant role for lagged unemployment in explaining current wages. In our empirical work using NLSY data, we investigate whether Kudlyak’s finding of implicit wage contracts is sensitive to Hagedorn and Manovskii’s critique.

4 The benchmark model

We begin by extending a standard business cycle model to allow for several real-world features of wage setting. Our benchmark model is a standard New Keynesian DSGE system built on the basic framework analyzed in CEE. We build on the baseline model by allowing for (i) endogenous variation in the composition of the workforce and (ii) differences between the allocative wage and the measured remitted payments to workers. We will spend more time describing our treatment of labor supply and wage setting and the mapping between the model variables and data because these are the non-standard features of the model. Many of the other mechanisms in the model are now common in the DSGE literature and the quantitative

New Keynesian literature and so we present them with relatively less detailed discussion.

4.1 Households

Consumers get utility from consumption and real money balances and get disutility from working. Let C_t be consumption of a nondurable good, let N_t be labor supplied at date- t and let M_t/P_t be real money balances held at date- t . Households act to maximize

$$E_t \left[\sum_{j=0}^{\infty} \beta^j \left\{ \frac{\sigma}{\sigma-1} [C_{t+j} - hC_{t+j-1}]^{\frac{\sigma-1}{\sigma}} - \phi \frac{\eta}{\eta+1} N_{t+j}^{\frac{\eta+1}{\eta}} + \Lambda \left(\frac{M_{t+j}}{P_{t+j}} \right) \right\} \right] \quad (1)$$

subject to the nominal budget constraint

$$P_t (C_t + I_t + b(u_t) K_t) + S_t + M_t = W_t N_t + R_t K_t u_t + S_{t-1} (1 + i_{t-1}) + \Pi_t + M_{t-1} \quad (2)$$

and the capital accumulation equation

$$K_{t+1} = K_t (1 - \delta) + F(I_t, I_{t-1}) \quad (3)$$

Here P_t is the nominal prices of the durable and the nondurable, W_t is the nominal wage rate and R_t is the nominal rental price of capital services, which is the product $K_t u_t$. Π_t denotes profits returned to the household through dividends. M_t is the supply of nominal money balances held at time t , S_t is nominal savings and i_t is the nominal interest rate. σ is the intertemporal elasticity of substitution, η is the Frisch labor supply elasticity, $\Lambda(\cdot)$ expresses the household's valuation for real money balances and $h \geq 0$ is a habit persistence term ($h > 0$ implies habit persistence in utility). The function $F(\cdot)$ is an investment adjustment cost function and $b(u_t)$ gives the resource cost of additional utilization per unit of physical capital. Following CEE, we assume that

$$F(I_t, I_{t-1}) = \left[1 - f \left(\frac{I_t}{I_{t-1}} \right) \right] I_t$$

with $f(1) = 1$, $f'(1) = 0$ and $f''(1) = \kappa$.

Households choose C_t , I_t , M_t , u_t and K_{t+1} to maximize (1) subject to (2) and (3). The determination of labor supply N_t is the key object of interest for this paper and we discuss this in greater detail below.

4.2 Firms and price setting

Following much of the New Keynesian literature we model the production and pricing component of the model as a two-stage process. Final goods are produced from a combination of intermediate goods. Final goods producers are competitive and have flexible prices. Intermediate goods firms are monopolistically competitive and change prices infrequently according to the Calvo mechanism.

4.2.1 Final goods producers

Final goods are produced from intermediates. Specifically, final output is given by the standard Dixit-Stiglitz aggregator

$$Y_t = \left[\int_0^1 y_t(s)^{\frac{\varepsilon-1}{\varepsilon}} ds \right]^{\frac{\varepsilon}{\varepsilon-1}}, \quad (4)$$

where $\varepsilon > 1$. Final goods producers are perfectly competitive and take the final goods price P_t and intermediate goods prices $p_t(s)$ as given. It is straightforward to show that demand for each intermediate good has the standard isoelastic form

$$y_t(s) = Y_t \left(\frac{p_t(s)}{P_t} \right)^{-\varepsilon}. \quad (5)$$

Competition among final goods producers ensures that the nominal price of the final good is a simple combination of the nominal prices of the intermediate goods used in production. Specifically,

$$P_t = \left[\int_0^1 p_t(s)^{1-\varepsilon} ds \right]^{\frac{1}{1-\varepsilon}}. \quad (6)$$

4.2.2 Intermediate goods producers

Intermediate goods are produced by monopolistically competitive firms who take the demand curve (5) as given when they set their prices. Each intermediate goods firm has a constant returns to scale production function

$$y_t(s) = Z_t k_t(s)^\alpha l_t(s)^{1-\alpha},$$

where $k_t(s)$, $l_t(s)$ and $y_t(s)$ denote capital, labor and output for intermediate producer s at time t . k_t is the quantity of capital services *inclusive of utilization*, and thus is not the firm-level equivalent of K_t , which is the stock of physical capital. Similarly, l_t is number

of standardized units of labor employed by the firm. That is, it is an index of the total labor input the firm derives from the potentially heterogenous workers it employs, expressed in a common numeraire, such as the number of high-school-educated workers. This concept of labor, which is relevant for productivity, should be distinguished from N_t , which is akin to total employment or the total number of hours worked by all persons, and is the object relevant for utility. Here Z_t is an aggregate productivity shock common to all firms. While the intermediate goods firms have some monopoly power in their output markets, they are competitive in the input markets, and take the nominal input prices W_t and R_t as given when making their decisions. Each period, firms choose their inputs to minimize costs. For any given level of production \bar{y} , the firm's cost-minimization problem is $\min_{l,k} Wl + Rk$ subject to $Zk^\alpha l^{1-\alpha} \geq \bar{y}$.

Because the production functions have constant returns to scale, and because capital and labor can flow freely across firms, firms choose the same capital-to-labor ratios. That is, for each intermediate producer s ,

$$\frac{k_t(s)}{l_t(s)} = \frac{K_t u_t}{L_t},$$

where we have used the market clearing conditions $\int k_t(s) ds = u_t K_t$ and $\int l_t(s) ds = L_t$. The nominal marginal cost of production, MC_t , for the intermediate goods producers is common to all firms (because the firms have constant returns to scale production functions). It can be shown that the date- t nominal marginal cost is

$$MC_t = \left(\frac{1}{\alpha}\right)^\alpha \left(\frac{1}{1-\alpha}\right)^{1-\alpha} \frac{W_t^{1-\alpha} R_t^\alpha}{Z_t}. \quad (7)$$

Price setting for each intermediate good producer is governed by a Calvo mechanism. Let θ_p be the probability that an intermediate goods producing firm cannot reset its price in a given period. Thus, each period, $1 - \theta_p$ firms reset their prices as they see fit. In many DSGE models, the firms that cannot reset their prices (i.e., those that don't get the Calvo draw) are assumed to reset their prices according to a backward-looking rule. CEE refer to this modelling device as "lagged inflation indexation" in the DSGE literature. To allow for inflation indexing, we would assume that the remaining θ_p firms set their prices according to the backward-looking rule $p_t(s) = p_{t-1}(s) (1 + \pi_{t-1})$ where

$$1 + \pi_t = \frac{P_t}{P_{t-1}}$$

is the gross nominal inflation rate. Without inflation indexing, firms that do not get the Calvo draw simply continue to charge the same nominal price they had at the beginning of the period.

Intermediate goods firms maximize the discounted value of profits for their shareholders (the households) and thus discount future nominal profits in period $t + j$ by the stochastic discount factor $\beta^j \lambda_{t+j}$ (technically, λ_t is the Lagrange multiplier associated with the nominal constraint (2)). The optimization problem for an intermediate goods firm is to choose a reset price p_t^* to maximize the objective

$$E_t \left[\sum_{j=0}^{\infty} (\theta_p)^j \left[\beta^j \lambda_{t+j} \left(p_t^* \prod_{s=0}^{j-1} (1 + \pi_{t+s}) - MC_{t+j} \right) \left(\left[\frac{p_t^* \prod_{s=0}^{j-1} (1 + \pi_{t+s})}{P_{t+j}} \right]^{-\varepsilon} Y_{t+j} \right) \right] \right]$$

where it is understood that $\prod_{s=0}^{-1} (1 + \pi_{t+s}) \equiv 1$. (The expression above, and those that follow, are written under the assumption that firms index their prices to lagged inflation as discussed above. The equations corresponding to the model without inflation indexing are the same except that the terms $(1 + \pi_{t+s})$ are all simply 1.)

Given the reset price p_t^* , and using (6), the price of the final good evolves according to

$$P_t = [\theta_p ([1 + \pi_{t-1}] P_{t-1})^{1-\varepsilon} + (1 - \theta_p) (p_t^*)^{1-\varepsilon}]^{\frac{1}{1-\varepsilon}}.$$

Well-known methods show that the optimal reset price together with the dynamic evolution of the aggregate price level imply that the model satisfies a hybrid New Keynesian Phillips curve of the form

$$\tilde{\pi}_t - \tilde{\pi}_{t-1} = \gamma^p \tilde{m}c_t + \beta E_t [\tilde{\pi}_{t+1} - \tilde{\pi}_t],$$

where $\gamma^p = \frac{(1-\theta_p\beta)(1-\theta_p)}{\theta_p}$ is the microeconomic rate of price adjustment.¹⁹ We use the notation \tilde{v}_t to denote the percent deviation of the variable v_t from its steady-state value \bar{v} . That is, $\tilde{v}_t = dv_t/\bar{v}$.

¹⁹We can also allow for the possibility for *partial* inflation indexing $p_t(s) = p_{t-1}(s) (1 + \pi_{t-1})^\omega$ with $\omega \in [0, 1]$. In this case, the implied Phillips curve is

$$\tilde{\pi}_t - \omega \tilde{\pi}_{t-1} = \gamma^p \tilde{m}c_t + \beta E_t [\tilde{\pi}_{t+1} - \omega \tilde{\pi}_t].$$

4.3 Labor supply and wage setting

The supply of labor features several mechanisms that are prominent in the empirical literature on labor supply and the measurement of wages. As in Erceg *et al.* (2000) and CEE, we allow for nominal wage rigidity in the model. In addition to nominal wage stickiness, we augment the model to include two new features: (i) endogenous composition bias and (ii) a difference between allocative wages and remitted wages. Both mechanisms influence the mapping between model predictions on the one hand and empirical measures of wages and labor supply on the other. To accommodate these mechanisms we treat the supply of labor as occurring in two separate stages within a period. We refer to these simply as stage 1 and stage 2.

In the first stage, the composition bias mechanism allocates workers with differential productivity to the market. This stage results in a single nominal wage paid for units of productivity-adjusted labor and an average wage for employed workers. We denote the wage for productivity-adjusted labor as W_t^1 , the average hourly wage for employed workers as \bar{W}_t^1 and the total supply of effective (productivity-adjusted) labor as L_t^1 where the superscript indicates that these variables are determined in stage 1.

In the second stage, an allocative wage is determined. The allocative wage is sticky and evolves according to a Calvo mechanism taking the stage 1 wage W_t^1 as the effective marginal cost of supplying units of effective labor. In addition to the allocative wage, which governs actual employment, the second stage also produces two separate observed wages: a new-hire wage W_t^{New} and a wage for all employed workers that corresponds to average hourly earnings W_t^{AHE} .

4.3.1 Composition bias

It is well-understood that the composition of the workforce changes systematically over the course of the business cycle. Typically, the labor force has a higher fraction of low-wage workers in booms than in recessions, making the average wage somewhat more countercyclical than the wage for a representative worker with fixed human capital characteristics, which is the concept of the wage in most macroeconomic models. To the extent that composition fluctuates over the cycle, the changing characteristics of the workforce automatically makes output per person appear more counter-cyclical than otherwise.

To introduce composition bias into the model, we imagine that labor varies by productivity.

Specifically, we assume that total actual *hours* of labor (the argument in the utility function (1)) is given by

$$N_t = \int_0^A n_t(a) \varphi(a) da. \quad (8)$$

Here a is an index of productivity and A is the maximum productivity of any individual in society. $\varphi(a)$ is the measure of the population with labor productivity a and $n_t(a)$ denotes hours worked per person with productivity a . For each type, $n_t(a) \in [0, 1]$. The total population is $\bar{N} = \int_0^A \varphi(a) da$. Each type is paid a nominal wage $w_t^1(a)$.

Workers supply labor to labor aggregating firms who in turn sell an effective labor aggregate at a wage W_t^1 . The labor aggregating firms' maximization problem is to hire different types of labor to maximize nominal profits.

$$\max_{n_t(a)} \left\{ W_t^1 \int_0^A a n_t(a) da - \int_0^A w_t^1(a) n_t(a) da \right\}$$

The labor aggregating firms' first order conditions for the choice of $n_t(a)$ requires

$$w_t^1(a) = W_t^1 a$$

for all a . That is, the individual's wage is a direct reflection of the worker's individual productivity.

