MANCHESTER 1824

The University of Manchester

Centre for Growth & Business Cycle Research

CGBCR

Discussion Paper Series

Coping with Uncertainty: Historical and Real-Time Estimates of the Natural Unemployment Rate and the UK Monetary Policy

By

George Chouliarakis

Centre for Growth and Business Cycle Research, Economic Studies, University of Manchester, Manchester, M13 9PL, UK

November 2009
Number 129

The Manchester School Vol 77 No. 4 July 2009
1463–6786 479–511

COPING WITH UNCERTAINTY: HISTORICAL AND REAL-TIME ESTIMATES OF THE NATURAL UNEMPLOYMENT RATE AND THE UK MONETARY POLICY*

by
GEORGE CHOULIARAKIS†
The University of Manchester

The paper derives and compares historical and real-time estimates of the UK natural unemployment rate and shows that real-time estimates are fraught with noise and should be treated with scepticism. A counterfactual exercise shows that, for most of the 1990s, the Bank of England tracked changes in the natural rate relatively successfully, albeit with some recognition lag which, at times, might have led to excessively cautious policy. A careful scrutiny of the minutes of the monetary policy committee meetings reveals that such 'cautiousness' should be taken as evidence of awareness of the real-time informational limitations that monetary policy is facing.

1 Introduction

It is widely accepted that, to attain the objective of low and stable inflation, the monetary policy authority should adjust its policy instrument in response to exogenous inflation shocks and to deviations of output and employment from their natural levels. In particular, when aggregate expenditure and employment fall below their natural levels, or when other disinflationary concerns appear, the central bank should consider adjusting interest rates downwards, whereas when aggregate expenditure and employment exceed their natural levels, or when other inflationary concerns emerge, the central bank should consider tightening monetary policy. In this setting, the natural levels of output and employment, and the associated natural rate of unemployment, are levels at which—absent supply shocks—inflation is stable. If the policy maker were able to obtain an accurate quantitative estimate of them, macroeconomic stabilization policy would be a more straightforward and predictable exercise.

Unfortunately, in practice, monetary policy making is never as straightforward because policy decisions are made under conditions of model,

* Manuscript received 27.5.07; final version received 9.9.08.
† I am grateful to Chandranath Amarasekara, Michael Artis, Ralf Becker, Monica Correa Lopez, Stephen Millard, Denise Osborn, participants at the International Conference of Economic Modeling, Sao Paulo, July 2007, the Royal Economic Society conference, University of Warwick, March 2008, and the referee for helpful comments and suggestions.

Published by Blackwell Publishing Ltd, 9600 Garsington Road, Oxford OX4 2DQ, UK, and 350 Main Street, Malden, MA 02148, USA.

sampling and data uncertainty. More particularly, policy makers may be uncertain about the specification of the model that describes the structural relations of the economy or about the quantitative strength and stability of those relations, e.g. the exact value of the model parameters. They may also be uncertain about the speed and pattern of the dynamic response of the economy to policy actions. Finally, they may be uncertain about the recent past and present state of the economy, because preliminary economic data, upon which policy makers are called to take decisions, are often subject to considerable measurement error and substantial *ex post* revisions. Indeed, uncertainty is so pervasive in monetary policy that Alan Greenspan refers to it as 'not just an important feature of the monetary policy landscape but as the defining characteristic of that landscape'.[1]

The problem and policy implications of uncertainty are especially pronounced in the measurement of output and unemployment gaps using only information available at the time of decision making. On one hand, these resource utilization indicators require the estimation of unobservable and time-varying parameters, such as the natural level of output and the natural rate of unemployment; thus their measurement is subject to model and sampling uncertainty. On the other, unlike historical estimates of output and unemployment gaps whose estimation is based on revised and largely noise-free data, real-time measures of aggregate resource utilization contain considerable noise and are often subject to significant revision. As such, they may provide a misleading picture of the state of the business cycle and lead macroeconomic stabilization to failure.[2]

A dramatic illustration of the perils of resource utilization mismeasurement can be found in the recent US and UK economic history. Specifically, a number of recent studies have convincingly argued that, at the root of the Great Inflation of the 1970s lies the policy makers' failure to take into account the inherent uncertainty surrounding real-time estimates, and the limits that this uncertainty imposes upon macroeconomic stabilization policies. In an important paper, Orphanides (2003b) forcefully shows that the surge in US inflation during the 1970s could be attributed, at least in part, to an unfortunate pursuit of activist policies in the face of output gap mismeasurement, the latter being the result of slow recognition of the productivity slowdown of the late 1960s and early 1970s. Likewise, Nelson and Nikolov (2003) show that monetary policy errors due to real-time output gap mismeasurement contributed 3–7 percentage points to average UK inflation in the 1970s and 1–5 percentage points to average UK inflation in the 1980s.

[1]See Greenspan (2004, p. 36).

[2]For a discussion of the problems of sampling and model uncertainty in the estimation of natural rates of unemployment, see Staiger *et al.* (1997a, 1997b) and Laubach (2001); for a discussion of the problem of data uncertainty, see Orphanides and Van Norten (2002) and Orphanides and Williams (2002).

The contribution of the Federal Reserve to the remarkable stability of the US economy during the second half of the 1990s offers yet another example of the importance of recognizing the real-time informational problem when setting policy. Blinder and Yellen (2001) describe the US monetary policy during 1996–99 as a policy of 'forbearance' and argue that 'one important question for macroeconomic historians is why the Fed chose to forebear even as the unemployment rate drifted down to levels that previously had been associated with accelerating inflation'.[3] Their answer suggests that the Greenspan Fed was able to detect early signs of a productivity surge by scrutinizing and successfully decoding a plethora of information, such as the stability of prices despite a significant rise in hourly labour compensation and profit margins or the unusual rise in high-tech capital goods orders. More importantly, the Fed was prepared to acknowledge that the likely improvement in trend productivity growth heightened the uncertainty surrounding preconceived estimates of the natural rate of unemployment. In the face of such uncertainty, the Federal Reserve refrained from responding aggressively to falling unemployment, as they would have done otherwise.[4] In Blinder and Yellen's words 'given increased uncertainty about structural shifts, Greenspan was able to persuade his committee not to attack inflation pre-emptively but rather to wait to shoot until they could see the whites of rising inflation's eyes. And that, of course, never happened.'[5] The benefits of a policy that was prepared to acknowledge imperfect knowledge and real-time informational limitations in the face of structural economic changes were non-trivial. Sinai (2004) compares the actual policy of 'forbearance' with a counterfactual simulation in which the Federal Reserve tightened monetary policy during 1996–99 by gradually raising the federal funds rate from $5\frac{1}{2}$ to $7\frac{3}{4}$ percentage points. He shows that the actual policy delivered stable inflation with an average annual unemployment rate of 4.8 per cent against a counterfactual average unemployment rate of 5.7 per cent.

With the advantage of hindsight, one can safely say that, during the second half of the 1990s, the UK economy also witnessed a substantial fall in its natural unemployment rate, although the origins of this fall may not have been as apparent as in the case of the USA.[6] In this context, it is worth asking how successfully the UK monetary policy authorities dealt with the increased

[3]See Blinder and Yellen (2001, p. 115). The term forbearance was coined to describe the patient behaviour of the US monetary policy authority in the face of historically low unemployment rates and fast growth.

[4]Greenspan (2004) clarifies that the Fed's stance during 1996–99 was 'motivated, in part, by the view that the evident structural economic changes rendered suspect, at best, the prevailing notion in the early 1990s of an elevated and reasonably stable non-accelerating inflation rate of unemployment'.

[5]See Blinder and Yellen (2001, p. 128).

[6]Note that the UK economy did not experience productivity growth acceleration.

uncertainty surrounding the value of the natural rate of unemployment at a period of structural economic change, and what kind of lessons one can draw from it. The goal of the paper is to address this question. More particularly, the paper intends to assess the extent of the problem of real-time data uncertainty that the UK policy makers might have faced while trying to judge the underlying state of the labour market, focusing especially on uncertainty surrounding real-time labour market quantity indicators and measures of resource utilization. It also intends to assess the cost of uncertainty by examining how different monetary policy would have been under the assumption of perfect information, i.e. if the UK monetary policy authorities had the wisdom of hindsight.

To this end, the paper aims first to derive alternative historical estimates of the natural rate of unemployment using a range of univariate and multivariate methods. These estimates are based on recent data releases and could not have been available to policy makers at the time decisions were made. Instead, they can be thought of as reflecting a retrospective or 'ultimate' state of knowledge about the supply potential of the UK economy and as such they will be used in this paper. Second, in light of the considerable dispersion of alternative estimates at any point in time, the paper uses pseudo out-of-sample forecasts to find out which of these statistical methods delivers natural rates that are relatively superior in understanding subsequent movements in inflation. Third, the paper uses the preferred estimate of the natural rate of unemployment in the context of a counterfactual exercise to assess whether perfect information of the economy's supply potential at the time of decision making might have led to a considerably different interest rate path. The results of the counterfactual exercise are then interpreted in the light of the minutes of the monetary policy committee meetings. Finally, the paper uses vintages of labour market quantity indicators and inflation rates to derive vintages of natural rates of unemployment and unemployment gaps as they would have been available to policy makers when decisions were made. The paper then examines whether these data contain reliable information that could have helped policy makers to gauge the 'true' underlying state of the labour market in real time. The findings are not encouraging. They suggest that real-time natural rates of unemployment and unemployment gaps contain substantial noise and should be treated with scepticism. The paper also finds that, since the adoption of inflation targeting, the Bank of England has tracked the secular changes of the natural rate of unemployment relatively successfully, albeit with some recognition lag which, at times, might have led to an excessively cautious policy. A careful examination of the minutes of the monetary policy committee meetings reveals that this 'cautiousness' should be interpreted as evidence of recognition of the real-time informational limitations that monetary policy is facing, especially at times of structural economic change.

The remainder of the paper is organized as follows. Section 2 derives alternative time-varying natural rates of unemployment for the UK during 1960–2003, using a range of univariate and multivariate methods, and discusses the results. Section 3 discriminates between competing methodologies by assessing the power of the different natural rate estimates to 'predict' inflation. Section 4 turns to the monetary policy of the Bank of England and the problem of uncertainty in real-time labour market quantity indicators and Section 5 concludes.

