Understanding the United Kingdom's Wageless Recovery*

Benjamin D. Nelson
Rokos Capital Management LLP and Centre for Macroeconomics

The recovery of 2013 that followed the United Kingdom's Great Recession featured a rapid fall in unemployment but stagnant wage growth. Did the wage Phillips curve break down? These dynamics have two main candidate explanations: declining labor frictions, meaning lower unemployment without increasing wage growth; or a demand recovery accompanied by weak productivity, meaning unemployment fell but equilibrium wage growth remained low. This paper investigates using an estimated New Keynesian model featuring unemployment. The data favor a mix of explanations, but with the balance of evidence favoring the second. A demand recovery reduced unemployment, but wages are likely to have remained weak mainly because of poor productivity.

JEL Codes: E23, E32.

1. Introduction

The United Kingdom's recovery following the global financial crisis featured two notable labor market "puzzles." From around the middle of 2013, unemployment fell rapidly—faster than would be implied by GDP growth—from 7.8 percent (2013:Q2) to 6.3 percent a year later, to around its pre-crisis average by the end of 2015, and to historically low levels by the end of 2017 (figure 1). At the same time, wage growth remained weak by past standards—weaker than would have been implied by the unemployment rate alone. Annual private-sector regular pay growth rose from around 1.0 percent in

*I am grateful to the editor and two anonymous referees for valuable comments and suggestions. All errors and omissions remain those of the author. The views expressed herein represent those of the author and not necessarily those of any employer. E-mail: benjamindnelson@hotmail.com.

Figure 1. Wage Growth Weak Throughout the Recovery While Unemployment Fell Rapidly from 2013

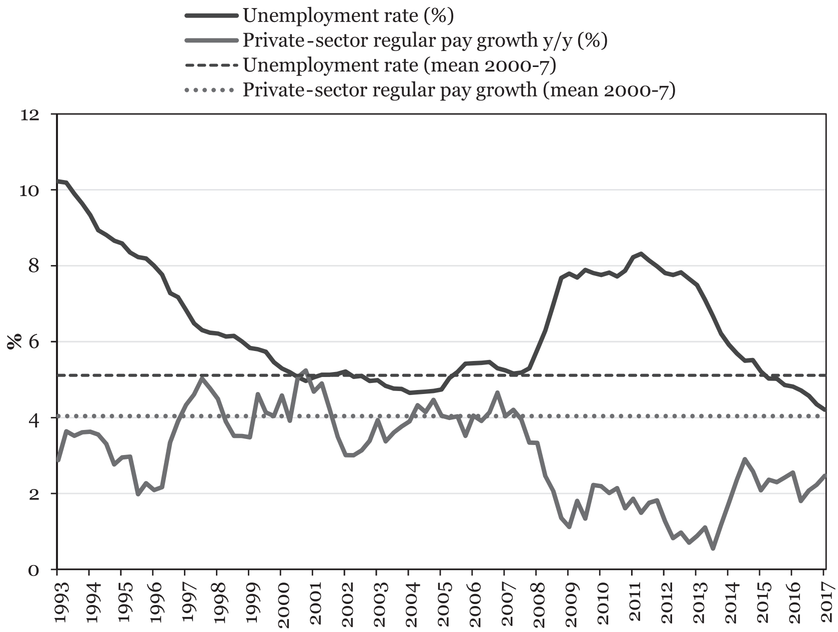

Source: Office for National Statistics.

2013 to 2.9 percent in 2015, before declining to 2.7 percent in 2016 and 2.3 percent in 2017. At the end of 2017, private-sector regular pay growth was around 2.6 percent, a shortfall of around 1.5 percentage points below its pre-crisis average, despite a return to low rates of unemployment.

Why did unemployment fall so rapidly without a corresponding rise in wage growth? There are two broad candidate explanations. First, it could be that declining labor market frictions meant unemployment could fall without upward pressure on wages. In other words, the natural rate of unemployment, having risen, could have taken a downward path in tandem with measured unemployment. The result would have been a lingering output gap and a lack of inflationary pressure. Alternatively, it could be that the fall in unemployment actually represented a strong recovery in aggregate demand and a closing output gap. In this case, weak wage growth would have to be explained by some other factor. A key candidate would be weak productivity. Reduced-form explanations, like simple Okun's law or wage Phillips curve relationships, are not much help in distinguishing these deeper causes. Yet shedding light on this issue has important implications, especially for monetary policy. If weak wage growth is thought to reflect poor productivity growth, this would have consequences that are less disinflationary, at the margin, than the case in which weak wage growth reflects large declines in the natural rate of unemployment, which would imply a continued margin of excess capacity in the economy, all else equal.

To investigate, I construct a New Keynesian model featuring unemployment, as in Galí (2011), and estimate it on six U.K. macroeconomic time series using data from the inflation-targeting era. In the baseline model, six structural shocks determine the dynamics of output, inflation, the policy rate, employment, unemployment, and wage growth. They suggest that the primary cause of the rapid fall in unemployment from 2013 was a recovery in demand—or, rather, a fading of past demand headwinds. This was accompanied by a more modest, though still notable, contribution from falling labor market frictions, and a falling natural rate of unemployment in particular. At the same time, wage growth remained weak primarily because of weak productivity, which also dragged on GDP. This squares the circle between rapidly falling unemployment given GDP growth on the one hand, and subdued wage growth given rapidly falling unemployment on the other.

Conditioning on the model's structural shocks reveals the shape of the United Kingdom's wage Phillips curve. The model illustrates how the search for such a relationship can be confounded when there are significant supply-side disturbances, particularly to price markups and to the natural rate of unemployment. So although the failure of wage growth to accelerate as unemployment fell quickly after 2013 seems, on the face of it, inconsistent with the wage Phillips curve, a sufficiently structural perspective suggests the relationship is, in fact, alive and well. This is worth emphasizing in the context of a search, post-crisis, for a stable Phillips-curve relationship in the face of what the International Monetary Fund (2013) described as the "missing disinflation."

We examine the robustness of the explanations for weak wage growth suggested by the baseline model to a series of alternative modeling assumptions. The first allows for a more flexible stochastic structure in the baseline model's wage Phillips curve to capture features of the wage data such as compositional effects. The second enriches the baseline model with financial frictions—a banking sector in particular. The third incorporates search-and-matching frictions into the model economy in place of the monopolistically competitive labor market embedded in the baseline model. Naturally, these variants offer different interpretations of the data. However, the model variant that best explains the data, as judged by its log data density, offers support for the interpretation of wage growth produced by the baseline model. Interestingly, the baseline unemployment-augmented New Keynesian model appears to fit the data more satisfactorily than the variants of the model that incorporate financial or search-and-matching frictions.

There has been a resurgence of interest in unemployment within the New Keynesian literature since the crisis (e.g., Gertler, Sala, and Trigari 2008; Blanchard and Galí 2010; and Galí, Smets, and Wouters 2012a, 2012b), building on advances made elsewhere to better understand labor market dynamics at the micro level (e.g., Mortensen and Pissarides 1994) and work to reconcile these approaches with the macro literature (e.g., Shimer 2005). Likewise, for obvious reasons, labor market dynamics remain center stage for monetary policymakers (e.g., Broadbent 2014; Carney 2014; Draghi 2014; and Yellen 2014). The present work is most closely related to Galí, Smets, and Wouters (2012a, 2012b), whose New Keynesian framework I employ, but also to the small literature which addresses formal estimation of New Keynesian models for the United Kingdom (e.g., Harrison and Oomen 2010; Burgess et al. 2013; and Facini, Millard, and Zanetti 2013).

Galí, Smets, and Wouters (2012a) is similar to the present study in that it is concerned with the structural drivers of macro and labor market dynamics in the wake of the Great Recession. Their focus, however, is on the United States, and with the earlier phase of the Great Recession experience, covering data only up to mid-2011. They find that the primary driver of the U.S. slow recovery was the absence of favorable shocks, possibly related to the zero lower bound, that were characteristic of more rapid recoveries following past (pre-1990) recessions. In contrast, I consider the United Kingdom's recovery phase, which began in earnest around 2013, following which I find that strong demand growth led the recovery in unemployment.

This work is also related to a small number of other DSGE studies of U.K. labor market dynamics during the Great Recession, though most are concerned with the onset of the crisis. Millard (2015) investigates the causes of the United Kingdom's recession, estimating a Smets-Wouters type model, extended to include open economy aspects and labor matching frictions, finding a key role for "demand" type shocks in explaining GDP in the initial phases of the crisis. The set of observables used in estimation differs from the present paper, which estimates a more compact model, and the sample ends at the end of 2013:Q4, before emergence of the "puzzles" with which I am concerned here. Also taking a search-and-matching DSGE approach, Faccini, Millard, and Zanetti (2013) use U.K. data to evaluate the role played by wage rigidities in U.K. inflation, finding only small effects. Pinter (2015) focuses on collateral channels, showing that the rise in U.K. unemployment was closely related to the collapse in house prices at the onset of the recession, via a collateral channel. Unlike these papers, the present study uses Galí's reinterpretation of the New Keynesian sticky wage model, à la Erceg, Henderson, and Levin's (2000), to relate macro and labor market dynamics, although as a robustness check we also consider a search-and-matching formulation. Recent non-DSGE studies of the United Kingdom's labor market during the crisis include Speigner (2014), who uses U.K. data to investigate the influence of long-term unemployment on U.K. wage growth, and Elsby and Smith (2010), who study the onset of the crisis in the United Kingdom, examining the opposite puzzle—the relatively small (given the size of the recession, and compared with the United States) rise in unemployment. They pointed to the prospective importance of a timely recovery in aggregate demand in averting persistently high unemployment. In relating weak wage growth to productivity, the present paper is also related to various studies of the United Kingdom's productivity puzzle (for which see, inter alia, Barnett et al. 2014).

The remainder of this paper proceeds as follows. Section 2 provides some motivating evidence. Section 3 outlines the baseline DSGE model. Section 4 describes estimation and the model's properties. Section 5 examines the labor market puzzles. Section 6 examines the robustness of the baseline model's conclusions to alternative modeling assumptions, including by allowing for financial and search-and-matching frictions. Section 7 concludes.

2. Motivating Evidence

Figure 1 shows the paths for private-sector regular pay growth and unemployment in the United Kingdom since the advent of inflation targeting, including during the United Kingdom's Great Recession. From around the middle of 2013, unemployment fell rapidly, from 7.8 percent (2013:Q2) to 6.7 percent a year later, falling to around its pre-crisis average by the end of 2015, and to historical lows by the end of 2017. At the same time, wage growth remained weak by past standards. Annual private-sector regular pay growth rose from around 1.0 percent in 2013 to 2.9 percent in 2015, before declining to 2.7 percent in 2016 and 2.3 percent in 2017. At the end of 2017, private-sector regular pay growth was around 2.6 percent, a shortfall of around 1.5 percentage points below its pre-crisis average, despite a return of low rates of unemployment.

A simple window on these puzzles can be provided by "Okun's law"—the reduced-form relationship between GDP growth and unemployment—and a simple wage Phillips curve—the reduced-form relationship between wage growth and unemployment. Figure 2 shows the fitted values from a simple Okun's law relationship. From 2013, the Okun-implied path for unemployment and the data diverge notably, to the extent that by the middle of 2014, unemployment had fallen over 1 percentage point more rapidly than expected given GDP growth outturns. Figure 3 shows the fitted values from a similarly simple reduced-form wage Phillips curve, derived from a regression of annual private-sector regular pay growth on unemployment. Just as unemployment was falling rapidly, wage growth was undershooting what a simple wage Phillips curve would have produced (figure 4). By the middle of 2014, wage growth was between 1.5 and 2 percentage points shy of its fitted values.

These simple reduced-form approaches highlight the basic labor market "puzzles," but they leave much unexplained, and they offer little by way of a structural explanation. For example, the concurrence of both negative Okun's law residuals and negative wage Phillips curve residuals, as shown in figure 5, is consistent with a

Figure 2. Unemployment Fell More Rapidly from 2013 Onwards Than a Simple Okun's Law Relationship Would Have Predicted

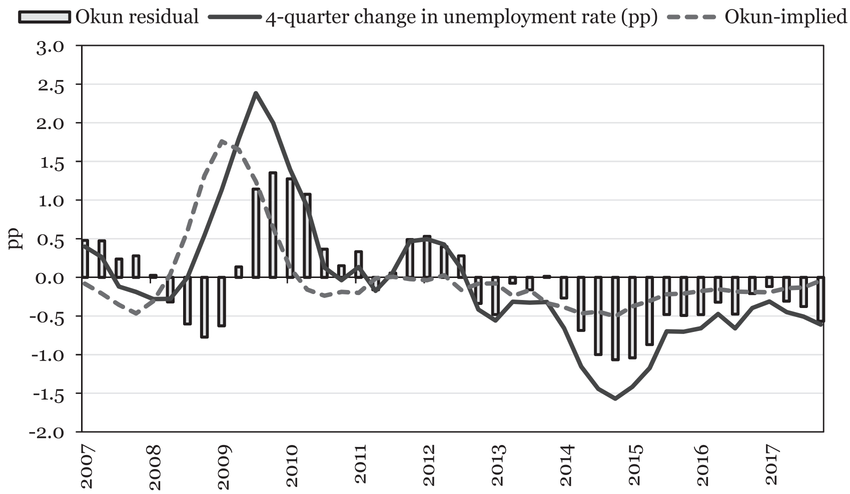

Notes: "Okun-implied" is based on the fitted values from a regression of the four-quarter change in unemployment on four-quarter GDP growth, 1993:Q4–2017:Q4. The resulting simple model suggests $u_t^{\Delta 4} = 0.302 - 0.246 y_t^{\Delta 4}$.

range of structural interpretations. One way to get a more structural perspective is to turn to a VAR model. These offer rich yet parsimonious descriptions of the data. Together with restrictions on the reduced-form variance-covariance matrix of the VAR's residuals, such a model can also help to make structural statements.

The simple VAR model I consider contains GDP, unemployment, and wage growth. This is the smallest possible model capable of distinguishing three key shocks of interest for the joint dynamics of wages and unemployment: productivity, demand, and labor frictions, or "u*" (natural rate of unemployment) shocks. The VAR is estimated using Bayesian methods, contains two lags, and is estimated over the inflation-targeting era in the United Kingdom.[1]

[1]To be consistent with the DSGE model I later estimate, I pre-filter the data in the following way: GDP enters as a quarterly log difference, unemployment as a level, and the wage index as a quarterly log difference. The priors for the BVAR are implemented as described in Banbura, Giannone, and Reichlin (2010).