Consider an increase in $n_t(a)$ from the perspective of the representative household. The utility impact of this increase is

$$\left[-\phi N_t^{\frac{1}{\eta}} + \lambda_t w_t^1(a) \right] \varphi(a) \times dn_t(a)$$

where λ_t is the shadow value of money payments to the representative household (i.e., λ_t is the Lagrange multiplier on the nominal budget constraint). If the term in brackets is positive, then it is optimal to set $n_t(a) = 1$. If the term in brackets is negative, then it is optimal to set $n_t(a) = 0$. Using $w_t^1(a) = W_t^1 a$ we can express the critical productivity cutoff \hat{a}_t as

$$\frac{\phi N_t^{\frac{1}{\eta}}}{\lambda_t W_t^1} = \hat{a}_t. \quad (9)$$

For any type $a > \hat{a}_t$ it is optimal to work full-time. Types $a < \hat{a}_t$ are out of the labor force.

Total employment is $N_t = \int_{\hat{a}_t}^A \varphi(a) da$ and total effective-productivity-adjusted labor is

$$L_t^1 = \int_{\hat{a}_t}^A a \varphi(a) da.$$

Except for two important differences, (9) is essentially a standard labor supply condition. First, \hat{a}_t is endogenous and co-varies negatively with aggregate employment N_t . Second, there is a difference between effective labor L_t^1 and measured hours of employment N_t .

The *average* wage \bar{W}_t^1 of employed workers in the first stage is simply the ratio of total wage payments to total hours of work, that is,

$$\bar{W}_t^1 = \frac{\int_0^A w_t^1(a) n_t(a) \varphi(a) da}{N_t} = \frac{L_t^1 W_t^1}{N_t}.$$

In contrast, the composition-adjusted wage from stage 1 is simply W_t^1 . Notice that the ratio of total hours worked to effective labor is equal to the ratio of the composition-adjusted wage to the average wage, $\frac{N}{L} = \frac{W}{\bar{W}}$. Using log-linear expressions for N_t and L_t , one can show that composition bias (the log difference between \bar{W}_t and W_t) satisfies

$$\widetilde{\bar{W}}_t - \widetilde{W}_t^1 = - \left[\frac{LN - 1}{LN} \right] \widetilde{N}_t \tag{10}$$

where we use the notation LN to denote the ratio of effective labor to measured hours worked L/N . Since the average wage exceeds the wage for the marginal worker (i.e., since $LN > 1$), composition bias imparts a negative comovement between the average wage and aggregate hours. In U.S. data, the cyclical variation in average real wages is negligible, while the composition-corrected wage is procyclical.

4.3.2 Allocative wage rigidity

In addition to the composition-bias mechanism presented above, the model features nominal wage rigidity, as in CEE. The wage block of their model is from Erceg, Henderson and Levin (2000). Like these earlier papers, we assume that wage rigidity applies directly to an “allocative wage,” by which we mean, the relevant wage for determining employment and work effort. Unlike these earlier papers, we allow the allocative wage to differ from the remitted wage that is readily observed in data. In addition to this allocative wage, the model produces a measured remitted wage that we discuss below.

We denote the allocative wage by X_t . The allocative wage adjusts sluggishly according to a Calvo mechanism. As we did in our treatment of composition bias, we assume that there is a labor aggregating firm that assembles an aggregate of labor “types.” This aggregating firm supplies effective labor to the productive firms at flow allocative wage X_t , but hires labor by type according to the type-specific allocative wages $x_t(s)$. The labor aggregate is given by a CES aggregate of labor types s ,

$$L_t^2 = \left[\int_0^1 l_t(s)^{\frac{\psi-1}{\psi}} ds \right]^{\frac{\psi}{\psi-1}},$$

where the superscript 2 refers to the fact that this labor supply is determined in stage 2. (Note, this treatment is essentially identical to our treatment of prices. As we did earlier, we let s be an index of different types though in this context s refers to a type of labor while before s was a type of intermediate good.) If the aggregating firm chooses to supply labor force L_t^2 , its demand for type s work is given by the isoelastic function,

$$l_t^2(s) = L_t^2 \left(\frac{x_t(s)}{X_t} \right)^{-\psi}.$$

The allocative wages $x_t(s)$ for each type s of labor are set by a monopolist in that type (similar to a union). The aggregate allocative wage X_t for units of the labor aggregate is a reflection of the type-specific allocative wages $x_{i,t}$

$$X_t = \left[\int_0^1 x_t(s)^{1-\psi} ds \right]^{\frac{1}{1-\psi}}.$$

Note that the labor market clearing condition implies that, up to a first-order approximation, the labor aggregate from stage 1 is equal to the resulting labor aggregate from stage 2 (i.e., $\tilde{L}_t^1 \approx \tilde{L}_t^2$).

As we did with the price setters, we assume that the type-specific wages are set according to a Calvo mechanism. The probability of adjusting a type-specific wage is $1 - \theta_w$ and the probability of not adjusting is θ_w . As we did with the price setters, we allow for the possibility of wage inflation indexing. In this case, wage setters who do not get the Calvo draw, instead follow the wage inflation indexing rule $x_t(s) = x_{t-1}(s)(1 + \pi_{t-1})$. Without wage inflation indexing, these wage setters would simply maintain the constant nominal allocative wage they had at the start of the period.

The union tries to maximize the present discounted value of wage markups $x_t(s) - W_t^1$. An extra dollar in period $t + j$ is worth $\beta^j \lambda_{t+j}$ to the household. Thus, a monopolist who has the option to set his wage at time t should choose a reset wage w_t^* to maximize

$$\max_{x_t^*} \left\{ E_t \left[\sum_{j=0}^{\infty} (\beta \theta_w)^j \lambda_{t+j} \left(x_t^* \prod_{s=0}^{j-1} (1 + \pi_{t+s}) - W_{t+j}^1 \right) L_{t+j}^2 \left(\frac{x_t^* \prod_{s=0}^{j-1} (1 + \pi_{t+s})}{X_{t+j}} \right)^{-\psi} \right] \right\}$$

where again is it understood that $\prod_{s=0}^{-1} (1 + \pi_{t+s}) = 1$. Given the reset wage x_t^* , the aggregate allocative wage evolves according to

$$X_t = \left[\theta_w \{(1 + \pi_{t-1}) X_{t-1}\}^{1-\psi} + (1 - \theta_w) (x_t^*)^{1-\psi} \right]^{\frac{1}{1-\psi}}.$$

4.3.3 Remitted wages

Our discussion highlights the difference between the *allocative* wage – the shadow wage X_t that governs work effort and employment – and the *measured* wage that governs the periodic payments from the employer to the workers. We assume that the remitted wage is a smoothed function of the allocative wage. Specifically, we assume that workers periodically renegotiate their contract terms (or separate from their current jobs and get new jobs with new terms). When wage contracts are renegotiated, the workers are given a new remitted wage. Let PDV_t be the expected present discounted value of future nominal allocative wages for a newly employed worker that resets the remitted wage with probability $s \in (0, 1]$. That is,

$$PDV_t = X_t + \beta (1 - s) E_t \left[\frac{\lambda_{t+1}}{\lambda_t} PDV_{t+1} \right] = E_t \left[\sum_{j=0}^{\infty} [\beta (1 - s)]^j \frac{\lambda_{t+j+1}}{\lambda_t} X_{t+j} \right].$$

Clearly PDV_t depends on the reset rate s (even though the reset rate plays no role in allocations). The measured remitted wage for new hires (or workers who newly renegotiated their contract) at date- t will be a smoothed version of the PDV. Specifically, we assume that the measured wage for new hires will solve

$$PDV_t = W_t^{New} E_t \left[\sum_{j=0}^{\infty} [\beta (1 - s)]^j \frac{\lambda_{t+j+1}}{\lambda_t} \right].$$

That is, W_t^{New} is a constant wage that will transfer the same expected amount to the workers given the reset rate s as they would receive by getting the time-varying aggregate allocative wage, X_t . For purposes of comparison with the data, W_t^{New} is the new-hire wage.

We can also track the average outstanding wage for all workers in the model. Let W_t^{AHE} be the average hourly earnings of all employed workers. By construction, the average outstanding wage at time t is the average wage for all workers that did not renegotiate together with the new-hire wage

$$W_t^{AHE} = W_{t-1}^{AHE} (1 - s) + H_t W_t^{New}$$

where $H_t = L_t - L_{t-1} (1 - s)$ denotes “new hires” which we interpret as all workers who are newly hired plus those who remain employed but receive new contract terms for their remitted wage.

It is worth mentioning some of the difference between the different wage concepts W_t^1 , \bar{W}_t^1 , W_t^{AHE} , W_t^{New} and X_t . One key difference between the wages in the first stage (W_t^1 and \bar{W}_t^1) and the wages in the second stage (W_t^{AHE} , W_t^{New} and X_t) is that the wages in the second stage have a (potentially, time-varying) wage markup. That is, in the non-stochastic steady state, $W^{AHE} = W^{New} = X = \frac{\psi}{\psi-1} W^1$. If the allocative wage is flexible, then the markup $\frac{\psi}{\psi-1}$ is constant even away from the steady state and in this case the dynamic behavior of the allocative wage and the stage 1 wage is the same (i.e., $\tilde{X}_t = \tilde{W}_t^1$). If there is no composition bias then the two stage 1 wages are the same, $W_t^1 = \bar{W}_t^1$. If the renegotiation rate $s = 1$ then all of the stage 2 wages are identical, $W_t^{AHE} = W_t^{New} = X_t$. In general all of the wages will differ.

4.4 Aggregate conditions and the steady-state

The goods market clearing condition is

$$Y_t = C_t + I_t + K_t a(u_t).$$

Although in principle there can be many different sources of uncertainty in the model, we focus our attention here on monetary shocks. We assume that monetary policy is described by a Taylor rule

$$\tilde{i}_t = (1 - \rho^i) \left[\phi_Y \frac{\tilde{Y}_t}{4} + \phi_\pi \tilde{\pi}_t \right] + \rho^i \tilde{i}_{t-1} + \varepsilon_t^i$$

Here, ϕ_Y and ϕ_π give the relative reaction of the monetary authority to output and inflation while ρ^i is an interest rate “smoothing” parameter. Here ε_t^i is a shock to the monetary authorities policy rule. We assume that ε_t^i is mean zero and i.i.d. over time.

4.4.1 Non-stochastic steady-state

We choose parameters to ensure that in the non-stochastic steady-state, $L = P = u = 1$. The steady-state markups are $\mu^p = \frac{\varepsilon}{\varepsilon-1}$ and $\mu^w = \frac{\psi}{\psi-1}$. We normalize the steady-state productivity cutoff to $\hat{a} = 1$. The steady-state nominal marginal cost is $MC = 1/\mu^p$. Since there is no inflation and no economic growth in the steady-state, $1 + r = 1 + i = \frac{1}{\beta}$. It is straightforward to show that the nominal rental price is $R = r + \delta$ and we must have $R = b'(1)$. Steady-state capital is

$$K = \left(\frac{\alpha MC}{R} \right)^{\frac{1}{1-\alpha}}.$$

The remaining details of the steady-state are standard and are therefore omitted.

4.4.2 Calibration

To compare the model to the data, as we do in Section 6, we will need to calibrate the parameters in the model. When possible, we adopt calibration settings based on conventional parameter values used for medium scale DSGE models. The discount factor β is set to 0.97 which implies a steady state annual real interest rate of three percent. We set both the Frisch elasticity η and the intertemporal elasticity of substitution σ to 1.00. Capital’s share, α , is set to 0.36. We set the type-specific elasticity of labor demand ψ to 21 which implies a 5 percent markup of the allocative wage over the base wage. We set the type-specific elasticity of product demand ε to 6 which implies a 20 percent markup of nominal price over nominal marginal cost. We set the Calvo parameters for wage and price adjustment, θ_w and θ_p , to 0.90 (quarterly). This implies that both wages and prices have an average duration of roughly 10 quarters. These durations are somewhat longer than most studies of microeconomic price adjustment data but are comparable to estimates from DSGE models.

Following Basu and Kimball (1997), we set the utilization elasticity $\frac{b''(1)}{b'(1)} = 1.00$. We set the investment adjustment cost parameter is set at $\kappa = 4.00$ and the habit persistence parameter $h = 0.65$. We allow firms to index their prices to past inflation as in CEE.