2 Historical Estimates of the Natural Rate of Unemployment of the UK

Below I will derive six alternative measures of the natural rate of unemployment of the UK since 1960. These measures correspond to alternative specifications of two basic statistical methodologies that have been used extensively in the literature, i.e. univariate filters and multivariate unobserved components models. A third way of estimating the time-varying natural rate of unemployment would rely on the estimation of a system of wage-setting and price-setting equations and would require the prior measurement of the factors that determine it, such as the degree of centralization of wage bargaining, the extent of coordination between employers and workers, or the degree of mismatch between workers and vacancies. Recent research has suggested that estimates of the natural rate based on such labour market models are sensitive to small changes in the specification and measurement of those variables; hence they are usually received with scepticism.[7] In light of this, I will not consider this approach here.

2.1 Univariate Filter Estimates of the Natural Rate

The natural rate of unemployment is here defined as the rate of unemployment at which, in the absence of supply shocks, inflation is stable. It is well known that, on the assumption of imperfections of information or in the presence of staggered wage and price setting, demand shocks have real short-run effects.[8] More particularly, positive demand shocks may reduce unemployment below the natural rate at the cost of rising inflation whereas negative demand shocks may increase unemployment above the natural rate at the expense of falling inflation. Yet, in the long run, once expectations have fully adjusted to changes in inflation, actual unemployment converges to its time-varying natural rate. It is therefore possible to consider the natural rate of unemployment as a slowly evolving trend from which actual unemployment may temporarily deviate, thus a trend that could be estimated using standard univariate filters. Univariate filters have been widely used in the

[7]For a discussion of the drawbacks of the structural approach, see Cassino and Thornton (2002).
[8]For a survey of the literature, see Romer (2006).

trend–cycle decomposition of a number of macroeconomic time series and possess the advantage of imposing very little structure on the problem at hand. Below, we will consider two versions of each of the two most popular univariate filters, i.e. the Hodrick and Prescott (1997) filter and the approximation of an ideal low-pass filter as described by Baxter and King (1999).

The Hodrick–Prescott (HP) filter views a given time series as the sum of a transitory cyclical component and a trend component that varies smoothly over time. An estimate of the latter is derived by minimizing the variance of the cyclical component subject to a penalty on the variability of the trend component. Naturally, the larger the value of the parameter that penalizes the variability of the trend component, the smoother is the path of this component and the greater is the standard deviation and the degree of persistence of the cyclical component. The value of the penalty (or smoothness) parameter is here given exogenously and Hodrick and Prescott recommend a value of $\lambda = 1600$ when using data with quarterly frequency. In estimating the natural rate of unemployment of the UK, I will consider two alternative implementations of the filter, one with the smoothness parameter recommended by Hodrick and Prescott and another with a smoothness parameter of $\lambda = 25{,}600$ which allows for greater persistence of cyclical unemployment and is closer in spirit to the approach advocated by Rotemberg (1999).

Baxter and King's approximation of the ideal low-pass filter is a two-sided moving average that 'passes through' the low-frequency trend component of an economic time series while suppressing components whose frequency is associated with cyclical or irregular variation. The application of the low-pass filter requires the specification of a cut-off frequency below which all frequencies are 'passed through'. Following the pioneering work of Burns and Mitchell (1946) on the US business cycle, Baxter and King adopt a cut-off frequency of $\pi/16$ but their low-pass filter is flexible enough to allow researchers to set the cut-off frequency they consider appropriate for their specific purpose. We will examine two alternative implementations of the low-pass filter, one with the cut-off frequency favoured by Baxter and King and another with a cut-off frequency of $\pi/30$, used by Staiger *et al.* (2001), i.e. a low-pass filter that admits time-series components with periodicity of 60 quarters or more. The underlying idea here is that, in the presence of unemployment persistence mechanisms, cyclical unemployment is more persistent than output deviations from trend.

2.2 Multivariate Unobserved Components Estimates of the Natural Rate

Unlike univariate methods of trend–cycle decomposition, which do not make use of information provided by economic models of the natural rate, the multivariate unobserved components method derives estimates of the natural rate based on the accelerationist Phillips curve. More particularly, the method uses the Kalman filter technique, a recursive procedure that, combined with the maximum likelihood estimation method, estimates the unobserved natural rate of unemployment and the slope coefficients of the accelerationist Phillips curve simultaneously. In this framework, the natural rate is modelled as a stochastic time-varying parameter, with restrictions imposed on its volatility to avoid implausible variation and make the estimated series conform to our understanding of the natural rate as a slowly moving variable. Below, I will obtain estimates of the UK time-varying natural rate of unemployment based on two alternative specifications of this method, a standard specification, i.e. a specification widely used in the estimation of unobserved components, and a hybrid specification introduced by Staiger *et al.* (2001).

2.2.1 The Standard Specification. The key components of the standard specification, from which the behaviour of the unobserved variable is elicited, are a model of inflation dynamics and the stochastic process that generates the unobserved variable. A few comments about these two components are worth making.

First, the standard model of inflation dynamics suggests that inflation is determined by inflation inertia, the unemployment gap and temporary supply shocks. Specifically, the model implies that at steady state, i.e. when actual unemployment is on the natural rate and there are no supply shocks, actual inflation is stable; yet, if unemployment deviates from its natural rate or if temporary supply shocks occur, changes in inflation will set in. In particular, if and for as long as unemployment stays above the natural rate, inflation will be falling, whereas if and for as long as unemployment stays below the natural rate, inflation will be rising. The model can be illustrated as follows:

$$\pi_t = a(L)\pi_{t-1} + b(L)(u_t - u_t^*) + c(L)s_t + e_t \tag{1}$$

where π denotes actual inflation, u denotes the observed unemployment rate, u^* denotes the unobserved time-varying natural unemployment rate, s denotes a vector of temporary supply shock variables, L is a polynomial in the lag operator, and the disturbance e is assumed to be independent identically distributed normal with zero mean and variance σ_e^2.

The inflation inertia term in equation (1) captures backward-looking expectations or sluggish price adjustment due to costs in obtaining and processing information on the part of price setters.[9]

The vector of temporary supply shock variables includes those supply shocks that might be reasonably expected to revert to zero over a horizon of one to two years, such as changes in real oil price deflators, real import price deflators or real effective exchange rates. Supply-side shocks of a more permanent nature, such as changes in labour market institutions, are thought to

[9] See Mankiw and Reis (2002).

affect sustainable employment and hence are captured by variations in the natural rate of unemployment.

Second, the stochastic process that determines how the natural rate varies over time is usually assumed to be a random walk, i.e.

$$u_t^* = u_{t-1}^* + v_t \tag{2}$$

where v_t is a random error with zero mean and variance σ_v^2.

Equations (1) and (2) contain all necessary economic information for the estimation of the time-varying natural rate of unemployment following the standard specification of the multivariate unobserved components method. The specification has been used by, among others, Gordon (1997, 1998) and Batini and Greenslade (2006) for the estimation of the natural rate of unemployment of the USA and UK, respectively.

2.2.2 The Hybrid Specification. The hybrid specification uses a two-step procedure to estimate the time-varying natural rate. First, a univariate trend of unemployment is derived from a low-pass filter and, second, the natural rate of unemployment is estimated using a multivariate model within which the univariate trend of unemployment serves as an input. To the extent that a considerable part of the time variation in the natural rate is likely to be reflected in changes in the univariate unemployment trend, the additional information contained in this trend can potentially contribute in delivering more precise estimates of the natural rate of unemployment.

More particularly, the hybrid specification defines the potential deviation of the natural rate from the univariate trend of unemployment as

$$\xi_t = u_t^* - u_t^{\mathrm{N}} \tag{3}$$

where, as above, u^* denotes the natural rate of unemployment and u^{N} denotes the univariate trend of unemployment.

Accordingly, the equation of inflation dynamics (1) can be rewritten as

$$\pi_t = \mu_t + a(L)\pi_{t-1} + b(L)(u_t - u_t^{\mathrm{N}}) + c(L)s_t + e_t \tag{4}$$

where μ_t $(= -b(L)\xi_t)$ is a time-varying intercept and all other variables are defined as in (1).

If the slope coefficients of equation (4) are stable, the time-varying intercept will simply reflect the departure of the natural rate from the univariate trend of unemployment and can be estimated by using the multivariate unobserved components method. To implement the procedure, the intercept drift is modelled as a stochastic time-varying parameter that follows a random walk, i.e.

$$\mu_t = \mu_{t-1} + \eta_t \tag{5}$$

where η_t is independent identically distributed $\mathrm{N}(0, \sigma_\eta^2)$.

Estimates of the time-varying natural rate can thus be obtained by combining the univariate trend and the estimated intercept drift, i.e.

$$u_t^* = u_t^{\mathrm{N}} - \left(\hat{\mu}_t / \hat{b}\right) \tag{6}$$

where $\hat{b}$ is the estimator of b and $\hat{\mu}$ is the estimator of μ.

The hybrid specification has been used by Staiger *et al.* (2001) for the estimation of the time-varying natural rate of unemployment of the USA.

2.2.3 Estimation Issues. The standard specification formed by equations (1) and (2) and the hybrid specification formed by equations (4) and (5) can be written in state–space form and estimated by maximum likelihood using the Kalman filter.[10] Yet, before proceeding, some issues pertaining to the estimation of the two multivariate unobserved components models are worth highlighting.