Figure 3. Despite Rapidly Falling Unemployment, Wage Growth Was Surprisingly Weak

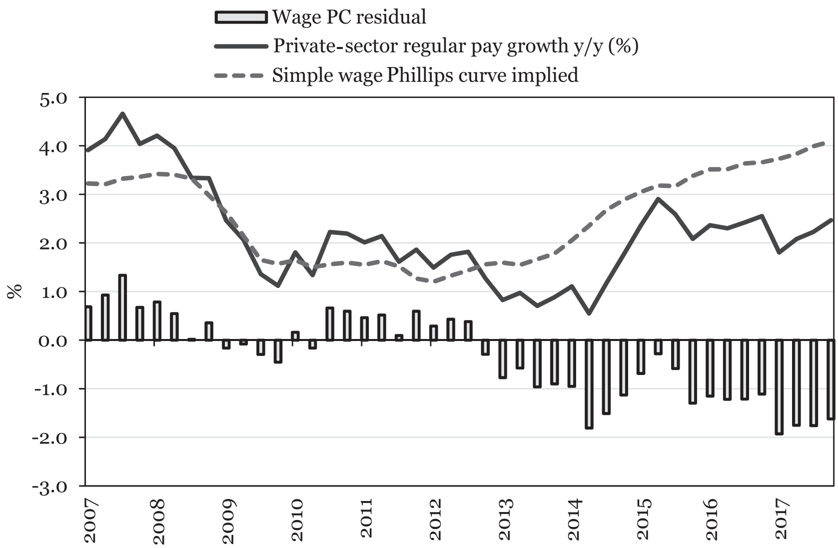

Notes: Based on a simple reduced-form relationship between annual private-sector regular pay growth and unemployment, 2001:Q1–2015:Q4. The simple regression suggests $\pi_t^{wann} = 7.05 - 0.70u_t$.

Figure 4. Significant Departure of Wage Growth from Levels Implied by a Simple Wage Phillips Curve

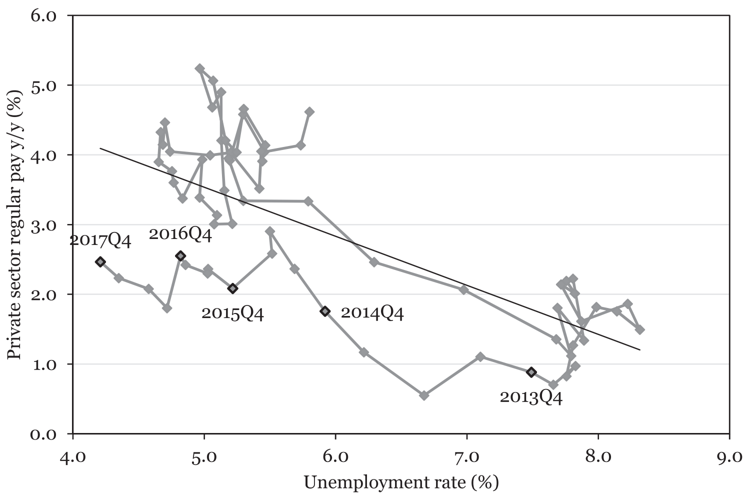

Figure 5. Okun's Law and Wage Phillips Curve Residuals Were Positively Correlated and Negative Over 2013–07

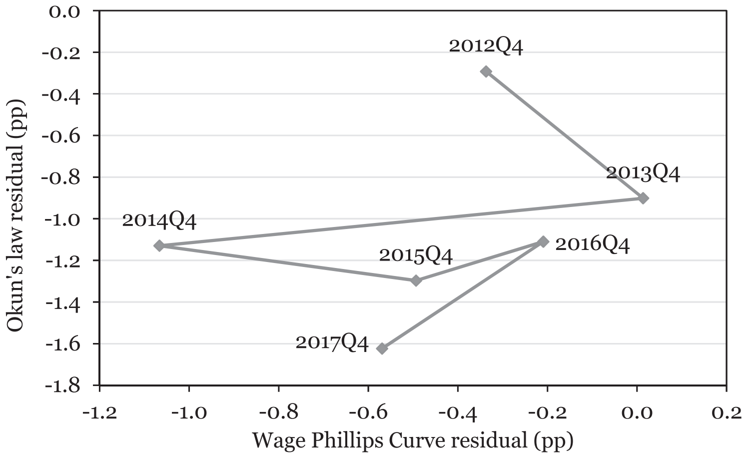

Notes: Residuals from simple wage Phillips curve and Okun's law relationships are described in figures 2 and 3.

Table 1. Sign Restrictions for the Three-Variable VAR

Shock	Variable		
	y	u	π^w
Demand	+	−	+
Labor Frictions (u*)	+	−	−
Productivity	+	+	+

Notes: "+" indicates a positive response of the variable in the column header and "–" a negative response. The sign restrictions are imposed on impact.

I identify the three structural shocks by using sign restrictions. To recover these shocks from the reduced-form covariance matrix, I look for rotations whose impulse responses satisfy the sign restrictions described in table 1. The restrictions are imposed for a single period (i.e., on impact).

Some discussion of these restrictions is in order. First, demand shocks raise output and lower unemployment, while raising wage growth. This is intuitive and is consistent with the DSGE model

I later consider. Second, shocks to labor frictions, which change the natural rate of unemployment ("u*"), raise output and lower unemployment but also lower wage growth. This pattern of comovement distinguishes these shocks from demand shocks, because they have opposite implications for short-term wage growth. Moreover, these restrictions uniquely identify these shocks within the set of impulse responses produced by the DSGE model that I later consider (i.e., no other shocks in the model produce this pattern of comovement). Finally, technology (productivity) shocks raise output and wage growth but also raise unemployment, as fewer labor inputs are needed per unit of output. This is reminiscent of Galí (1999), who finds that positive technology shocks typically lower hours across G-7 economies, and is also consistent with the DSGE model estimated below in which productivity shocks are the *only* shocks for which output, unemployment, and wage growth co-move positively.[2]

The impulse responses satisfying these restrictions are shown in figure 6. Shocks to the natural rate are estimated to have persistent effects on unemployment and output growth, but relatively transient effects on wage inflation. Productivity shocks have more persistent effects on wage growth, while the impact on unemployment quickly becomes insignificant. Finally, the demand shock has significant and persistent effects on all three variables.

Using these estimated shocks, figure 7 provides a VAR-based narrative for the fall in unemployment and subdued wage growth from 2013. The VAR explains around 2 percentage points of the 3.2 percentage point decline in unemployment by the end of 2017 with a fall in the natural rate of unemployment. The remainder of the decline in unemployment is explained largely by stronger demand, which explains around 1 percentage point of the decline. What of the explanation for weak wage growth? Weak demand and lower u* exert significant drags on wage growth between 2013 and 2015, while weak productivity plays an increasingly significant role from 2014 onwards, accounting for essentially all of the shortfall of wage growth below trend by the end of 2017.

The VAR can be used to construct a simple measure of the natural rate of unemployment, which corresponds to the estimated

[2]Labor supply shocks also generate positive co-movement between unemployment and output, but reduce wage growth, unlike productivity shocks.

Figure 6. Impulse Responses from Three-Variable VAR

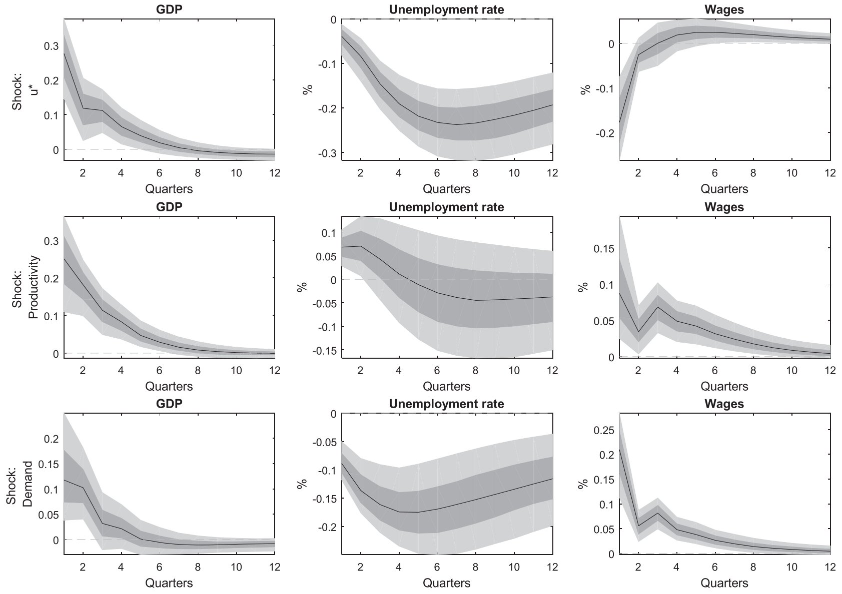

Notes: Rows contain structural shocks and columns contain variables. "u*" = natural rate of unemployment shock; "Productivity" = productivity shock; "Demand" = demand shock. Shaded bands show the {16th, 84th} and {32nd, 68th} percentiles of the set of identified impulse responses.

Figure 7. Three-Variable VAR Explanation of Unemployment and Wage Growth Since 2013:Q2

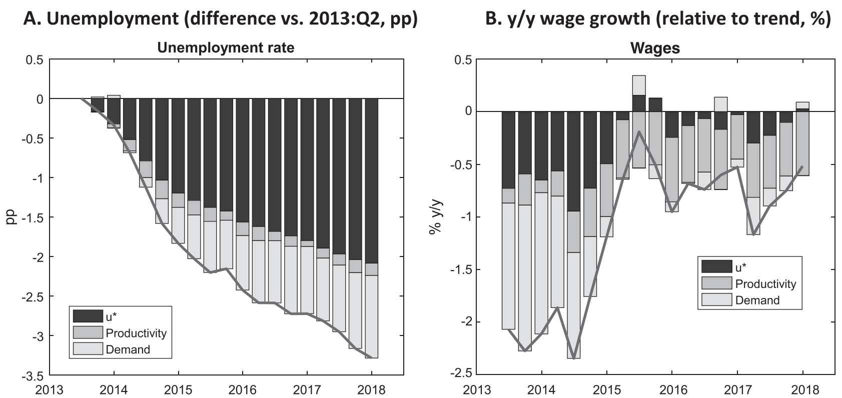

Figure 8. VAR-Based Estimate of the Natural Rate of Unemployment

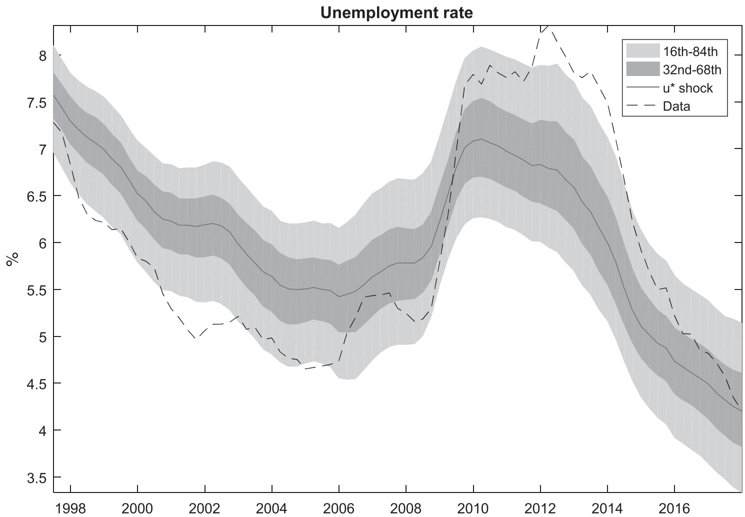

Notes: Dashed line is data; solid line is the median estimate of the contribution of identified u* shocks to the unemployment rate; shaded regions indicate the {16th, 84th} and {32nd, 68th} credible set of contributions of the u* shock to unemployment.

contribution of labor frictions (u*) shocks to the observed unemployment rate over the sample period. Figure 8 shows this implied path, together with the credible set of contributions generated across accepted rotations of the VAR's variance-covariance matrix within the 16th and 84th percentiles. It suggests that the start of the crisis coincided with a worsening of labor market frictions, causing a rise in the natural rate of unemployment of around 1 percentage point. After the initial onset of the crisis, however, the natural rate is estimated to have fallen gradually from around 7 percent in 2009 to around 6.5 percent at the end of 2012, before falling rapidly from 2013 to 2014, eventually reading around 4.2 percent at the end of 2017.

3. A Benchmark Model

The VAR evidence is suggestive but can be given a richer structural underpinning by using a model with theoretically coherent foundations. This is helped by recent advances in New Keynesian macroeconomics that have helped embed unemployment in otherwise standard, and empirically successful, New Keynesian models. In this section I begin by setting out the key elements of such a model. I use a New Keynesian framework, similar to Smets and Wouters (2007), extended to include unemployment as in Galí (2011) and Galí, Smets, and Wouters (2012b). The model features staggered nominal wage setting, as in Erceg, Henderson, and Levin (2000). The key innovation relative to this approach, as introduced by Galí (2011), is to offer a reinterpretation of the wage markup that results from imperfect competition in the labor market and, in particular, from the "differentiated varieties" assumption applied to types of labor. In the present setup, this markup is reinterpreted as the natural rate of unemployment, and it is an unemployment gap—a gap between the unemployment rate and its natural rate—that enters the wage Phillips curve. This means unemployment can be included as an additional observable in the model and, importantly for our purposes, (inefficient) shocks to the natural rate of unemployment can be distinguished from other (efficient) shocks to labor supply via the different implications these shocks have for the dynamics of unemployment. The model set out in the next few subsections is the "benchmark" model.[3] We examine the robustness of the conclusions drawn from this benchmark model in subsequent sections, where we add financial frictions and a richer description of labor market frictions in turn.

3.1 The Model's Linearized Conditions

This section briefly summarizes the model's linearized conditions, where lowercase letters denote log-deviations from steady state, beginning with the demand side of the economy.

3.1.1 Demand

Aggregate demand y_t comprises consumption c_t and investment in_t:

$$y_t = \frac{C}{Y} c_t + \left(1 - \frac{C}{Y}\right) in_t, \tag{1}$$

[3]A full derivation, which is standard, appears in the appendix.

where C/Y is the steady-state share of consumption in income. Consumption dynamics are governed by a consumption Euler equation; the resulting "dynamic IS curve" relates demand to its past and future values, and negatively to the ex ante real interest rate:

$$c_t = \frac{h}{1+h}c_{t-1} + \frac{1}{1+h}E_t c_{t+1} - \frac{1-h}{1+h}\left(i_t - E_t\pi_{t+1}\right) + \varepsilon_t^c, \tag{2}$$

where i_t is the nominal interest rate, π_t is goods price inflation, and ε_t^c is a shock to households' consumption demand.[4] Parameter h governs the strength of consumption habits.