In addition to the standard parameters discussed above, the model also requires values for the parameters that govern composition bias and the remitted wage. There are three

key parameters that govern these mechanisms: the renegotiation hazard s , the steady-state ratio of effective labor to total hours worked LN and the density of types at the steady state productivity cutoff $\varphi(1)$. For our baseline setting we assume that neither of these mechanisms is operative and thus we set $s = 1.00$ (so the remitted wage is equal to the allocative wage), $LN = 1.00$ (so there is no difference in average productivity per hour) and $\varphi(1) = \infty$ (so hours can be varied without changing the productivity of the marginal worker). This baseline specification is thus essentially equivalent to a standard medium-scale sticky-price/sticky-wage DSGE model. When we introduce composition bias and infrequent resetting of remitted wages we set $s = 0.21$ following Barattieri *et al.* (2014), $LN = 2.0$ and $\varphi(1) = 2$.

5 Empirical measures of real wages

Empirically, the cyclical nature of the real wage is potentially influenced by several different features. Cyclical variations in the composition of employed workers has been emphasized as an important component of variation in real wages (see, Solon, Barsky and Parker (1994), and Elsby, Shin and Solon (2016)). Even after correcting for compositional changes however, it is difficult to speak unambiguously about a single concept of “the real wage.” As emphasized by Haefke *et al.* (2013), wages of newly hired workers appear to be much more cyclical than the wages of workers who are continually employed. Beaudry and DiNardo (1991, 1995) argue that wage payments are shaped by implicit agreements between employers and employees and thus the remitted wage at a given point in time provides at best an incomplete measure of the worker’s compensation. Similarly, Kudlyak (2014) finds that wage paths of workers hired at various points of the business cycle exhibit great differences in present value suggesting that the theoretical concerns articulated by Becker (1962) and Barro (1977) regarding implicit long-term contracts have empirical as well as theoretical merit.

In this section we will examine micro data on real wages to attempt to assess whether available evidence can provide insight into how various measures of real wage payments move over the business cycle and also whether these wage measures react to monetary policy shocks. We empirically quantify the separate contributions of composition bias, variations in the new-hire wage, and variations in the present value of wage commitments.

5.1 Background

In the model in the previous section, there are several objects that map to measured wages, but only one, which we have called X_t , is allocative. Unfortunately this allocative wage is not directly measured in the data. In principle, this allocative wage can be uncovered by differencing measures of the present value at two points in time. This difference is what Kudlyak (2014) calls the *user cost of labor* (UCL).²⁰ Specifically,

$$UCL_t = PDV_t - \beta(1-s)PDV_{t+1} \approx X_t. \quad (11)$$

The UCL_t is only approximately equal to the allocative wage since the calculation above ignores the expectations operator and the stochastic discount factor. Kudlyak finds that unlike the average wage, the UCL is highly procyclical, even more so than the wage of new hires.

Below we will construct measures of both the UCL and the new-hire wage. Both the UCL and the new-hire wage are difficult objects to measure, since one needs to observe individual workers over time. The two panel data sets for the U.S., the NLSY and the PSID, both have relatively small samples and limited sample periods. Furthermore, the data are annual, which is not ideal for business-cycle analysis. A benefit of using individual level panel data is that such data can correct for composition bias. Thus, in their early analyses of composition bias over the business cycle, Bils (1985) uses the NLSY and Solon, Barsky and Parker (1994) use the PSID. The recent papers by Haefke, Sonntag and van Rens (2013) and Elsby, Shin and Solon (2016) instead use CPS data. The drawback to using the CPS is that since it is not a true panel, one cannot remove the effects of *unobserved* individual effects from the wages. On the other hand, the CPS has the advantage that it provides a large and nationally-representative sample, and continuous monthly data going through the Great Recession period and its aftermath.

5.2 NLSY

For the purposes of the analysis in this chapter, we focus on wage data from the National Longitudinal Survey of Youth 1979 (NLSY). The NLSY is an unbalanced panel of workers initially interviewed in 1979 and then, if possible, interviewed every subsequent year until 1994

²⁰The term is used as an analogy to the “user cost of capital” under adjustment costs, in which case the decision to add an extra unit of capital is a dynamic decision with long-term consequences.

and every second year after 1994 (see below). The initial sample included 12,686 individuals born between 1957 and 1964. The birth years of the individuals in the data are distributed roughly uniformly over the years between 1957 and 1964. At the time of the initial survey in 1979 these individuals were all between 14 and 21 years of age. The initial respondents consisted of 6,403 males and 6,283 females.

While there are many advantages to using NLSY data, there are some disadvantages as well. Chief among these disadvantages is the fact that, due to the nature of the survey's construction, the sample in the NLSY "ages" systematically with the passage of time. This immediately means that the average wage of employed workers in the NLSY should not be directly compared to the average hourly earnings wage series constructed from NIPA data which presumably reflects wage payments to all individuals employed at any given point in time. Moreover, there is only a small amount of age variation in any single year. Thus, while we can in principle control for age in our wage regressions, the age (or more accurately "experience") coefficients will be difficult to distinguish from the growth of average wages over time.

Our data includes all data from the first interview in 1979 up to 2013. Because the NLSY was modified to a biennial survey starting in 1994, we drop all of the odd years between 1994 and 2012.²¹ For our analysis here, we focus exclusively on men. Thus our sample consists of the 6,403 men interviewed initially in 1979 and then followed until 2012. Although the NLSY is not a representative sample of the U.S. population, the survey provides a yearly cross-sectional weight variable that can be used to make the sample comparable to that year's population.²²

The NLSY reports wage information for up to five jobs each year. Our sample technically includes data from jobs in 1978 even though the first interviews were done in 1979. (The 1978

²¹While the NLSY does ask the respondent to remember information for the previous years after it made the transition to biennial surveys, the wage series and responses appear to be systematically different for the odd years.

²²Of the whole initial interview sample, roughly half (6,111) comprised what the NLSY refers to as a "representative sample" of the non-institutionalized working-age population born between 1957 and 1964. In addition to the representative sample, the NLSY also collected data on a "disadvantaged sample" consisting of 5,295 individuals who identified as Hispanic, Latino, Black and economically disadvantaged respondents. Finally, the NLSY includes 1,280 respondents who are representative of the population serving in the armed forces. This latter sample is referred to as the "military" sample. Both the disadvantaged sample and the military sample were severely cut back or eliminated entirely from the NLSY in 1984 and again in 1990. We keep all males in each of the three sub-groups (the representative sample, the military sample and the disadvantaged sample). We then use the cross sectional weights to convert the NLSY data to an overall representative sample. Note, we do not use the longitudinal weights that are included in the NLSY. The longitudinal weights are intended to produce a representative panel over the entire period.

data come from interviews that were done early in 1979 and so pertained to jobs in 1978.) We focus on the “hourly rate of pay” variable that is constructed by the NLSY. Respondents are asked for the most convenient way to report their total earnings.²³ They could report pay per hour, per day, per week, per month, or per year. In every case, the reported statistic is then converted to an hourly pay rate based on a measure of the respondents typical hours worked. The resulting hourly rate of pay includes tips, overtime pay, and bonuses but is computed before any deductions. To construct real wages, we deflate the hourly rate of pay with a price index. We considered two separate price deflators in our analysis: the Consumer Price Index and the Implicit Price Deflator for the Non-farm Business Sector.^{24,25} The analysis for the two separate price indices were quite similar overall. Since our focus is on intertemporal labor demand from the firms’ perspective (i.e., the real product wage), we focus on the real wage measures using the deflator for the non-farm business sector in the discussion below.

In addition to the information on wages, the NLSY includes information on the industry of the jobs and whether the jobs are covered by a union. We do not include union status in our analysis because the union variables included in the Employer History Roster exhibit an unusual change following the 1994 change from annual to biennial coverage in the NLSY. (See the technical appendix for more discussion).

WAGE REGRESSIONS

We begin by describing how we construct the various measures of real wage series from the NLSY data. Given the available data, as described above, we run regressions of the following form:

$$\ln w_{t,\tau}^i = c + \alpha^i + \zeta t + \Psi X_t^i + \sum_{d_0=1}^T \sum_{d=d_0}^T \chi_{d_0,d} D_{d_0,d}^i + \varepsilon_t^i. \quad (12)$$

This is the basic empirical specification considered in Kudlyak (2014). Here $w_{t,\tau}^i$ is the real wage for individual i at time t who was hired at time τ . This regression provides a best linear prediction of the log real wage at time t of a worker i , who started his job in period τ . In its most general form, this wage regression allows for a time trend, demographic and industry controls (included in X_t^i), individual fixed effects (the α^i coefficients) and time effects that

²³QES-71A in the 2012 survey.

²⁴We used the Consumer Price Index for All Urban Consumers: All Items and the Non-farm Business Sector Implicit Price Deflator. Both variables are seasonally adjusted and are available from the FRED database as CPIAUCSL and IPDNBS.

²⁵We exclude wage rates less than 1 dollar per hour and above 100 dollars per hour measured in 1979 dollars. This restriction led to the elimination of 2,894 wage-year observations. This censoring at 1 dollar and 100 dollars is the same censoring used by the BLS when it uses NLSY data.

depend on two periods: when the individual began work at his current job and the current date. The additional covariates in the X_t^i matrix are the individual's experience at time t (and experience squared), tenure at time t (and tenure squared), schooling completed, and industry fixed effects. Experience is defined as the maximum of (Age – 6 – years of schooling) and 0. The dummy variables $D_{d_0,d}^i$ take the value 1 if $d_0 = \tau$ and $d = t$ and 0 otherwise.

The χ coefficients are particularly important for interpretation of the new-hire wage series and the user-cost series that we emphasize below. At time t , all workers who began work at their current job at date- τ get an additional adjustment to their predicted wage given by the coefficient $\chi_{\tau,t}$. These adjustments imply that workers who begin at date- τ experience an expected strip of log wage realizations given by $\{\hat{\chi}_{\tau,\tau}, \hat{\chi}_{\tau+1,\tau}, \hat{\chi}_{\tau+2,\tau}, \dots, \hat{\chi}_{\tau+j,\tau}, \dots\}$ etc. These dummy variables thus adjust for vintages of hired workers, where the vintage is defined by when the worker was hired in addition to the current calendar date. Notice that the variable $\hat{\chi}_{\tau,\tau}$ reflects the wages of a newly hired worker (i.e., the date- τ wage of a worker hired at date- τ). In the estimation, we truncate the χ strips at seven years (including year 0).²⁶

This specification can also be used to calculate composition adjusted wages. For instance, if we restrict the $\chi_{\tau+j,\tau}$ coefficients to be zero then the resulting specification gives a predicted wage that adjusts for both observed changes in workforce composition (by including the X_t^i variables) and unobserved workforce composition (by including the individual fixed effects α^i), but does not allow for vintage effects on the wage. Adding the χ dummy variables allows us to recover composition-adjusted wages with vintage effects. For example, the coefficient on $\chi_{\tau,\tau}$ tells us whether a newly hired worker receives a wage increase or reduction relative to workers hired in previous years, controlling for any differences in human capital between new hires and other workers.

AVERAGE HOURLY EARNINGS AND NEW-HIRE WAGES

Before we consider our measures of the user cost of labor, we first examine average hourly earnings. We consider two measures. The first is a measure taken from the BLS. The BLS reports a measure of compensation per hour for the non-farm business sector. We then deflate this measure by the implicit price deflator.²⁷ We refer to this measure as AHE-BLS.

Our second measure of average hourly earnings is constructed from the NLSY data. We refer to this measure as AHE-NLSY. This wage series is constructed by first running the

²⁶More precisely, we include all of the dummy variables in the estimation of (12); however, following Kudlyak (2014), we use only seven χ estimates when we calculate the user cost of labor.

²⁷The variables used in this calculation are from the FRED database. We use Nonfarm Business Sector: Compensation Per Hour (COMPENFB) and the Nonfarm Business Sector: Implicit Price Deflator (IPBNBS).

simplified version of regression (12)

$$\ln w_t^i = c + \Psi X_t^i + \sum_{d=1}^T \omega_d D_d^i + \varepsilon_t^i,$$

where D_d^i is a time dummy variable (D_d^i takes the value 1 if $d = t$ and zero otherwise). For the NLSY measure of average hourly earnings, the controls X_t^i include only experience and experience squared. Because the experience variable is defined as $\max\{\text{age} - 6 - \text{schooling}\}$ this is close to being a control for age and age squared. The estimated time fixed effects $\hat{\omega}_t$ are then an estimated time series of average hourly earnings. Note that, because the NLSY is based on a fixed set of individuals who were entering the workforce in the late 1970s, it is crucially important to include controls for age in this measure. If we did not include experience and experience squared then the sample would systematically age and this would impart a systematic aging component to the wage measures.

