Inflation and real disequilibria

Mark S Astley

and

Tony Yates

Monetary Analysis, Bank of England, London, EC2R 8AH
E-mail: mark.astley@bankofengland.co.uk
tony.yates@bankofengland.co.uk

The views expressed in this paper are not necessarily those of the Bank of England, nor of the MPC, but those of the authors alone. We would like to thank Hasan Bakhshi for his insightful thinking on this work. We would also like to thank Paul Fisher, Neal Hatch, Nigel Jenkinson and participants at a Johannesburg Economic Society seminar hosted by RAU in July 1998.

Issued by the Bank of England, London EC2R 8AH to which requests for individual copies should be addressed: envelopes should be marked for the attention of the Publications Group. (Telephone: 020 7601 4030.) Working Papers are also available from the Bank's Internet site at http:\\www.bankofengland.co.uk/wplist.htm.

ISSN 1368-5562

Contents

Abstract

This paper outlines some problems with the methods often used to construct measures of real ‘disequilibria’ or ‘gaps’ (eg the output gap), and to examine their relation to inflation. It then offers a structural vector autoregression alternative, which we use to construct estimates of output, unemployment and capacity utilisation gaps. We construct our gap estimates by summing the effects of particular structural shocks on output etc – where the shocks are identified using long-run restrictions derived from theory. Our approach has four main advantages over other methods. First, it uses economics rather than statistics to construct the gaps. Second, the estimates are not contingent upon particular assumptions about the structure of the economy. Third, it does not impose a rigid causal chain running from gaps to inflation. Fourth, it allows us to construct several gaps and examine their relation to inflation in a single framework – so our three gaps are internally consistent and can be used to make inferences about the structure of the economy.

1. Introduction

How should we construct measures of real ‘disequilibria’ or ‘gaps’? How should we model their relation to inflation?

This paper offers some comments on the methods used to answer those questions and suggests an alternative technique. Specifically, we estimate a structural vector autoregression model to construct several ‘gap’ measures and, jointly, to examine their relation to inflation. This allows us to identify the unobservable aggregate demand and aggregate supply shocks that underlie movements in inflation and real variables – including the unobservable ‘gap’ components of the latter. Our identification scheme elaborates on one basic consensual idea: that in the short run nominal and real variables may be related, but in the long run they are not. Our approach means that the output, unemployment and capacity utilisation gaps we construct are internally consistent, and follow from the short-run non-neutralities that cause them to be potentially related to nominal variables. We use these three gaps to demonstrate that a point well known in theory — the three gaps will only be identical under very restrictive circumstances — is also apparent in the data.

The paper is offered as much as a methodological comment as a source of new ‘gap’ estimates. So we spend some time —in Section 2 —discussing what we see as the shortfalls of other approaches to constructing ‘gaps’ and modelling their relation to inflation. Section 3 presents the results of our SVAR analysis. Section 4 concludes.

2. Motivation: some comments on previous work on inflation and real disequilibria

It is still common to see empirical models that look like this:

$$p_t = A(L)^n p^e_{t,t-n} + B(L)^n p_{t-n} + C(L)^n (y_{t-n} - y^*_{t-n}) \qquad \mathbf{(1)}$$

where $p_t = p_t - p_{t-1}$, is the difference between the log price levels at t and t-1

$p^e_{t,t-n}$ is the expectation of inflation at t given information at $t-n$

$y_{t-n} - y^*_{t-n}$ is output (or unemployment or capacity utilisation) minus its long run level, both at $t-n$

and L is a lag operator and A,B,C are vectors of parameters. Equation **(1)** — a reduced-form Phillips curve — says that the current inflation rate is related to some combination of lags of inflation or lags of inflation expectations, (with A and B constrained to sum to unity, so that real variables are neutral in the long run), and lags of some real variable. This variable is often called a ‘gap’ or a ‘real disequilibrium’. We will interpret this variable as trying to measure ‘movements in real variables due to demand shocks’ (which our priors tell us should be transitory) as distinct from ‘movements in real variables due to supply shocks’ (which our priors tell us should be permanent). Perhaps because they are convenient, models like **(1)** are still popular in the literature. But our surmise of the current state of knowledge is that they could be improved on in a number of different ways. The remainder of this section elaborates on this point.

2.1 Which Phillips curve? One reduced form for all structures?

The first point to note is that **(1)** is a reduced form that is a good approximation to some ‘structural’ economic models, but not all. And without other evidence to guide the choice between them, the empirical researcher is left in the dark. The following are the Phillips curves that drop out of some popular economic models (Appendix A provides a discussion of the parameters):(1)

(Rotemberg, 1982)

$$p_t = p^e_{t+1,t} + \frac{b}{c}(y_t - y^*_t) - \frac{e_t}{c} \qquad \textbf{(2)}$$

(Calvo, 1983)

$$p_t = p^t_{t+1,t} + \frac{g^2 b}{(1-g)}(y - y^*) + \frac{g^2}{(1-g)}e_t \qquad \textbf{(3)}$$

(Taylor, 1979, 1980)

$$p_t = p^e_{t+1,t} + s - b(u_{t-1} + u_t + u^e_{t,t-1} + u^e_{t+1,t}) + 2(e_t + e_{t-1}) + (p^e_{t,t-1} - p_t) \qquad \textbf{(4)}$$

(Fuhrer and Moore, 1995)

$$\Delta p_t = \Delta p^e_{t+1,t} + s' - b'(u_t + u_{t-1}) + 2(e'_t + e'_{t-1}) + (p^e_{t,t-1} - p_t) \qquad \textbf{(5)}$$

(Layard et al ,1991)

$$p_t = p^e_{t+1,t} - \frac{b''}{l_p + l_w}(u_t - u^*_t) \qquad \textbf{(6)}$$

---

(1) See Roberts (1995, 1997) for the derivation of these equations.

The main point to note is that Phillips curves **(2)-(6)** are tantalisingly similar, but <u>not</u> identical. Specifically, some assume nominal rigidity in prices, others nominal wage rigidity, and some both. Some invoke models of staggered contracts, some do not. Some are written in terms of output, and some in terms of unemployment. Some have moving averages in the shocks, some do not. Some have rates of change of inflation as the dependent variable, while some have its level. Even if **(2)-(6)** were an exhaustive list of Phillips curves derived from competing structural models, we would be in trouble, forced to rely on what Sims (1980) called 'incredible' identification restrictions to estimate our model. In fact the problem is more acute since our list is far from exhaustive, and it is also not a list of resolutely 'structural' models. The aggregate supply components of Fuhrer and Moore (1995), Taylor (1979,1980), Calvo (1983), Rotemberg (1982) and Layard *et al* (1991) are all, in their own way, 'reduced forms' and they are often at odds with fully optimal behaviour by rational agents.(2) For example, the key features of the nominal frictions that determine the slope of the short-run aggregate supply curve — the degrees of time-dependence and staggering — are both treated as exogenous: time and regime invariant.

We argue for an encompassing empirical framework that relies on a minimum of the most uncontroversial identification restrictions; and can therefore potentially shed some light on the empirical performance of the competing theories behind equations **(2)-(6)**.

2.2 Identifying demand and supply shocks: statistics versus economics

Equation **(1)** requires construction of a measure of the output (or some other) gap. Several different methods have been used to uncover this unobservable component of the time series of output etc.

The first general approach is to draw some trend line through observed output/unemployment/capacity utilisation. This could be a straight-line trend (see, eg, Fisher *et al* (1996)), a moving average or a Hodrick-Prescott (hereafter HP) filter (Fisher *et al* (1996), Clark *et al* (1996), Turner (1995), and many more). These methods embody some particular, and, in our view, peculiar assumptions about the behaviour of demand and supply. The straight-line trend approach attributes all movements in output around the trend to demand shocks, forcing all supply shocks to be positive (since output rises over time).(3) A moving average or HP filter allocates movements in output proportionately to demand and supply: it also induces a positive correlation between

---

(2) Ascari (1997) shows that the Taylor model is a particularly bad reduced form for a dynamic general equilibrium model with wage stickiness.