Because of "flow" adjustment costs, the investment Euler equation is also inertial and reads

$$in_t = \frac{1}{1+\beta}in_{t-1} + \left(1 - \frac{1}{1+\beta}\right)E_t in_{t+1} + \frac{1}{1+\beta}\omega q_t^k, \tag{3}$$

where q_t^k is the real price of capital goods (in terms of final goods) and ω parameterizes the adjustment costs in investment.

3.1.2 Supply

On the supply side, firms' production technology takes the form of the production function:

$$y_t = (1-\alpha)\,n_t + \alpha k_{t-1} + \varepsilon_t^a, \tag{4}$$

where α is the capital share of income, n_t is employment, k_t is physical capital, and ε_t^a is the disturbance to total factor productivity. Capital evolves according to

$$k_t = (1-\delta)\,k_{t-1} + \delta in_t, \tag{5}$$

where δ governs the capital depreciation rate, and firms demand labor and capital such that

$$m_t = w_t - (y_t - n_t)\,, \tag{6}$$

$$m_t = z_t^k - (y_t - k_{t-1})\,, \tag{7}$$

[4] Formally, ε_t^c is a shock to households' intertemporal preferences.

where m_t is real marginal cost, w_t is the real wage, and z_t^k is the real marginal product of capital. As noted above, because goods prices are reset only periodically, with the remaining set of prices indexed to inflation, inflation is governed by the price Phillips curve:

$$\pi_t = \gamma_p \pi_{t-1} + \beta E_t \left(\pi_{t+1} - \gamma_p \pi_t\right) + \lambda_p m_t + \varepsilon_t^p, \qquad (8)$$

where ε_t^p is the shock to firms' natural price markups, and where γ_p governs the extent of inflation inertia due to indexation. The parameter λ_p governs the strength of the relationship between real marginal cost and goods price inflation, and reflects the underlying stickiness of goods prices.

3.1.3 Labor Market

Households' labor supply schedule is

$$w_t = s_t + \varphi l_t + \varepsilon_t^l, \qquad (9)$$

where φ is the inverse Frisch elasticity of labor supply, l_t is desired employment, and ε_t^l is the labor supply shock. The variable s_t capturing the size of the wealth effect on labor supply is given by

$$s_t = (1-\nu)\, s_{t-1} + \nu \left(\frac{1}{1-h} c_t - \frac{h}{1-h} c_{t-1}\right), \qquad (10)$$

where the final term is the marginal utility of consumption (in the presence of habits and under log utility), and parameter ν governs the strength of the wealth effect on labor supply.

The unemployment rate is given by

$$u_t = l_t - n_t, \qquad (11)$$

which is higher the greater is the wage markup in the labor market. Wage inflation is determined by the wage Phillips curve:

$$\pi_{w,t} = \gamma_w \pi_{w,t-1} + \beta E_t \left(\pi_{w,t+1} - \gamma_w \pi_{w,t}\right) + \lambda_w \left(u_t - u_t^*\right), \qquad (12)$$

where π_w is nominal wage inflation, γ_w governs the degree of wage indexation, λ_w is the relationship between unemployment and wage inflation, and u_t^* is the natural rate of unemployment. Another way to view $u_t - u_t^*$ is as the *unemployment gap*; when it is positive, there is upward pressure on wage growth, and when it is negative, there is downward pressure.

3.1.4 Monetary Policy, Returns, and Shocks

Finally, monetary policy is described by a Taylor rule:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i)\left(\phi_\pi \pi_t + \phi_y \left(y_t - y_t^*\right)\right) + \varepsilon_t^i, \qquad (13)$$

in which y_t^* is the natural level of output and where ε_t^i is a monetary policy shock. Parameter ρ_i governs inertia in the monetary policy rule, and $\{\phi_\pi, \phi_y\}$ the response of the nominal interest rate to inflation and the output gap, respectively. The model is closed with an arbitrage relationship between capital and bonds, $E_t r_{t+1}^k = r_t = i_t - E_t \pi_{t+1}$, where r_t^k is the return on capital given by

$$r_t^k - q_{t-1}^k = (1 - \beta (1 - \delta))\, z_t^k + \beta (1 - \delta)\, q_t^k. \qquad (14)$$

In summary, the model features six shocks: to consumption demand, monetary policy, technology, labor supply, the natural rate of unemployment, and the price markup. These are all AR(1) processes which follow:

$$\varepsilon_t^j = \rho_j \varepsilon_{t-1}^j + u_t^j, \quad u_t^j \sim N\left(0, \sigma_j^2\right), \quad j = \{c, i, a, l, u, mp\}, \qquad (15)$$

(where the persistence of the monetary policy shock is restricted to be zero).

4. Estimating the Benchmark Model

Six aggregate time series are used in estimation: GDP, core CPI inflation,[5] the quarterly policy rate adjusted for the effects of the Bank of England's asset purchases ("Bank Rate"), nominal wage inflation, employment, and unemployment. To correct for the presence of the zero bound, I use a measure of the policy rate adjusted

[5]This is CPI inflation excluding food, energy, alcohol, and tobacco. I use core inflation in order to focus on the structural links between the labor market and inflationary pressure, excluding those due to volatile items like energy and food. Equally, I use a CPI-based measure as the measure of inflation because the Bank of England's inflation objective is cast in terms of the CPI, not, say, the GDP or other expenditure component deflator. Besides this, there are also well-known measurement issues around GDP deflator inflation, including the challenges faced in measuring the relative price of public-sector output. For these reasons I prefer to use the CPI as a measure of prices.

for the impact of the MPC's asset purchases—a so-called shadow rate, in the spirit of Wu and Xia (2016). This uses methods to identify the macroeconomic impact of the MPC's asset purchase program, and converts these into a "shadow" path for the policy rate with equivalent macroeconomic effects (see Joyce, Tong, and Woods 2011). The sample begins with the advent of inflation targeting in the United Kingdom and runs from 1993:Q1 to 2017:Q4. Measurement equations are defined relating the model variables to their empirical counterparts; GDP, employment, core inflation, and wage inflation enter as demeaned quarterly log-differences, while Bank Rate and unemployment enter in levels deviations from steady state, taken to be 5 percent for each, around their sample averages. Five MCMC chains of 100,000 draws each are simulated to explore the model's posterior distribution.

Six of the model's parameters are fixed prior to estimation. These are the discount factor, β, which is set to give an annualized real interest rate of 3 percent, around the pre-crisis average rate; the capital share, α, set to 1/3; the price elasticity, set to 5, implying a markup of 1.25; the wage markup, set to imply a steady-state unemployment rate of 5 percent, around the average in the United Kingdom over the period of central bank independence; the share of consumption in output, which, given the simple demand side modeled here, is 0.75, per the U.K. data; and the depreciation rate on capital, δ, set at 2.5 percent.

The remaining parameters are estimated. Table 2 contains a detailed description of the priors used to do this. The prior for habit formation in consumption is centered on 1/2. Some studies for the United Kingdom have estimated relatively low habit parameters on U.K. data (e.g., Harrison and Oomen 2010; Burgess et al. 2013) compared with similar models fitted to euro-area and U.S. data (Smets and Wouters 2003, 2007); however, the present model has a simpler demand side than these papers, giving little guide for the appropriate choice in this context. In that light, 1/2 seems reasonable. The labor supply elasticity is estimated to be relatively low in the United Kingdom by Harrison and Oomen (2010) (an inverse Frisch elasticity of 1/0.142 at the mean), but high by Faccini, Millard, and Zanetti (2013) (an inverse Frisch elasticity of 1.64 at the mean), and closer to 2 by Burgess et al. (2013). Once again, the propagation mechanisms in these models all differ in important ways from the

Table 2. Priors and Parameter Estimates

	Prior			Posterior		
Parameter	**Dist.[a]**	**Mean**	**St. Dev.**	**Mean**	**Lower**	**Upper**
ρ_a	B	0.500	0.075	0.974	0.965	0.982
ρ_c	B	0.500	0.075	0.711	0.649	0.772
ρ_{mp}	B	0.500	0.075	0.295	0.209	0.377
ρ_l	B	0.500	0.075	0.819	0.746	0.895
ρ_u	B	0.500	0.075	0.600	0.501	0.702
h	B	0.500	0.100	0.849	0.805	0.892
v	B	0.500	0.200	0.028	0.022	0.034
φ	IG	3.000	0.250	3.679	3.182	4.145
ω	IG	1.000	0.200	0.625	0.510	0.738
λ_p	IG	0.050	0.005	0.041	0.036	0.047
λ_w	IG	0.050	0.005	0.054	0.045	0.063
γ_p	B	0.500	0.150	0.183	0.064	0.302
γ_w	B	0.500	0.150	0.121	0.060	0.182
ρ_i	B	0.500	0.150	0.873	0.842	0.905
ϕ_π	IG	0.500	0.100	0.646	0.424	0.863
ϕ_y	IG	0.125	0.100	0.135	0.058	0.209
Shock St. Devs.						
Technology (a)	IG	0.01	0.01	0.006	0.005	0.007
Monetary (i)	IG	0.01	0.01	0.003	0.002	0.003
Demand (c)	IG	0.01	0.01	0.002	0.002	0.003
Markup (mp)	IG	0.01	0.01	0.002	0.002	0.002
Labor Supply (l)	IG	0.01	0.01	0.009	0.008	0.011
Unemployment (u)	IG	0.01	0.01	0.044	0.034	0.054

[a]Prior distribution: B = beta; IG = inverse gamma.

present paper, including the fact that the present model allows for a variable wealth effect on labor supply. Together, these considerations also suggest an intermediate choice of prior, which we set at 3, together with a prior for the wealth effect on labor supply of 1/2, in line with Galí, Smets, and Wouters (2012b). We choose the priors for the wage and price Phillips curves directly and to be consistent with wages and prices resetting on average around once every three to four quarters, given the discount factor, and we center the price and wage indexation parameters at 1/2. The priors for the policy rule are chosen to be standard and are set to give responses to inflation and the output gap of 1.5 and 0.125, respectively. All priors governing the shock persistence are set to 1/2, and their standard deviations are set to 1 percent.

The posterior mean estimates are contained in the final three columns of table 2. Relative to the prior, consumption habits (h) are estimated to be larger, and closer to the estimates obtained for U.S. and euro-area model, than some past estimates for the United Kingdom. The wealth effect on labor supply (ν) is weak, consistent with the estimates for the United States in Galí, Smets, and Wouters (2012b). This means that most of the effects of demand on the labor market come not through changes in households' work incentives (which shift labor supply) but through the effects on labor demand. The inverse Frisch elasticity (φ) is estimated to be around 3.7, indicating a relatively inelastic labor supply response to real consumption wages. The price Phillips curve is estimated to be flatter than the wage Phillips curve, implying, perhaps counterintuitively, goods prices that are stickier than nominal wages. Finally, the monetary policy rule contains responses to inflation and the output gap which are close to the prior values, whereas policy smoothing is much higher (0.87) than under the prior mean.

Figure 9 gives a snapshot of the dynamic properties of the model by showing the estimated impulse response functions. The rows of the figure contain the different structural shocks, indicated by the labels at the start of each row, while the columns contain the estimated responses of the different observable variables. A couple of features are worth noting.

First, the technology shock is the only shock that results in a negative co-movement of output and employment on impact. In this case, a positive technology shock means firms can produce the same

Figure 9. Estimated Impulse Response Functions—Baseline Model

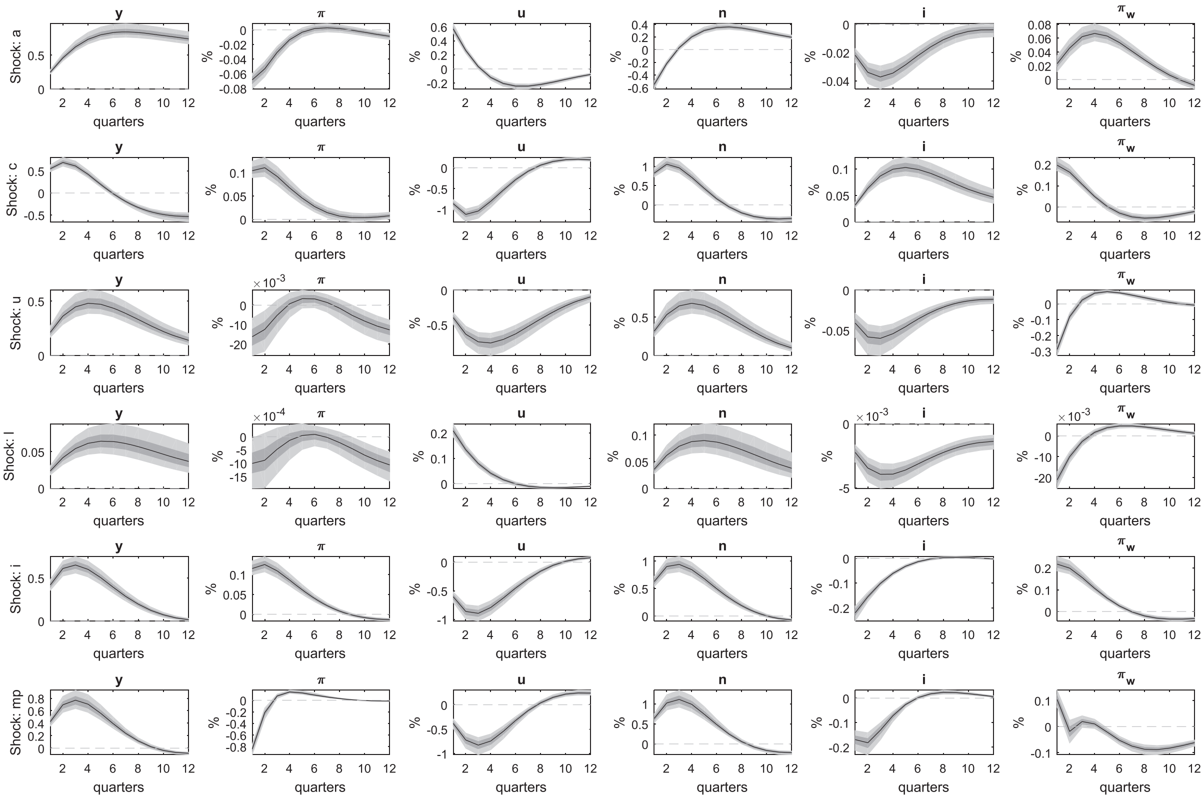

Notes: Rows contain structural shocks; columns contain variables. “a” = technology shock; “c” = demand shock; “i” = monetary policy shock; “l” = labor supply shock; “u” = natural rate of unemployment shock; “mp” = price markup shock.

output with fewer labor inputs. As a result, employment falls. In part, this is due to features of the model that build persistence into demand, including habits in consumption and investment adjustment costs. These features mean that when total factor productivity (TFP) rises, demand does not respond "immediately"; as a consequence, firms have a lesser need for labor. In contrast, for the remaining shocks employment is procyclical. This feature limits the extent to which technology shocks alone can explain the macroeconomic fluctuations in the data, a point emphasized by Galí (1999).