To construct the new-hire wage, we return to the original regression specification (12). As noted, the new-hire wage series corresponds to the estimated coefficients $\hat{\chi}_{t,t}$. We include all of the available demographic controls in X_t^i and we also include individual fixed effects in the regression.²⁸

CALCULATING THE USER COST OF LABOR

We base our calculation of the user cost of labor (UCL) on equation (11). To calculate the user cost, we need to calculate a forecast of the present value of wage payments for a worker hired at date- t and the present value of wage payments for a worker hired at date- $t + 1$. For an individual hired at date- t and still employed at date- $t + j$, we construct the predicted value of the log real wage $\widehat{\ln w_{t,t+j}}$. We can then calculate the implied present value of compensation as

$$\widehat{PDV}_t = \sum_{j=0}^{\infty} \beta^j (1 - s)^j \exp \left\{ \widehat{\ln w_{t,t+j}} \right\}.$$

Note that in addition to requiring a sequence of predicted log wages $\{\widehat{\ln w_{t,t+j}}\}_{j=0}^{\infty}$, this calculation requires a separation rate s and a discount factor β .

To construct the projected wage payments $\widehat{\ln w_{t,\tau}}$, we consider the anticipated wage payments for a firm that hires an ‘‘average worker’’ at date- t . As the employment relationship

²⁸Our method for constructing the new-hire wage differs from that in Kudlyak (2014), who simply examines the wages for workers hired in the current year. Our procedure creates a wage series for new hires correcting for composition, in parallel with our construction of the user cost of labor.

continues, our measure of the worker’s experience and our measure of the worker’s tenure both increase. We assume that the initial experience is fixed at the sample average of 11.72 years and we set the initial tenure variable to 0.5 years (this implicitly assumes that a worker who reports being newly hired at his current job at the time of the interview was hired 6 months earlier). We set the worker’s schooling to 12.57 years, again the sample average in the NLSY. Then, based on (12), at date- τ , a worker hired at date- $t \leq \tau$ has a projected log wage

$$\widehat{\ln w_{t,\tau}} = \hat{c} + \hat{\zeta}\tau + \hat{\Psi}\bar{X}_{\tau-t} + \hat{\chi}_{\tau,t} \quad (13)$$

where $\bar{X}_{\tau-t}$ are demographic controls for the “average worker” (i.e, schooling = 12.57, experience = $11.72 + \tau - t$ and tenure = $0.5 + \tau - t$).

For the separation rate s , we follow Kudlyak (2014) who uses a monthly separation rate of 0.0295. This figure is based on calculations of the average separation rate in the NLSY. We then convert this monthly separation rate into an annual separation rate by setting $s = 1 - (1 - 0.0295)^{12} = 0.3019$. The NLSY figure might be somewhat low relative to other datasets. The separation rate from the JOLTS dataset is closer to 0.035. The annual discount factor is set to 0.97. Note that our calculation of the present value of wage payments is truncated at seven years (including the initial year). Given the high observed separation rates in the data, this truncation has a relatively small effect on the present value.

5.3 The cyclicity of real labor compensation

We are now in a position to examine the cyclical behavior of real wages. Tables 3 and 4 report cyclicity estimates for six different measures of log real wages. For each measure of real wages, we regress the calculated wage series on an indicator of the business cycle (and a time trend and a constant). Table 3 examines the cyclicity of real wages with respect to the HP filtered unemployment rate. We use HP filtered unemployment rather than the unemployment rate in levels because the average unemployment rate changes substantially over the time period for the NLSY.²⁹ Thus, the coefficients reported are semi-elasticities: the percent change in a real wage measure in response to a one percentage-point deviation of unemployment relative to its trend. The sample for columns 1 - 5 consists of 25 data points

²⁹The HP filtered unemployment rate is taken from monthly data from 1985 to 2016. To avoid the well known endpoint problem in HP filtering, we add 120 months of predicted unemployment rates taken from an estimated AR(6) to the end of the sample. We then HP filter the padded series using a smoothing parameter of 500,000. The regression uses annual averages of the monthly HP deviations.

from 1979 to 2012, dropping the odd years between 1994 and 2012 (see Section 3.2). To construct the UCL, we need to impute values of wages for the odd years between 1994 and 2012. The final user cost series itself ends in 2007 because we require seven subsequent wage observations to calculate the value of the UCL in year t (again see Section 3.2 for details).

Columns 1-4 report results for average hourly earnings. For the BLS wage series, AHE-BLS, the coefficient on the unemployment rate is -0.507: real average hourly earnings fall by roughly 0.5 percent for each percentage point increase in the cyclical component of the unemployment rate. Columns 2-4 report results for our constructed measure of average hourly earnings from the NLSY data, AHE-NLSY. As noted in our discussion above, the dependent variables in Columns 2-4 are estimated time fixed effects from regressions of individual wages on the listed set of controls. The columns differ according to the number of controls included in the regression. Column 2 includes only experience and experience squared; column 3 adds industry fixed effects, job tenure and schooling; column 4 includes all of the aforementioned controls and adds individual fixed effects. The NLSY sample exhibits greater cyclical for all of the measures of average hourly earnings, and the cyclical rises with the number of controls for worker characteristics. We interpret this finding as being supportive of the basic composition-bias effect emphasized by Bils (1985) and Solon *et al.* (1994). Typically, as we add more controls for worker heterogeneity, the point estimate of the cyclical rises (though note, the standard errors are high enough that we cannot say with any certainty that any one of these measures is clearly more or less cyclical than any other).

Column 5 reports results for the new-hire wage. The point estimate for the cyclical coefficient is -0.698 , so a one percentage point increase in the cyclical component of unemployment corresponds to a 0.7 percent reduction in the real new-hire wage. By itself, the point estimate seems to be at odds with the findings in Haefke *et al.* (2013) who reported that in CPS data, the wages of newly hired workers appeared substantially more cyclical than average hourly earnings. We should note that while our point estimates do not indicate greater cyclical of the new-hire wage, the estimates are quite noisy and admit a range of interpretations.

Column 6 reports results for the user cost of labor (UCL). Our measure of the UCL exhibits *much* greater cyclical than either the composition-adjusted wage or the new-hire wage series. In Table 3, the cyclical estimate is -5.818 indicating that for every one percentage-point increase in the cyclical component of unemployment, the real user cost of labor falls by almost 6 percent (!).

The estimates in Table 3 are robust to alternate measures of the business cycle. Table 4 reports estimates for the same dependent variables as those in Table 3, but uses HP filtered GDP as the indicator of the business cycle instead of the unemployment rate. Again, average hourly earnings seem to be only moderately cyclical. When HP filtered GDP is above trend by one percent, AHE-BLS is above trend by only 0.311 percent. By contrast, holding the set of workers fixed in the NLSY and controlling for observed and unobserved heterogeneity increases this estimate to 1.165 percent. The point estimate of the cyclical component of the new-hire wage is more cyclical. The point estimate is a rise of roughly 1.3 percent for every one percent change in the cyclical component of GDP. Finally, as before, the UCL is the most cyclical wage measure. For each percent increase in GDP above trend, the UCL rises by approximately 3.1 percent.

What these results seem to suggest is that both composition bias and implicit contracting play important roles in shaping the wage payments made to workers over the business cycle. Quantitatively, controlling for composition (by including individual fixed effects and controls for observed worker differences in the wage regressions) increases wage cyclicalities by perhaps as much as a factor of two relative to a group of workers without such controls. The effects of implicit contracting and wage-smoothing seem to be even greater than the effects of composition bias. According to our calculations, the user cost of labor has a cyclicalities that is, in some cases, about six times greater than the log real wages of the base group. Since average payments are less cyclical than the user cost, workers hired in bad times are paid a wage greater than their user cost. In return, the workers expect to receive fewer and smaller wage increases over their employment spell.

Our findings (which are consistent with the results in Kudlyak (2014)) seem to corroborate the results in Beaudry and DiNardo (1991, 1995), who argued that current wage payments seem to be tied to past labor market conditions. In that paper, the authors showed that the maximum unemployment rate during a job spell and the unemployment rate that prevailed when the worker was hired both have a significant influence on current wage payments. The specification above, which we have adapted from Kudlyak's work, is a more general econometric specification than the one in Beaudry and DiNardo, but implicit contracts still appear to play an important role in shaping wage payments.

5.4 Wage responses to monetary shocks

Almost all of the literature on wage cyclicality examines the response of real wages to a cyclical indicator, typically the unemployment rate. However, the monetary business-cycle literature has also emphasized the importance of replicating estimated impulse response functions to identified shocks – most often monetary shocks. The modern literature on estimating the effects of monetary policy shocks using VARs began with Bernanke and Blinder (1992). Here we follow the approach in CEE, since we ultimately want to make comparisons between empirical and theoretical impulse responses to monetary shocks.

To implement the VAR procedure, we first extend our annual real wage measures from the NLSY to a longer quarterly series using the Chow-Lin procedure. The extension to quarterly data is important for the validity of the identifying assumptions commonly used in the VAR literature. The identification assumptions invoked are plausible in quarterly observations but this plausibility becomes strained if the data are sampled at an annual frequency. The Chow-Lin method uses the annual data to estimate the relationship between the annual wage measures constructed above and other variables that are available at a quarterly frequency. The variables used in the Chow-Lin procedure are Real Gross Domestic Product, Real Hourly Compensation in the Nonfarm Business sector, the Civilian Unemployment Rate, and Total Nonfarm Payrolls for All Employees.³⁰ The resulting interpolated series distributes the annual measure to the corresponding quarters (thus, the annual averages of the constructed quarterly series equal the original annual measures). We then extend the series by projecting the missing data to periods outside the years 1979-2012 covered by the NLSY. We first regress the interpolated quarterly wage measures on the variables in the Chow-Lin procedure above. We then use the OLS estimates to form estimates \hat{w}_t for time periods earlier than 1979 and later than 2012. Figure 1 plots quarterly average hourly earnings, the new-hire wage, and the user cost of labor for the period 1965:1 to 2015:3. For each series, a separate linear trend was removed prior to plotting. Each series is in log points and is plotted so that the mean of each series is centered at 1.00.

The impulse response functions to monetary policy shocks are constructed following the approach recommended by CEE. We include the same variables, in the same Choleski ordering as in the original CEE specification. In order, the variables are real output, real consumption expenditure, the price level, real investment spending, real average hourly compensation,

³⁰All variables are in logs except for the unemployment rate which is entered in levels.

average labor productivity, the federal funds rate, real corporate profits, and the growth rate of the money supply (M2). Following Bernanke and Blinder (1992), the innovation to the funds rate is identified as a structural shock to monetary policy. Notice that by assumption, none of the variables in the first block (output, consumption, the price level, investment, compensation and labor productivity) responds contemporaneously to a shock to monetary policy. In contrast, both corporate profits and the growth rate of M2 respond contemporaneously to monetary shocks. Our approach is to extend the CEE specification by appending a single additional variable – an additional wage measure – to the second block of variables. Thus, our augmented VAR introduces a wage which is allowed to respond contemporaneously to monetary shocks. However, we add the restriction that monetary policy does not respond contemporaneously to shocks to the new wage measure. This restriction is sufficient to identify the impulse response of the wage measure to a monetary policy shock.

We do not want to allow the new wage measures to influence the identified monetary policy shocks. That is, we wish to ensure that the identified shocks remain the same as we change our measure of wages in the VAR. This first consideration implies that the new wage measures should be excluded from the dynamic equations governing the variables originally included in the CEE specification. It also suggests that we should order the new wage series last so that these new measures will respond to the other variables but the other variables — in particular the federal funds rate — will not respond to the alternate wage measures.³¹ One consequence of ordering the new wage measures last is that they respond contemporaneously to a monetary policy shock. In CEE’s original specification their measure of the wage, average hourly earnings, comes before the federal funds rate, and may respond to monetary policy shocks only with a one-quarter lag. Thus, our procedure treats the new wage measures differently from average hourly earnings, but only for the first quarter after a monetary policy shock.³²

As in CEE, consumption, investment and corporate profits come from the BEA’s NIPA tables. Unlike CEE, our measures of output, the price level, employee compensation and labor productivity are only for the Non-farm Business sector. Our decision to use the Non-farm business sector is motivated primarily by our belief that the non-farm business sector is a

³¹As in CEE, the federal funds rate is assumed to respond to average hourly compensation. Thus, we *always* include real hourly compensation in the VAR. This ensures that our identification assumptions match those in CEE.

³²We experimented with other identification schemes, including ones that constrain the contemporaneous responses of the new wage measures to zero, symmetrically with average hourly earnings. The results we report below are qualitatively robust to all the variants we tried.

better match to models of infrequent price-adjustment by firms that are trying to maximize profits. Excluded industries (such as utilities and government production) likely do not set prices optimally the way most macroeconomic models posit. All variables are in log levels, except for the federal funds rate which is in levels and M2 which is in log differences. All variables were downloaded from the Federal Reserve Bank of St. Louis FRED web database.