(3) Negative supply shocks can be allowed for by combining the straight-line trend with some other variable that might proxy shocks to the production possibility frontier (eg commodity prices).

the two.(4) Yet there is no reason why demand and supply shocks should always — or ever — behave in this way. These methods base the decomposition of output (or whatever) on statistics rather than economics, and suffer as a result.(5)

A second approach is to construct the supply component of output (or some other real variable) from the trend levels of its determinants. In the case of output this is often called the ‘production function’ approach(6) and involves using some economics to specify a production technology to weight up the trend levels of supply factors.(7) But it also involves some well-known difficulties. First, it relies on finding accurate data on the determinants of output (or whatever) that often do not exist(8) (eg the capital stock, or the determinants of unemployment). Second, we need to be confident in the technology used to weight up trend levels of factors: should we, for example, use Cobb-Douglas or CES to construct trend output? A third — and more fundamental — problem, which has received little attention, is that the production function approach leaves us with the problem of decomposing movements in factors into trend and cyclical components. The solution often adopted is to filter the series in one of the ways described above, the drawbacks of which prompted the appeal to the production function approach in the first place.

A third method is to use a Kalman filter to uncover the unobserved, supply-induced movements in output or unemployment. Kuttner (1994) applied this approach to US output; King *et al* (1995), Staiger *et al* (1996) and Gordon (1997) applied it to US unemployment; and Saleheen(1998) applied it to UK output and unemployment. This method does use economics to identify demand and supply: the estimates of potential output are constructed so that they have no long-run effect on inflation. But identification is achieved by imposing — in addition to plausible long-run restrictions — a particular form on the dynamics of the Phillips curve. This conflicts with our preference for minimalism in identification articulated in Section 2.1.

---

(4) This is because even if a rise in output is due entirely to a positive supply shock the moving average method will indicate that there has been a positive supply shock and a positive demand shock — with the mix depending on the memory of the moving average. The HP filter also suffers from the well-known ‘end-points’ problem. And Cogley and Nason (1995) argued that it can generate spurious cycles.

(5) A potential additional problem in the case of unemployment is that the filter will have to deal with the fact that theory suggests that some types of supply shocks — in particular technology shocks — as well as demand shocks will only have temporary effects on unemployment. But demand and supply shocks, of course, have opposite effects on inflation. This adds a further caveat to the interpretation of Phillips curves based upon unemployment gaps.

(6) See *inter alia* Giorno *et al* (1995), Fisher *et al* (1996).

(7) Likewise the natural rate of unemployment is estimated by regressing observed unemployment on a collection of the likely determinants of the natural rate — such as replacement ratios and union density (see Adams and Coe (*op cit*), Manning (1993)).

(8) See, for example, Bean (1994) and Manning (1993).

A fourth way of extracting the supply component of output (or whatever) — and one that also incorporates economics — is to use the multivariate Beveridge-Nelson decomposition (MBND hereafter) technique developed by Evans and Reichlin (1994) to decompose output into its trend and cyclical components.(9) This method has been used by, amongst others, Barrell and Sefton (1995). But it also suffers from problems. For example, Lippi and Reichlin (1994) argued that the MBND underestimates the proportion of output movements due to supply shocks. This is because the MBND assumes that the entire impact of a supply shock on GDP (or whatever) occurs immediately (because the permanent — supply — component is assumed to be a random walk). In reality, however, productivity shocks that come about because of innovation may take time to have their full effect on output — as the new technology diffuses through the economy. The MBND will mistakenly allocate such gradual effects to demand shocks — inducing an upward bias in the average size and variability of the estimated output gap.

We would not pretend that the SVAR method we present below is perfect — we will discuss some of its drawbacks later — but it does address some of the criticisms outlined above. In particular, it uses some fairly uncontroversial economics to construct the output gaps; it does not require gathering data on unobservables; and it allows the data to determine how demand and supply shocks diffuse over time.

2.3 Endogeneity: real disequilibria do not cause inflation or vice versa

Our third broad comment on methods like those in equation **(1)** is that they assume that causality runs from the real disequilibria to inflation. We view that assumption as incorrect because nominal (price) rigidities are one of the reasons why exogenous shocks generate real disequilibria in the first place. So it seems wrong to argue that such real disequilibria generate, with a lag, the (shock induced) price movements that occur once those nominal rigidities stop binding. We prefer a method that treats both inflation and real disequilibria as endogenous variables, whose movements are driven by exogenous shocks. The following quotes from Irving Fisher (1926) show just how spectacularly unoriginal we are in arguing for this approach:

> '.... what the figures show is… … a genuine and straightforward causal relationship; that the ups and downs of employment are the effects, in large measure, of the rises and falls of prices, due in turn to... money and credit...... Of course, this relationship might conceivably not be causal ... both might be conceivably caused by some third

---

(9) In brief, the MBND formulates trend output growth as the long-run output forecast derived from a VECM. Because the importance of the cyclical component asymptotes to zero as we forecast further into the future, long-run forecasts only pick up the trend component. This is the multivariate generalisation of the method proposed by Beveridge and Nelson (1981).

> influence. Or it might be conceived that price-change simply represents a forecast of good or bad business.'

In addition to suggesting that the 'Phillips curve' should more properly be renamed the 'Fisher Curve'(10) those extracts illustrate our point that we could just as well write regressions **(1)–(6)** with the output terms on the left-hand side, although this typically is not done. Treating real disequilibria and inflation as endogenous variables requires, of course, a systems approach. There are examples of systems approaches based on the Kalman filter, but these often fall short of allowing for full simultaneity.(11) Our SVAR method models both real disequilibria and inflation as endogenous variables — their movements being caused by the underlying demand and supply shocks hitting the economy.

2.4 Co-movements in real disequilibria: incredible identification and a missed opportunity

Our final comment on previous work is that it has usually — though not universally — only considered one measure of real disequilibria at a time. We argue that this constitutes an act of 'incredible' identification and a missed opportunity to learn something about the economy. (In contrast in Section 3 we model simultaneously output, unemployment and capacity gaps, and inflation).

To explain, suppose that we chose to investigate the data with a model like **(3)**, which Roberts (1995) derives from Calvo (1983). This Phillips curve, which is written in terms of output disequilibria, is based on sticky product market prices. If the 'real world' is characterised by sticky product prices and there are no other nominal frictions (e.g. in the labour market) then we will obtain unbiased coefficient estimates (other problems aside). But we will be in trouble if it is wages rather than prices that are sticky: we will then need some other assumptions about production technology, the form of nominal inertia in labour markets — which we may or may not believe — to convert the true unemployment-propagated Phillips curve into the output space format actually estimated. Similarly, if prices and wages are both sticky, then we might be able to generate models that give us the same functional form as **(3)** —which assumes that

---

(10) Donner and McCallum (1972) examine the history of the Phillips curve.

(11) In particular, Apel and Jansson (1997) formulate a system composed of a Phillips curve, an Okun's law equation (relating changes in the output gap to changes in the unemployment gap) and assumed stochastic processes for potential output and the natural rate of unemployment. But that system imposes a restrictive path for the transmission of demand shocks — running from demand shocks to inflation via the unemployment gap in the Phillips curve, and then from the unemployment gap to the output gap via the Okun's law equation — that we would suggest has no theoretical underpinnings.

only prices are sticky; but then again we might not. In order to measure the potentially differing temporary responses of output and unemployment to an underlying demand shock, we need to develop a model that is as agnostic as possible about the source (product or labour market?) and form (staggered or not?) of nominal rigidities. If we do not develop such a model, choosing either output *or* unemployment-based models of inflation constitutes an act of 'incredible identification'. In short, we are arguing that the practice popularised by Okun (1962) of using a rule of thumb to convert changes in output into changes in unemployment is deficient.(12)

Why would estimating a 'single disequilibrium' model amount to a missed opportunity? Simply because the co-movements of gaps in output, unemployment and capacity utilisation can tell us something about the economy. If it were impossible to substitute one factor for another, and there were only one source of nominal friction (and capital adjustment costs), there would be(13) an exact mapping between movements in output, unemployment and capacity utilisation caused by demand shocks. In other, more realistic worlds, however, there will be no such mapping. To develop the Calvo-based example articulated above, if there are sticky product prices, but only *real* inertia in labour markets, then nominal demand shocks will generate output and capacity utilisation movements, but will have no effect on unemployment. In other words, although this model indicates that demand shock will generate identical (zero) long-run movements in output and unemployment, the predicted short-run movements differ dramatically. The response of different real variables to supply shocks may also be revealing: in some models the same technology shocks that generate *permanent* movements in output generate *temporary* movements in unemployment (eg in Layard *et al* 1991), while in others they do not.

We are not original in arguing for the simultaneous modelling of different real disequilibria — Adams and Coe (1990) and Apel and Jansson (1997) have models that include output and unemployment gaps. But those papers also embody assumptions about product and labour market rigidities, and about the direction of propagation of shocks (from the labour to the product market or *vice-versa*?) that, while ensuring a kind of internal consistency, preclude them from exploiting the two advantages of gap 'multiplicity': agnostic identification, and inference from the co-movement between the gaps in response to different shocks. We think that the SVAR approach we implement is better designed for this purpose.