First, to ensure the long-run neutrality of inflation with respect to unemployment, one has to impose on (1) and (4) the restriction that the sum of coefficients on lagged inflation equals unity. Although the estimated sum of coefficients on lagged inflation, over the period under consideration, approaches unity, a meaningful calculation of the natural rate requires that the sum should be constrained to be exactly equal to unity. Following Staiger *et al.* (1997a) I impose this restriction by casting and estimating the two models in first differences—rather than levels—of inflation. Second, the vector of supply shock variables includes the log change of real import price and real oil price deflators, shocks that are expected to have only a transitory effect on domestic inflation. The supply shock variables are standardized so that they have a zero mean and a standard deviation of one. Third, a number of studies have noted that inflation may be affected by whether unemployment is rising or falling—not just by the level of unemployment relative to the natural rate. To allow for, and measure directly, possible speed limit effects, the unemployment gap variable enters the model in level as well as in difference form. Fourth, during the late 1960s and 1970s, there have been several attempts to contain inflation by means of incomes policies. To capture their potential effect, a number of dummy variables are included. In particular, following Henry (1981) and Flanagan *et al.* (1983), I include intercept dummies for two instances of statutory incomes policies: the 12-month freeze on wages and prices of the Wilson administration during 1966Q3–1967Q2 and the social contract of the Wilson and Callaghan administrations during 1975Q3–1977Q2. I also include intercept dummies to capture the possibility of 'catch-up' wage explosions following the relaxation of incomes policies. The periods covered by the latter dummies are of the same length as the

[10]See Hamilton (1994).

periods for which the incomes policies were in operation. A final specification issue concerns the selection of lags in the polynomials $a(L)$, $b(L)$ and $c(L)$. All variables enter initially with eight lags and are then tested down on the basis of significance. Given these choices, the standard and hybrid specifications of the unobserved components method are, respectively,

$$\Delta\pi_t = a(L)\,\Delta\pi_{t-1} + b'(u_t - u_t^*) + b''\Delta u_t + c(L)\,s_t + d_t + e_t \tag{7}$$

where $e_t \sim \mathrm{N}(0, \sigma_e^2)$,

$$u_t^* = u_{t-1}^* + v_t \tag{8}$$

where $v_t \sim \mathrm{N}(0, \sigma_v^2)$, and

$$\Delta\pi_t = \mu_t + a(L)\,\Delta\pi_{t-1} + b'(u_t - u_t^{\mathrm{N}}) + b''\Delta u_t + c(L)\,s_t + d_t + e_t \tag{9}$$

where $e_t \sim \mathrm{N}(0, \sigma_e^2)$,

$$\mu_t = \mu_{t-1} + \eta_t \tag{10}$$

where $\eta_t \sim \mathrm{N}(0, \sigma_\eta^2)$.

Before the Kalman filter procedure starts, the vector of parameters, including the time-varying natural rate of unemployment, needs to be initialized. Initial parameter values for the slope coefficients of equations (7) and (9) were obtained from their ordinary least squares estimation, in which the time-varying natural rate was proxied by a univariate unemployment trend.[11] Also, the initial condition of the state variable in the standard specification of the model is set equal to the first observation of the univariate unemployment trend while the initial condition of the state variable in the hybrid specification of the model is set equal to zero. It is worth noting that the unobserved components model estimates of the natural unemployment rate are robust to alternative parameter starting values and initial conditions.

To implement the procedure, one needs to restrict the value of the signal-to-noise ratio. The signal-to-noise ratio is the ratio of the variance of the error term of the stochastic process that drives the natural rate to the variance of the error term in the inflation dynamics equation and determines the smoothness of the resulting time-varying natural rate series. In particular, a high signal-to-noise ratio implies that a large part of the residual variation in the inflation equation is soaked up by the natural rate; hence the natural rate series is quite volatile. On the other extreme, a signal-to-noise ratio equal to zero implies a constant natural rate. In line with previous studies, I experiment with alternative values of the signal-to-noise ratio, all of which permit the natural rate to move about as much as it likes without though exhibiting high frequency variation.

[11]Low-pass filter with cut-off frequency $\pi/30$.

2.3 Empirical Results

Estimates of the natural rate of unemployment are obtained using UK quarterly data from 1960Q1 to 2003Q4. Columns 3 and 4 of Table 1 report results obtained from the estimation of the standard and hybrid specifications of the unobserved components model respectively. In both cases the signal-to-noise ratio is set to 0.09, an entirely standard value in the estimation of natural unemployment rates, yet below I consider the robustness of the estimated natural rates to alternative values of the ratio. The following findings are worth highlighting. First, all regressors are signed as expected with the unemployment gap having a significant negative effect and changes in real import price inflation a significant positive effect on changes in domestic price inflation. The unemployment gap coefficients, in particular, range between –0.17 and –0.30, an entirely plausible range implying that the impact effect of a 1 per cent fall in the unemployment gap is a rise in annualized inflation of 0.68–1.20 per cent. In the presence of real import price inflation, real oil price inflation is statistically insignificant and therefore not included, a rather expected result given that changes in real import price inflation already capture the effect of oil price shocks on domestic inflation. Furthermore, the

Table 1
Estimation of Inflation Dynamics Equations Using the Kalman Filter, UK 1960Q1–2003Q4

Dependent variable: $\Delta\pi$			
Regressors	*Lags*	*Standard specification*	*Hybrid specification*
$\Delta\pi$	1	–0.63 (0.07)	–0.65 (0.07)
	2	–0.60 (0.09)	–0.64 (0.09)
	3	–0.58 (0.08)	–0.62 (0.09)
	4	–0.19 (0.11)	–0.25 (0.11)
	5	–0.15 (0.01)	–0.20 (0.11)
Unemployment gap	1	–0.17 (0.08)	–0.30 (0.11)
ΔUnemployment	1	–0.34 (0.27)	–0.27 (0.28)
ΔRMP inflation	1	0.08 (0.06)	0.08 (0.06)
	2	0.18 (0.06)	0.17 (0.06)
	3	0.18 (0.07)	0.17 (0.07)
	4	0.19 (0.06)	0.18 (0.06)
	5	0.21 (0.05)	0.21 (0.05)
	6	0.12 (0.06)	0.13 (0.06)
	7	0.15 (0.05)	0.15 (0.05)
	8	0.17 (0.04)	0.17 (0.04)
Incomes policy 1966/67		–0.31 (0.46)	–0.37 (0.45)
Incomes policy 1975/77		–0.40 (0.24)	–0.32 (0.24)
Dummy 1		–3.38 (0.44)	–3.37 (0.52)
Dummy 2		1.64 (0.61)	1.67 (0.58)
Log likelihood		606.8	609.1

Notes: Standard errors in parentheses; Δ is the first difference operator; RMP denotes real import prices; Dummy 1 takes the value of –1 for 1979Q3 and 1 for 1980Q3, zero otherwise; Dummy 2 takes the value of 1 for 1974Q4 and 1975Q1, zero otherwise.

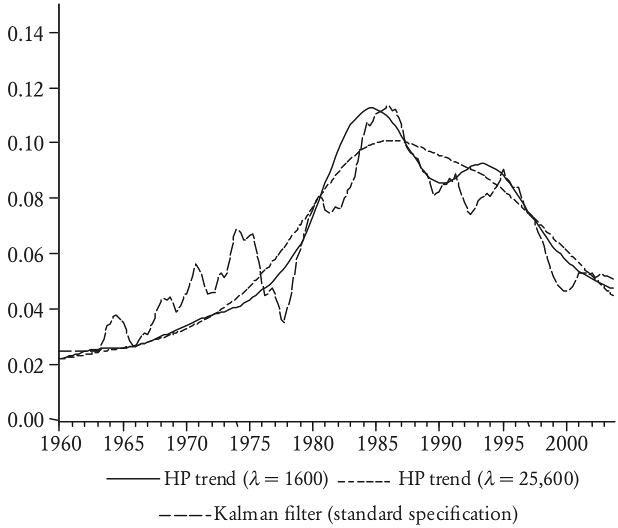

FIG. 1 Historical Estimates of Time-varying Natural Rates, UK 1960Q1–2003Q4

speed limit effect is correctly signed but not statistically significant, perhaps suggesting that the unemployment persistence mechanisms are rather weak. The incomes policy dummies are correctly signed but not significant while the 'catch-up' wage explosion dummies were wrongly signed and insignificant and therefore are not reported.

Figures 1 and 2 show the time profiles of time-varying natural rates estimated using alternative univariate and multivariate methods. Regardless of the method, all measures reveal the same overall pattern. The natural rate of unemployment was low during the first half of the 1960s floating around 3 per cent. It drifted upwards from the late 1960s to the mid-1980s stabilizing towards the end of the decade around 8 per cent. It declined steadily during the second half of the 1990s and, at the turn of the millennium, it was standing close to 5 per cent. The causes of the rise and fall of the natural rate have been discussed extensively elsewhere. The rise of the natural rate is often ascribed to the interaction of adverse macroeconomic shocks and labour market institutions that propagate those shocks over time, such as the intermediate degree of centralization of wage bargaining, the absence of coordination between employers and trade unions, and the bias of the system of unemployment insurance in favour of passive and against active support.[12] Likewise, the fall of the natural rate is ascribed to the benign macroeconomic environment of the past 15 years coupled with the implementation of some key labour market reforms.[13]

[12]See Layard *et al.* (2005).
[13]For a detailed discussion see Nickell (2001).

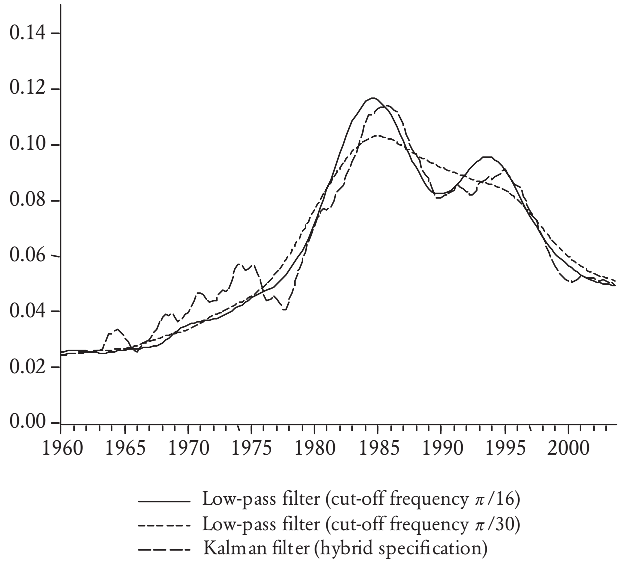

FIG. 2 Historical Estimates of Time-varying Natural Rates, UK 1960Q1–2003Q4

The bell-shaped pattern of the natural rate is robust to a number of specification changes, some of which are especially worth reporting. First, I re-estimate the unobserved components model using alternative values of the signal-to-noise ratio. In particular, I set the ratio to 0.04 and 0.16, i.e. I allow the quarterly variation of the time-varying natural rate to equal 20 and 40 per cent respectively of the residual variation of the inflation dynamics model. The resulting estimates of the time-varying natural rates, together with the original estimate, are reported in Fig. 3(a). As expected, the estimated series differ somewhat between themselves—lower signal-to-noise ratios produce relatively flatter estimates—yet the differences are very small and the overall pattern does not change.