To estimate the VAR we use the same sample used by CEE, namely 1965:3 to 1995:3. We do so in order to make it easy to compare our results to the ones in this benchmark paper.³³ We experimented with other sample periods, including extending the sample forward to 2007. The extended sample would allow us to estimate the VAR using a larger data set, while stopping short of the zero-lower-bound period in which the identifying assumptions do not apply. Unfortunately, we found that the impulse responses reported by CEE change significantly when the later data are added, in ways that are difficult to interpret in light of the underlying theory. (For example, the well-known “price puzzle” is clearly apparent in the extended sample.) Since our main objective is to explore how our novel real wage measures responded to a well-known identified shock, we chose to estimate the VAR over the original CEE sample. However, we believe that the instability of the impulse responses to monetary shocks over different sample periods—reminiscent of the results in Hanson (2004)—is worthy of investigation in its own right.³⁴

Our focus in this chapter is on the responses of the various wage measures constructed above. For each “new” measure of the real wage, we estimate a different VAR system to recover the response of the wage to the monetary shock. For example, we estimate the system separately for the new-hire wage and for the user cost. Given the VAR structure discussed below, the new variable does not affect the responses of the original CEE variables in any way. As a result, we report the impulse responses of the baseline set of variables to a monetary policy shock in one block, since these do not change as we change the additional wage variable added to the CEE specification.

Figure 2 shows the reaction of the standard macroeconomic variables included in the CEE system to an identified increase in the federal funds rate of 50 basis points. (To save space, the figure omits the responses for corporate profits and for the growth rate of M2.) Each

³³Even though our sample is identical with the sample in CEE, differences will arise because of data revisions and also because we use NFB output and prices rather than GDP and the GDP deflator.

³⁴In line with our findings, in her chapter in this *Handbook*, Ramey (2016) also emphasizes that the original CEE estimates “do not hold up well in later samples.” She concludes that “the most likely reason for the breakdown [...] in the later sample is simply that we can no longer identify monetary policy shocks well.”

panel reports the impulse response of a single variable. The units of all variables are reported in percentage points (i.e., 1.00 corresponds to one percent). The dotted lines correspond to 1 standard deviation error bands. The shock leads to a reduction in output of roughly 0.25 percent, a reduction in non-durable consumption of slightly less than 0.2 percent and a reduction in investment of nearly 0.5 percent. Note that measured productivity also declines, suggesting that unobserved factor utilization contributes to the decline in production.

Figure 3 reports the impulse responses for the three wage measures we constructed above. The left side panel reports the reaction of average hourly earnings (AHE-BLS). This wage measure barely reacts to the shock. In the center panel, the new-hire wage falls by substantially more. After roughly a year and a half, the new-hire wage has fallen by more than 0.5 percent. The right side panel shows the user cost of labor. The UCL falls even more than the new-hire wage and remains relatively low even more than two years after the shock.

We found similar results when estimating impulse response functions over the time period 1979:4-2007:4. The beginning of this alternative sample corresponds to the beginning of Paul Volcker’s chairmanship of the Federal Reserve, but also has the benefit of excluding any backward projection of the user cost series. In general, the impulse responses from the main block of variables are more muted but take longer to return to trend. Despite this difference, the user cost series still has a peak response near 0.75 percent. The new-hire wages oscillate rapidly, but reach a similar peak response after a similar lag. We conclude that the results are qualitatively unchanged over this shorter sample period that overlaps significantly with Kudlyak’s data sample and is also the period when the “modern” era of U.S. monetary policy may be said to have begun.

5.5 Extension: controlling for match quality

As discussed in Section 3, Hagedorn and Manovskii (2013) argue that much of the observed history dependence of current wages can be understood by appealing to labor search when workers face a job ladder. In Hagedorn and Manovskii’s model, the match quality of a job can be proxied by including the cumulative labor market “tightness” during the employment cycle in the wage regression. Labor market tightness is the ratio of vacancies to unemployment. Intuitively, during an employment cycle, a worker gradually climbs up the match-quality ladder. How fast he or she climbs is determined by current aggregate labor-market tightness. Ultimately, how high the person gets is given by *cumulative* labor-market tightness over the

employment cycle. In this section we extend the results above to include Hagedorn and Manovskii’s proposed measure of labor market tightness, to see whether the results we have reported are robust to the inclusion of this variable.

To implement Hagedorn and Manovskii’s proposed correction, we use the NLSY’s weekly arrays to classify each respondent’s work history into employment cycles. An employment cycle begins when a person finds a job and exits involuntary unemployment. The employment cycle spans the full length of time employed, even if a worker switches employers, as well as *voluntary* spells of unemployment. The employment cycle ends through involuntary unemployment or voluntary unemployment that turns involuntary if the person cannot find a job within two months of voluntarily entering unemployment. The NLSY survey asks individuals why they left their last job, and we use this information to determine whether unemployment is voluntary or involuntary.³⁵

We then calculate the sum of labor market tightness for each job cycle for each individual and include the resulting variable in the individual wage regression (12) as an additional control.^{36,37} Formally, let $\xi_t = \frac{v_t}{u_t}$ denote labor market tightness at date t . Then, for an individual i , currently in an employment spell that began at date $J(i)$, we calculate the sum of the individual’s labor market tightness as $\omega_t^i = \sum_{s=t-J(i)}^t \xi_s$. We then re-estimate (12) and include ω_t^i in the vector of controls X_t^i .

Finally, we modify the prediction equation by assuming that firms hire an individual with average characteristics (as before) and with a fixed average duration of an employment cycle in the NLSY, $\bar{J} = 3.24$ years. That is, we form the projections $\widehat{\ln w_{t,\tau}}$ by including the variable $\bar{x}_t = \sum_{s=t-\bar{J}}^t \xi_s$ in the estimated equation (13). Unlike the variable x_t^i , which varies across workers depending on when their employment spell began, the variable \bar{x}_t exhibits no cross-sectional variation. However, since aggregate labor market tightness ξ_t changes over time, \bar{x}_t

³⁵A person leaves a job involuntarily if he or she is fired, laid off or if the plant closes. If the person voluntarily quits to look for a new job (etc.) and finds a new job within two months, the employment cycle is assumed to continue. If the person voluntarily leaves but it takes longer than two months to find a job, the employment cycle ends and the person falls off the job ladder.

³⁶To construct our measure of labor market tightness, we use the help wanted index calculated by Barnichon (2010).

³⁷Hagedorn and Manovskii actually use two separate controls. They control for cumulative labor market tightness both during a job spell (the variable q^{HM} in their 2013 paper) and also prior to starting a job if the worker either makes a job-to-job transition or starts from a period of voluntary unemployment (the variable q^{EH}). Because (12) includes an arbitrary set of fixed effects $\chi_{\tau,t}$ for the current job spell, the first adjustment (the q^{HM} variable) is already included in our baseline specification. Thus in our analysis we confine our attention to the second adjustment, which controls only for cumulative labor market tightness *prior* to the start of a job.

has time-series variation which is included in the forecasts of $\widehat{\ln w_{t,\tau}}$.

Tables 5 and 6 report results for the cyclicality of the various wage measures after we control for unobserved idiosyncratic match quality. Notice that the cyclicality estimates do decline somewhat, though the overall cyclicality of the wage series is only moderately changed. In particular, the high cyclicality of the new-hire wage and the user cost of labor remain. Similar results are found in response to monetary shocks. Repeating the steps above for the wage measures including the control for match quality gives impulse response functions that are close to the response functions we saw earlier. Figure 4 compares the impulse response of the new-hire wage and the user cost of labor with and without the control for match quality. As shown in the figure, the impulse responses are almost indistinguishable.

6 Comparing the model and the data

We now relate the empirical evidence we have presented regarding the cyclicality of real wages to the predictions of New Keynesian models of the class developed in Section 4. To build intuition, we begin by examining the models without either composition bias or infrequent renegotiation of wage remittances. This requires that we set the renegotiation rate s and the ratio of effective labor to total hours L/N both to 1.00, and the inverse density of types $\varphi(1)^{-1} = 0$. Since s is the renegotiation rate of long-term wage contracts, $s = 1$ implies that remitted wages are changed in every period and thus move in lock step with changes in the (sticky) allocative wage. Since the remitted wage is equal to the allocative wage in this specification, there are no implicit wage contracts in the economy. The assumption of no composition bias implies that the marginal and average worker supplies the same number of efficiency units of labor per observed hour of work. We examine the responses of this baseline New Keynesian model to a monetary shock under the assumption of sticky prices only, sticky wages only, and both sticky prices and wages. The results will guide us regarding the features to add to the model in order to bring the model's predictions closer to the observed wage data.

We set the Calvo parameter for price adjustment to $\theta_p = 0.90$ (quarterly) implying that prices adjust on average once every 10 quarters, or once every two and a half years. We do the same for the initial sticky wage specification, so $\theta_w = 0.90$. While these calibrations are somewhat high relative to the micro data on the average frequency of price and wage changes, they are in line with many DSGE estimates and the implied model impulse response functions have enough persistence for their computed trajectories to be comparable with

the empirical impulse responses. The DSGE model also features traditional mechanisms considered by business cycle analysts to better match the dynamic effects of monetary shocks on output. Specifically, the model features investment adjustment costs, habit persistence in consumption, variable capital utilization, price and wage indexation, and increasing returns to scale in production. The parameter values used are reported in Table 2, and are roughly in line with prevailing estimates in the literature.

Figure 5 shows impulse responses of this baseline model to a 25 basis point shock to the central bank’s policy rate (a shock to the Taylor rule). Since our main interest is in comparing the model responses of different wage measures to the corresponding empirical responses, we show the model and data responses of average hourly earnings (AHE), the new-hire wage (NHW) and the user cost of labor (UCL), as well as output. We reproduce the data responses that appeared in Figure 3 together with responses from the baseline model with only price rigidity, only wage rigidity, and with both types of nominal inertia. We find that, as one might expect, following the increase in the interest rate, all three types of nominal rigidity cause output to decline. Output falls significantly more in the model with only sticky prices than in the model with just sticky wages; of course, it falls further in the model with both nominal frictions. Although the baseline models reproduce the hump-shaped output response observed in the data, the trough of output comes 2-3 quarters earlier in the models than in the data. That is, the models need additional persistence mechanisms or stronger persistence mechanisms to match the estimates.

One of the main findings in the empirical section is that different wage measures behave differently over the business cycle and in response to monetary shocks. In the baseline model, since there is just a single wage (or more precisely, since the average wage, the new hire wage, and the allocative wage are all identical), the model is completely incapable of matching the differential patterns of wages in the data. We see that the impulse responses for the three concepts of the wage are identical in each model. As usual, the UCL in the model is the allocative wage, but with constant wage negotiation and no composition bias, AHE and the NHW are identical to the UCL. Thus, in this set of models there is a single unambiguous wage response to a monetary shock.

In the sticky-price model with flexible wages, the wage declines sharply (thus, it is “pro-cyclical” in the sense that it moves in the same direction as output does following the monetary shock). The wage decline qualitatively matches the responses of the NHW and UCL, but does not match the fact that AHE responds much less. On the other hand, the model with sticky

wages and flexible prices shows a mild *countercyclical* response of the wage to a monetary shock, for the same reasons that Keynes’s (1936) model in the *General Theory* predicted high real wages in recessions. Finally, note that the model with equal (and high) price and wage rigidity shows that the real wage is basically acyclical in the wake of a monetary policy shock. One might summarize the three models by saying that in the model with price rigidity firms are off their labor demand curves but workers are on their labor supply curves, so real wages are procyclical. The situation is reversed in the model with sticky wages, so real wages are countercyclical. (However, the assumption of variable capital utilization flattens the labor demand curve significantly, so the degree of countercyclicity is modest.) Finally, with both sticky wages and prices, both workers and firms are off their notional supply and demand curves in the labor market, and the real wage has no clear cyclical pattern. Note that the wage response in this variant of the model is qualitatively consistent with the empirical response of AHE. Hence, it is clear why modelers who interpret AHE as the allocative wage have been led towards models with both wage and price rigidity, as in CEE.

Starting from the baseline model above, we now consider the effects of implicit contracting and composition bias on model predictions for our three wage measures. In the following discussion, we consider models with sticky prices and flexible wages only. We do this both to conserve space and also because the sticky wage models typically have simulated wage paths that are either sharply counterfactual (i.e., there are sharp increases in wage payments following a negative monetary shock) or wage paths that are acyclical which, while matching the observed behavior of average hourly earnings, fail to match the responses of new hire wages and the UCL.