---

(12) Prachowny (1993) outlines some other criticisms of the Okun (1962) method.

(13) For shocks small enough not to warrant an adjustment in the capital stock immediately.

3 An SVAR investigation of inflation and real disequilibria

We now implement an SVAR approach to construct measures of the output, unemployment and capacity utilisation gaps and examine their relation to inflation. Several previous papers have used SVARs to estimate real disequilibria. In particular, Sterne and Bayoumi (1995), DeSerres *et al* (1995), Funke (1997), St Amant and Van Norden (1997) and Thomas, Dhar and Pain (1998) estimated SVAR-derived output gaps.(14) Similarly, Roberts (1993) and Dolado and Salido (1996) present SVAR-derived unemployment gaps. Our value-added is to study the behaviour of several real disequilibria, including the previously undiscussed capacity utilisation gap,(15) in the same system.

3.1 The data

The first stage in the SVAR method is to estimate a reduced-form VAR. In our case we estimated a five variable system including: the first difference of the (log) oil price (Δoil); the first difference of (log) RPIX inflation(16) (Δp); the first difference of (log) real GDP (Δy); the first difference of unemployment (Δu); and the first difference of the CBI capacity utilisation series (Δcu). Appendix B details the variable definitions and sources. Our sample period runs from 1974 Q1 to 1998 Q2. On the basis of the results of sequential LR tests and Akaike/Schwartz criteria (available on request) we included three lags in the first-stage VAR. ADF tests (again available on request) and the variable plots of Chart 1 indicate that all of these transformations were sufficient to induce stationarity in each of the series.(17) As we explain in Section 3.2, this allows us — in the second stage of the method — to impose restrictions on the long-run impact of structural shocks on the integrals of those variables (eg the level of real GDP), which we plot in Chart 2.

---

(14) An output gap is also implicit, though not discussed, in the bivariate Quah and Vahey (1995) system.

(15) An exception is Gordon (1999), who estimates a time series of the non accelerating inflation rate of capacity utilisation (NAIRCU) in a Phillips curve model – and hence (implicitly) a capacity utilisation gap.

(16) That is, the second difference of the log price level.

(17) Our finding that the level of capacity utilisation is I(1) is the most suspect — because the series is bounded between 0 and 100. Likewise the non-stationarity of the unemployment rate and inflation are subject to controversy. These findings probably reflect small sample factors. That said, similar findings are implied for the United States in Fuhrer and Moore (1995), and for the United Kingdom in Quah and Vahey (1995) and Westaway (1997), for example.

Chart 1: Variables entering the first-stage VAR

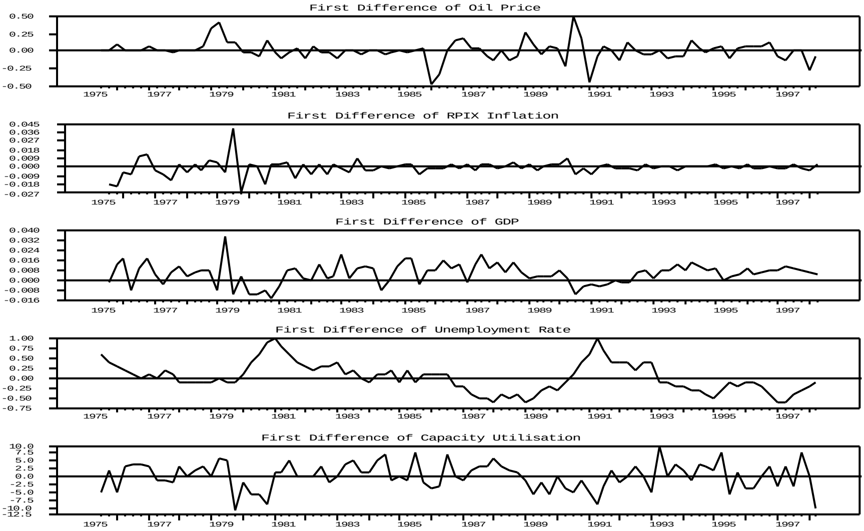

Chart 2: Variables upon which identifying restrictions are imposed

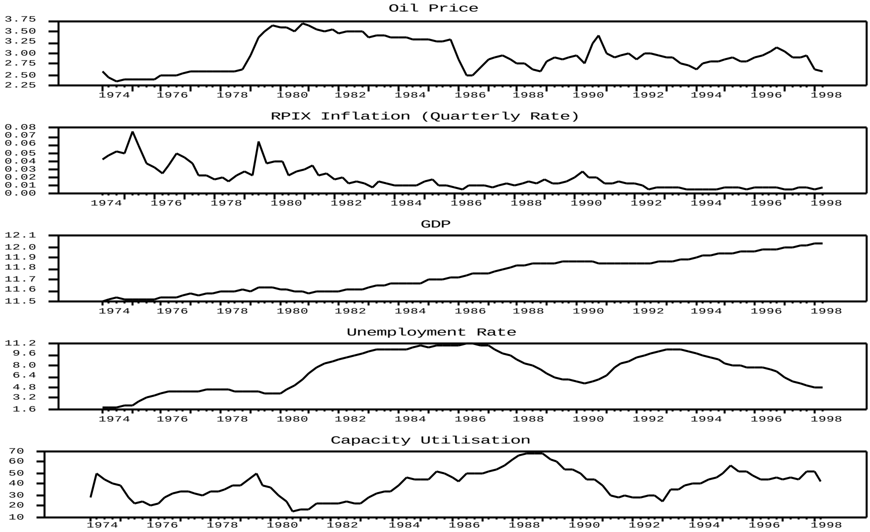

The focus of this paper is on the co-movements of inflation with different measures of real gaps — hence our inclusion of inflation, GDP, unemployment and capacity utilisation in the VAR. Capacity utilisation is included to take account of changes in the intensity with which the factors of production are used. But it is worth noting that the (CBI measured) capacity utilisation series we use relates only to the manufacturing sector, while several of the other variables in the system are whole-economy measures. This is an unavoidable imperfection, as the CBI series is the only one available over a long period in the United Kingdom.(18) We use RPIX inflation — retail price inflation excluding mortgage interest payments — because of its role as the UK inflation target (from 1992 onwards). We include the oil price because we think that it will be useful to distinguish between oil *shocks* and other types of supply shocks. (For example, we might expect the dynamics of output to be different following a positive supply shock that comes from innovation compared to the response to a fall in the price of oil that makes existing production technologies cheaper).

3.2 Identification

We use this system to identify the effects of five structural shocks, some of which we will use to create our gap measures. The five shocks are: oil shocks (ε^{oil}) — reflecting changes in the demand and supply conditions in the global oil market; IS shocks (ε^{IS}) — movements in the preferences of home/foreign consumers, fiscal policy etc.; technology shocks (ε^{TECH}) — shifts in the productivity of the factors of production; unemployment (or natural rate of unemployment) shocks (ε^{UNEM}) — reflecting changes in the factors (union bargaining strength, efficiency wage considerations etc) that drive long-run unemployment movements; and LM shocks (ε^{LM}) — shifts in money demand, supply or velocity.

To move from the reduced-form VAR representation to the structural moving average representation — achieve exact identification — it is necessary, in a system including n variables, to find $n(n\text{-}1)/2$ *theory derived* restrictions.(19) In our five variable system this means we need ten restrictions. Following Blanchard and Quah (1989) we use exclusively *long-run* restrictions — specifying that a particular shock has no long-run effect on the *level* of a particular endogenous variable. We use long-run restrictions because the theory underpinning them is often less controversial than that underlying contemporaneous restrictions.(20) In our case, the long-run restrictions are consistent

---

(18) The British Chambers of Commerce (BCC) produce separate capacity utilisation series for the manufacturing and service sectors. Unfortunately, however, they only start in 1989.

(19) A total of n^2 restrictions have to be imposed to achieve exact identification. The usual assumptions of orthogonality and unit variance of the structural shocks provide $n(n+1)/2$ restrictions (15 in our case). See Blanchard and Quah (1989) or Bank of England (1999) for more detailed explanations of identification procedures.

(20) Our identification scheme does not include any cointegrating relationships. This reflects the fact that theory indicates that no meaningful cointegration vectors exist. And the cointegration tests support that prior: while they provided some (but not conclusive) statistical evidence of one cointegrating vector, that vector had no economic content. The approach of King, Plosser, Stock and Watson (1991) should be applied if meaningful cointegrating relationships are present.

with all the reduced-form Phillips curves sketched in Section 2. We leave the short-run responses completely unrestricted because it allows us to be agnostic about issues such as the form, degree and duration of nominal price/wage or capital rigidities.