A second sensitivity test is to consider the robustness of the estimated time-varying natural rates to alternative definitions of price inflation. To do so, I re-estimate the unobserved components model using inflation rates based alternatively on the retail price and consumer price indices. The resulting natural rate estimates, shown in Fig. 3(b), are broadly similar to the original one. A third specification check is to consider the robustness of the results to alternative measures of unemployment. I do so by replacing the UK labour force survey (LFS) unemployment rate with the OECD standardized unemployment rate. Figure 3(c) shows that the resulting natural rate estimates remain very close to the original one. A fourth sensitivity test is to use an alternative model of how the natural rate varies over time. In particular,

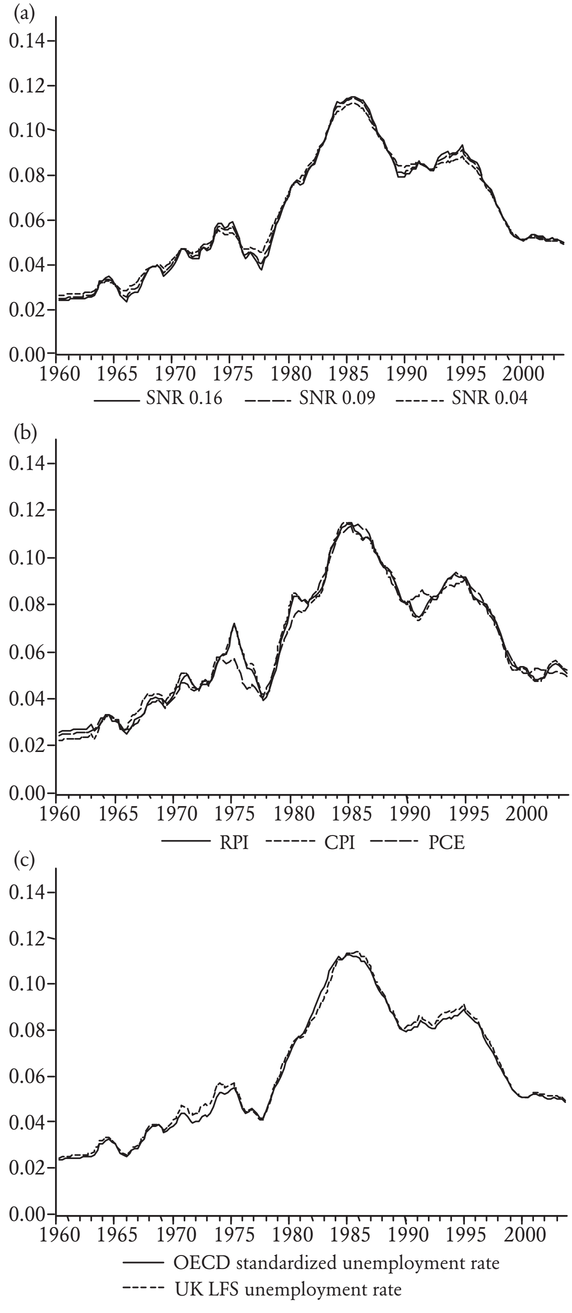

Fig. 3 Sensitivity Analysis, Kalman Filter Estimates of Natural Rates: Hybrid Specification
SNR, signal-to-noise ratio; RPI, retail price index; CPI, consumer price index; PCE, final consumption expenditure deflator.

TABLE 2
HISTORICAL ESTIMATES OF THE NATURAL UNEMPLOYMENT RATE OF THE UK

Method	*1964Q1*	*1974Q1*	*1984Q1*	*1994Q1*	*2000Q1*
HP trend ($\lambda = 1600$)	2.6	4.1	11.2	9.2	5.7
HP trend ($\lambda = 25{,}600$)	2.5	4.4	9.8	8.6	6.1
Low-pass filter (cut-off frequency $\pi/16$)	2.5	4.2	11.5	9.5	5.7
Low-pass filter (cut-off frequency $\pi/30$)	2.6	4.3	10.2	8.5	6
Kalman filter (standard specification)	3.4	6.9	10.3	8.1	4.6
Kalman filter (hybrid specification)	3.2	5.7	10.8	8.8	5.1
Memoranda					
Median of estimates	2.6	4.3	10.5	8.7	5.7
Range of estimates	2.5–3.4	4.1–6.9	9.8–11.5	8.1–9.5	4.6–6.1
Actual unemployment rate	2.9	3.6	11.8	9.9	5.8

I model the change in the natural unemployment rate as a persistent stationary process whereby, following a shock, the natural rate converges to its new steady state gradually, i.e.

$$\Delta u_t^* = \phi \Delta u_{t-1}^* + \upsilon_t \tag{11}$$

where $0 \leq \phi < 1$ and υ_t is a random error with zero mean and variance σ_υ^2. I re-estimate the standard specification of the unobserved components model using the above stochastic process as state equation. A Wald test on the ϕ coefficient fails to reject the null hypothesis that the coefficient equals zero, i.e. it fails to reject the null of a random walk process.[14]

Table 2 reports estimates of the natural rate, at various points in time, and shows that differences across methods can be wide.

Differences across methods are also reflected clearly in the time profile of unemployment gaps. Figures 4 and 5 report unemployment gaps constructed on the basis of the six alternative estimates of the natural rates derived above.

The unemployment gap series are positive during the early 1980s and early 1990s, i.e. in the aftermath of negative demand shocks, and negative during the early 1970s and late 1980s, i.e. during periods of excess demand such as the years of the Barber and Lawson booms. Yet, despite similarities in the overall pattern, the differences between the unemployment gap series are quite striking, demonstrating that, even with the advantage of hindsight, there remains significant uncertainty about the past state of the economy. In principle, because multivariate methods bring additional information to bear on the decomposition of trend and cycle, one can argue that they provide more accurate estimates of natural rates assuming that the underlying model is correctly specified. Yet, in practice, given uncertainty about model specification, especially uncertainty regarding the proper modelling of the behaviour of the unobserved component, the advantage of multivariate methods

[14]Wald chi-square statistic 2.4 (p value 0.12).

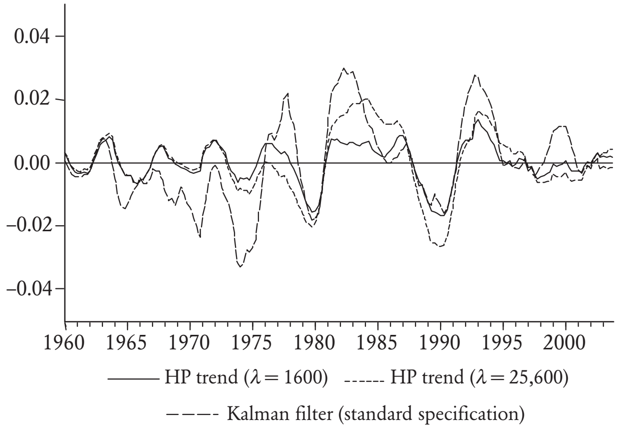

Fig. 4 Unemployment Gaps

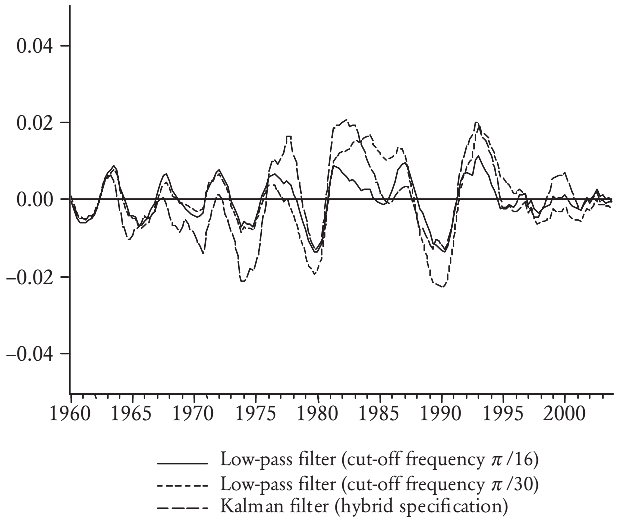

Fig. 5 Unemployment Gaps

may be largely illusory. One way to deal with this uncertainty is to discriminate between the natural unemployment estimates and find out which one is more accurate. It is to this that I now turn.

3 Assessment of Alternative Natural Rate Estimates

The aim of this section is to determine which of the above methods delivers more accurate natural rate estimates. As the natural rate of unemployment, and the associated unemployment gap, is an important input in models of inflation dynamics, a good way to do so is on the basis of the power of alternative natural rate estimates to understand subsequent movements in inflation. For this purpose, I generate pseudo out-of-sample forecasts of changes in inflation using alternative natural rate estimates and compare their forecasting performance using their root mean squared forecast error. I also consider pseudo out-of-sample forecasts of changes in inflation from a model of inflation dynamics with a constant natural rate which serve as a benchmark.

As observed differences in root mean squared errors across equations are not necessarily statistically significant, it is important to use a formal statistical procedure to test them. I do so using the test of equality of predictive accuracy proposed by Diebold and Mariano (1995). The Diebold–Mariano test procedure is designed to test the null hypothesis of equality of predictive accuracy by considering the mean loss differential that is constructed from the mean squared errors of pairs of competing models. Under the null hypothesis, this mean, suitably normalized, has an asymptotic standard normal distribution. For small sample sizes, Diebold and Mariano recommend the use of finite-sample tests of predictive accuracy to complement theirs. In line with their recommendation, I also use a standard sign test.[15]

Column 2 of Table 3 reports the root mean squared errors obtained from one-step-ahead pseudo out-of-sample inflation forecasts. The root mean squared errors are based on forecasts using the same equation (7) but each of them corresponds to a different measure of the natural rate of unemployment. *Prima facie*, the results suggest that models based on Kalman filter estimates of the natural rate outperform those based on univariate methods. They also suggest that models based on univariate methods fail to outperform even models with constant natural rates. Columns 3 and 5 of Table 3 examine to what extent the superior forecasting performance of models based on Kalman filter estimates of the natural rates is statistically relevant. More particularly, column 3 reports the Diebold–Mariano test statistic and column 5 the sign test statistic derived from comparisons of all models against the inflation dynamics model whose natural rate is estimated using the hybrid specification of the multivariate method. Overall, the superior performance of the latter against models based on univariate method estimates of the natural rate is highly statistically significant whereas its

[15]Noting that, for small and moderate-sized samples, the Diebold–Mariano test is over-sized, Harvey *et al.* (1997) propose a modification in order to improve the test's finite-sample performance. The main modification involves using an approximately unbiased estimator of the variance of the mean of squared error differences. Also, they propose comparing their modified Diebold–Mariano statistic with critical values from the Student's t distribution with $n - 1$ degrees of freedom rather than from the standard normal distribution. The results of the Harvey, Leybourne and Newbold test were only marginally different from those of the Diebold–Mariano test and thus I do not report them.