Second, "demand" and monetary policy shocks are distinguished only by the responses of the policy rate. Positive demand shocks entail the policy rate rising to curtail inflationary pressures whereas, clearly, expansionary monetary shocks entail a cut in the policy rate. Third, compare the responses of the labor supply (fourth row) and u shocks (fifth row). Expansions of labor supply and reductions in labor frictions both entail increases in output, falls in inflation, reductions in the policy rate, and rising employment. However, they have opposite implications for unemployment. When labor frictions abate, the expansion that follows entails a reduction in unemployment. When labor supply increases, by contrast, the expansion entails higher unemployment. As a result, the "u*" shock is the main disturbance for which there is a material positive co-movement of wage inflation and unemployment.[6] With the exception of the price markup shock, for which this pattern is also just visible, the remaining disturbances in the economy generate opposing co-movement of wages and unemployment.

5. The United Kingdom's Great Recession and Its Aftermath

5.1 Explaining the Dynamics of the Crisis and the Recovery

This section turns to the model's interpretation of U.K. macroeconomic dynamics during the Great Recession and recovery. Begin with figure 10, which shows the shock-based decomposition of the six observable series used in estimation. It is clear from the top-left panel

[6] This co-movement is also observed for the technology shock, but the response of wage inflation relative to unemployment is much more muted in this case.

Figure 10. The Dynamics of the Great Recession and Its Aftermath

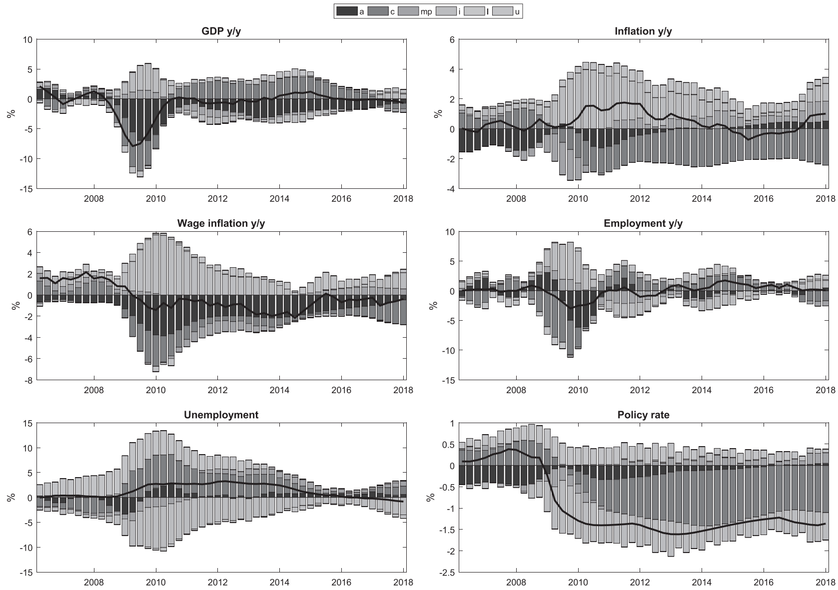

Notes: "a" = technology shock; "c" = demand shock; "i" = monetary policy shock; "l" = labor supply shock; "u" = natural rate of unemployment shock; "mp" = price markup shock. Data are shown relative to their sample means. The policy rate is quarterly.

that the United Kingdom faced severe negative shocks at the onset of the crisis, centered on both weak demand and weak productivity. The effects of the negative demand shock on output and employment were somewhat ameliorated, though not completely offset, by stimulative monetary policy. Indeed, the policy rate is estimated to have fallen faster than would have been expected given the average Taylor rule of the sample. This stimulus is estimated to have been broadly sufficient to offset the impact of negative demand shocks on inflation, which rose not only due to this but also due to a series of price markup shocks from around 2010 onwards.

Now turn to the labor market. As noted, despite significant stimulus, annual employment growth fell to –3 percent (relative to trend) in 2009, and unemployment rose some 3 percentage points above its past average. It is notable in particular that a large proportion of the initial rise in unemployment reflected net demand weakness. It is also notable, though, that the United Kingdom appears to have been operating at below its natural rate of unemployment for some time prior to the crisis. In fact, the flexible-price output gap is estimated to have been positive and between +3 and +5 percent in the lead-up to the crisis. As the recovery became embedded, the natural rate of unemployment is estimated to have fallen back.

In that context, the model's explanation for the fall in wage growth is quite striking. In particular, annual wage growth is almost completely explained by the contribution of productivity shocks following the crisis episode and into the recovery. Alongside this effect, demand weakness and monetary policy largely offset one another and, curiously perhaps, wage growth was supported by labor market friction, or u*, shocks. It is notable that over the period during which the effects of these shocks began to wane for unemployment, they were also pushing *up* on wage growth.

What lies behind this finding? Part of the explanation lies in the dynamics of u* shocks themselves; aside from the price markup shock, the wage markup or u* shock is the least persistent of the model's remaining macroeconomic disturbances. The dampening effect on wage growth of a given fall in unemployment due to lower u* is, therefore, relatively short-lived. The other part of the explanation relates to the impact of declines in u* on firms' incentives to invest. In particular, once the initial negative effect on wage growth fades, a countervailing positive effect on wage growth arises via the positive impact that lower labor market frictions have on aggregate demand, investment, and hence the real wage. Consider the following: as u* falls, firms take on more labor, raising the marginal product of capital and so the incentive to invest; this, eventually, boosts the real wage, a dynamic illustrated by the impulse responses shown in figure 9. This makes it relatively challenging for lower u* shocks to generate a sustained decline in wage growth at the baseline model's estimated posterior mode.

Turn next to the "puzzle" period—from around 2013. Here, unemployment fell around 2.5 percentage points, with growth only modestly above trend, but wage growth did not accelerate very markedly at all. Figure 11 shows the explanation for the fall in unemployment and the absence of a wage pickup. The model suggests

Figure 11. The Fall in Unemployment from 2013:Q2 Largely Reflected a Recovery in Demand While Weak Wage Growth Reflected Weak Productivity

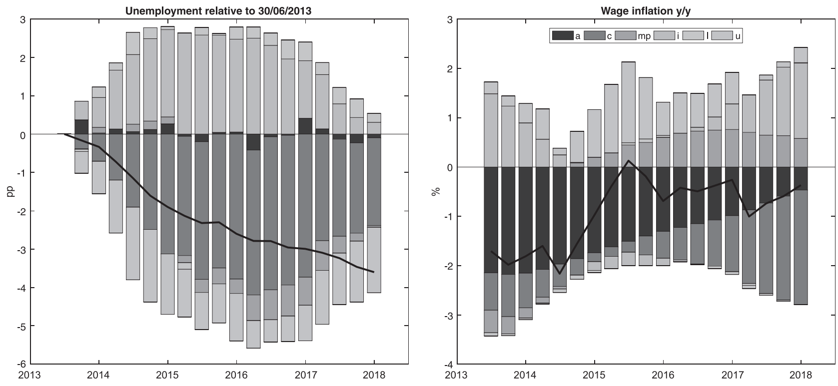

Notes: "a" = technology shock; "c" = demand shock; "i" = monetary policy shock; "l" = labor supply shock; "u" = natural rate of unemployment shock; "mp" = price markup shock.

unemployment fell largely because demand headwinds fell back, net of the effect of fading monetary policy stimulus. By the end of 2015, this accounted for the bulk of the fall, with some additional contribution from lower labor frictions. The prominent role for positive demand developments is consistent with the cyclical recovery in job vacancies observed in other data over this period, which are strongly procyclical in general. At the same time, although demand headwinds were fading, headline growth was held back by a continued drag from weak productivity, which meant unemployment could fall rapidly despite the prevailing GDP growth rate. These estimates contributions from weak productivity are corroborated by other estimates. For example, the year-over-year growth rate of the model's fitted TFP shock series is 69 percent correlated with top-down estimates of TFP produced by the United Kingdom's Office for National Statistics.

And what of wage growth? This continued to be largely explained by weak productivity, although the (seemingly temporary) acceleration between mid-2014 and mid-2015, and its fall back thereafter, is explained by a rise and then a fall in the contribution from labor frictions—the u* shock. Over the broad sweep of the recovery, then, weak productivity appears to have dragged on wage growth, with demand and monetary policy effects broadly netting off, and with volatility being induced via variation in labor frictions.

One important point to emphasize is that the model allows for a stationary TFP process. An alternative is to allow for non-stationary technology, and to investigate the effects of shocks to the *growth rate* of technology on the dynamics of wages, unemployment, and other macroeconomic variables. We experimented with this case on a version of the model calibrated at the mean of the posterior parameter values estimated above, comparing the responses of the model to the estimated (stationary) TFP shock with an alternative, temporary shock to the growth rate of TFP. While the precise quantitative magnitudes of the responses differ under these two cases, qualitatively we found strong similarities in the short-run responses of the output gap, inflation, policy rate, employment, unemployment, and wage growth to persistent stationary TFP shocks on the one hand and temporary TFP growth rate shocks on the other. In particular, in both cases, negative technology shocks generate reductions in unemployment and persistently weak nominal wage inflation. The qualitative similarity of these two shocks appears in part due to the presence of realistic model features, including habit persistence and investment adjustment costs, which add inertia to demand in the short run, as discussed above.[7]

5.2 Whither the Wage Phillips Curve?

Does the absence of a clear co-movement between wage growth and unemployment over this period mean the wage Phillips curve is no

[7]Notwithstanding this point, the longer-run dynamics of employment and wage inflation are likely to be quite sensitive to the persistence of the process governing technology. In the stationary technology case, following a negative shock, wage inflation eventually overshoots its equilibrium level as the productivity level eventually recovers to its baseline value, and the real wage "gap" resulting from weak nominal wage growth in the short term is closed. In contrast, if the productivity process does not result in a recovery in the level of TFP, and hence the real wage, this nominal overshooting need not occur. If anything, then, estimating a model with stationary as opposed to non-stationary TFP may therefore understate the extent to which weak TFP dragged on wage growth.

Figure 12. Conditional U.K. Wage Phillips Curve Is Alive and Well

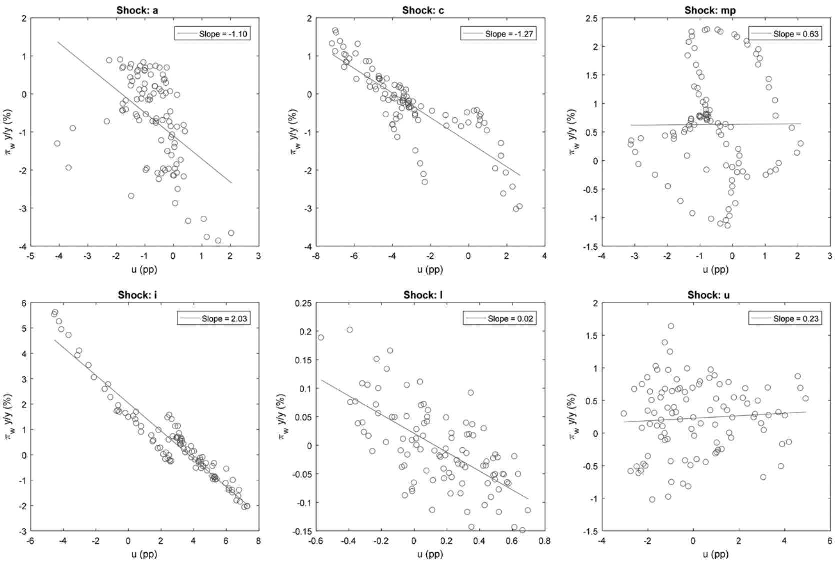

Notes: "a" = technology shock; "c" = demand shock; "i" = monetary policy shock; "l" = labor supply shock; "u" = natural rate of unemployment shock; "mp" = price markup shock.

more? To the contrary. We can use the model to describe the *structural* wage Phillips curve in the United Kingdom—which is not so much a single relationship as a set of conditional relationships which vary with the underlying structural disturbance.

Figure 12 shows the conditional relationship between wage inflation and unemployment for different structural disturbances. Here, each panel plots a given shock's contribution to wage inflation and unemployment, and the axis scales are common to each shock to give an indication of relative magnitudes. For example, demand shocks (middle panel; top row) trace a clear downward relation between wage inflation and unemployment; so do monetary policy shocks. The covariation of wages and unemployment is especially small for labor supply shocks. And finally, shocks to u* and to price markups have traced broadly flat, even upward-sloping, relationships between wage inflation and unemployment, and are both responsible for a reasonable degree of variation in the unemployment rate. These shocks would otherwise confound the search for a stable wage Phillips curve. Figure 12 shows that, conditional on the right shocks, the Phillips curve is alive and well.

6. Robustness

6.1 Specification of the Wage Phillips Curve

In this section, we investigate the robustness of the baseline results to different specifications of the wage Phillips curve. One concern with the baseline specification of the wage Phillips curve is that the shock term, which in theory captures variation in the natural rate of unemployment (proportional to the natural wage markup), must capture lots of features of the wage data, including the possibility of transitory factors affecting wage dynamics, such as residual seasonal patterns, or other high-frequency variation in wage growth due, for example, to sample variability.[8] Arguably, these factors could swamp low-frequency variation in wages due to change in the gap between unemployment and its natural counterpart. To allow for this possibility, we estimate a variant of the baseline model in which the wage Phillips curve (12) is modified to be

$$\pi_{w,t} = \gamma_w \pi_{w,t-1} + \beta E_t \left(\pi_{w,t+1} - \gamma_w \pi_{w,t}\right) + \lambda_w \left(u_t - u_t^*\right) + u_t^w, \tag{16}$$

where $u_t^w \sim \mathcal{N}(0, \sigma_w^2)$ is a purely transitory wage shock capturing the features of the data described above.

The resulting parameter estimates are shown in table 3, column 2. The mean of the posterior estimates for most parameters remain broadly similar, with the possible exception of the coefficient on the output gap in the monetary policy rule, which rises from 0.135 in the baseline model to 0.587 in the modified model. Likewise, the shock processes are estimated to have properties as in the baseline model, with the obvious exception of the standard deviation of the

[8]This could include compositional effects—changes in the composition of characteristics of the underlying workforce—which can have consequences for measured wage changes. In the U.K. context, see, inter alia, Broadbent (2015).