We begin by examining the role of implicit contracting. Starting with the baseline model above, we consider the effects of gradually reducing the parameter s from its initial value of 1.00. When $s < 1$, the *remitted* wage is changed infrequently even though the allocative wage (the UCL) is fully flexible, since in this model we assume no wage rigidity. While the UCL is free to react to changing economic conditions, other measures of the wage—AHE and the NHW—change by substantially less than the UCL. The results are shown in Figure 6. Note that the results for $s = 1$ reproduce the sticky-price impulse responses of the previous figure. As discussed in Section 3, Barattieri *et al.* (2014) find that $s = 0.21$ is the approximate frequency of changes in remitted wages observed in micro wage data. We include $s = 0.50$ as an intermediate case. Note that the three impulse responses are identical for the UCL – the allocative wage is unaffected by the value of s . However, $s < 0$ implies the existence of

implicit contracts, which makes the three wage measures differ in their responses to monetary policy shocks. Particularly interesting is the result for the measured value of $s = 0.21$. For this value of s , the allocative wage falls sharply, the wage for new hires falls less, and average hourly earnings fall only slightly. The pattern of wage responses for the three wage measures relative to the output response bears a strong qualitative resemblance to the empirical impulse responses. This observation leads us to conclude tentatively that the evidence suggests that a model with sticky prices, flexible wages, and a significant role for implicit wage contracts has the best chance of matching the data.

Figure 7 shows the impulse response function for the model with composition bias effects. Endogenous composition adjustment has two separate effects on the responses of wages to monetary (and non-monetary) shocks. First, by introducing a difference between the average labor compensation of employed workers and the labor compensation for the “marginal” workers, composition adjustments cause average hourly earnings to be less responsive than the user cost of labor which holds labor force composition fixed. The magnitude of this differential is given by the ratio $(LN - 1)/LN$ in equation (10) where $LN \geq 1$ is the steady state ratio of the effective labor supply to hours worked or equivalently, LN is the ratio of average labor compensation to wages paid to marginal workers. In a model without endogenous composition adjustment, this ratio is 1 and there are no effects of composition bias on measured wages. If $LN > 1$ then average wages move by less than the user cost.

The second effect of compositional changes is that the effective labor supply elasticity in such an environment is strictly less than the individual labor supply elasticities. The reason for this is that expanding employment means hiring workers who are increasingly less productive. The magnitude of this effect is governed by the inverse density of types at the productivity cutoff $\varphi(1)^{-1}$. In a typical model in which all workers are the same, $\varphi(1) = \infty$ (i.e., there is a mass point at the common productivity 1) and thus $\varphi(1)^{-1} = 0$. If the density of types at the cutoff is smooth however, $\varphi(1)^{-1}$ is greater than 0 indicating that hiring more workers requires lowering the marginal productivity. If $\varphi(1)^{-1}$ is large then expanding the workforce requires tolerating much lower productivity workers and thus the effective labor supply elasticity is substantially lower.

We consider three different model specifications in Figure 7. First we report the response for the standard model (the thin solid line) without composition bias. The dashed line reports the impulse responses for a model with an “intermediate” degree of composition bias. For this specification we consider a case with $LN = 2.0$ (so the average worker is paid twice as

much as the marginal worker) and $\varphi^{-1}(1) = 2.00$. The dotted line reports the responses for a “high” degree of composition bias in which $LN = \varphi^{-1}(1) = 4.00$. The figure displays both of the effects mentioned above. Notice in particular that the specifications with composition bias feature notably sharper reductions in wages. This is because, in the sticky-price environment, output, and thus labor, is effectively demand determined. Given demand, the firms simply hire or fire as many workers as necessary to increase or decrease production. Since the effective labor supply elasticity is reduced by the compositional adjustments, the wages must fall by more. Also, notice that the reduction in average hourly earnings is less than the decline in the user cost. This is a direct consequence of $LN > 1$: exactly the effect highlighted in Solon, Barsky and Parker (1994).

Figure 8 considers the baseline model with both implicit wage smoothing contracts and a modest amount of composition bias. For this simulation, we set $s = 0.21$ as suggested by Barattieri *et al.* (2014) and we adopt the intermediate composition bias specification, $LN = \varphi(1)^{-1} = 2.0$.

As we did in Figure 5, Figure 8 shows the impulse responses under the assumption of pure sticky prices (solid line), pure sticky wages (dashed line) and a specification with both sticky prices and sticky wages (dotted line). Based on this simulation, the sticky price specification seems to outperform both of the other cases. In the sticky wage case and the specification with both types of nominal rigidities, average hourly earnings rise noticeably following the monetary contraction, while the sticky price model is surprisingly close to the actual point estimates from the SVAR impulse response, which show a slight decline. While still not matching the shape of the dynamic responses to the new hire wage or the user cost of labor, the model responses are quantitatively close. The model predicts a maximum decline in the new hire wage of roughly 0.50 percent and a maximum decline in the user cost of labor of 0.70 percent. In comparison, the SVAR estimates a decline in the new hire wage of roughly 0.70 percent and a decline in the user cost of almost 0.80 percent. The simulated responses for the other two model specifications display much smaller movements of either the new hire wage or the user cost (in all cases, the peak declines are less than 0.10 percent).

The main problem with the impulse responses of the sticky-price model relative to the data is the lack of persistence of output following a monetary contraction. Output in the model attains its maximum response three quarters after the shock, as opposed to six quarters in the data. The new hire wage and the user cost of labor hit their troughs at the same time as output, meaning that they display a phase shift relative to the empirical impulse responses.

Thus, we conjecture that a more persistent behavior of real variables after a monetary shock would bring the model and the data into closer alignment.

At the same time, it is easy to see why the model lacks persistence. Models such as CEE use sticky wages as an important persistence mechanism, but as Figure 8 shows, allocative wage stickiness is inconsistent with the observed behavior of the UCL following a monetary contraction. Thus, our preferred model specification with sticky prices and flexible wages is missing one of the key propagation mechanisms featured in many standard medium-scale DSGE models in the literature. An important agenda for future research is to find new propagation mechanisms for sticky-price models to replace the assumption of sticky allocative wages, which seems to be fundamentally inconsistent with the data.

7 Conclusion

Recent empirical studies suggest that the cyclicity of real wages is greater than conventional wisdom would suggest. The literature emphasizes two reasons for this enhanced cyclicity. First, endogenous changes in the composition of the workforce mechanically causes average hourly wage payments to understate the change in wages relative to wage changes holding workforce composition fixed. Second, there are indications that the allocative wage – the wage that governs hours worked and that firms internalize when making production and pricing decisions – may not equal the contemporaneous remitted wage. In particular, firms and workers may well have an implicit understanding that the remitted wage will be a smoothed version of the expected allocative wage. By estimating the expected present value of wage payments, one can construct a “user cost of labor,” which should measure the underlying allocative wage.

In this chapter, we have reproduced and extended the key empirical results in Kudlyak’s (2014) work. Our empirical analysis confirms her calculations for the cyclicity of the allocative wage. In addition, the NLSY also allows us to decompose the cyclical response of wages by controlling separately for compositional changes and at the same time controlling for the wage smoothing effects of implicit contracts. The data suggest that while compositional changes contribute significantly to the dynamics of average hourly earnings, the effects of implicit contracts and wage smoothing are even greater.

Using the annual estimates of the user cost of labor and the new hire wages from the NLSY as a starting point, we extend the estimated series to a quarterly series and include the

extended data in a structural VAR for the purposes of studying the reactions of real wages to monetary shocks. The estimated structural VAR suggests that the user cost of labor and new hire wages both decline sharply following a contractionary monetary shock. In contrast, average hourly earnings—the usual measure of the wage in macroeconomic research—barely respond to such shocks. Our model, if extended to allow for wage smoothing within long-term worker-firm associations, can match the fact that average hourly earnings and even the wage paid to new hires are significantly less cyclical than the allocative user cost of labor.

The differential reactions of these wage measures present two key challenges for prevailing New Keynesian models of the monetary transmission mechanism. First, in most New Keynesian models, there is no conceptual difference between the allocative wage, the remitted wage, and average hourly earnings. Thus, at a basic level, these models will not be able to match the empirical findings we study. Second, to the extent that New Keynesian models include a prominent role for sticky nominal wages, the models typically predict that either the allocative wage will counterfactually rise in the wake of an increase in the policy rate or that wages will not respond at all. Neither of these predictions would match our basic finding that allocative wages appear to fall sharply after a monetary tightening.

Analysis of a medium-scale DSGE model suggests that successful models will emphasize price rigidity rather than wage rigidity. In addition, to match the estimates in this chapter, such models will likely allow for a relatively smooth remitted wage which is not allocative. The wage data, therefore, favor the “old New Keynesian economics”—the early models of Rotemberg (1982), Ball and Romer (1989), and Kimball (1995), which all assumed competitive labor markets and flexible allocative wages—rather than the “new New Keynesian economics” with allocative prices and wages both sticky, as in Blanchard and Kiyotaki (1987), Erceg, Henderson and Levin (2000), Smets and Wouters (2007) and CEE.

We conclude by suggesting that research is needed on two fronts, one theoretical and one empirical. The theoretical challenge arising from these results is the need to explain the observed persistence of the real effects of monetary shocks without being able to rely on wage stickiness to temper the response of marginal cost to a monetary shock. Our estimates suggest that real marginal cost, properly computed, is strongly procyclical.³⁸ This conclusion, if correct, casts doubt on the main persistence mechanism of medium-scale New Keynesian

³⁸This agrees with some, although not all, of the literature on the countercyclicality of the price markup, which is of course just the inverse of real marginal cost. See for example, Rotemberg and Woodford (1999) and Bils, Klenow and Malin (2014).

models, such as CEE and Smets and Wouters (2007), which generally rely on assumptions that make marginal cost acyclical. Thus, the problem facing monetary economics can be restated as: Why do prices behave sluggishly, even though wages and hence marginal costs are strongly procyclical?

This question was the focus of a major research program several decades earlier. Relative to the research that took place in the 1980s, new observations on firm-level prices and quantities, and the desire to have persistence mechanisms that are consistent with both time-dependent and state-dependent pricing models, impose additional constraints on the proposed solutions.³⁹ To use the language of Ball and Romer (1990), the search for “real rigidities” in price setting must arrive at a satisfactory conclusion in order to make models of the monetary transmission mechanism consistent with recent observations of the data.

The empirical challenge is to extend the measurement of real wage cyclicality to other data sets and other countries. It would be particularly interesting to compare the results reported here with similar calculations for the major continental European economies, or for the Euro area as a whole. The labor market is one area where economists have argued that the differences between the U.S. and Europe are most pronounced. Galí (2016) follows in this tradition by incorporating “hysteresis” effects into a standard New Keynesian framework, arguing that this modification is necessary to match the greater persistence of the unemployment rate in Europe. Yet this change, where the unemployed exert little downward pressure on wages, makes real wages even less sensitive to the business cycle than they are in the standard New Keynesian models developed to explain U.S. macro data. Do European micro data indicate that there is enormously more allocative wage rigidity in Europe than in the U.S.? Finding out whether the answer is yes or no is clearly of first-order importance for understanding cyclical fluctuations in these two major economies.

³⁹See, for example, Klenow and Willis (2016), Dotsey and King (2005), and Nakamura and Steinsson (2010). Leahy (2011) and Nakamura and Steinsson (2013, Section 12) provide insightful discussions based on Ball and Romer (1990).

References

- Abraham, Katharine G., and John C. Haltiwanger. 1995. "Real Wages and the Business Cycle." *Journal of Economic Literature* 33(3): 1215–1264.
- Akerlof, George, William R. Dickens, and George Perry. 1996. "The Macroeconomics of Low Inflation." *Brookings Papers on Economic Activity* 27(1): 1–76.
- Ball, Laurence, and David Romer. 1989. "The Equilibrium and Optimal Timing of Price Changes." *Review of Economic Studies* 56(2): 179–198.
- Ball, Laurence, and David Romer. 1990. "Real Rigidities and the Non-neutrality of Money." *Review of Economic Studies* 57(2): 183–203.
- Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk. 2014. "Some Evidence on the Importance of Sticky Wages." *American Economic Journal: Macroeconomics* 6(1): 70–101.
- Barnichon, Regis. 2010. "Building a Composite Help-Wanted Index." *Economics Letters* 109(3): 175–178.
- Barro, Robert J. 1977. "Long-term Contracting, Sticky Prices, and Monetary Policy." *Journal of Monetary Economics* 3(3): 305–316.
- Barro, Robert J., and Robert G. King. 1984. "Time-separable Preferences and Intertemporal-substitution Models of Business Cycles." *Quarterly Journal of Economics* 99(4): 817–839.
- Basu, Susanto, and Brent Bundick. 2017. "Uncertainty Shocks in a Model of Effective Demand." *Econometrica*, forthcoming.
- Basu, Susanto, and Miles S. Kimball. 1997. "Cyclical Productivity with Unobserved Input Variation." Working Paper 5915. National Bureau of Economic Research.
- Basu, Susanto, and Alan M. Taylor. 1999. "Business Cycles in International Historical Perspective." *Journal of Economic Perspectives* 13(2): 45–68.
- Beaudry, Paul, and John DiNardo. 1991. "The Effect of Implicit Contracts on the Movement of Wages Over the Business Cycle: Evidence from Micro Data." *Journal of Political Economy* 99(4): 665–688.
- Beaudry, Paul, and John DiNardo. 1995. "Is the Behavior of Hours Worked Consistent with Implicit Contract Theory?" *Quarterly Journal of Economics* 110(3): 743–768.