Because the first-stage reduced-form VARs are estimated in first differences of the variables, imposing restrictions on the *level* of the variables amounts to restricting the cumulated first difference impulse responses. If we denote the matrix of these long run multipliers as $C(1)$,(21) the restriction that shock j has zero long-run effect on the *level* of endogenous variable i requires that $C_{ij}(1)=0$ be imposed. Using this notation our system can be represented in matrix form as follows:

$$\begin{bmatrix} Oil\,\mathrm{Pr}ice \\ Inflation \\ GDP \\ Unemp \\ CapUtil \end{bmatrix} = \begin{bmatrix} C_{11}(1) & \mathbf{C_{12}(1)} & \mathbf{C_{13}(1)} & \mathbf{C_{14}(1)} & \mathbf{C_{15}(1)} \\ C_{21}(1) & \mathbf{C_{22}(1)} & \mathbf{C_{23}(1)} & \mathbf{C_{24}(1)} & C_{25}(1) \\ C_{31}(1) & \mathbf{C_{32}(1)} & C_{33}(1) & C_{34}(1) & \mathbf{C_{35}(1)} \\ C_{41}(1) & C_{42}(1) & C_{43}(1) & C_{44}(1) & C_{45}(1) \\ C_{51}(1) & C_{52}(1) & C_{53}(1) & \mathbf{C_{54}(1)} & C_{55}(1) \end{bmatrix} \begin{bmatrix} e^{OIL} \\ e^{IS} \\ e^{TECH} \\ e^{UNEM} \\ e^{LM} \end{bmatrix} \qquad \mathbf{(7)}$$

Table 1 shows that theory offers us 16 possible long-run restrictions in our system; in the table an ‘LR’ denotes the fact that theory suggests that the shock (in the first column) has zero long-run effect on the level of a particular endogenous variable. (22)

We discuss these restrictions below. As only ten theory-based restrictions are required to exactly identify the system, unless we are to estimate an overidentified system, which we do not do in this paper, we have to choose to leave six of the long-run responses unrestricted. We determine our choice by a combination of (i) imposing the type of restrictions that have been most frequently implemented in previous papers (reflecting their uncontentious nature); and (ii) ensuring that we can distinguish all the shocks from each other. We did experiment with several alternative systems ; the one we present below appears to have the highest economic content, by a long way, of the alternatives available.(23)

---

(21) If we let lower case c’s denote the effect of a shock on the first difference of a variable and the upper case C’s denote the equivalent effect on the level of a variable (which is the accumulation of the first difference effects) then $C(1)=\sum_{n=0}^{\infty} c_n$

(22) It is important to note that our method does not specify a particular horizon at which the identifying restrictions are imposed. Rather we allow for the restrictions biting up to an infinite horizon. In practice, however, they are usually imposed far sooner than this.

(23) In particular, we obtained implausible impulse responses and forecast error variance decompositions when we transferred the restrictions on output onto unemployment. But broadly similar results to those presented below were sometimes obtained under some alternative identification schemes (for example, when we switched the IS-inflation restriction with an LM-unemployment restriction).

Table 1: Possible Identifying Restrictions

	Effect on Variable:				
Shock	**Oil Price**	**Inflation**	**GDP**	**Unemploy**	**Cap. Util**
Oil		*LR*			*LR*
IS	**LR**	**LR**	**LR**	*LR*	
Technology	**LR**	**LR**		*LR*	
Unemploy	**LR**	**LR**			**LR**
LM	**LR**		**LR**	*LR*	*LR*

Key: A bold LR denotes a restriction used to exactly identify the system; an italicised, plain text LR denotes a possible overidentifying restriction.

The ten restrictions underlying our exactly identified system (which appear in bold in equation **(7)** and Table 1) are:

$$C_{12}(1) = C_{13}(1) = C_{14}(1) = C_{15}(1) = C_{22}(1) = C_{23}(1) = C_{24}(1) = C_{32}(1) = C_{35}(1) = C_{54}(1) = 0$$

We restrict IS shocks, technology shocks, unemployment shocks and LM shocks each to have no long-run effect on the oil price ($C_{12}(1) = C_{13}(1) = C_{14}(1) = C_{15}(1) = 0$). These restrictions reflect the fact that the world oil price is determined by world oil market demand and supply factors.(24)

We restrict IS shocks, technology shocks and unemployment shocks each to have no long-run effect on the (RPIX) inflation rate ($C_{22}(1) = C_{23}(1) = C_{24}(1) = 0$). These restrictions embody the view that inflation is a monetary phenomenon in the long-run — determined purely by LM shocks (which Roberts(1995) labels as 'target inflation' shocks). These type of restrictions can, of course, be traced back to Friedman (1968) and Phelps (1968) and are present in the models described by equations **(2)-(6)**.

The two restrictions that IS and LM shocks both have no long-run effect on the level of GDP ($C_{32}(1) = C_{35}(1) = 0$) imply that the long-run aggregate supply curve is vertical. So output movements reflect the three supply shocks (oil, technology and unemployment) in the long run (represented by the $C_{31}(1)$, $C_{33}(1)$ and $C_{34}(1)$ coefficients). This type of restriction is, of course, familiar from Blanchard and Quah (*op cit*) and has been used by a large number of other papers.

Finally, we restrict unemployment shocks to have no long-run effect on capacity utilisation ($C_{54}(1) = 0$). This is motivated by thinking of capacity utilisation being determined in the long run by adjustment costs and demand uncertainty (see eg

---

(24) Our specification that the four non-oil shocks have no effect on oil prices presumes that the United Kingdom is a 'small' country in the global market for oil. These restrictions would be less valid if we were to apply our method to data for a large country such as the United States. This type of restriction has previously been used by DeSerres *et al* (1995) and Bjornland (1997).

Greenwood *et al* (1988) and Finn (1996)) rather than the union power, efficiency wage (etc) factors that are embodied by unemployment shocks. This restriction is also consistent with the Joyce and Wren-Lewis (1991) findings.

This leaves six possible overidentifying restrictions:

- Oil shocks should have no long-run effect on inflation (they only affect the price *level*) — implying that $C_{21}(1) = 0$.

- IS shocks, LM shocks and technology shocks should have no long-run effect on unemployment — ie $C_{42}(1) = C_{45}(1) = C_{43}(1) = 0$. The first two of those are analogous to the equivalent restrictions on output. The final one is proved in Layard, Jackman and Nickell (1991)(25) and implemented in Bean (1992). By leaving these responses unrestricted we allow for the *possibility* of Blanchard and Summers (1986) hysteresis-type effects being uncovered in the data.

- Oil shocks and LM shocks should have no long-run effect on capacity utilisation: $C_{51}(1) = C_{55}(1) = 0$. These restrictions are based upon interpreting the capital adjustment costs and demand uncertainty factors that Greenwood *et al* (*op cit*) and Finn (*op cit*) argue should underlie long-run CU movements as being most closely related, respectively, to the technology and real demand (IS) shocks we identify. This interpretation is, however, open to question.

3.3 Results

Since we don't engage in formal overidentification tests, we will begin by discussing what we might call 'informal' overidentification — examining if the forecast error variance decompositions (FEVDs) and the impulse responses accord with our economic priors. We also examine how close the unimposed overidentifying restrictions are to holding in our exactly identified system. And we discuss what inferences (if any) we can make about the underlying economic structure in the United Kingdom, before going on to construct the measures of real disequilibria.

3.3.1 Forecast error variance decompositions (FEVDs)

Our first act of informal overidentification is to examine the FEVDs. These tell us which shocks were the primary *sources* of movement in the endogenous variables over the sample period. In particular, FEVDs tell us the proportion of the forecast error variance

---

(25) This restriction does, however, rest on the assumption that the coefficients on productivity shocks are the same in Layard *et al*'s (*op cit*) price-setting and wage-setting schedules.

of each endogenous variable at different forecast horizons attributable to each of the shocks.