Table 3. Posterior Mean Parameter Estimates for Baseline and Variants of the Baseline Model

Parameter	(1)	(2)	(3)	(4)
ρ_a	0.974	0.979	0.960	0.776
ρ_c	0.711	0.737	0.706	0.684
ρ_{mp}	0.295	0.325	0.281	0.515
ρ_l	0.819	0.832	0.862	0.498
ρ_u	0.600	0.543	0.722	0.965
h	0.849	0.808	0.844	0.901
v	0.028	0.027	0.032	N/A
φ	3.679	3.933	3.751	3.335
ω	0.625	0.667	0.518	0.702
λ_p	0.041	0.039	0.040	0.039
λ_w	0.054	0.046	0.066	0.012
γ_p	0.183	0.169	0.129	0.564
γ_w	0.121	0.158	0.119	0.272
ρ_i	0.873	0.897	0.901	0.905
ϕ_π	0.646	0.602	0.573	0.526
ϕ_y	0.135	0.587	0.183	0.087
σ_b	N/A	N/A	0.795	N/A
N_b/S_b	N/A	N/A	0.036	N/A
Shock St. Devs.				
Technology (a)	0.006	0.006	0.006	0.041
Monetary (i)	0.003	0.002	0.003	0.002
Demand (c)	0.002	0.002	0.002	0.002
Markup (mp)	0.002	0.002	0.002	0.003
Labor Supply (l)	0.009	0.010	0.010	0.013
Unemployment (u)	0.044	0.011	0.032	0.023
Wage (w)	N/A	0.005	N/A	N/A
Marginal Density	2,147.4	2,151.4	2,116.0	1,769.9

Note: Column numbers refer to the following models: (1) = baseline; (2) = (1) + transitory shock in wage Phillips curve; (3) = financial frictions model; (4) = search-and-matching frictions model.

shock driving the natural rate of unemployment, which drops by 75 percent. This is fully expected given the change in the stochastic specification of the wage Phillips curve. Overall, the model fit as judged by the model's marginal density improves slightly, with the

Bayes factor being exp(4), indicating positive evidence in favor of the modified model.[9]

What does this mean for the dynamics driving wage growth? Figure 13, panel A shows the baseline wage decomposition, while panel B shows the decomposition resulting from the modified wage Phillips curve. Clearly idiosyncratic wage shocks explain a portion of the overall variance of wage growth in this model. Interestingly, some of the buoyant wage growth in the pre-crisis years is now explained by positive wage shocks, giving a lesser role to demand shocks vis-à-vis the baseline model. In addition, the period of weak wage growth from 2013 onwards accompanying the rapid fall in unemployment is now explained, to some extent, by negative wage shocks, and, interestingly, an additional drag from contractionary monetary policy shocks, whereas the baseline specification suggests monetary policy was supportive of wage growth in this period. These estimated wage shocks could capture compositional effects in the wage data.[10] The biggest driver of sub-trend wage growth remains productivity, however. That said, the effect in 2013 is roughly halved, and the persistence of the drag from productivity is notably lessened. The modified specification also reduces the extent to which shocks to the natural rate of unemployment are estimated to have increased wage growth over the period, relative to the baseline model.

6.2 Specification of the Model Structure

This section explores two extensions to the benchmark model designed to elaborate on two important issues. The first arises given the experience of the financial crisis and the Great Recession that followed, and seeks to add more realistic financial frictions to the benchmark model, and a financial intermediary sector in particular. This is a potentially powerful propagation mechanism affecting the

[9] See, e.g., Kass and Raftery (1995) and An and Schorfheide (2007).

[10] The correlation between these shocks' contributions to wage growth and the Bank of England's estimates of total wage compositional effects is around 33 percent between 2008 and 2017. Of the Bank of England's identified total compositional effects, the highest correlation with the wage shocks identified in our model is with those effects relating to changes in age, tenure, gender, region of residence, whether working full time, and whether in public-sector employment, where the correlation rises to 48 percent. To make this comparison, we use compositional effects published in Bank of England (2018, chapter 4).

Figure 13. Wage Growth Decompositions

A. Baseline Specification

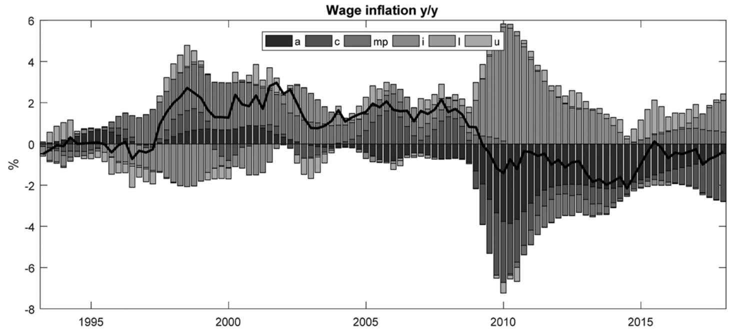

B. Baseline Specification + Transitory Shock in Wage Phillips Curve

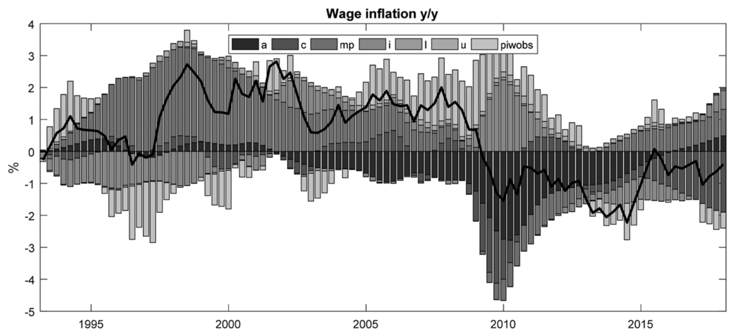

C. Baseline Specification + Financial Frictions

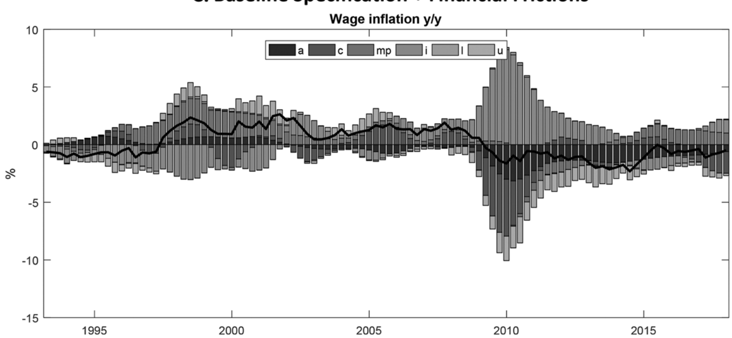

(continued)

Figure 13. (Continued)

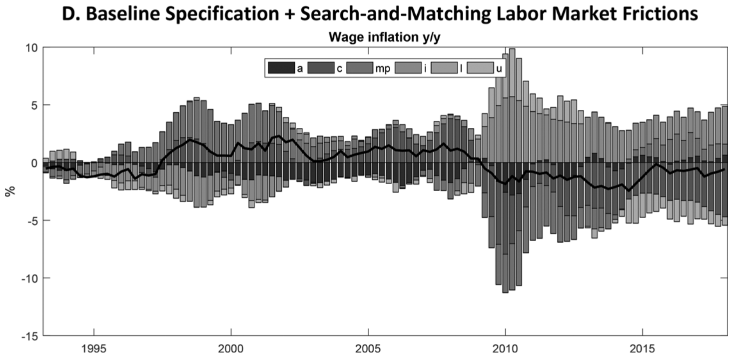

Notes: "a" = technology shock; "c" = demand shock; "i" = monetary policy shock; "l" = labor supply shock; "u" = natural rate of unemployment shock; "m" = price markup shock; "piwobs" = transitory wage Phillips curve shock.

transmission of macroeconomic disturbances through to the labor market. Equally, to the extent that financial frictions create a "bottleneck" between the funds provided by households and those available to firms, such frictions may have the effect of raising investment adjustment costs. A less elastic response of investment to macro disturbances may mean wages get a smaller boost from higher labor productivity following reductions in u*, giving this shock a better chance at explaining low wage growth and falling unemployment. Hence, incorporating financial frictions is not only interesting in its own right, but also has the potential to interact in important ways with labor market shocks by changing the way these disturbances propagate through the economy. This extension is considered in section 6.2.1. The second arises out of the benchmark model's relatively parsimonious description of the labor market. In particular, it lacks any role for labor market search-and-matching frictions that have been studied extensively in the literature and are regarded as giving rise to a more "realistic" wage-setting curve. This extension to the benchmark model is considered in section 6.2.2.

6.2.1 Financial Frictions

In this section we introduce more realistic financial frictions. In particular, rather than assuming that households can obtain claims on the physical capital stock directly, we instead assume households deposit all their savings with a commercial bank. Commercial banks, in turn, combine their own net worth with deposit funding to make loans to final goods firms. Finally, final goods firms use the loans obtained from banks to finance the physical capital they use in production. This description of financial intermediation follows Gertler and Kiyotaki (2010), and Gertler and Karadi (2011). The result is an endogenous spread between the return on one-period nominal bonds, or bank deposits, and the return on capital; the spread, in turn, varies with the net worth of the financial intermediary sector, generating a financial accelerator, à la Bernanke, Gertler, and Gilchrist (1999). The following section describes these modifications to the benchmark model in more detail.[11]

Modifications to the Benchmark Model. In the model with financial frictions, a continuum of competitive banks indexed by $i \in [0, 1]$ raise deposits $D_t(i)$ from households, combining these with their own net worth $N_t^b(i)$, to make loans to final goods firms $S_t^b(i)$. Final goods firms use bank loans to finance the capital stock; due to competition, the return on a bank's assets is equal to the real return on physical capital.

In this setting the household budget constraint is

$$P_t C_t + D_t = (1 + i_{t-1})\, D_{t-1} + \int_0^1 W_t(i) N_t(i) di + J_t + J_t^b, \quad (17)$$

where J_t^b are profits returned to the household lump sum from the banking sector. In turn, bank i's balance sheet reads

$$Q_t^k S_t^b\,(i) = D_t\,(i) + N_t^b(i). \quad (18)$$

Banks pay the real risk-free return R_t on their deposits. As such, the profit of the bank is $R_t^k Q_{t-1}^k S_{t-1}^k(i) - R_{t-1} D_{t-1}(i)$. After gathering returns, each period the bank pays out a fraction $1 - \sigma_b \in [0, 1]$ of its profits as a dividend to the bank's ultimate owners, households. Equivalently, with probability $1 - \sigma_b$, the manager of bank i receives a signal to "exit," in which case she returns the bank's net worth to the owning household lump sum. The franchise value of the bank

[11] A full derivation is in the appendix.

is therefore $V_0^b(i) \equiv E_0 \sum_{t=0}^{\infty} \Lambda_{0,t}[(1-\sigma_b) N_t^b(i) + \sigma_b V_t^b(i)]$. Banks are operated by managers subject to incentive frictions. The financial friction takes the following form. Each period, the manager of bank i can abscond with a fraction $\theta_b \in [0,1]$ of the bank's assets (e.g., by paying herself an excessive bonus, by engaging in vanity projects, etc.). It is incentive compatible for the bank manager not to do so if the franchise value of the bank $V_t^b(i)$ exceeds the proceeds of malfeasance, or

$$V_t^b(i) \geq \theta_b Q_t^k S_t^b(i). \tag{19}$$

The bank manager's problem is therefore to choose the scale of her balance sheet, $S_t^b(i)$, so as to maximize her bank's franchise value, but subject to the incentive compatibility constraint (17). The full solution to the banker's problem is given in the appendix. The key equilibrium condition emerging from this problem is the following:

$$\nu_t^k - \nu_t^d = \theta_b - \frac{\nu_t^d}{\ell_t}, \tag{20}$$

where $\ell_t(i) \equiv \frac{Q_t^k S_t^b(i)}{N_t^b(i)}$ is the bank's *leverage* (at market value), which by symmetry is the same for all i, and $v_t^j, j = k, d$ are the marginal returns to the bank of loans (when $j = k$) and deposits (when $j = d$). As a result, the term $v_t^k - v_t^d$ is closely related to the spread between the return on capital, R_t^k, and the risk-free interest rate, R_t. Expression (18) makes this spread a function of the severity of the incentive friction, θ_b, the marginal return on deposits, ν_t^d, and banking-sector leverage, ℓ_t. The financial accelerator mechanism then operates as follows. When asset prices (the price of capital) fall, the value of bank net worth, or equity capital, also falls, reflecting the lower value of the bank's claims on the physical capital stock. Lower net worth raises mark-to-market leverage, which via equation (18) causes a cutback in credit supply and so a rise in credit spreads. This endogenous contraction in credit supply in turn reduces investment by firms, lowering the demand for capital, and so further lowering the price of capital, amplifying the impact of initial disturbance on asset prices, investment, and so demand. This mechanism also operates in reverse, amplifying positive shocks that cause asset prices to rise.

Estimation Results: Financial Frictions. Estimation employs the same observables as in the baseline model, and all model priors are identical to the baseline, with the exception of two that are new and are related to the financial frictions. These are σ_b, 1 minus the "dividend rate" for banks, and N_b/S_b, the bank's steady-state capital ratio, which is related to the severity of financial frictions in this model.[12] The priors for these parameters are centered at 0.90 and 0.10, respectively, indicating a 10 percent quarterly dividend rate and a 10 percent equity capital ratio. The remaining additional banking parameters are governed by our calibration of the steady-state credit spread, which we set at 200 basis points on an annualized basis.

The parameter estimates are contained in table 3, column 3. Relative to the baseline model, a few changes stand out. First, the forcing process for the natural rate of unemployment shock is estimated to be more persistent. Second, the elasticity of investment with respect to the price of capital, ω, falls, indicating slightly higher investment adjustment costs. Third, the wage Phillips curve steepens a little. Finally, the standard deviations of the forcing processes are broadly unchanged, with the exception of the unemployment shock, whose standard deviation falls by around 25 percent. The bottom row in table 3 contains the marginal density of the model. Here it is apparent that the fit of the model is inferior to the baseline model, despite the extra richness arising from the financial frictions, with a Bayes factor of exp(31.4), which, somewhat surprisingly, provides strong support for the baseline model over the financial frictions model.