- Becker, Gary S. 1962. "Investment in Human Capital: A Theoretical Analysis." *Journal of Political Economy* 70(5): 9–49.
- Ben S. Bernanke, Alan S. Blinder. 1992. "The Federal Funds Rate and the Channels of Monetary Transmission." *American Economic Review* 82(4): 901–921.
- Bernanke, Ben S., and Kevin Carey. 1996. "Nominal Wage Stickiness and Aggregate Supply in the Great Depression." *Quarterly Journal of Economics* 111(3): 853–883.
- Bewley, Truman F. 1999. *Why Wages Don't Fall During a Recession*. Cambridge, MA: Harvard University Press.
- Bils, Mark. 1985. "Real Wages over the Business Cycle: Evidence from Panel Data." *Journal of Political Economy* 93(4): 666–689.
- Bils, Mark, and Peter J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112(5): 947–985.
- Bils, Mark, Peter J. Klenow, and Benjamin A. Malin. 2014. "Resurrecting the Role of the Product Market Wedge in Recessions." Working Paper 20555. National Bureau of Economic Research.
- Blanchard, Olivier Jean, and Nobuhiro Kiyotaki. 1987. "Monopolistic Competition and the Effects of Aggregate Demand." *American Economic Review* 77(4): 647–666.
- Card, David, and Dean Hyslop. 1997. "Does Inflation 'Grease the Wheels of the Labor Market'?" In *Reducing Inflation: Motivation and Strategy*, eds. Christina D. Romer and David H. Romer. Chicago: University of Chicago Press.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 1999. "Monetary Policy Shocks: What Have We Learned and to What End?" In *Handbook of Macroeconomics*, eds. John B. Taylor and Michael Woodford, vol. 1, chap. 2, 65–148. Elsevier, 1st ed.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113(1): 1–45.
- Daly, Mary C., and Bart Hobijn. 2016. "The Intensive and Extensive Margins of Real Wage Adjustment." Federal Reserve Bank of San Francisco Working Paper 2016-04. <http://www.frbsf.org/economic-research/files/wp2016-04.pdf>.
- Dotsey, Michael, and Robert G. King. 2005. "Implications of State-dependent Pricing for Dynamic Macroeconomic Models." *Journal of Monetary Economics* 52(1): 213–242.

- Dunlop, John T. 1938. “The Movement of Real and Money Wage Rates.” *Economic Journal* 48(191): 413–434.
- Eichengreen, Barry, and Jeffrey Sachs. 1985. “Exchange Rates and Economic Recovery in the 1930s.” *Journal of Economic History* 45: 925–946.
- Elsby, Michael W. L. 2009. “Evaluating the Economic Significance of Downward Nominal Wage Rigidity.” *Journal of Monetary Economics* 56(2): 154–169.
- Elsby, Michael W. L., Donggyun Shin, and Gary Solon. 2016. “Wage Adjustment in the Great Recession and Other Downturns: Evidence from the United States and Great Britain.” *Journal of Labor Economics* 34(S1): S249–S291.
- Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin. 2000. “Optimal Monetary Policy with Staggered Wage and Price Contracts.” *Journal of Monetary Economics* 46(2): 281–313.
- Fleischman, Charles A. 1999. “The Causes of Business Cycles and the Cyclicity of Real Wages.” Finance and Economics Discussion Series 1999–53. Board of Governors of the Federal Reserve System.
- Friedman, Milton, and Anna Jacobson Schwartz. 1963. *A Monetary History of the United States*. Princeton: Princeton University Press.
- Galí, Jordi. 2013. “Notes for a New Guide to Keynes (I): Wages, Aggregate Demand, and Employment.” *Journal of the European Economic Association* 11(5): 973–1003.
- Galí, Jordi. 2016. “Insider–Outsider Labor Markets, Hysteresis and Monetary Policy” (working paper, http://crei.cat/people/gali/mp_hysteresis_jan2016.pdf).
- Galí, Jordi, and Mark Gertler. 1999. “Inflation Dynamics: A Structural Econometric Analysis.” *Journal of Monetary Economics* 44(2): 195–222.
- Galí, Jordi, Mark Gertler, and J. David López-Salido. 2007. “Markups, Gaps, and the Welfare Costs of Business Fluctuations.” *Review of Economics and Statistics* 89(1): 44–59.
- Geary, Patrick T., and John Kennan. 1982. “The Employment-Real Wage Relationship: An International Study.” *Journal of Political Economy* 90(4): 854–871.
- Gertler, Mark, and Antonella Trigari. 2009. “Unemployment Fluctuations with Staggered Nash Wage Bargaining.” *Journal of Political Economy* 117(1): 38–86.

- Gottschalk, Peter. 2005. "Downward Nominal-wage Flexibility: Real or Measurement Error?" *Review of Economics and Statistics* 87(3): 556–568.
- Gottschalk, Peter, and Minh Huynh. 2010. "Are Earnings Inequality and Mobility Overstated? The Impact of Nonclassical Measurement Error." *Review of Economics and Statistics* 92(2): 302–315.
- Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman. 1988. "Investment, Capacity Utilization, and the Real Business Cycle." *American Economic Review* 78(3): 402–417.
- Haefke, Christian, Marcus Sonntag, and Thijs van Rens. 2013. "Wage Rigidity and Job Creation." *Journal of Monetary Economics* 60(8): 887–899.
- Hagedorn, Marcus, and Iourii Manovskii. 2013. "Job Selection and Wages over the Business Cycle." *American Economic Review* 103(2): 771–803.
- Hall, Robert E. 2005. "Employment Fluctuations with Equilibrium Wage Stickiness." *American Economic Review* 95(1): 50–65.
- Hall, Robert E., and Paul R. Milgrom. 2008. "The Limited Influence of Unemployment on the Wage Bargain." *American Economic Review* 98(4): 1653–74.
- Hanes, Christopher. 1993. "The Development of Nominal Wage Rigidity in the Late 19th Century." *American Economic Review* 83(4): 732–756.
- Hanes, Christopher. 1996. "Changes in the Cyclical Behavior of Real Wage Rates, 1870–1990." *Journal of Economic History* 56(4): 837–861.
- Hanes, Christopher, and John A. James. 2003. "Wage Adjustment under Low Inflation: Evidence from U.S. History." *American Economic Review* 93(4): 1414–1424.
- Hansen, Gary D. 1985. "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics* 16(3): 309–327.
- Hanson, Michael S. 2004. "The 'Price Puzzle' Reconsidered." *Journal of Monetary Economics* 51(7): 1385–1413.
- Huang, Kevin X. D., Zheng Liu, and Louis Phaneuf. 2004. "Why Does the Cyclical Behavior of Real Wages Change over Time?" *American Economic Review* 94(4): 836–856.
- Hume, David. (1742) 1987. *Essays, Moral, Political, and Literary*. Indianapolis, IN: Library of Economics and Liberty. [Online] available from <http://www.econlib.org/library/LFBooks/Hume/hmMPL.html>; accessed 9 March 2016; Internet.

- Kahn, Shulamit. 1997. “Evidence of Nominal Wage Stickiness from Microdata.” *American Economic Review* 87(5): 993–1008.
- Keynes, John M. 1936. *The General Theory of Employment, Interest, and Money*. London: Macmillan.
- Kimball, Miles S. 1995. “The Quantitative Analytics of the Basic Neomonetarist Model.” *Journal of Money, Credit and Banking* 27(4): 1241–1277.
- King, Robert, and Sergio Rebelo. 1999. “Resuscitating Real Business Cycles.” In *Handbook of Macroeconomics*, eds. J. B. Taylor and M. Woodford, vol. 1, Part B, chap. 14, 927–1007. Elsevier, 1st ed.
- Klenow, Peter J., and Jonathan L. Willis. 2016. “Real Rigidities and Nominal Price Changes.” Working Paper, available <http://klenow.com/RealRigidities.pdf>.
- Kudlyak, Marianna. 2014. “The Cyclicalities of the User Cost of Labor.” *Journal of Monetary Economics* 68: 53–67.
- Le Bihan, Hervé, Jérémie Montornès, and Thomas Heckel. 2012. “Sticky Wages: Evidence from Quarterly Microeconomic Data.” *American Economic Journal: Macroeconomics* 4(3): 1–32.
- Leahy, John. 2011. “A Survey of New Keynesian Theories of Aggregate Supply and Their Relation to Industrial Organization.” *Journal of Money, Credit and Banking* 43: 87–110.
- Lebow, David E., Raven E. Saks, and Beth Anne Wilson. 1999. “Downward Nominal Wage Rigidity: Evidence from the Employment Cost Index.” Finance and Economic Discussion Series Working Paper 1999-31. <http://www.federalreserve.gov/pubs/feds/1999/199931/199931pap.pdf>.
- Lünnemann, Patrick, and Ladislav Wintr. 2009. “Wages Are Flexible, Aren’t They? Evidence from Monthly Micro Wage Data.” European Central Bank (ECB) Working Paper 1074.
- McLaughlin, Kenneth J. 1994. “Rigid Wages?” *Journal of Monetary Economics* 34(3): 383–414.
- Nakamura, Emi, and Jón Steinsson. 2010. “Monetary Non-neutrality in a Multisector Menu Cost Model.” *Quarterly Journal of Economics* 125(3): 961–1013.
- Nakamura, Emi, and Jón Steinsson. 2013. “Price Rigidity: Microeconomic Evidence and Macroeconomic Implications.” *Annual Review of Economics* 5(1): 133–163.

- Nekarda, Christopher J., and Valerie A. Ramey. 2013. “The Cyclical Behavior of the Price-Cost Markup.” Working Paper 19099. National Bureau of Economic Research.
- Pencavel, John. 2015. “Keynesian Controversies on Wages.” *Economic Journal* 125(583): 295–349.
- Pissarides, Christopher A. 2009. “The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?” *Econometrica* 77(5): 1339–1369.
- Ramey, Valerie A. 2016. “Macroeconomic Shocks and Their Propagation.” In *Handbook of Macroeconomics*, eds. John B. Taylor and Harald Uhlig, vol. 2, 71–162. Elsevier.
- Ravenna, Federico, and Carl E. Walsh. 2008. “Vacancies, Unemployment, and the Phillips Curve.” *European Economic Review* 52(8): 1494–1521.
- Rogerson, Richard. 1988. “Indivisible Labor, Lotteries and Equilibrium.” *Journal of Monetary Economics* 21(1): 3–16.
- Romer, Christina D., and David H. Romer. 1989. “Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz.” *NBER Macroeconomics Annual* 4: 121–184.
- Romer, Christina D., and David H. Romer. 2004. “A New Measure of Monetary Shocks: Derivation and Implications.” *American Economic Review* 94(4): 1055–1084.
- Rotemberg, Julio J. 1982. “Monopolistic Price Adjustment and Aggregate Output.” *Review of Economic Studies* 49(4): 517–531.
- Rotemberg, Julio J., and Michael Woodford. 1991. “Markups and the Business Cycle.” *NBER Macroeconomics Annual* 6: 63–140.
- Rotemberg, Julio J., and Michael Woodford. 1999. “The Cyclical Behavior of Prices and Costs.” In *Handbook of Macroeconomics*, eds. John B. Taylor and Michael Woodford, vol. 1, Part B, chap. 16, 1051–1135. Elsevier.
- Sbordone, Argia M. 2002. “Prices and Unit Labor Costs: A New Test of Price Stickiness.” *Journal of Monetary Economics* 49(2): 265–292.
- Schmitt-Grohé, Stephanie, and Martín Uribe. 2013. “Downward Nominal Wage Rigidity and the Case for Temporary Inflation in the Eurozone.” *Journal of Economic Perspectives* 27(3): 193–212.
- Shimer, Robert. 2005. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies.” *American Economic Review* 95(1): 25–49.