Chart 3 presents our FEVD results; the structural shocks appear on the vertical axis and the endogenous variables on the horizontal axis. (So, for example, the relative importance of unemployment shocks in GDP movements is presented in the graph in the fourth row, third column of the chart.) The vertical axis of each of the graphs in the chart represents the proportion of the endogenous variable movement that the shock accounts for — ranging, of course, from zero to 1. The horizontal axis of each of the graphs represents the forecast horizon of interest. To stress the fact that the identifying restrictions are not constrained to be imposed at any finite horizon we plot the FEVDs up to 100 quarters (25 years) out. We present the FEVDs on the *levels* of the endogenous variables in the chart because these correspond most closely to the restrictions we have in mind from our theory.(26)

Our FEVD results are generally consistent with our theoretical priors. This suggests that our system has at least some economic content — the shocks that we label 'LM' shocks, for example, seem to play the kind of roles we would expect. Our FEVDs do, however, throw up a couple of puzzles.

Our theoretical priors suggest that nominal demand (LM) shocks should not contribute significantly to long-run movements in real variables. And this turns out to be the case. In particular, we find that LM shocks account for only 1.9% of unemployment and 3.8% of capacity utilisation movements in the long run (the $C_{35}(1)=0$ identifying restriction imposes that prior on output). Interestingly, LM shocks have little effect on real variables in the short run either. This finding — which is supportive of the 'real business cycle' explanation for output fluctuations — is actually quite a common one.(27)

We also expect that real demand ('IS') shocks should make minor contributions to long-run real variable movements. Our result again generally confirm this prior. The exception is that we find that IS shocks account for 40% of capacity utilisation movements in the long run. This suggests that our Section 3.2 interpretation of IS shocks as loosely proxying demand uncertainty factors may be reasonable.

---

(26) The FEVDs of the first differences of variables produced broadly similar results to the levels ones. This need not necessarily be the case. This is because the levels (first difference) FEVD of a variable are based upon non-linear transformations of the levels (first differences) impulse responses of that variable to each of the shocks.

(27) Blanchard and Quah (*op cit*), Clarida and Gali (1994), Quah and Vahey (1995) Funke (1995) and Astley and Garratt (1998) all uncover a similar finding.

Our priors suggest that inflation movements should be explained by purely nominal shocks in the long run, but by a combination of nominal and real shocks in the short run. Again such a pattern is apparent in our results (although the long-run result is *partly* imposed.(28) The main puzzle here is that our results show oil shocks contributing almost as much to long-run inflation movements as to short-run ones (13%).

We expect that technology shocks should be an important determinant of long-run output movements. And our results accord with that reasoning — technology shocks account for almost 90% of long-run GDP movements (this plausible result was again only partly imposed(29)). And our finding that technology shocks account for a large proportion (58%) of long-run capacity utilisation suggests that our section 3.2 interpretation of technology shocks as proxying for capital adjustment costs may be reasonable. But our finding that technology shocks are the main determinant of long-run unemployment movements (accounting for 56% of them) is hard to rationalise — Layard, Jackman and Nickell (*op cit*) predict that such shocks should have zero long-run effect. We return to this puzzle below.

---

(28) Our identifying restrictions mean that only LM and oil shocks can affect long-run inflation movements. But is it reassuring that we find that LM shocks rather than oil shocks underlie most of the long-run inflation movements.

(29) Our identifying restrictions mean that technology, oil and unemployment shocks are able to affect GDP in the long run.

Chart 3: Forecast error variance decomposition of variables

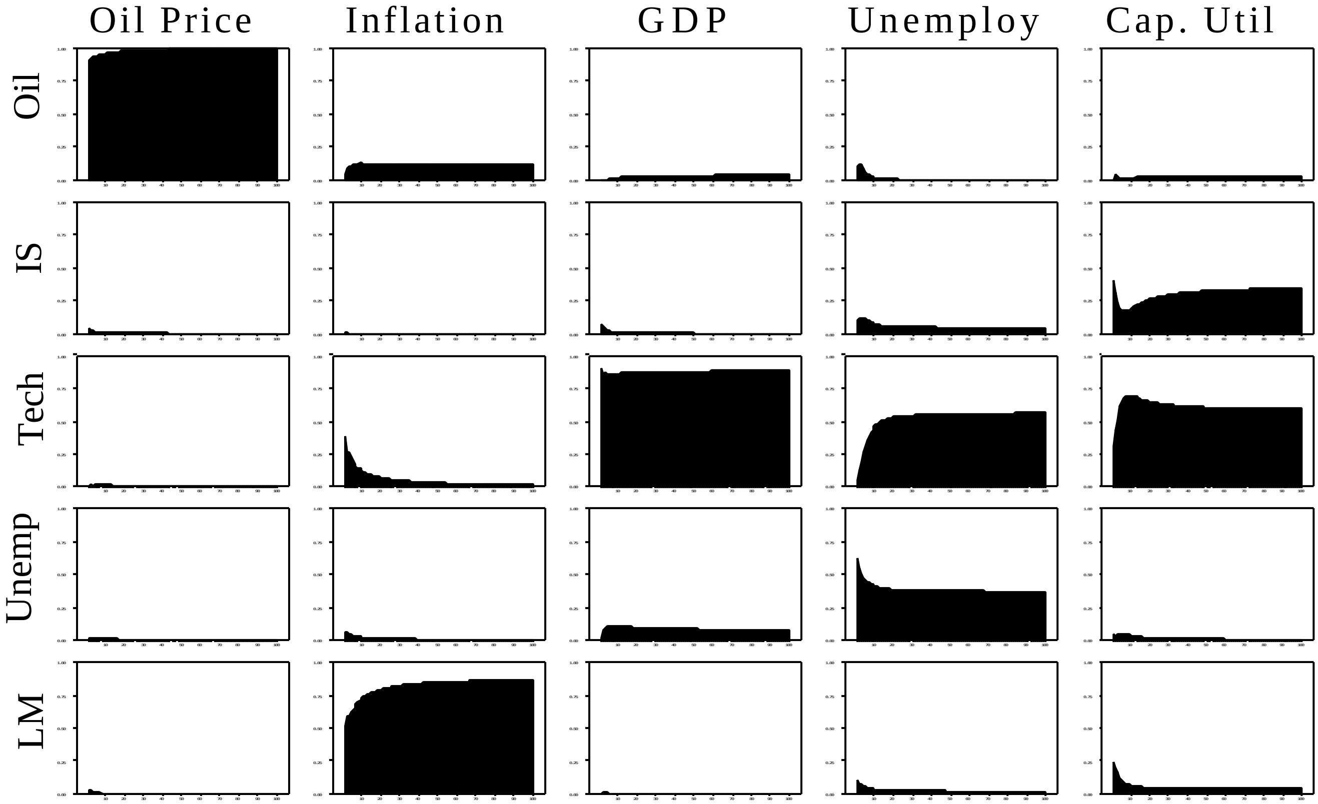

3.3.2 Impulse responses

Our second informal check on the economic content of our system is to examine the impulse responses of the endogenous variables to the structural shocks. Chart 4 presents the estimated impulse responses of the *levels* of the variables (following a unit innovation in each of the structural shocks). The format of the chart is identical to that used to display the FEVD results (see Chart 3). (So, for example, the effect of an IS shock on GDP is presented in the graph in the second row, third column of the chart.) In contrast to the FEVD results, here we only plot the impulse responses for the first forty quarters (ie ten years). This is because the impulse responses have usually roughly ‘levelled off’ by then and this is arguably the period we are most interested in.

Table 2 summarises what theory predicts about the effects of shocks on the endogenous variables. Most of the expected responses are intuitive, and will be discussed below. The only exceptions are the response of unemployment to technology shocks and the response of capacity utilisation to unemployment shocks. Technology shocks have, as Bean (1992) discusses, an ambiguous effect on unemployment. On the one hand the rise in the productivity of labour associated with a positive technology shock raises the demand for labour and hence tends to reduce unemployment.(30) On the other hand, this increased productivity means that fewer workers will be required to produce a given quantity of output — tending to raise unemployment. Turning to the effects of unemployment shocks on capacity utilisation, the predicted negative effect is based upon the assumption that factors of production are less than perfectly substitutable.

Table 2: Expected responses of variables to (positive) shocks

			Effect on variable:		
Shock:	**Oil price**	**Inflation**	**GDP**	**Unemploy**	**Cap. Util**
Oil	↑	↑ (LR=0)	↓	↑	↓ (LR=0)
IS	**? (LR=0)**	**↑ (LR=0)**	**↑ (LR=0)**	↓ (LR=0)	↑
Technology	**? (LR=0)**	**↓ (LR=0)**	↑	**?** (LR=0)	↑
Unemploy	**? (LR=0)**	**↑ (LR=0)**	↓	↑	**↓ (LR=0)**
LM	**? (LR=0)**	↑	**↑ (LR=0)**	↓ (LR=0)	↑ (LR=0)

Key: ↑ = increase; ↓ = decrease; (LR=0) = response expected to go to zero in long run (imposed if in bold); ? = no priors about expected direction of response.