TABLE 3
FORECASTING PERFORMANCE OF ALTERNATIVE NATURAL RATE ESTIMATES

		Diebold–Mariano statistic		*Sign test statistic*	
		Benchmark		*Benchmark*	
Method	*RMSE*	*Hybrid*	*Constant*	*Hybrid*	*Constant*
HP trend ($\lambda = 1600$)	0.294	3.063	–0.188	1.633	–1.225
		(0.002)	(0.850)	(0.102)	(0.221)
HP trend ($\lambda = 25600$)	0.294	3.047	–0.201	2.041	–0.816
		(0.002)	(0.840)	(0.041)	(0.414)
Low-pass filter (cut-off frequency $\pi/16$)	0.286	3.195	–0.697	1.633	–1.633
		(0.001)	(0.486)	(0.102)	(0.102)
Low-pass filter (cut-off frequency $\pi/30$)	0.300	3.039	0.524	2.858	0.408
		(0.002)	(0.600)	(0.004)	(0.683)
Kalman filter (standard specification)	0.223	–1.117	–3.022	–1.225	–2.449
		(0.264)	(0.002)	(0.221)	(0.014)
Kalman filter (hybrid specification)	0.231	—	–2.906	—	–2.449
		—	(0.004)	—	(0.014)
Constant natural rate	0.296	2.906	—	2.449	—
		(0.004)	—	(0.014)	—

Notes: Root mean squared errors (RMSEs) based on one-step-ahead pseudo out-of-sample inflation forecasts; estimation sample 1960Q1–1997Q4; forecast sample 1998Q1–2003Q4; *Benchmark* refers to the model against which all other models are compared; p values in parentheses.

performance against the model based on natural rate estimates using the standard specification of the multivariate method is insignificant. Thus, the results confirm the superior forecasting performance of the models based on the Kalman filter method. Columns 4 and 6 of Table 3 report the Diebold–Mariano and sign test statistics respectively derived from comparisons of all models against the inflation dynamics model with a constant natural rate. The results confirm that the hypothesis of equality of predictive accuracy between models based on univariate estimates of the natural rate and the model with a constant natural rate cannot be rejected.

In a similar exercise, using US data, Orphanides and Williams (2002) failed to detect differences in the forecasting performance across models with different estimates of the natural rate and concluded that 'inflation forecasting accuracy is virtually identical across the specifications that include the unemployment gap'. The above results suggest that this is not true for the UK. In particular, the results show that one cannot reject the hypothesis that models of inflation dynamics based on multivariate estimates of the time-varying natural rates have superior forecasting performance to those based on univariate estimates of natural rates or those based on constant natural rates. To the extent that the power to 'predict' inflation is considered as one of the key features of the natural rate of unemployment, the multivariate estimates of the natural rate are clearly superior.

4 A Taylor Rule with a Time-varying Natural Rate

Starting with Taylor (1993) a substantial literature argues that the US monetary policy of the past 20 years can be well explained by a simple interest rate reaction function where interest rates respond to deviations of inflation from its target as well as to aggregate slack. Ball and Tchaidze (2002) and Blinder and Reis (2005), among others, show that a simple Taylor rule based on a time-varying natural rate can track closely the US monetary policy of the 1990s, suggesting that the policy makers' 'forbearance' in face of falling unemployment and a booming economy reflected simply a real-time awareness of the steady decline of the natural rate of unemployment. To the extent that 'forbearance' in face of falling unemployment has been a key component of the remarkable economic record of the USA during the 1990s, the American Federal Reserve's early recognition of the steady decline of the natural rate should claim part of the credit.[16]

As shown above, the UK economy too witnessed a significant and sustained decline in its natural rate of unemployment during the second half of the 1990s. Indeed, in retrospect one can say with confidence that towards the end of the decade the UK economy could sustain rates of unemployment not seen since the early 1970s. How successful have the UK monetary policy authorities been in recognizing early, as the US Federal Reserve did, the structural changes that led to the steady fall of the natural rate? Or, to put it differently, how different would the interest rate path have been if policy makers had perfect information about the natural rate, i.e. if they knew in real time what we learnt in hindsight?

Recent evidence suggests that, since the adoption of inflation targeting, a simple interest rate reaction function comes close to mimic the behaviour of the Bank. Nelson (2000) estimated interest rate reaction functions for each UK monetary policy regime since 1972 and concluded that monetary policy during inflation targeting is characterized by a reaction function with slope coefficients not dissimilar to those proposed by Taylor. If so, a simple Taylor rule based on historical estimates of the time-varying natural rate could describe how policy makers would have behaved if they had observed in real time the actual changes in the natural rate. It can thus serve as benchmark against which some preliminary answers to the above questions can be drawn.

One caveat is in order. The actual interest rate might deviate from the one prescribed by the above Taylor rule for reasons other than real-time uncertainty about the true state of the economy, e.g. instability in the slope

[16]Ball and Tchaidze (2002) and Blinder and Reis (2005) argue that the ability of the American Federal Reserve to recognize changes in the natural rate of unemployment was due to an early appreciation of the emerging new economy and the ensuing acceleration of productivity growth. For the potential effect of productivity growth on the natural rate of unemployment see Ball and Moffitt (2001) and Ball and Mankiw (2003).

coefficients of the rule. Any indication therefore of systematic deviations between the two should be interpreted carefully and corroborated by additional evidence. I do so by scrutinizing the minutes of the monetary policy committee meetings and speeches of members of the committee. Before proceeding, some issues pertaining to the calibration of the natural-rate-based Taylor rule are discussed.

The interest rate reaction function that I consider is a standard Taylor rule whereby the short-term nominal interest rate responds positively to deviations of actual inflation from target and to the output gap, i.e.

$$i_t = \pi_t + \omega_1(\pi_t - \pi^*) + \omega_2(y_t - y_t^*) + r_t^* \tag{12}$$

or using the unemployment gap as a measure of economic slack

$$i_t = \pi_t + \omega_1(\pi_t - \pi^*) - \frac{\omega_2}{\phi}(u_t - u_t^*) + r_t^* \tag{13}$$

where i denotes the short-term nominal interest rate, π^* denotes the inflation target of the Bank, y denotes the log of real output and y^* its potential level, r^* denotes the natural rate of interest, ω_1 and ω_2 are slope coefficients and ϕ is Okun's coefficient.

In effect, the standard Taylor rule postulates that when the economy is at steady state and actual inflation equals target inflation, the real interest rate equals the natural interest rate. The values of the slope coefficients ω_1 and ω_2 that Taylor recommends are both equal to 0.5 reflecting the observation that the opportunity cost of more inflation stability increases sharply when the fluctuations in inflation are smaller than those of output and, conversely, the opportunity cost of more output stability increases sharply when the fluctuations in output are smaller than those of inflation.

Before proceeding to derive the interest rate path that is consistent with our preferred estimate of the time-varying natural rate, one has to have an estimate of Okun's coefficient and an estimate of the natural rate of interest which, like the natural rate of unemployment, is unobservable and time-varying. To estimate Okun's coefficient, I first derive estimates of potential output by applying a low-pass filter that admits cycles of periodicity of eight years or longer, on the basis of which I calculate the output gap. Figure 6 shows the resulting series. The series behave as one would expect registering a recessionary gap in the first half of the decade and being close to zero thereafter.

Then I estimate a dynamic version of Okun's equation, as described by Gordon (1984) and Weber (1995). Table 4 reports the results. It is worth noting that the estimate of the long-run Okun's coefficient that I obtain is close to the entirely standard value of –0.5.

Modern New Keynesian models of monetary policy define the natural interest rate as the rate of interest at which the output gap converges to zero.

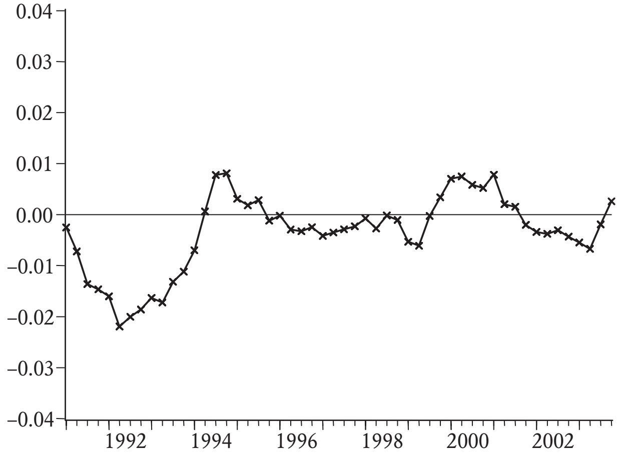

FIG. 6 The UK Output Gap, 1991Q1–2003Q4

TABLE 4
ORDINARY LEAST SQUARES ESTIMATION OF OKUN'S EQUATION

Dependent variable: Unemployment gap		
Regressors	*Lags*	
Unemployment gap	1–2	0.851 (0.048)
Output gap	1–2	–0.078 (0.045)
Implied long-run Okun's coefficient		–0.526 (0.214)
R^2		0.95
SER		0.001
LM		3.36 (0.50)

Notes: Standard errors in parentheses; Breusch–Godfrey Lagrange multiplier (LM) statistic for serial correlation up to order 4 (*p* value in parentheses).

In the spirit of this definition, I estimate the natural interest rate using the following multivariate unobserved components model:

$$q_t = a(L)q_{t-1} + b(L)(r_t - r_t^*) + \varepsilon_t \tag{14}$$

and

$$r_t^* = r_{t-1}^* + \eta_t \tag{15}$$

where q denotes the output gap, r the short-term real interest rate, ε is a random error with zero mean and variance σ_ε^2, and η is a random error with zero mean and variance σ_η^2.