- Sigurdsson, Jósef, and Rannveig Sigurdardottir. 2016. “Time-dependent or State-dependent Wage-setting? Evidence from Periods of Macroeconomic Instability.” *Journal of Monetary Economics* 78: 50–66.
- Smets, Frank, and Rafael Wouters. 2007. “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach.” *American Economic Review* 97(3): 586–606.
- Solon, Gary, Robert Barsky, and Jonathan A. Parker. 1994. “Measuring the Cyclical-ity of Real Wages: How Important is Composition Bias.” *Quarterly Journal of Economics* 109(1): 1–25.
- Solow, Robert M. 1979. “Another Possible Source of Wage Stickiness.” *Journal of Macroeconomics* 1(1): 79–82.
- Stock, James H., and Mark W. Watson. 1999. “Business Cycle Fluctuations in US macroeconomic Time Series.” In *Handbook of Macroeconomics*, eds. John B. Taylor and Michael Woodford, vol. 1, Part A, chap. 1, 3–64. Elsevier.
- Stockman, Alan. 1983. “Aggregation Bias and the Cyclical Behavior of Real Wages.” University of Rochester, Economics Department, Mimeo.
- Sumner, Scott, and Stephen Silver. 1989. “Real Wages, Employment, and the Phillips Curve.” *Journal of Political Economy* 97(3): 706–720.
- Tarshis, Lorie. 1939. “Changes in Real and Money Wages.” *Economic Journal* 49(193): 150–154.
- Taylor, J.B. 2016. “The Staying Power of Staggered Wage and Price Setting Models in Macroeconomics.” In *Handbook of Macroeconomics*, eds. John B. Taylor and Harald Uhlig, vol. 2, 2009–2042. Elsevier.
- Tobin, James. 1972. “Inflation and Unemployment.” *American Economic Review* 62(1/2): 1–18.
- Walsh, Carl E. 2003. “Labor Market Search and Monetary Shocks.” In *Dynamic Macroeconomic Analysis*, eds. Sumru Altug, Jagjit S. Chadha, and Charles Nolan, 451–486. Cambridge, UK: Cambridge University Press.
- Wolpin, Kenneth I. 1992. “The Determinants of Black–White Differences in Early Employment Careers: Search, Layoffs, Quits, and Endogenous Wage Growth.” *Journal of Political Economy* 100(3): 535–560.

Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.

Table 1. Wage Change Frequency in SIPP Data

	Hourly Workers [BBG (2014)]	Salaried Workers
Reported	0.565	0.721
Adjusted	0.120	0.061
Adjusted + Correction [GH (2010)]	0.211	0.209
Number of obs.	17,148	21,947

Table 2. Parameters for New Keynesian DSGE Model

Parameter	Value
Discount factor, annual (β)	0.97
Intertemporal elasticity of substitution (σ)	1.00
Frisch labor supply elasticity (η)	1.00
Depreciation rate, annual (δ)	0.10
Capital's share (α)	0.36
Type-specific labor elasticity (ψ)	21.00
Type-specific product elasticity (ε)	6.00
Average duration of prices, quarterly $((1 - \theta_p)^{-1})$	10.00
Average duration of wages, quarterly $((1 - \theta_w)^{-1})$	10.00
Inflation indexing	Yes
Marginal cost of capital utilization $(b''(1)/b'(1))$	1.00
Investment adjustment cost (κ)	4.00
Habit weight (h)	0.65
Ratio of effective labor to total hours (LN)	2.00
Inverse density at unit productivity $(\varphi(1)^{-1})$	2.00
Renegotiation rate, quarterly (s)	0.21

Table 3. Real Wage Cyclicalities: Unemployment Rate

	AHE-BLS	AHE-NLSY			New Hire	UCL
	(1)	Base (2)	Controls (3)	Controls, FEs (4)	(5)	(6)
HP-filtered						
Unemployment Rate	-0.507 (0.471)	-0.976 (1.530)	-1.185 (1.507)	-1.328 (1.623)	-0.698 (1.822)	-5.818 (2.079)
Observations	34	25	25	25	25	27

Notes: OLS standard errors are in parentheses. Coefficients are multiplied by 100.

Table 4. Real Wage Cyclicalilty: GDP

	AHE-BLS	AHE-NLSY			New Hire	UCL
	(1)	Base (2)	Controls (3)	Controls, FEs (4)	(5)	(6)
HP-filtered GDP	0.311 (0.353)	0.984 (1.093)	0.960 (1.082)	1.165 (1.161)	1.325 (1.287)	3.122 (1.351)
Observations	34	25	25	25	25	27

Notes: OLS standard errors are in parentheses.

Table 5. Real Wage Cyclicalilty Controlling for Match Quality: Unemployment Rate

	AHE-BLS	AHE-NLSY			New Hire	UCL
	(1)	Base (2)	Controls (3)	Controls, FEs (4)	(5)	(6)
HP-filtered Unemployment Rate	-0.507 (0.471)	-1.039 (1.833)	-1.092 (1.729)	-1.294 (1.764)	-0.691 (1.851)	-4.773 (2.049)
Observations	34	25	25	25	25	27

Notes: OLS standard errors are in parentheses. Coefficients are multiplied by 100.

Table 6. Real Wage Cyclicalilty Controlling for Match Quality: GDP

	AHE-BLS	AHE-NLSY			New Hire	UCL
	(1)	Base (2)	Controls (3)	Controls, FEs (4)	(5)	(6)
HP-filtered GDP	0.311 (0.353)	1.000 (1.682)	0.844 (1.568)	1.069 (1.631)	1.244 (1.311)	2.284 (1.336)
Observations	34	25	25	25	25	27

Notes: OLS standard errors are in parentheses.




Fig. 1. Measures of the real wage.

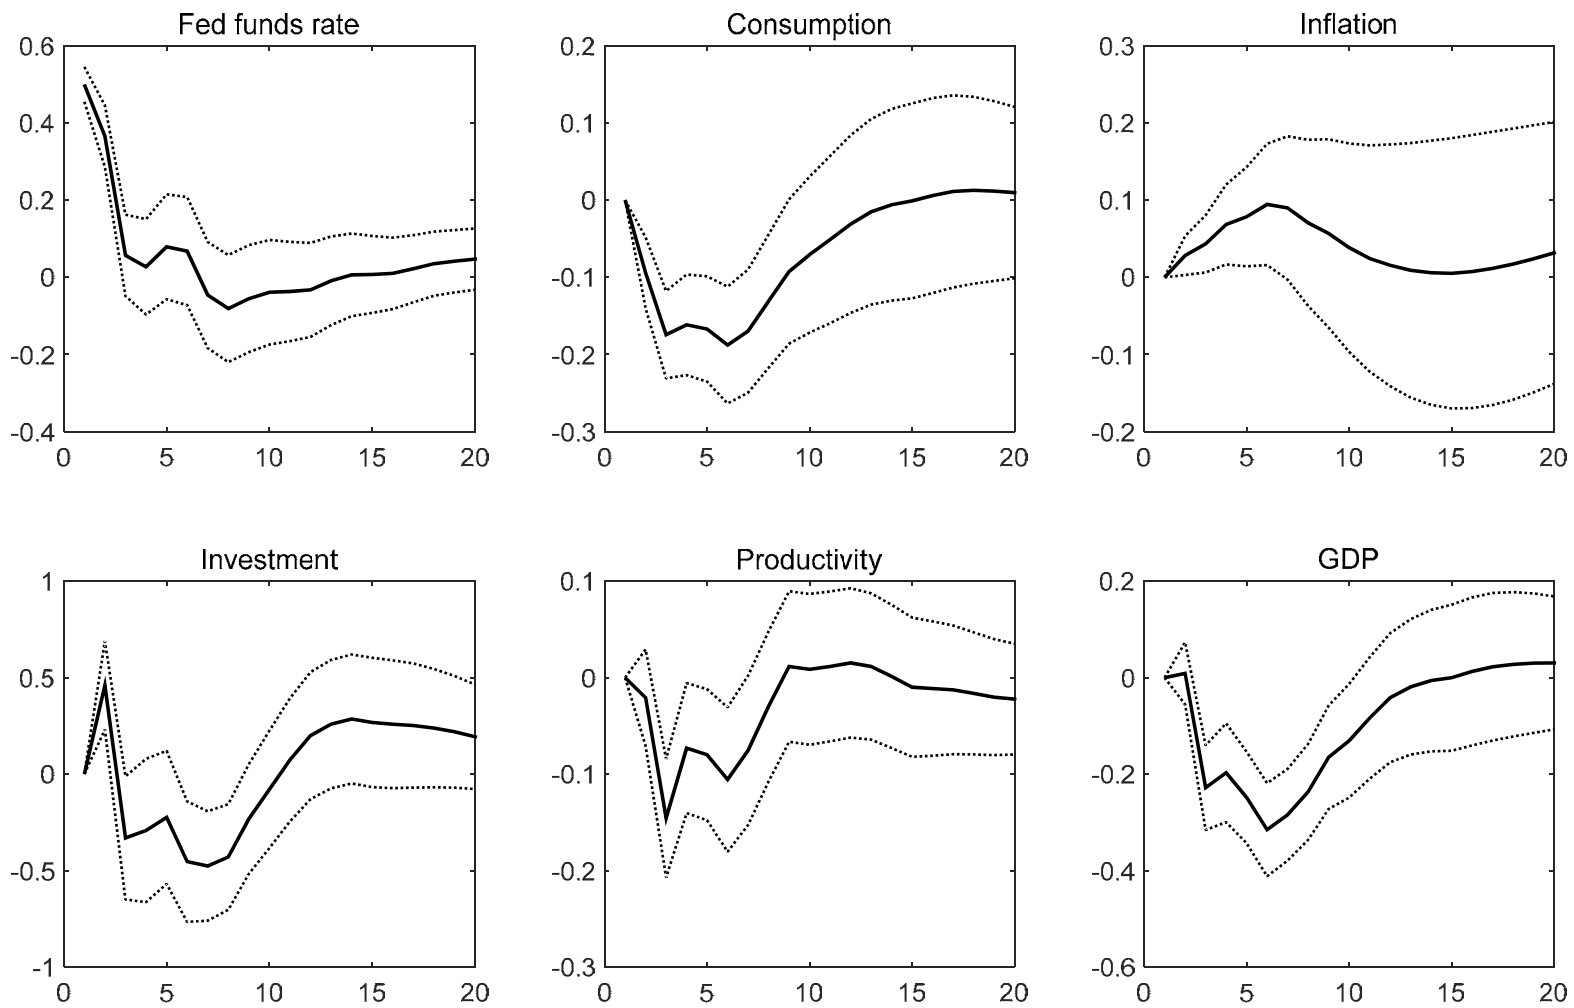


Fig. 2. Impulse responses to an identified monetary contraction: standard variables.



Fig. 3. Impulse responses to an identified monetary contraction: real wage measures.




Fig. 4. Impulse responses controlling for match quality.




Fig. 5. Wage dynamics in Baseline New Keynesian models.

Notes: Each panel reports the estimated impulse responses (heavy line) and model impulse responses to a 25 basis point shock to the federal funds rate. The thin solid line displays the response of the baseline New Keynesian model with sticky prices but flexible wages. The dashed line displays the response of the model with sticky wages but flexible prices, and the dotted line displays the responses of the model with both sticky prices and sticky wages.




Fig. 6. Implicit contracts in sticky price models.

Notes: Each panel reports the estimated impulse responses and model impulse responses to a 25 basis point shock to the federal funds rate. The different lines correspond to different parameter values. In all cases, prices are sticky but wages are flexible. The thin solid line displays the responses of the baseline model when remitted wages are reset with a quarterly probability (s) of 0.21 (roughly every 15 months). The dashed line displays results for the model with $s = 0.50$ and the dotted line displays results when $s = 1.00$ (continual adjustment of the remitted wage).




Fig. 7. Wage dynamics and composition bias in sticky price models.

Notes: Each panel reports the estimated impulse responses and model impulse responses to a 25 basis point shock to the federal funds rate. The different lines correspond to different parameter values. In all cases, prices are sticky but wages are flexible. The thin solid line displays the responses of the model with no composition bias ($LN = 1.00$ and $\varphi^{-1}(1) = 0.0001$). The dashed line displays results for an intermediate level composition bias ($LN = 2.00$ and $\varphi^{-1}(1) = 2.00$). The dotted line displays results for high composition bias ($LN = 4.00$ and $\varphi^{-1}(1) = 4.00$).




Fig. 8. Real wage dynamics in the modified model.

Notes: Each panel reports the estimated impulse responses (heavy line) and model impulse responses to a 25 basis point shock to the federal funds rate. The parameter values for the model are set according to Table 2. We use $s = 0.21$ and the intermediate level of composition adjustment ($LN = 2.00$ and $\varphi^{-1}(1) = 2.00$). The thin solid line displays the response of a model with sticky prices but flexible wages. The dashed line displays the response of a model with sticky wages but flexible prices, and the dotted line displays the responses of a model with both sticky prices and sticky wages.