It is apparent that the estimated responses presented in Chart 4 are remarkably similar to the priors of Table 2. In particular, only one of the twenty (non-oil) impulse responses we are interested in is incorrectly signed. (Although several of the

---

(30) Bean (1992) argues that this effect will dominate if ‘the elasticity of the marginal product of labour (in efficiency units) is less than unity in absolute value’.

responses that are correctly signed overall are incorrectly signed at a few horizons.(31)) This tells us that our identification scheme is plausible and has generated shocks which we can interpret in the way suggested by the labels (IS, technology etc) that we have given them. And it implies that the FEVD results of Chart 3 are not just statistical artefacts.

For example, our priors suggest that all three types of supply shocks (oil, technology and unemployment) should generate a negative (positive) co-movement between inflation and GDP (unemployment). And this is precisely the pattern apparent in Chart 4. Similarly, we expect that the two demand shocks (IS and LM) will generate a positive (negative) co-movement between inflation and GDP (unemployment). And this is precisely what we uncover.

It is also apparent, however, that several of the responses that theory suggests should asymptote to zero, end up not doing so when we leave them unconstrained in our exactly identified system. The most obvious examples are the persistent effects that IS, technology and LM shocks have on unemployment.(32) We consider those 'puzzles' further in the next section. There are several other puzzling aspects of the impulse responses. First, some of the very short-run responses to oil shocks are surprising. For example, unemployment is lower for about two years after a positive oil shock — suggesting, if we take the impulse response at face value, that firms substitute into labour-intensive methods of production when the cost of running machines increases. This is not completely implausible, but, we think, unlikely — since the short-run costs of adjusting factors will probably be higher for capital than for labour. Second, the response of capacity utilisation to IS shocks is counter-intuitive. This could, among other things, indicate problems with the capacity utilisation data.

---

(31) Combining these two sets of problems means that our system would 'score' around 18 out of a possible 20 impulse responses being correctly signed and interpretable.

(32) Similarly, the effects of neither oil nor LM shocks on capacity utilisation asymptote to zero. But we view this as less of a puzzle than the unemployment responses — because our theoretical priors are weaker about which shocks should drive long-run capacity utilisation movements.

Chart 4: Impulse responses of levels of variables

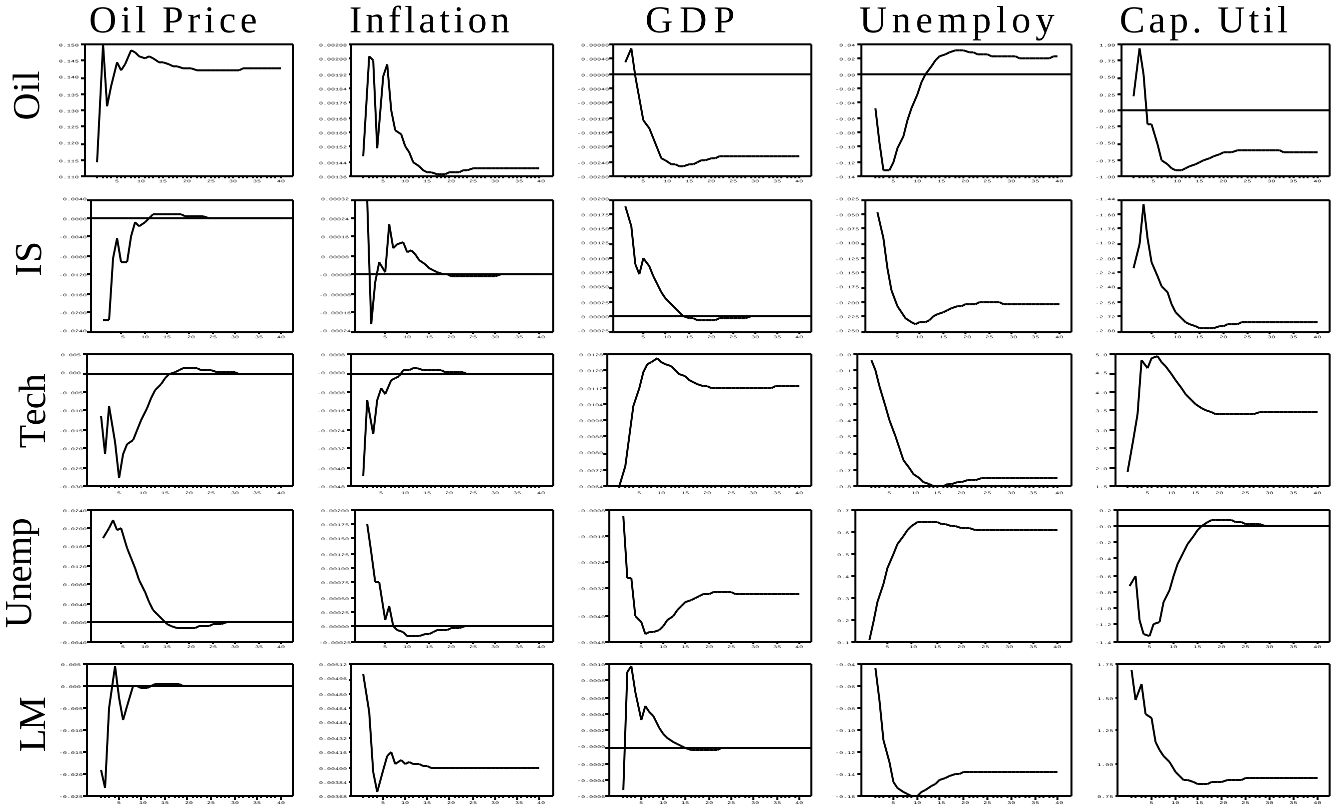

3.4 Making economic inferences from the impulse responses

One of our motivations for considering different measures of real disequilibria was to enable us to make some kind of inference about the structure of the economy from the comovement of gaps in response to particular types of shock. We do this directly by examining selected impulse responses.

We first return to the finding that technology, IS and LM shocks all have persistent effects on unemployment. These findings are contrary to the assumptions built into most conventional models (like Bean (1992)), but could indicate weak hysteresis effects along the lines of Blanchard and Summers (1986). So we can either query the identification scheme or use our results to shed light on competing theories.

This pattern is most pronounced for technology shocks (because our FEVD results reveal that technology shocks account for a large proportion of long-run unemployment movements). This casts doubt on those models of unemployment that are technology-neutral at all horizons (those based upon real inertia in labour markets). And it indicates that the Layard *et al* (1991) assumption that productivity shocks have the same effect on their wage-setting and price-setting schedules does not hold. Alternatively, of course, this result could cast doubt on our having correctly identified technology shocks. But several other aspects of our results suggest that this is not the case. First, our result that positive technology shocks generate falls in unemployment is in line with the Bean (1992) findings. Second, technology shocks appear to play plausible roles in sub-period endogenous variable movements (see below).

Our finding that *real* demand (IS) shocks persistently affect unemployment, although accounting for a small proportion (4.3%) of long-run unemployment movements, could be telling us many things. First, it points to weak Blanchard and Summers (1986) type hysteresis. Second, it lends some support to models where unemployment is determined by real factors, but where the adjustment back to equilibrium (in the long run independent of labour demand) is very slow. Third, the result could again be a statistical artefact, reflecting the fact that we have left this response unconstrained. What our results definitely do show, however, is that we can reject the hypothesis of the economy operating on a ‘backward-L’ labour supply curve (where labour supply is horizontal at the level of benefits, and vertical at the total labour supply). This is because such a theory requires that output and unemployment respond in a similar manner to IS shocks (see the Section 2.4 example) — which they clearly do not in our results.

Our finding that nominal demand (LM) shocks affect unemployment suggests that the labour market is characterised by some degree of nominal rigidity: unemployment cannot be a purely ‘real’ phenomenon if in the short run, it responds to nominal shocks.(33)

The estimated impulse responses also show that that LM shocks — Roberts’ (1993) ‘target inflation’ shocks — affect output in the short run. This is consistent with the Fuhrer and Moore (1995) model (Phillips curve **(5)** above, written in terms of rates of change of inflation) which emphasises that, contrary to the Taylor (1979, 1980) model, inflation is costly to reduce — even under rational expectations and credible policy announcements. This result is, of course, also consistent with models with non-rational expectations and incredible policy announcements. Against this, our finding that inflation itself responds quite speedily to a variety of shocks is consistent with models with only limited degrees of nominal frictions.