The basic specification and methodology are close to that used by Laubach and Williams (2003) but, unlike them, I assume that the natural rate

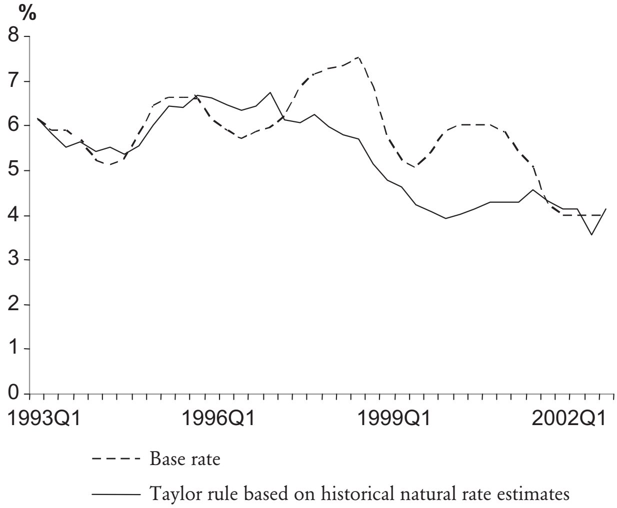

FIG. 7 The Base Rate versus the Time-varying Natural Rate Taylor Rule

of interest follows a random walk. To ensure that the estimated natural rate of interest exhibits low frequency variation, I restrict the signal-to-noise ratio ($\sigma_\eta^2/\sigma_\varepsilon^2$) to 0.01. The estimated natural rate of interest hovers around 3.8 per cent for the first three years of the inflation-targeting regime but declines gradually thereafter. At the end of 2002, the natural rate of interest stood close to 2.5 per cent.

Having retrieved estimates of Okun's coefficient and the natural rate of interest, I proceed to consider how different the interest rate path would have been if policy makers had perfect knowledge of the natural rate of unemployment in real time. I will do so in the context of the natural-rate-based Taylor rule (13) where the historical estimate of the time-varying natural rate is derived from the hybrid specification of the unobserved components model. I measure inflation using the retail price index excluding mortgage payments and I set the inflation target to 2.5 per cent, respectively, the preferred measure of inflation and the inflation target of the Bank until the end of 2003.

Figure 7 plots the actual interest rate against the interest rate prescribed by the natural-rate-based Taylor rule. Two salient features of the graph are worth noting. First, the two series follow broadly the same path. In particular, both series fluctuate around 6 per cent during the first three years of inflation targeting and both decline thereafter, converging to an interest rate of 4 per cent towards the early years of the new millennium. The behaviour of the actual and calibrated interest rates suggests that, overall, the UK monetary policy authorities have tracked the changes in the natural rate of unemployment reasonably well. Second, the actual and calibrated interest rates move very closely together during periods at which the natural rate remains stable whereas they temporarily deviate during periods at which the natural rate declines. In particular, during the first three years of inflation targeting, while the natural rate was roughly constant, the actual and calibrated interest rates move closely together, whereas, during the second half of the 1990s, while the natural rate was declining, the two series temporarily deviate. This pattern points towards the presence of a perception lag regarding changes in the natural rate, suggesting that in periods of sustained changes of the natural rate it may be taking some time for policy makers to process the new information and update their model parameters accordingly.

To substantiate this hypothesis, I examine the monetary policy debate, as reflected in the minutes of the monetary policy committee meetings and public speeches of the members of the committee. I do so by focusing, in particular, on the periods from 1997Q2 to 1998Q3 and 1999Q3 to 2001Q1, both periods during which the actual base rate increased while the interest rate prescribed by the Taylor rule suggested a somewhat looser monetary policy stance.

Early in the first period, the committee identified an interesting pattern in the data: although quantity indicators of the labour market, such as the unemployment claimant count, vacancies and skilled labour shortages, pointed consistently towards a tight labour market with considerable upside risks to the domestic component of inflation, earnings growth exhibited remarkable stability, rising on average at an annual rate of 4–4.5 per cent, a rate consistent with the Bank's inflation target. This 'benign conjunction of strong quantities and modest earnings growth' was new and puzzling. In principle, it could signify either that the natural rate of unemployment was lower than thought or that the substantial real appreciation of sterling since 1996 reduced temporarily the wedge between the real consumption and the real product wage. Both explanations would show in the data as reduced pressure on earnings growth for any given level of unemployment, although in the first case this would be a permanent whereas in the second a temporary effect.[17] In the minutes of the monetary policy committee meeting of July 1997, the committee noted:

> It remained surprising that earnings growth had not increased further during the last few years given the recorded increases in employment and falls in unemployment. It appeared possible that the rate of unemployment compatible with a stable rate of inflation was lower than had earlier been thought. But it was also possible that there could be a sudden sharp increase in earnings growth, as there had been in the late 1980s, when earnings growth had risen after what had seemed at the time a surprisingly long period of stability. (Minutes of the monetary policy committee meeting, July 1997, §45)

[17]Note that this analysis is consistent with a standard open economy imperfect competition model. For details, see Layard *et al.* (2005).

In subsequent meetings, the committee explicitly acknowledged the theoretical possibility that the natural rate of unemployment might have declined but remained uncertain about the likelihood of this outcome and dismissed taking any action upon it. While waiting to see how long-lasting the stability of earnings growth would be, the committee attached much higher weight to the probability that, following the expected reversion of sterling towards its equilibrium value, inflation would reignite. The minutes of the meeting of September 1997 noted:

> The Committee concluded that uncertainties remained about how much further tightening of the labour market could be tolerated without generating upward pressure on wages. It would be unwise in the current state of knowledge to take a strong view about the level of the natural rate of unemployment but it remained essential to monitor closely wage settlements and average earnings. (Minutes of the monetary policy committee meeting, September 1997, §21)

In subsequent meetings, and as the underlying trend in earnings growth remained stable, the possibility of a decline in the natural rate of unemployment came under closer scrutiny. In the minutes of April 1998, and for the first time, the committee discussed extensively the determinants of the natural rate and considered alternative reasons for its possible decline. On balance, the committee acknowledged the presence of a perception lag and invited further scrutiny of the natural rate, while also raising concerns about trends in the labour market quantity indicators. In particular, the minutes noted:

> Taken together, the quantity signals appeared to indicate a tightening of the labour market and some surveys suggested that the tightening would continue. The problem was in judging where the natural rate of unemployment lay. On one view, the subdued behaviour of earnings was encouraging evidence that the natural rate of unemployment might be lower than had previously been thought. (Minutes of the monetary policy committee meeting, April 1998, §42)

In subsequent meetings, and in particular from mid-2000 onwards, as the expected correction of the real exchange rate did not materialize while the benign combination of falling unemployment with subdued earnings growth continued, the minutes reveal a greater confidence in revising downwards the perceived estimates of the natural rate and gradually coming to terms with a new period of lower sustainable unemployment. The minutes of the meetings of November 2000 noted:

> More generally, it was possible that the sustainable rate of unemployment had, for a variety of reasons, improved by more than the Committee had already allowed for in its recent forecasts. A further adjustment had been made by the Committee in the latest Inflation Report projections. (Minutes of the monetary policy committee meeting, November 2000, §20).

Overall, the evidence from the minutes of the monetary policy committee meetings bears out the view that the temporary deviation of the actual base rate from the natural-rate-based Taylor rule is, at least partly, due to perception lags regarding the changing natural rate of unemployment. To a large extent, the monetary policy committee could not possibly have avoided them. Their presence reflects the gradual process of making sense of structural economic changes under conditions of uncertainty. A good deal of uncertainty stems from the fact that, unlike the econometrician or economic historian who attempts to estimate the natural rate retrospectively, policy makers lack the benefit of hindsight. They have to judge the state of the labour market, and take decisions, on the basis of real-time estimates of economic aggregates, i.e. estimates derived from a subset of survey responses that are often subject to considerable measurement error and sizeable *ex post* revisions. The perils of real-time data mismeasurement for macroeconomic stabilization policies, especially in relation to real-time output gap estimates, have been explored extensively. In an important paper, Orphanides and Van Norten (2002) use real-time US output data to construct measures of real-time output gaps, and show that *ex post* revisions of real-time output gap estimates are often of the same order of magnitude as the estimated gap itself. In another paper, Orphanides and Van Norten (2005) use real-time US output data to show that, although historical output gap estimates appear to be useful for predicting inflation, real-time output gap measures do not perform nearly as well.

In this context, it is worth asking to what extent data uncertainty hindered the monetary policy committee's real-time assessment of the state of the UK labour market or, to put the question differently, to what extent real-time estimates of unemployment gaps could have provided timely and useful information about the state of the UK labour market, thus reducing the length of perception lags. To answer this question, one has to look at estimates of natural rates of unemployment as they would have been obtained by the Bank staff at the time of decision making. I derive such estimates by estimating a bivariate unobserved components model, using successive releases of labour market quantity indicators and price deflators.[18] In particular, I estimate vintages of natural rates of unemployment using vintages of unemployment rates and final consumption expenditure deflators, as published in successive issues of *Labour Market Trends* and *Quarterly National Accounts*, 1998Q3–2007Q2. Each vintage of natural rate estimates corresponds to a pair of unemployment rate and personal consumption expenditure deflator vintages; thus the state variables and parameters of the unobserved components model are estimated as many times as the number of vintages, using only data available at each point in time. Having constructed

[18]The unobserved components specification is $\Delta\pi_t = a(L)\Delta\pi_{t-1} + b(L)(u_t - u_t^*) + \varepsilon_t$ (signal equation) and $u_t^* = u_{t-1}^* + v_t$ (state equation) where Δ is the difference operator, L is the lag operator, π denotes the final consumption expenditure inflation, u denotes the unemployment rate, and ε and v are serially uncorrelated zero mean disturbances. The natural rate of unemployment is the Kalman smoother estimate of u^*.

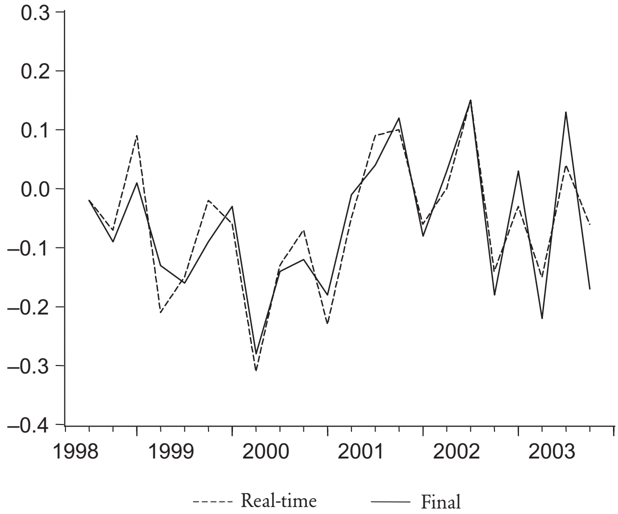

FIG. 8 Real-time versus Final Estimates of Unemployment Rates (Changes)

vintages of natural rate estimates, I construct vintages of unemployment gaps by subtracting each vintage of natural rate estimates from the corresponding vintage of unemployment rates. Finally, I use the latest estimate from each of the unemployment rate, natural rate and unemployment gap vintages to construct respectively real-time unemployment rate, natural rate and unemployment gap series. These real-time series can be thought of as reflecting the most up-to-date information available to the Bank staff at each point in time.