Overall, the propagation of shocks is relatively modestly affected by the presence of financial frictions. As judged from the impulse response functions, the transmission of technology shocks, for example, is little changed. Other shocks, including monetary, u* and labor supply shocks, exhibit some amplification, however, as would be expected in the presence of a financial accelerator mechanism. The modest degree of amplification is a corollary of relatively contained

[12]In particular, a higher value of θ_b—indicating worse financial frictions—requires the bank to be funded with more equity in equilibrium, reflecting a more severe agency problem between the banker and the household, such that N_b/S_b is higher.

movements in bond spreads in the estimated model.[13] This may be the result of not having included bond spreads within the set of observables.

Despite this inferior fit, for completeness we report the wage decomposition from this model in figure 13, panel C. Consistent with the intuition above, together with the changes in the estimated parameter vector noted previously, shocks to the natural rate of unemployment now exhibit a persistent drag on wage growth throughout the post-crisis period. Together with negative demand shocks, these are broadly offset by monetary policy shocks, which continue to be estimated to have boosted wage growth substantially. The residual wage weakness continues to be explained by weak productivity, as in the baseline model. So despite the inferior fit, the additional frictions implied by the financial frictions model do help resolve the puzzling absence of dampening effects on wages from labor market shocks exhibited by the baseline model.

6.2.2 Labor Market Frictions

In this section, we consider an alternative formulation of labor market frictions, switching the "monopolistic competition in the labor market" formulation employed in the baseline model for a formulation in the Diamond-Mortensen-Pissarides search-and-matching tradition, specifically one following Thomas (2008). In this framework, unemployment arises as a result of search-and-matching frictions in the labor market, rather than as a result of inefficiently high wage markups as in the baseline model.

In the model with search-and-matching frictions, the household's utility function is

$$\mathcal{U} \equiv E_t \sum_{t=0}^{\infty} \beta^t \mathcal{W}_t, \tag{21}$$

[13] For example, when evaluated at the posterior mean of the parameter distribution, the standard deviation of bond spreads implied by the model is around one-half that of GDP.

$$\mathcal{W}_t \equiv \log\left(C_t - hC_{t-1}\right) - \frac{\exp\left(u_t^l\right)}{1+\phi}\chi_h \int_0^1 H_t\left(i\right)^{1+\varphi} N_t\left(i\right) di \\ - b\exp\left(u_t^u\right) N_t, \tag{22}$$

where $N_t(i)$ denotes the number of workers employed in firm i, $H_t(i)$ are hours per worker, and b is the marginal disutility of working, subject to random variation u_t^u. This shock replaces the "u*" shock from the baseline model, although its role is similar in that variation in u_t^u causes fluctuations in workers' desired employment rate, implying, in turn, shocks to the "natural" rate of unemployment in this variant of the model. The budget constraint remains unchanged vis-à-vis the baseline model. The labor force remains normalized to unity.

Each period, the probability of job separation is λ. The remaining fraction $1-\lambda$ of employed people remains in employment the following period. In addition, there is a probability $p(\tilde{\theta}_t)$ that some of the $1 - N_t = U_t$ unemployed fraction of the population enters employment. This probability depends on labor market tightness, $\tilde{\theta}_t$, which is defined as the ratio of job vacancies posted by firms, V_t, and the unemployed, U_t, such that $\tilde{\theta}_t \equiv V_t/U_t$. A matching function $m(.)$ defines the efficiency of this process, such that

$$p\left(\tilde{\theta}_t\right) \equiv \frac{m\left(V_t, U_t\right)}{U_t}, \tag{23}$$

where $p(\tilde{\theta}_t)$ is increasing in $\tilde{\theta}_t$. The *matching rate* is $q\left(\tilde{\theta}_t\right) \equiv m(V_t, U_t)/V_t$, and $q(.)$ and $p(.)$ are therefore related according to $p\left(\tilde{\theta}_t\right) = \tilde{\theta}_t q(\tilde{\theta}_t)$. We then write the law of motion for employment as

$$N_t = \left(1-\lambda\right) N_{t-1} + p\left(\tilde{\theta}_t\right) U_t. \tag{24}$$

The first element of labor market equilibrium is derived from this problem, namely, the value to the household of an addition job. This is given by $\mathcal{S}_t^w\left(i\right) \equiv \frac{\partial \mathcal{W}_t}{\partial N_t(i)}\frac{1}{U_{c,t}}$, where $U_{c,t}$ is the marginal utility of consumption.

The second element comes from the firm side. In particular, intermediate goods production entails the operation of the following production technology, at firm i, given by $Y_t\left(i\right) = \exp\left(u_t^a\right)\left(H_t\left(i\right) N_t\left(i\right)\right)^{1-\alpha} K_t\left(i\right)^{\alpha}$. Firm i posts vacancies

$V_t(i)$, and doing so entails a utility cost to the firm's management of $C(V_t(i))$, where $C(.)$ is convex. Specifically, following Thomas (2008),

$$C\left(V_t\left(i\right)\right) = \frac{1}{U_{c,t}} \frac{\chi_v}{1+\psi_v} \left(\frac{V_t\left(i\right)}{N_t\left(i\right)}\right)^{1+\psi_v} N_t\left(i\right). \tag{25}$$

The stock of employment at firm i evolves according to $N_{t+1}\left(i\right) = \left(1-\lambda\right) N_t + q\left(\tilde{\theta}_t\right) V_t(i)$. The firm's problem is to maximize the discounted flow of future profits subject to the vacancy posting cost and to the law of motion for employment by choice of vacancies, employment, and capital. The resulting marginal value of employment $\mathcal{S}_t^f\left(i\right)$ to the firm can be used to derive the job creation curve, which relates the firm's vacancy rate $Z_t^v(i) \equiv V_t(i)/N_t(i)$ to the marginal product of labor, the vacancy posting cost, and—given the law of motion for employment—the future vacancy rate, $Z_{t+1}^v(i)$. Nominal wages are set on a staggered basis. When not reoptimized, wages are indexed to past outturns. When reoptimized, firms and workers engage in Nash bargaining, which results in the joint surplus to workers and firms of an employment match being divided according to

$$\mathcal{S}_t^f\left(i\right) = \xi\left[\mathcal{S}_t^f\left(i\right) + \mathcal{S}_t^w\left(i\right)\right], \tag{26}$$

where $\xi \in [0,1]$ is the bargaining power of firms. This condition results in an expression for the Nash bargained real wage, and the gap between this Nash wage and real wage outturns becomes the key argument on the right-hand side of the wage Phillips curve (see Thomas 2008), replacing the unemployment gap in the baseline model.

Finally, bargaining over hours is assumed to be efficient, maximizing the joint surplus of workers and firms. This results in an expression equating the marginal revenue product of labor with the marginal rate of substitution between hours and consumption.

The remainder of the model is left unchanged vis-à-vis the baseline, with the exception that the labor supply block no longer contains a role for the parameter ν, which previously controlled the strength of the wealth effect on labor supply.

In addition to the new equations tracking the stock of employment, labor market tightness, the unemployment rate, and the vacancy rate, the matching problem generates two key additional relationships together with a modification to the wage Phillips curve. In log-linear terms, the wage Phillips curve becomes

$$\pi_{w,t} - \gamma_w \pi_{w,t-1} = \beta \left(1 - p\left(\tilde{\theta}\right) - \lambda\right) E_t \left(\pi_{w,t+1} - \gamma_w \pi_{w,t}\right) + \lambda_w \left(w_t^n - w_t\right), \quad (27)$$

where w_t^n is the Nash-bargained wage described above, and $w_t = w_{t-1} + \pi_{w,t} - \pi_t$ is the real wage. In addition, the labor market block includes a linearized Nash wage condition, which relates the Nash-bargained wage to the marginal disutility of labor supply, the marginal product of labor, and exogenous variation in u_t^u, households' disutility of employment; and a linearized job creation curve, which relates vacancy posting positively to the marginal product of labor and negatively to the cost of vacancy posting. A complete set of linearized equations is available in the appendix.

Estimation Results: Labor Market Frictions. The priors for estimation are the same as in the baseline model, but with two exceptions. First, the parameter ν no longer enters the model, and so is dropped. Second, because the Nash wage gap rather than the unemployment gap enters the wage Phillips curve, and the former has a greater volatility than the latter, we adjust the prior on its slope downwards so that the response of wages and unemployment to demand shocks broadly matches that in the baseline model. Finally, a number of new steady-state relationships are introduced in the labor frictions model that do not exist in the baseline model. We calibrate these to U.K. data. In particular, we assume a steady-state job-finding rate of 30 percent, which matches the U.K. pre-crisis average; a steady-state unemployment rate of 5 percent; a labor share of income of 67 percent; a vacancy posting cost of 1 percent of GDP, following Thomas (2008); and equal bargaining power between firms and workers. The steady-state employment rate, capital share of income, and disutility of employment follow as a result of these parameter settings.

The parameter estimates for the model with labor market matching frictions are shown in table 3, column 4. Relative to the baseline model reported in column 1, the productivity process is estimated to be less persistent, although its volatility is much increased. Some of the supply-side persistence previously estimated to originate in technology shocks is replaced by labor market shocks—the "u*" shock in particular, whose persistence is estimated to be much greater in this variant of the model, but whose volatility falls. Estimates of the other structural parameters remain broadly similar to the baseline model, with the main exceptions being the estimated degrees of persistence in the nominal wage and inflation processes (γ_w and γ_p), estimates of which approximately double. The investment elasticity rises—unlike in the model with financial frictions—indicating lower investment adjustment costs.

The shock-based decomposition of wage growth is shown in figure 13, panel D. Compared with the baseline model, the role of technology shocks in explaining wage weakness is much reduced. In their place, demand shocks and reductions in households' disutility from employment—the stand-in for "u*" shocks in this model—make substantial contributions to explaining post-crisis wage weakness. This appears to be for two reasons. First, the estimated size and persistence of technology shocks is lower in the labor frictions model. Second, and more importantly, conditional on technology shocks, the co-movement of wage inflation and output changes from positive in the baseline model to negative in the labor frictions model. This is because of greater sluggishness in the adjustment of labor market quantities in the labor frictions model compared with the baseline model. This means the disinflationary effects of positive technology shocks appear in labor "prices" (i.e., nominal wages) to a greater extent than they appear in labor "quantities" (i.e., unemployment and employment) when there are search-and-matching frictions. Thus, negative technology shocks tend to boost nominal wage inflation in the labor frictions model, whereas they drag on wage inflation in the baseline model.

However, the fit of the search-and-matching model is substantially inferior to the preceding variants, as indicated by the marginal data densities shown in the final row of table 3. Compared with the baseline model, for example, the marginal data density of the search-and-matching model is around 378 points lower. As a result, the Bayes factor indicates very strong support in favor of the baseline model over the model with matching frictions, suggesting that, at the margin, the data favor an explanation for weak wage growth that leans heavily on weak productivity outturns.

7. Concluding Remarks

As U.K. unemployment fell strongly from the middle of 2013, faster than implied by GDP growth, why did wage growth fail to pick up? Did the wage Phillips curve break down? There are two (broad) candidate explanations: first, that the fall in unemployment was accompanied by the concomitant fall in the natural rate of unemployment, such that weak wage growth reflected a lingering output gap; second, that the fall in unemployment in fact reflected a strong recovery in demand, with weak wage growth being explained by some other factors, perhaps weak productivity, which also dragged on demand. This paper suggests the data favor a mix of explanations, but with the balance of evidence favoring the latter explanation: a strong abatement of demand headwinds lowered unemployment, but weak productivity growth kept wage growth low and meant unemployment could fall more quickly than implied by GDP growth. These conclusions are subject to uncertainty, and necessarily depend on a number of modeling choices. For example, the baseline model suggesting these conclusions has been kept as parsimonious as possible. However, richer models including financial frictions or search-and-matching labor market frictions either suggest similar conclusions or seem to fit the data less well.

Appendix

Benchmark Model

Households

Households of unit measure each contain a continuum of members represented by the unit square, indexed by the pair $(i, j) \in [0, 1]^2$, where dimension i represents the type of labor service the individual supplies, and where j indexes his or her disutility from work. The disutility of labor is given by j^{φ}. $N_t(i)$ denotes the fraction of a particular household's members that supplies labor of type i, so that for the household as a whole, the disutility of labor is $\int_0^1 \int_0^{N_t(i)} j^{\varphi} dj\, di = (1+\varphi)^{-1} \int_0^1 N_t\,(i)^{1+\varphi}\, di$. In addition to choosing the fraction of the household that supplies labor (described further below), the household chooses consumption to maximize

$$\mathcal{U} \equiv E_t \sum_{t=0}^{\infty} \beta^t \exp(\tilde{u}_t^c) \Bigg[\log(C_t - hC_{t-1}) - \frac{\exp(u_t^l)}{1+\varphi} X_t \int_0^1 N_t(i)^{1+\varphi} di \Bigg], \tag{28}$$

in which E_t is the mathematical expectations operator, C_t denotes the household's consumption of a CES bundle of differentiated final goods varieties, β determines the rate of time preference, h determines the strength of habits in consumption, and φ determines the elasticity of labor supply. Variable u_t^l is a shock to the marginal disutility of working, which we refer to as a labor supply shock, and which follows an exogenous AR(1) process, $u_t^l = \rho_l u_{t-1}^l + \varepsilon_t^l$, $\varepsilon_t^l \sim N(0, \sigma_l^2)$. Similarly, the variable $\tilde{u}_t^c$ is a disturbance to intertemporal preferences, a linear scaling of which u_t^c also follows an AR(1) process, $u_t^c = \rho_c u_{t-1}^c + \varepsilon_t^c \varepsilon_t^c \sim N(0, \sigma_c^2)$.[14] Finally, variable X_t is introduced to allow for a flexible parameterization of the wealth effect on labor supply. In particular,

$$X_t \equiv S_t \left(C_t - hC_{t-1}\right)^{-1}, \tag{29}$$

$$S_t = S_{t-1}^{1-\nu} \left(C_t - hC_{t-1}\right)^{\nu}. \tag{30}$$

Maximization of $\mathcal{U}$ is subject to the following flow budget constraint:

$$P_t C_t + D_t + Q_t^k S_t^k = (1 + i_{t-1}) D_{t-1} + R_t^k Q_{t-1}^k S_{t-1}^k + \int_0^1 W_t(i) N_t(i) di + J_t, \tag{31}$$

where P_t is the aggregate price level; D_t are savings in the form of one-period risk-free bonds, paying the nominal interest rate $1 + i_t$; S_t^k are claims on the physical capital stock with price Q_t^k in terms of final goods, which yield a return R_t^k; $W_t(i)$ is the nominal wage of labor type i; and J_t are profits returned to the household lump sum from firms in the economy. The familiar optimality conditions

[14] The rescaling ensures the intertemporal preference shock enters the consumption Euler equation with a unit coefficient. The rescaling is $u_t^c \equiv \frac{1-h}{1+h}(1-\rho_c)\tilde{u}_t^c$, where ρ_c is the AR(1) coefficient in the process governing $\tilde{u}_t^c$.

resulting from this problem include the optimal consumption plan, which satisfies

$$E_t \beta \frac{\exp(\tilde{u}_{t+1}^c)}{\exp(\tilde{u}_t^c)} \frac{U_{c,t+1}}{U_{c,t}} \frac{1+i_t}{\Pi_{t+1}} = 1, \tag{32}$$

where $\Pi_t \equiv P_t/P_{t-1}$ is the gross CPI inflation rate and $U_{c,t} \equiv (C_t - hC_{t-1})^{-1}$ is the marginal utility of consumption, together with

$$E_t \beta \frac{\exp(\tilde{u}_{t+1}^c)}{\exp(\tilde{u}_t^c)} \frac{U_{c,t+1}}{U_{c,t}} \frac{1}{\Pi_{t+1}} \left[(1+i_t) - R_{t+1}^k\right] = 0, \tag{33}$$

which equates returns on bonds and capital. The intertemporal disturbance in the consumption Euler equation can be seen as driving a wedge between the ex ante real interest rate and the realized real interest rate. This plays the role of a demand shock in the model. An alternative interpretation is that it serves as a stand-in for shocks originating in the financial sector that drive a wedge between the interest rate paid or received by households, and the risk-free interest rate implemented by the monetary authority.