Our results also show that demand (IS and LM) shocks affect capacity utilisation. The rise in capacity utilisation generated by a positive LM shock, together with the concomitant rise in GDP, points to the existence of costs of adjusting capital. The counterintuitive response of capacity utilisation to IS shocks, however, casts doubt on the ability of capacity utilisation data to deliver economically interpretable responses, so we would not want to place too much weight on this.

A summary inference from our impulse responses, and by far the most important one to take away from this paper, is that the impulse responses of our different real variables (unemployment, output and capacity utilisation) to a given shock (demand or supply, nominal or real) are different. So, of course, the response of our different measures of the real disequilibria to given shocks will also be different. In other words, our results indicate that the very restrictive assumptions (outlined in Section 2.4)(34) necessary for the movements in different real variables following a given shock to be highly correlated do not hold. Granted, this finding might not surprise an informed reader, but it emphasises the merits of considering different real gaps in a consistent framework, and of not moving — Okun Style — seamlessly from talking in terms of ‘output gaps’ to talking in terms of ‘unemployment gaps’.

3.5 SVAR estimates of real disequilibria

This section presents our estimates of the time series of the output gap, the unemployment gap and the capacity utilisation gap. We also present our accompanying estimates of potential output and the natural rate of unemployment. Our

---

(33) Though the impulse responses show that LM shocks have persistent effects on unemployment, the FEVDs show that LM shocks account for a very small proportion (1.9%) of long-run unemployment movements.

(34) No factor substitution, only one source of nominal friction, capital adjustment costs.

gaps are going to be constructed by adding up the effect of different shocks on the endogenous variables. So the first task is to decide which shocks underlie the gap components of each of our real variables. Our decisions are based largely on theoretical considerations. In particular, each of our gaps is composed of the effects of those shocks that theory suggests should have no long-run effect on that variable. This means that, contrary to some approaches, the unemployment and capacity utilisation gaps include both demand and supply shocks. It also means that both those gaps include shocks whose long-run effects do not asymptote to zero in our results. This theoretically based approach is, of course, least well defined for the capacity utilisation — simply because, as previously discussed, the theory is a little vague!

This approach means that our output gap reflects the sum of the effects of IS and LM shocks on output. Similarly, the unemployment gap is the sum of the effects of IS, technology and LM shocks on unemployment. Our inclusion of technology shocks in the unemployment gap is, of course, the most important example of including a shock whose long-run effects do not asymptote to zero in our results. Finally, our capacity utilisation gap is the sum of the effects of oil, unemployment and LM shocks on capacity utilisation.

Since we cannot uncover information about the shocks before the sample period begins, all of the gaps are constrained to equal zero initially. Unless it happened to be the case that there were no shocks currently working through the real variables at the start of the sample, this will mean that our estimates of the gaps should not be taken seriously around that period. Just how much of a problem this is will vary from gap to gap, depending on how long it takes the relevant shocks to have their full effects on output, unemployment or capacity utilisation. We take comfort from the fact that when we experimented with the start-date, the end-of-sample estimates of the gaps were unaffected. This ‘initial equilibrium’ problem affects the results of previous SVAR output gap papers — eg DeSerres *et al* (1995), Sterne and Bayoumi (1995), Funke (1997) and St Amant and Van Norden (1997) — but it does not seem to have been discussed before.

Chart 5 plots the five structural shocks which are used to generate the gap series. We do not comment on them in isolation here, but we will refer to them as we discuss each of the gaps.

Chart 5: Structural shocks derived from SVAR

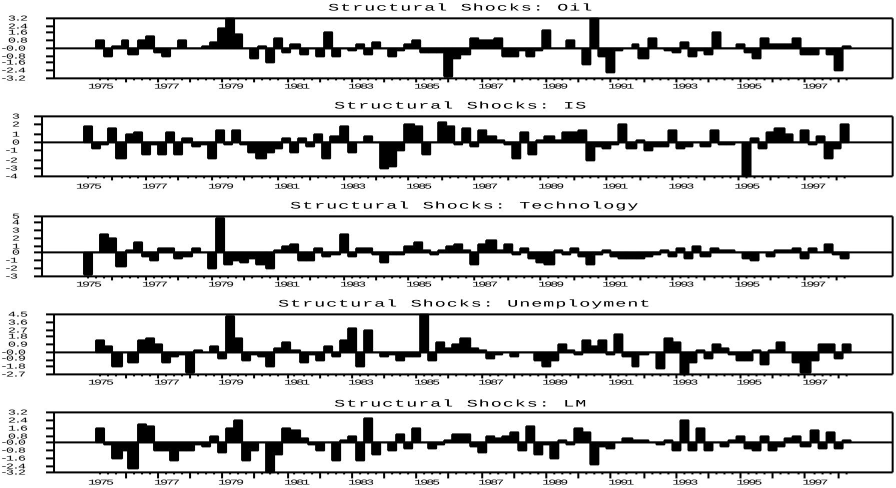

3.5.1 The SVAR output gap

We calculate the output gap (y_t^{gap}) as follows:

$$y_t^{gap} = \sum_j C_{32}(j)e^{IS}{}_{(t-j)} + \sum_j C_{35}(j)e^{LM}{}_{(t-j)}$$

where: e^{IS} and e^{LM} represent the IS and LM shocks respectively; C_{32} and C_{35} represent the impulse responses of (the level of) output to IS and LM shocks respectively (both of which are restricted to zero in the long run).

So potential output (y_t^{pot}) is calculated as:

$$y_t^{pot} = m + \sum_j C_{31}(j)e^{OIL}{}_{(t-j)} + \sum_j C_{33}(j)e^{TECH}{}_{(t-j)} + \sum_j C_{34}(j)e^{UNEM}{}_{(t-j)}$$

Potential output therefore equals the cumulated effects of the three supply shocks (technology, unemployment and oil) on output, plus the constant in the VAR, m. This is, of course, equal to actual output minus the output gap (as $y_t = y_t^{gap} + y_t^{pot}$); potential output is the level of real GDP that occurs when the effects of the demand disturbances on output have dissipated. Our practise of allocating the VARs unconditional forecasts (or constants) to the trend component follows Sterne and Bayoumi (1995). It reflects our view that we do not expect gaps to have constants in them, since they are motivated by thinking of shocks that have temporary (not constant) effects on real variables.

Chart 6 presents our estimate of the time series of the output gap. It has several intuitive features. First, a large positive output gap (output above potential) is apparent in the mid to late 1980s — when the UK economy was in a ‘boom’ period. Our system attributes this positive output gap to a series of positive IS shocks and, to a lesser extent, a series of positive LM shocks that it uncovers prior to and around this period (see Chart 5). We think that the positive IS shocks can be linked to the loose fiscal policy around this period. Likewise we view the positive LM shocks as reflecting loose monetary conditions around this period — given that financial liberalisation had unleashed excess liquidity. The second intuitive feature is that we can see that there are negative output gaps apparent around what we might from independent sources call the ‘recessions’ of the early 1980s and (to a lesser extent) the early 1990s. We interpret these as reflecting the tight monetary or fiscal policy at the time — which our results pick up as negative LM and IS shocks around those periods.

Chart 6: SVAR derived output gap

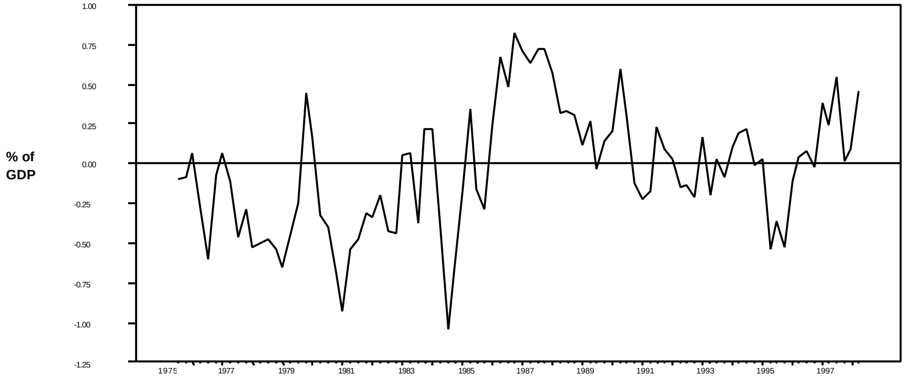

Our estimates show that a positive output gap was present from the start of 1997 until the end of our sample period (1998 Q2) — which was mainly due to the impact of positive IS shocks in 1996 and 1997 (see Chart 5). We think of those as reflecting several factors that raised (liquid) personal sector wealth — tax cuts,(35) equity price rises and building society demutualisations(36) — which were associated with a strong

---

(35) Two tax changes were implemented in 1996 — the basic tax rate was cut by 1% (to 24%) and the lower rate (20%) band was raised by £700. Together these changes were worth around £2 billion.