Figures 8–10 compare the real-time unemployment rates, natural rates of unemployment and unemployment gaps during 1998Q2–2004Q1 with their 'final' estimates. A 'final' estimate is here defined as the estimate published 12 quarters after the first data release. Although one has to acknowledge that 'final' is an elusive concept when it comes to issues of macroeconomic measurement, it is not unreasonable to assume that, within a 12-quarter window, revisions filter out most of the noise contained in preliminary data. The difference between the 'final' and real-time estimates of each series gives the extent of data revision, or real-time measurement error, at each point in time. Figure 8, in particular, compares changes in real-time against changes in 'final' unemployment rates. Simple visual inspection shows that the real-time series tracks closely the 'final' one suggesting that the size of real-time measurement error is small.

Figure 9 compares changes in real-time against changes in 'final' natural rates of unemployment. The real-time estimates are a good deal more volatile than the 'final' ones, pointing towards the presence of substantial noise in the real-time data.

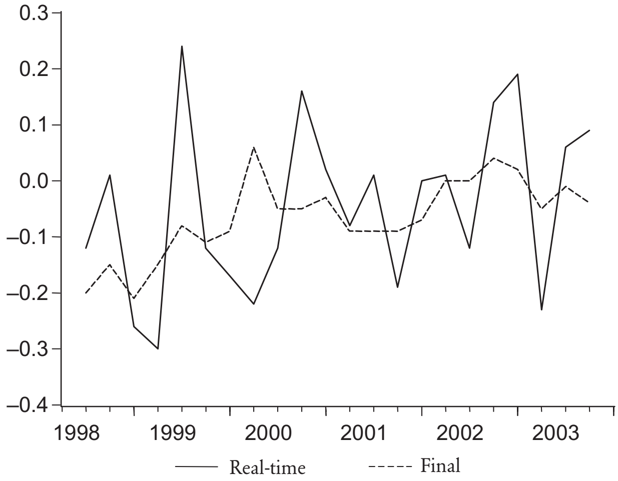

FIG. 9 Real-time versus Final Estimates of Natural Rates of Unemployment (Changes)

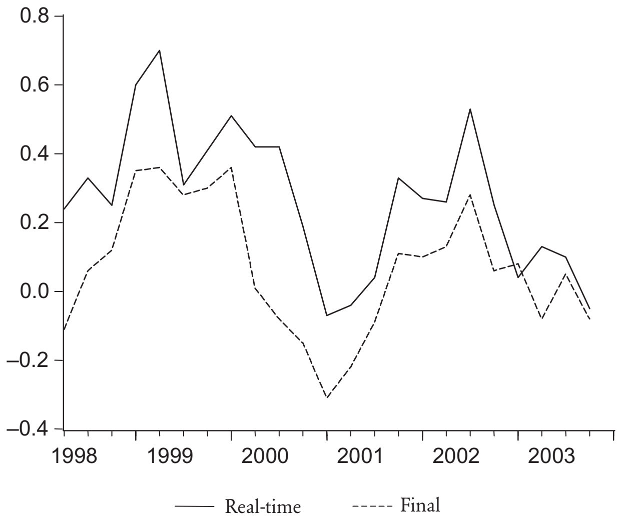

FIG. 10 Real-time versus Final Estimates of Unemployment Gaps (Levels)

Figure 10 compares levels of real-time against levels of 'final' unemployment gap estimates. A number of features are readily discernible. The two graphs seem to fluctuate around different means, suggesting that the real-time unemployment gap persistently underestimates the tightness of the UK labour market. In addition, the real-time series is more volatile than the 'final' one, suggesting that real-time unemployment gaps may not contain reliable information about the 'true' state of the labour market.

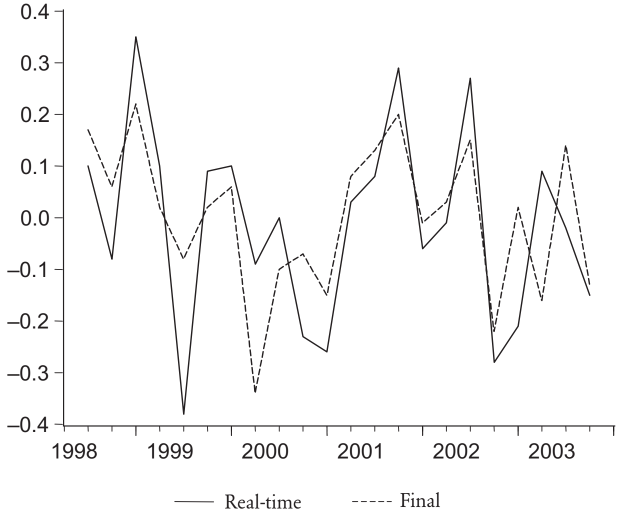

FIG. 11 Real-time versus Final Estimates of Unemployment Gaps (Changes)

Orphanides (2003a) and Walsh (2003) noted that, in the presence of persistent upward or downward revisions in (output) gap series, changes—rather than levels—of (output) gaps can help avoid a good deal of the real-time measurement problem and, thus, provide more reliable indicators of the underlying state of the economy. Figure 11 compares changes of real-time against changes of 'final' unemployment gaps. Although, as one would expect, the persistent underestimation of the tightness of the UK labour market disappears, the real-time series remains somewhat more volatile than the 'final' one, again pointing towards the presence of noise in the real-time data.

Table 5 reports summary reliability statistics of real-time labour market quantities. Column 2 (COR) presents the correlation between the real-time and 'final' series of each of the labour market quantity indicators. Columns 3 (NSR1) and 4 (NSR2) provide two proxies for the noise-to-signal ratio in the real-time estimates. NSR1 denotes the ratio of the root mean square of the revision to the standard deviation of the 'final' estimate while NSR2 denotes the ratio of the standard deviation of the total revision to the standard deviation of the 'final' estimate. The two proxies differ in that the numerator of the former contains the squared bias, and thus also captures the presence of persistent positive or negative revisions. Column 5 (OP-SIGN) reports the frequency with which real-time estimates misclassify the sign of their 'final' counterparts.

Several results are worth noting. First, the noise-to-signal ratios of changes in real-time natural rates of unemployment are sizeable, confirming

TABLE 5
MEASURES OF RELIABILITY OF REAL-TIME LABOUR MARKET QUANTITIES

	COR	*NSR1*	*NSR2*	*OP-SIGN*
Unemployment growth	0.84	0.53	0.54	0.09
ΔUnemployment rate	0.89	0.45	0.46	0.05
ΔNatural rate of unemployment	0.39	1.98	2.01	0.45
Unemployment gap	0.79	1.25	0.68	0.22
ΔUnemployment gap	0.67	0.96	0.67	0.23

Notes: Vintages 1998Q3–2004Q1; Δ is the first difference operator; COR denotes the correlation of the real-time and final estimates; NSR1, a noise-to-signal ratio proxy, denotes the ratio of the root mean square of the revision to the standard deviation of the final estimate; NSR2, a noise-to-signal ratio proxy, denotes the ratio of the standard deviation of the revision to that of the final estimate; OP-SIGN denotes the frequency with which the real-time and final estimates have opposite signs.

that real-time measures of the natural unemployment rates are not reliable indicators of the 'true' state of the labour market. Likewise, the noise-to-signal ratios of the levels of unemployment gaps are substantial and indicate the presence of a non-trivial squared bias. The signal-to-noise ratios of the changes in unemployment gaps are somewhat lower but still high enough to render the real-time estimates unreliable indicators of the 'true' state of the labour market.[19] The noise-to-signal ratios of unemployment growth and the change in unemployment rates are considerably lower. The contrast between the low noise in real-time LFS-based measures and the high noise in real-time natural unemployment rates and unemployment gaps indicates that the unreliability of the latter is not due to sizeable *ex post* data revisions in labour market data but to revisions in private consumption expenditure deflators or to the end-of-sample problem of the estimator or a combination of the two. The unreliability of real-time natural unemployment rates and unemployment gaps is borne out by the other two summary statistics. The results of column 5, in particular, show that real-time estimates of natural rates get the quarterly change in the 'final' natural rate of unemployment wrong 45 per cent of the time while real-time levels and/or changes of unemployment gaps give the wrong sign about 22–23 per cent of the time.

Overall, although labour market data are not subject to significant *ex post* revisions, real-time estimates of natural unemployment rates and unemployment gaps contain substantial noise and misclassify frequently the sign of their 'final' counterparts. Moreover, real-time estimates of the level of the unemployment gap appear to have underestimated persistently the tightness of the UK labour market, demonstrating that real-time information should only be taken into account with a great deal of scepticism. Orphanides (2003a) showed that naïve policy that downplays the real-time informational limitations may become a source of instability. If so, and until significant

[19]Orphanides and Van Norten (2002) report noise-to-signal ratios of similar size for various measures of US output gaps.

progress in forecasting *ex post* revisions takes place, perception lags can be thought of as the price of macroeconomic stability in face of data uncertainty.

5 Conclusion

The main findings of the paper are as follows. First, historical estimates show that the natural rate of unemployment has followed a bell-shaped pattern, rising from 3 per cent in the early 1960s to an average of 8 per cent in the 1980s before falling back to an average of 5 per cent around the turn of the century. The causes of the rise and fall of the natural rate have been discussed extensively elsewhere, yet there still remain several questions to be answered, especially regarding its welcome decline during the past decade. Second, the multivariate unobserved components method delivers unambiguously superior natural rate estimates, at least in so far as the usefulness of the natural rate to understand subsequent movements in inflation is what matters. Third, real-time estimates of the natural rate of unemployment and unemployment gaps are not reliable indicators of the 'true' underlying state of the labour market. Real-time estimates of unemployment gaps during 1998–2003, in particular, have persistently underestimated the 'true' extent of labour market tightness and naïve reliance on them could have led to macroeconomic instability. These pessimistic findings about real-time unemployment gaps mirror earlier results on the unreliability of output gaps and confirm that, especially at times of heightened uncertainty, policy makers should consider their real-time resource utilization indicators with the appropriate degree of scepticism and act cautiously. Finally, a counterfactual exercise and the minutes of the monetary policy committee meetings indicate that the UK monetary policy authority was able to track the changes in the natural rate of unemployment relatively successfully, albeit with some lag of recognition. This perception lag stems, at least in part, from the committee's interpretation of the new and puzzling 'conjunction of strong quantities and modest earnings growth' of the second half of the 1990s as primarily the result of a temporary appreciation of the real exchange rate that would give way to overheating as soon as the real exchange rate reverted towards its preconceived equilibrium. The possibility that the natural rate of unemployment might be lower than the committee had allowed for emerged gradually in the minutes of the committee meetings and as a result of the expected depreciation failing to materialize. In light of the significant uncertainty surrounding real-time labour market indicators and measures of resource utilization, the 'cautious' stance of the monetary policy committee in the face of a new and benign conjunction should be taken to reflect awareness of the severe real-time informational limitations and an indication of the progress that monetary policy has achieved in the past 30 years.