Regarding labor supply, a household worker supplying labor type i will find it worthwhile to supply labor if $U_{c,t} \times W_t(i)/P_t \geq \exp\left(u_t^l\right) X_t j^{\varphi}$, such that the marginal supplier of labor $L_t(i)$ is

$$U_{c,t} \frac{W_t(i)}{P_t} = \exp\left(u_t^l\right) X_t L_t\left(i\right)^{\varphi}. \tag{34}$$

This is the labor supply schedule for labor type i. From this and the definitions of X_t and S_t above, we see that as $\nu \rightarrow 1$, the standard wealth effect on labor supply operates fully: as consumption rises, the marginal utility of consumption falls and the labor supply curve shifts inwards. Allowing for $\nu < 1$ means this effect can be reduced, however, and as $\nu \rightarrow 0$, the wealth effect on labor supply disappears. In this case, labor supply only rises in response to higher real wages or due to shocks to the disutility of working.

Production

Goods producers $f \in [0,1]$ generate differentiated-products final goods using the production function $Y_t\left(f\right) = \exp\left(u_t^a\right) N_t\left(f\right)^{1-\alpha}$

$K_{t-1}(f)^{\alpha}$, where $N_t(f)$ is a bundle of labor types employed, defined by

$$N_t(f) \equiv \left[\int_0^1 N_t(f,i)^{\frac{\varepsilon_{w,t}-1}{\varepsilon_{w,t}}} di\right]^{\frac{\varepsilon_{w,t}}{\varepsilon_{w,t}-1}}, \tag{35}$$

where $\varepsilon_{w,t}$ controls the elasticity of substitution between types of labor, which is subject to exogenous variation (described below), and where $K_{t-1}(f)$ is capital installed in period $t-1$ available for use in production in period t. The variable u_t^a follows an exogenous AR(1) process $u_t^a = \rho_a u_{t-1}^a + \varepsilon_t^a, \varepsilon_t^a \sim N(0, \sigma_a^2)$ and captures exogenous variation in total factor productivity, which is assumed to be stationary. Labor demand for labor type i is

$$N_t(f,i) = \left(\frac{W_t(i)}{W_t}\right)^{-\varepsilon_{w,t}} N_t(f), \tag{36}$$

where $W_t \equiv \left[\int_0^1 W_t(i)^{1-\varepsilon_{w,t}} di\right]^{\frac{1}{1-\varepsilon_{w,t}}}$, such that $\int_0^1 W_t(i) N_t(f, i) di = W_t N_t(f)$, allowing the firm's profit function to be written,

$$\mathcal{P}_t(f) \equiv (1+\tau) P_t(f) Y_t(f) - W_t N_t(f) - Z_t^k K_{t-1}(f), \tag{37}$$

where τ is a production subsidy offsetting the steady-state distortion associated with monopolistic competition in the goods market and Z_t^k is the return on capital. $P_t(f)$ is the firm's selling price, and it faces a demand curve for its product given by

$$Y_t(f) = \left(\frac{P_t(f)}{P_t}\right)^{-\varepsilon_{p,t}} Y_t, \tag{38}$$

where $\varepsilon_{p,t}$ is the elasticity of substitution between varieties of final good, which is subject to exogenous variation (described below). Firm f resets its price only with probability $1-\theta_p$ each period; when it is not reset, it is indexed to past inflation with intensity γ_p. The firm's problem is to choose its reset price to solve

$$\max_{P_t^*(f)} E_t \sum_{i=0}^{\infty} \theta_p^i \Lambda_{t,t+i} \left[(1+\tau) P_t^*(f) \prod_{k=0}^{i} \Pi_{t+k+i}^{\gamma_p} - M_{t+i}(f)\right] Y_{t+i|t}(f), \tag{39}$$

where $Y_{t+i|t}(f) = (P_t^*(f)/P_{t+i})^{-\varepsilon} Y_{t+i}$ is demand in period $t+i$ conditional on having reset price to $P_t^*(f)$ in period t, $\Lambda_{t,t+i} \equiv \beta U_{c,t+i}/U_{c,t}$ is the household's stochastic discount factor used to value future profits, and $M_{t+i}(f)$ is the firm's nominal marginal cost. This, in turn, is related to factor demands according to

$$M_t(f)(1-\alpha)\frac{Y_t(f)}{N_t(f)} = W_t, \tag{40}$$

$$M_t(f)\,\alpha\frac{Y_t(f)}{N_t(f)} = Z_t^k, \tag{41}$$

while the first-order condition for the firm's optimal pricing decision is

$$E_t \sum_{i=0}^{\infty} \theta_p^i \Lambda_{t,t+i}\left[(1+\tau)P_t^*(f)\prod_{k=0}^{i}\Pi_{t+k+i}^{\gamma_p} - \frac{\varepsilon_{p,t}}{\varepsilon_{p,t}-1}M_{t+i}(f)\right] \times Y_{t+i|t}(f) = 0. \tag{42}$$

Observe that in steady state, $(1+\tau)P^*(f) = \frac{\varepsilon_p}{\varepsilon_p-1}M(f)$, such that a subsidy that is increasing in the price markup $1+\tau = \frac{\varepsilon_p}{\varepsilon_p-1} \equiv \mu_p$ ensures that pricing is efficient in steady state. Letting the price markup $\mu_{p,t} \equiv \frac{\varepsilon_{p,t}}{\varepsilon_{p,t}-1}$ be subject to exogenous variation that is linear in logs captures exogenous variation in firms' pricing power which can generate increases or decreases in inflationary pressure. In particular, let $\log \mu_{p,t} \equiv u_t^p = \rho_p u_{t-1}^p + \varepsilon_t^p$,$\varepsilon_t^p \sim N(0,\sigma_p^2)$, such that by symmetry across firms, the price Phillips curve reads, in log-linear terms, as

$$\pi_t - \gamma_p\pi_{t-1} = \beta E_t(\pi_{t+1} - \gamma_p\pi_t) + \lambda_p m_t + u_t^p, \tag{43}$$

where m_t is log *real* marginal cost.

Wage Setting

Wage bargaining takes place on behalf of workers by recruitment agencies, which reset wages only periodically. In any period, the probability of wage renegotiation is $1-\theta_w$, which is undertaken by recruitment agencies with the objective of maximizing

$$\max_{W_t^*} E_t \sum_{k=0}^{\infty} \left(\beta\theta_w\right)^k U(C_{t+k|t}, N_{t+k|t}), \tag{44}$$

where $U(.,.)$ is the period utility function for workers, where $X_{t+k|t}, X = \{C, N\}$, denotes consumption or employment in period $t+k$ whose wage was reset in period t at W_t^*. The problem is to choose the optimal reset wage subject to the household's budget constraint and the sequence of labor demands from firms given above, resulting in

$$\sum_{k=0}^{\infty} \left(\beta\theta_w\right)^k E_t \left[N_{t+k|t} U_{c,t+k|t} \frac{W_t^*}{P_{t+k}} + \frac{\varepsilon_{w,t}}{\varepsilon_{w,t}-1} U_{n,t+k|t} \right] = 0, \tag{45}$$

where $U_{n,t}$ is the marginal utility of supplying extra labor. Defining the marginal rate of substitution between consumption and leisure $MRS_{t+k|t} \equiv -U_{n,t+k|t}/U_{c,t+k|t} = -\exp\left(u_t^l\right) S_t N_{t+k|t}^{\phi}$, this becomes

$$\sum_{k=0}^{\infty} \left(\beta\theta_w\right)^k E_t \left[N_{t+k|t} U_{c,t+k|t} \left(\frac{W_t^*}{P_{t+k}} - \frac{\varepsilon_{w,t}}{\varepsilon_{w,t}-1} MRS_{t+k|t} \right) \right] = 0, \tag{46}$$

an expression determining the wage markup. In steady state, the real wage is related to the marginal rate of substitution between consumption and labor according to $W^*/P = \frac{\varepsilon_w}{\varepsilon_w - 1} MRS$, where $\mu_w \equiv \frac{\varepsilon_w}{\varepsilon_w - 1}$ is the steady-state wage markup. Around the steady state, we allow the log wage markup to undergo exogenous variation, described further below. Further, wages that are not reoptimized are indexed to past wage inflation: $W_{t+k|t} = W_{t+k-1|t} \Pi_{w,t-1}^{\gamma_w} \Pi^{1-\gamma_w}$, where $\Pi_{w,t}$ is wage inflation, and such that the aggregate wage index evolves according to

$$W_t^{1-\varepsilon_w} = \theta_w \left(W_{t-1} \Pi_{w,t-1}^{\gamma_w} \Pi^{1-\gamma_w} \right)^{1-\varepsilon_w} + (1-\theta_w) \left(W_t^*\right)^{1-\varepsilon_w}. \tag{47}$$

Unemployment and the Wage Phillips Curve

The following is Galí's (2011) reinterpretation of the staggered-wage-setting process described by Erceg, Henderson, and Levin (2000) in terms of a wage Phillips curve à la Phillips (1958). Letting lowercases denote log-deviations, the unemployment rate is $u_t = l_t - n_t$, i.e., the excess of labor supply over labor demand. In log-linear terms, the labor supply condition derived above can be written $w_t - p_t = s_t + \varphi l_t + u_t^l$. Similarly, the optimal wage-setting condition gives a log-linear expression for the wage markup of $\mu_t^w = w_t - p_t - mrs_t$, in which $mrs_t = s_t + \varphi n_t + u_t^l$. Combining these gives the wage markup as

$$\mu_t^w = \varphi\left(l_t - n_t\right) = \varphi u_t, \tag{48}$$

i.e., the wage markup is directly proportional to the unemployment rate. Next, note that the wage-setting problem above results in the following expression for nominal wage inflation:

$$\pi_{w,t} - \gamma_w \pi_{w,t-1} = \beta E_t\left(\pi_{w,t+1} - \gamma_w \pi_{w,t}\right) - \tilde{\lambda}_w\left(\mu_t^w - \mu_t^{w*}\right), \tag{49}$$

in which μ_t^{w*} is the natural wage markup and $\tilde{\lambda}_w \equiv (1 - \beta\theta_w)(1 - \theta_w)/\theta_w(1 + \varepsilon_w \varphi)$. Using the expression for the wage markup, we can then write the wage Phillips curve as

$$\pi_{w,t} - \gamma_w \pi_{w,t-1} = \beta E_t\left(\pi_{w,t+1} - \gamma_w \pi_{w,t}\right) + \lambda_w\left(u_t - u_t^*\right), \tag{50}$$

where $\lambda_w \equiv \varphi\tilde{\lambda}_w$ such that $u_t^* \equiv \mu_t^{w*}/\varphi$ is the natural rate of unemployment, which we assume follows an exogenous AR(1) process, $u_t^* = \rho_u u_{t-1}^* + \varepsilon_t^{u*}$, $\varepsilon_t^{u*} \sim N(0, \sigma_u^2)$.

Capital Goods Producers

Capital goods are produced by perfectly competitive firms that transform final goods into capital goods, subject to "flow" adjustment costs. Such producers solve

$$\max_{\{I_t\}_{t=0}^{\infty}} E_t \sum_{i=0}^{\infty} \Lambda_{t,t+i}\left[\left(Q_t^k - 1\right) I_{t+i} - f\left(\frac{I_{t+i}}{I_{t+i-1}}\right) I_{t+i}\right], \tag{51}$$

where Q_t^k is the real price of capital goods in terms of final goods, and the function $f(.)$ captures investment adjustment costs, which satisfies $f'(1) = 0$, $f''(1) > 0$. The first-order condition for investment is

$$Q_t^k = 1 + f\left(\frac{I_t}{I_{t-1}}\right) + f'\left(\frac{I_t}{I_{t-1}}\right)\frac{I_t}{I_{t-1}} - E_t\Lambda_{t,t+1}f'\left(\frac{I_{t+1}}{I_t}\right)\left(\frac{I_{t+1}}{I_t}\right)^2, \quad (52)$$

which determines the dynamics of Tobin's Q. Finally, as a result of investment activity, the capital stock evolves according to $K_t = (1-\delta)K_{t-1} + I_t$, where δ is the depreciation rate of physical capital.

Market Clearing and Monetary Policy

Finally, the goods market must clear:

$$Y_t\Delta_t = C_t + \left(1 + f\left(\frac{I_t}{I_{t-1}}\right)\right)I_t, \quad (53)$$

where Δ_t is an index of price dispersion across producers; households' claims on the physical capital stock must equal the total stock, $S_t^k = K_t$; and monetary policy seeks to stabilize inflation and the output gap according to the rule

$$\frac{1+i_t}{1+i} = \left(\frac{1+i_{t-1}}{1+i}\right)^{\rho_i}\left(\Pi_t^{\phi_\pi}\left(\frac{Y_t}{Y_t^*}\right)^{\phi_y}\right)^{1-\rho_i}\exp(\varepsilon_t^i), \quad (54)$$

where Y_t^* is the level of output that would prevail under flexible prices, and where $\varepsilon_t^i \sim N(0, \sigma_i^2)$ is a shock to the monetary policy rule.