Appendix

Data Definitions and Sources

Prices: implicit final consumption expenditure deflator. Source: Office for National Statistics (Code: (ABJQ+HAYE)/(ABJR+HAYO)).

Prices: retail prices index excluding mortgage interest payments (RPIX). Source: Office for National Statistics (Code: CHMK).

Prices: consumer price index (all items). Source: International Monetary Fund (IMF) International Financial Statistics (Code: 11264ZF).

Import Prices: implicit import price deflator. Source: Office for National Statistics (Code: IKBI/IKBL).

Oil Prices: quarterly average crude price (US$ per barrel). Source: IMF International Financial Statistics (Code: 00176AAZZF).

Exchange Rate: quarterly average US$/GBP market rate. Source: IMF International Financial Statistics (Code: I 12RHZF).

Unemployment: LFS unemployment rate, 1971Q1–2003Q4. Source: Office for National Statistics (Code: MGSX). The series is extrapolated backwards using the OECD Labour Force Statistics unemployment rate.

Unemployment: OECD standardized unemployment rate, 1970Q1–2003Q4. The series is extrapolated backwards using the OECD Labour Force Statistics unemployment rate.

Aggregate Output: GDP at market prices (chained volume index). Source: Office for National Statistics (Code: YBEZ).

Short-term Interest Rate: quarterly average base rate. Source: Bank of England (Code: IUQABEDR).

Incomes Policy Dummies: 1 for periods of incomes policies; 0 otherwise. Incomes policy dummies cover the statutory wage freeze of 1966:3–1967:3 and the social contract of 1975:3–1977:2 (for details see Flanagan *et al.* (1983) *Unionism, Economic Stabilization, and Incomes Policies*, Washington, DC, The Brookings Institution).

Real-time Implicit Final Consumption Expenditure Deflator. Source: Office for National Statistics, Quarterly National Accounts, Various Issues (Code: (ABJQ+HAYE)/(ABJR+HAYO)). Available through the Bank of England's real-time database (Bank of England Code: PCNADEF.Q)

Real-time Unemployment. Source: Office for National Statistics, Labour Market Trends, Various Issues, Table A1 (Code: MGSC.Q). Available through the Bank of England's real-time database.

Real-time Labour Force. Source: Office for National Statistics, Labour Market Trends, Various Issues, Table Al (Code: MGSF.Q). Available through the Bank of England's real-time database.

The minutes of the monetary policy committee meetings of the Bank of England are published as an appendix to the quarterly Inflation Report of the Bank; they are also available on the Bank's website.

References

Ball, L. and Mankiw, G. (2003). 'The NAIRU in Theory and Practice', *Journal of Economic Perspectives*, Vol. 16, pp. 115–136.

Ball, L. and Moffitt, R. (2001). 'Productivity Growth and the Phillips Curve', in A. Krueger and R. Solow (eds), *The Roaring Nineties: Can Full Employment be Sustained?*, New York, The Russell Sage Foundation and the Century Foundation Press.

Ball, L. and Tchaidze, R. (2002). 'The Fed and the New Economy', *American Economic Review*, Vol. 92, pp. 108–114.

Batini, N. and Greenslade, J. V. (2006). 'Measuring the UK Short-run NAIRU', *Oxford Economic Papers*, Vol. 58, pp. 28–49.

Baxter, M. and King, R. (1999). 'Measuring Business Cycles: Approximate Band-pass Filters for Economic Time Series', *Review of Economics and Statistics*, Vol. 81, pp. 575–593.

Blinder, A. and Reis, R. (2005). 'Understanding the Greenspan Standard', in *The Greenspan Era: Lessons for the Future: a Symposium Sponsored by the Federal Reserve Bank of Kansas City*, Kansas City, MO, Federal Reserve Bank of Kansas City.

Blinder, A. and Yellen, J. (2001). 'The Fabulous Decade: Macroeconomic Lessons from the 1990s', in A. Krueger and R. Solow (eds), *The Roaring Nineties: Can Full Employment be Sustained?*, New York, The Russell Sage Foundation and the Century Foundation Press.

Burns, A. and Mitchell, W. (1946). *Measuring Business Cycles*, New York, National Bureau of Economic Research.

Cassino, V. and Thornton, R. (2002). 'Reduced-form Estimates of the Natural Rate of Unemployment Based on Structural Variables', *Bank of England Working Paper 153*.

Diebold, F. and Mariano, R. (1995). 'Comparing Predictive Accuracy', *Journal of Business and Economic Statistics*, Vol. 13, pp. 253–263.

Flanagan, R., Soskice, D. and Ulman, L. (1983). *Unionism, Economic Stabilization and Income Policies*, Washington, DC, The Brookings Institution.

Gordon, R. (1984). 'Unemployment and Potential Output in the 1980s', *Brookings Papers on Economic Activity*, Vol. II, pp. 537–564.

Gordon, R. (1997). 'The Time-varying NAIRU and its Implications for Economic Policy', *Journal of Economic Perspectives*, Vol. 11, pp. 1–32.

Gordon, R. (1998). 'Foundations of the Goldilocks Economy: Supply Shocks and the Time-varying NAIRU', *Brookings Papers on Economic Activity*, Vol. II, pp. 297–346.

Greenspan, A. (2004). 'Risk and Uncertainty in Monetary Policy', *American Economic Review*, Vol. 94, pp. 33–40.

Hamilton, J. (1994). *Time Series Analysis*, Princeton, NJ, Princeton University Press.

Harvey, D., Leybourne, S. and Newbold, P. (1997). 'Testing the Equality of Prediction Mean Squared Errors', *International Journal of Forecasting*, Vol. 13, pp. 281–291.

Henry, S. G. B. (1981). 'Incomes Policy and Aggregate Pay', in J. Fallick and R. Elliott (eds), *Incomes Policies, Inflation and Relative Pay*, London, George Allen & Unwin.

Hodrick, R. and Prescott, E. (1997). 'Postwar US Business Cycles: an Empirical Investigation', *Journal of Money, Credit and Banking*, Vol. 29, pp. 1–16.

Laubach, T. (2001). 'Measuring the NAIRU: Evidence from Seven Economies', *Review of Economics and Statistics*, Vol. 83, pp. 218–231.

Laubach, T. and Williams, J. (2003). 'Measuring the Natural Rate of Interest', *Review of Economics and Statistics*, Vol. 85, pp. 1063–1070.

Layard, R., Nickell, S. and Jackman, R. (2005). *Unemployment: Macroeconomic Performance and the Labour Market*, 2nd edn, Oxford, Oxford University Press.

Mankiw, G. and Reis, R. (2002). 'Sticky Information versus Sticky Prices: a Proposal to Replace the New Keynesian Phillips Curve', *Quarterly Journal of Economics*, Vol. 117, pp. 1295–1328.

Nelson, E. (2000). 'UK Monetary Policy 1972–1997: a Guide Using Taylor Rules', *Bank of England Working Paper 120*.

Nelson, E. and Nikolov, K. (2003). 'UK Inflation in the 1970s and 1980s: the Role of Output Gap Mismeasurement', *Journal of Economics and Business*, Vol. 55, pp. 353–370.

Nickell, S. (2001). 'Fundamental Changes in the UK Labour Market', *Oxford Bulletin of Economics and Statistics*, Vol. 63, pp. 715–736.

Orphanides, A. (2003a). 'Monetary Policy Evaluation with Noisy Information', *Journal of Monetary Economics*, Vol. 50, pp. 605–631.

Orphanides, A. (2003b). 'The Quest for Prosperity without Inflation', *Journal of Monetary Economics*, Vol. 50, pp. 633–663.

Orphanides, A. and Van Norten, S. (2002). 'The Unreliability of Output-gap Estimates in Real Time', *Review of Economics and Statistics*, Vol. 84, pp. 569–583.

Orphanides, A. and Van Norten, S. (2005). 'The Reliability of Inflation Forecasts Based on Output Gap Estimates in Real Time', *Journal of Money, Credit and Banking*, Vol. 37, pp. 583–601.

Orphanides, A. and Williams, J. (2002). 'Robust Monetary Policy Rules with Unknown Natural Rates', *Brookings Papers on Economic Activity*, Vol. II, pp. 63–118.

Romer, D. (2006). *Advanced Macroeconomics*, 3rd edn, New York, McGraw-Hill.

Rotemberg, J. (1999). 'A Heuristic Method for Extracting Smooth Trends from Economic Time Series', *NBER Working Paper 7439*.

Sinai, A. (2004). 'Preemption, Changing Structure, and US Monetary Policy', *American Economic Review*, Vol. 94, pp. 49–52.

Staiger, D., Stock, J. and Watson, M. (1997a). 'How Precise are Estimates of the Natural Rate of Unemployment?', in C. Romer and D. Romer (eds), *Reducing Inflation: Motivation and Strategy,* NBER Research Studies in Business Cycles, Chicago, IL, University of Chicago Press.

Staiger, D., Stock, J. and Watson, M. (1997b). 'The NAIRU, Unemployment and Monetary Policy', *Journal of Economic Perspectives*, Vol. 11, pp. 33–49.

Staiger, D., Stock, J. and Watson, M. (2001). 'Prices, Wages, and the US NAIRU in the 1990s', in A. Krueger and R. Solow (eds), *The Roaring Nineties: Can Full Employment be Sustained?*, New York, The Russell Sage Foundation and the Century Foundation Press.

Taylor, J. (1993). 'Discretion vs Policy Rules in Practice', *Carnegie-Rochester Conference Series on Public Policy*, Vol. 39, pp. 195–214.

Walsh, C. (2003). 'Minding the Speed Limit', *FRBSF Economic Letter*, May.

Weber, C. (1995). 'Cyclical Output, Cyclical Unemployment, and Okun's Coefficient: a New Approach', *Journal of Applied Econometrics*, Vol. 10, pp. 433–445.