Model with Banking and Financial Frictions

Derivation

In the model with financial frictions, the banker's problem is to maximize $V_0^b(i) \equiv E_0\sum_{t=0}^{\infty}\Lambda_{0,t}[(1-\sigma_b)N_t^b(i) + \sigma_b V_t^b(i)]$ subject to the balance sheet constraint, $Q_t^k S_t^b(i) = D_t(i) + N_t^b(i)$, and the incentive compatibility constraint, $V_t^b(i) \geq \theta_b Q_t^k S_t^b(i)$. We form the Lagrangean:

$$L^b(i) = \left(1+\lambda_t^b\right)V_t^b(i) - \lambda_t^b(i)\,\theta_b Q_t^k S_t^b(i).$$

We guess that the value of the bank takes the following form: $V_t^b(i) = \nu_t^k Q_t^k S_t^b(i) - \nu_t^d D_t(i)$. Using the balance sheet constraint, we then get

$$V_t^b(i) = \nu_t^k Q_t^k S_t^b(i) - \nu_t^d \left[Q_t^k S_t^b(i) - N_t^b(i)\right],$$

such that the first-order condition for lending is

$$\left(1+\lambda_t^b\right)\left(\nu_t^k - \nu_t^d\right) = \lambda_t^b \theta_b,$$

which relates the lending spread to the tightness of the incentive compatibility constraint, measured by λ_t^b. We assume the constraint binds in the neighborhood of the steady state we study, in which case we also have

$$\left(\nu_t^k - \nu_t^d\right) Q_t^k S_t^b(i) + \nu_t^d N_t^b(i) = \theta_b Q_t^k S_t^b(i),$$

which relates the scale of the balance sheet to the size of "divertible" assets; this expression, in effect, determines the bank's leverage. Note that this can be rearranged to give

$$Q_t^k S_t^b(i) = \frac{\nu_t^d}{\theta_b - \left(\nu_t^k - \nu_t^d\right)} N_t^b(i).$$

So $V_t^b = \left(\nu_t^k - \nu_t^d\right) Q_t^k S_t^b(i) + \nu_t^d N_t^b(i) = \frac{\theta_b}{\theta_b - \left(\nu_t^k - \nu_t^d\right)} \nu_t^d N_t^b(i)$, or

$$V_t^b = \nu_t^d \left(1+\lambda_t^b\right) N_t^b(i).$$

Then the bank's value can be written as a linear function of its net worth:

$$\begin{aligned} V_t^b(i) &= E_t \Lambda_{t,t+1} \left[(1-\sigma_b) N_{t+1}^b(i) + \sigma_b V_{t+1}^b(i)\right] \\ &= E_t \Lambda_{t,t+1} \Omega_{t+1}^b N_{t+1}^b(i), \end{aligned}$$

where

$$\Omega_{t+1}^b \equiv 1 - \sigma_b + \sigma_b \frac{\theta_b}{\theta_b - \left(\nu_{t+1}^k - \nu_{t+1}^d\right)} \nu_{t+1}^d.$$

Note that this related to the Lagrange multiplier λ_t^b because

$$\lambda_t^b = \frac{\nu_t^k - \nu_t^d}{\theta_b - \left(\nu_t^k - \nu_t^d\right)}.$$

So

$$\Omega_{t+1}^b = 1 - \sigma_b + \sigma_b \nu_{t+1}^d \left(1 + \lambda_{t+1}^b\right).$$

Then, because $N_{t+1}^b(i) = R_{t+1}^k Q_t^k S_t^b(i) - R_t D_t(i)$, we have

$$\nu_t^k Q_t^k S_t^b\left(i\right) - \nu_t^d D_t\left(i\right) = E_t \Lambda_{t,t+1} \Omega_{t+1}^b [R_{t+1}^k Q_t^k S_t^b(i) - R_t D_t(i)]$$

such that equating coefficients yields

$$\nu_t^k = E_t \Lambda_{t,t+1} \Omega_{t+1}^b R_{t+1}^k$$
$$\nu_t^d = E_t \Lambda_{t,t+1} \Omega_{t+1}^b R_t.$$

Finally, we assume exiting bankers are replaced each period with a fraction ξ_b of last period's bankers' gross returns, such that, in aggregate,

$$N_{t+1}^b = \left(\sigma_b + \xi_b\right) R_{t+1}^k Q_t^k S_t^b - \sigma_b R_t D_t.$$

Moreover, bank loans finance the physical capital stock, such that $Q_t^k S_t^b = Q_t^k K_t$.

In summary, then, the complete set of banking equations is

$$\left(1 + \lambda_t^b\right)\left(\nu_t^k - \nu_t^d\right) = \lambda_t^b \theta_b$$
$$(\nu_t^k - \nu_t^d) Q_t^k S_t^b\left(i\right) + \nu_t^d N_t^b\left(i\right) = \theta_b Q_t^k S_t^b(i)$$
$$\Omega_t^b = 1 - \sigma_b + \sigma_b \nu_t^d \left(1 + \lambda_t^b\right)$$
$$\nu_t^k = E_t \Lambda_{t,t+1} \Omega_{t+1}^b R_{t+1}^k$$
$$\nu_t^d = E_t \Lambda_{t,t+1} \Omega_{t+1}^b R_t$$
$$N_{t+1}^b = (\sigma_b + \xi_b) R_{t+1}^k Q_t^k S_t^b - \sigma_b R_t D_t$$
$$Q_t^k S_t^b = D_t + N_t^b$$
$$Q_t^k S_t^b = Q_t^k K_t.$$

This represents eight new equations in seven new variables: $(S^b, D, N^b, \nu^k, \nu^d, \lambda^b, \Omega^b)$. The remaining amendment to the baseline model is to drop the bond-equity arbitrage condition, which is replaced by the equation for the credit spread provided by the banking-sector block.

Steady State

Compared with the frictionless model, the key difference to the model steady state in the presence of banking frictions is the presence of a steady-state credit spread, represented by $R^k - R > 0$, the excess return on capital over the risk-free interest rate. We treat the following as calibration targets, with the remaining parameters being chosen so as to be consistent with these targets in equilibrium: the steady-state credit spread, $R^k - R$; bank leverage, S^b/N^b; and the exit rate of bankers, or, equivalently, the bank's dividend rate, σ_b.

From the net worth accumulation equation, $N^b = (\sigma_b + \xi_b) R^k S^b - \sigma_b RD$, we get

$$\xi_b = \frac{\frac{N^b}{S^b} + \sigma_b \left[R \left(1 - \frac{N^b}{S^b} \right) - R^k \right]}{R^k},$$

since $1 - \frac{N^b}{S^b} = \frac{D}{S^b}$.

From the spread equation,

$$\left(1 + \lambda^b\right) \beta \Omega^b \left(R^k - R\right) = \lambda^b \theta_b.$$

From the balance sheet equation,

$$\theta_b = \beta \Omega^b \left[\left(R^k - R\right) + R \frac{N^b}{S^b} \right],$$

which gives us θ_b when we know Ω^b. So dividing one by the other,

$$\lambda^b = \frac{R^k - R}{R} \frac{1}{\frac{N^b}{S^b}}.$$

From the discount factor, $\Omega^b = 1 - \sigma_b + \sigma_b \beta \Omega^b R \left(1 + \lambda^b\right)$, since $v^d = \beta \Omega^b R = \Omega^b$, so $\Omega^b = \frac{1 - \sigma_b}{1 - \sigma_b (1 + \lambda^b)}$.

This can be plugged into the expression above to learn θ_b.

Linearized Banking Equations

Letting hats and lowercase variables denote log-deviations from steady state, we have

$$\left(1+\lambda^b\right)\left(v^k\ \hat{v}_t^k - v^d\ \hat{v}_t^d\right) + \left(v^k - v^d\right)\lambda^b\ \hat{\lambda}_t^b = \theta_b \lambda^b\ \hat{\lambda}_t^b$$

$$\left(v^k - v^d\right)\left(q_t^k + s_t^b\right) + \left(v^k\ \hat{v}_t^k - v^d\ \hat{v}_t^d\right) + v^d \frac{N^b}{S^b}\left(\hat{v}_t^d + n_t^b\right) = \theta_b (q_t^k + s_t^b)$$

$$\Omega^b\ \hat{\Omega}_t^b = \sigma_b v^d \hat{\lambda}^b\ \lambda_t^b + \sigma_b v^d \left(1+\lambda^b\right)\hat{v}_t^d$$

$$\hat{v}_t^k = \hat{\Lambda}_{t,t+1} + \hat{\Omega}_{t+1}^b + r_{t+1}^k$$

$$\hat{v}_t^d = \hat{\Lambda}_{t,t+1} + \hat{\Omega}_{t+1}^b + r_t$$

$$\frac{N^b}{S^b} n_t^b = \left(\sigma_b + \xi_b\right) R^k \left(r_t^k + q_{t-1}^k + s_{t-1}^b\right) - \sigma_b R \frac{D}{S^b}(r_{t-1} + d_{t-1})$$

$$q_t^k + s_t^b = \frac{D}{S^b} d_t + \frac{N^b}{S^b} n_t^b$$

$$s_t^b = k_t.$$

References

An, S., and F. Schorfheide. 2007. "Bayesian Analysis of DSGE Models." *Econometric Reviews* 26 (2–4): 113–72.

Banbura, M., D. Giannone, and L. Reichlin. 2010. "Large Bayesian Vector Auto Regressions." *Journal of Applied Econometrics* 25 (1): 71–92.

Bank of England. 2018. *Inflation Report*, February. Available at https://www.bankofengland.co.uk/inflation-report/2018/february-2018.

Barnett, A., S. Batten, A. Chiu, J. Franklin, and M. Sebastia-Barriel. 2014. "The U.K. Productivity Puzzle." *Quarterly Bulletin* (Bank of England) Q2: 114–28.

Bernanke, B. S., M. Gertler, and S. Gilchrist. 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." In *Handbook of Macroeconomics*, Edition 1, Vol. 1C, ed. J. B. Taylor and M. Woodford, 1341–93 (chapter 21). Elsevier.

Blanchard, O., and J. Galí. 2010. "Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment." *American Economic Journal: Macroeconomics* 2 (2): 1–30.

Broadbent, B. 2014. "Unemployment and the Conduct of Monetary Policy in the U.K." Speech given at the Federal Reserve Bank of Kansas City's Annual Economic Symposium, Jackson Hole, Wyoming, August 23.

———. 2015. "Compositional Shifts in the Labor Market." Speech given at the Bank of England's "Understanding the Great Recession: From Micro to Macro" Conference, London, September 23.

Burgess, S., E. Fernandez-Corugedo, C. Groth, R. Harrison, F. Monti, K. Theodoridis, and M. Waldron. 2013. "The Bank of England's Forecasting Platform: COMPASS, MAPS, EASE and the Suite of Models." Working Paper No. 471, Bank of England.

Carney, M. 2014. Speech given at the 146th Annual Trades Union Congress, Liverpool, September 9.

Draghi, M. 2014. "Unemployment in the Euro Area." Speech given at the Federal Reserve Bank of Kansas City's Annual Economic Symposium, Jackson Hole, Wyoming, August 22.

Elsby, M. W. L., and J. C. Smith. 2010. "The Great Recession in the U.K. Labor Market: A Transatlantic Perspective." *National Institute Economic Review* 214 (1): R26–R37.

Erceg, C. J., D. W. Henderson, and A. T. Levin. 2000. "Optimal Monetary Policy with Staggered Wage and Price Contracts." *Journal of Monetary Economics* 46 (2): 281–313.

Faccini, R., S. Millard, and F. Zanetti. 2013. "Wage Rigidities in an Estimated Dynamic, Stochastic, General Equilibrium Model of the U.K. Labor Market." *Manchester School* 81 (S1): 66–99.

Galí, J. 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89 (1): 249–71.

———. 2011. "The Return of the Wage Phillips Curve." *Journal of the European Economic Association* 9 (3): 436–61.

Galí, J., F. Smets, and R. Wouters. 2012a. "Slow Recoveries: A Structural Interpretation." *Journal of Money, Credit and Banking* 44 (s2): 9–30.

———. 2012b. "Unemployment in an Estimated New Keynesian Model." In *NBER Macroeconomics Annual 2011*, Vol. 26, ed.

D. Acemoglu and M. Woodford, 329–60. University of Chicago Press.

Gertler, M., and P. Karadi. 2011. "A Model of Unconventional Monetary Policy." *Journal of Monetary Economics* 58 (1): 17–34.

Gertler, M., and N. Kiyotaki. 2010. "Financial Intermediation and Credit Policy in Business Cycle Analysis." In *Handbook of Monetary Economics*, Edition 1, Vol. 3, ed. B. M. Friedman and M. Woodford, 547–99 (chapter 11). Elsevier.

Gertler, M., L. Sala, and A. Trigari. 2008. "An Estimated Monetary DSGE Model with Unemployment and Staggered Nominal Wage Bargaining." *Journal of Money, Credit and Banking* 40 (8): 1713–64.

Harrison, R., and Ö. Oomen. 2010. "Evaluating and Estimating a DSGE Model for the United Kingdom." Working Paper No. 380, Bank of England.

International Monetary Fund. 2013. "The Dog That Didn't Bark: Has Inflation Been Muzzled or Was It Just Sleeping?" In *World Economic Outlook: Hopes, Realities, Risks* (April). Washington, DC: International Monetary Fund.

Joyce, M., M. Tong, and R. Woods. 2011. "The United Kingdom's Quantitative Easing Policy: Design, Operation and Impact." *Quarterly Bulletin* (Bank of England) Q3: 200–212.

Kass, R. E., and A. E. Raftery. 1995. "Bayes Factors." *Journal of the American Statistical Association* 90 (430): 773–95.

Millard, S. 2015. "The Great Recession and the U.K. Labor Market." Working Paper No. 566 (November), Bank of England.

Mortensen, D., and C. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies* 61 (3): 397–415.

Phillips, A. W. H. 1958. "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957." *Economica* 25 (100): 283–99.

Pinter, G. 2015. "House Prices and Job Losses." Working Paper No. 569 (December), Bank of England.

Shimer, R. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95 (1): 25–49.

Smets, F., and R. Wouters. 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association* 1 (5): 1123–75.

———. 2007. “Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach.” *American Economic Review* 97 (3): 586–606.

Speigner, B. 2014. “Long-term Unemployment and Convexity in the Phillips Curve.” Working Paper No. 519, Bank of England.

Thomas, C. 2008. “Search and Matching Frictions and Optimal Monetary Policy.” *Journal of Monetary Economics* 55 (5): 936–56.

Wu, J. C., and F. D. Xia. 2016. “Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound.” *Journal of Money, Credit and Banking* 48 (2–3): 253–91.

Yellen, J. 2014. “Labor Market Dynamics and Monetary Policy.” Speech given at the Federal Reserve Bank of Kansas City’s Annual Economic Symposium, Jackson Hole, Wyoming, August 22.