(36) The demutualised building societies were, of course, previously owned by their members. So the main effect of the demutualisation was to increase the liquidity of this existing wealth.

upturn in consumer expenditure.(37) Notably, however, our estimates indicate that this end-sample positive output gap was smaller than its mid/late 1980s equivalent.

Our estimates of the output gap are small in relation to the observed GDP movements. Chart 7, which plots observed GDP and our measure of potential GDP (as defined above) makes this point. This is, of course, simply a corollary of FEVD finding that demand (IS and LM) shocks are unimportant sources of GDP fluctuations relative to supply shocks.

Chart 7: Actual and potential GDP

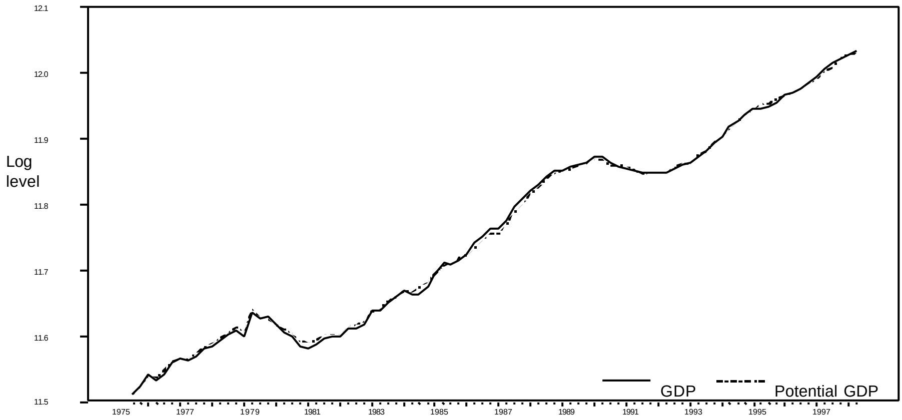

Chart 7 implies that potential output actually falls over some periods. Some of this is undoubtedly due in part to the impact of negative technology shocks — which are apparent in Chart 6 above — or what have elsewhere been called 'negative Solow residuals', which researchers have been uncomfortable about (since if taken literally it implies the possibility of uninventing a machine). But we also have other negative supply-side shocks — unemployment and oil shocks — that can more plausibly be thought of as shifting the production possibility frontier inwards. So the fact that potential output falls at times should not give cause for concern.

How do our estimates compare with alternative methods of estimating the output gaps? We consider two alternatives: the difference between GDP and Hodrick Prescott filtered GDP — Chart 8; and the output gap that Fisher *et al* (1996) derived using a 'production function' approach — Chart 9. Because Hodrick Prescott filter results are well known to be sensitive to the value of the smoothing parameter (λ) Chart 10 reports output gaps generated using three λ values — 100, 1600 (often the 'default' choice)

---

(37) Annual consumption growth rose from under 2% at the end of 1995 to 5% by the end of 1997.

and 5000. A lower λ constrains the HP filtered series to follow actual GDP movements more closely — and so will generate smaller output gaps.

The most obvious similarity between our SVAR generated output gap series and the two alternatives is that all three series point to a large positive output gap in the mid/late 1980s. But there are several differences:

- The SVAR-derived output gap is generally smaller than that obtained using the three other approaches. For example, the SVAR estimates of the late 1980s output gap are around a third of the size of that generated by the HP filters. Funke (1997) uncovered a similar finding, without explaining it. We think that the explanation lies in the fact that straight line and moving average filters attribute an arbitrarily large proportion of observed output movements to demand shocks (see Section 2.2).

- Negative output gaps in the early 1980s and (especially) the early 1990s are more apparent in the HP filter and production function approaches than they were in our SVAR-derived series.

- The HP filter and production function approaches both point to positive output gaps in the mid to late 1970s, whereas our SVAR points to a negative output gap.

Chart 8: Output gaps generated by HP filter

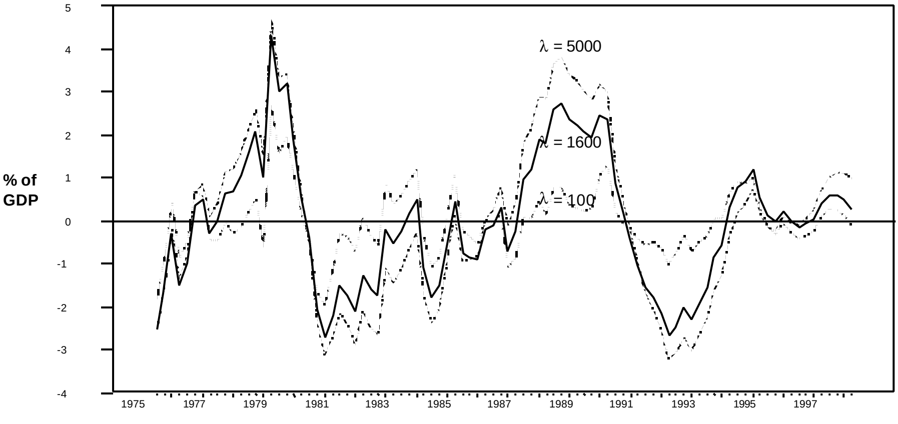

Chart 9: Output gap derived from production function approach

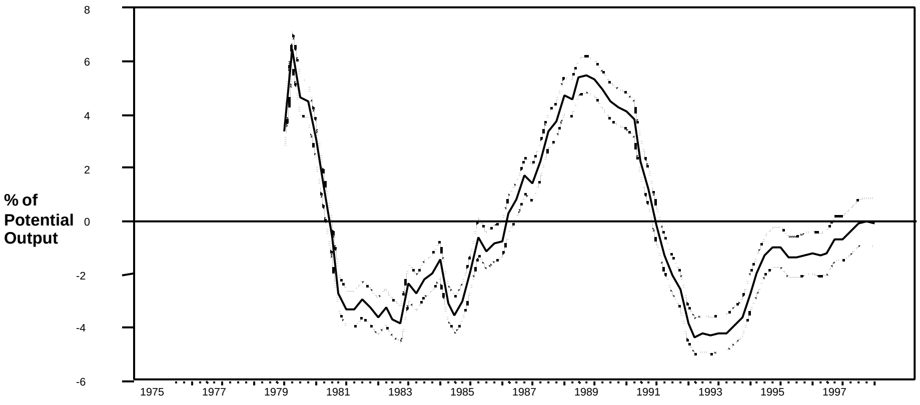

Key: Dark line denotes point estimate, dashed lines the 75% confidence intervals

3.5.2 The SVAR unemployment gap

We use an analogous method to calculate our unemployment gap time series (u_t^{gap}):

$$u_t^{gap} = \sum_j C_{42}(j)e^{IS}{}_{(t-j)} + \sum_j C_{43}(j)e^{TECH}{}_{(t-j)} + \sum_j C_{45}(j)e^{LM}{}_{(t-j)}$$

Or, in words, the sum of the effects of those shocks that theory indicates should have a temporary effect on unemployment. Here we need to bear in mind that our results indicate that technology shocks (and, to a considerable lesser extent, IS shocks) have permanent effects on unemployment.

Our estimate of the time series of the natural rate of unemployment (u_t^{nat}) is also constructed in an analogous fashion to our potential output series. In particular, it is the difference between measured unemployment and our estimate of the unemployment gap — so it represents the cumulated effects on unemployment of the shocks that theory indicates will permanently affect unemployment (unemployment and oil shocks), plus the constant in the VAR (α):

$$u_t^{nat} = a + \sum_j C_{41}(j)e^{OIL}{}_{(t-j)} + \sum_j C_{44}(j)e^{UNEM}{}_{(t-j)}$$

Chart 10 presents our estimate of the time series profile of the unemployment gap. Several things are apparent. First, the unemployment gaps are very persistent. This, of course, reflects the persistent effects(38) that our estimates suggest that IS, LM and (especially) technology shocks have on unemployment. Second, the unemployment gaps are considerably larger than their output gap equivalents. This means, of course, that the measured unemployment rate diverges quite markedly from the natural rate of unemployment (see Chart 11) and, because of the persistence of the unemployment gap series, does so for substantial periods.

Cromb (1993) and Coulton and Cromb (1994) survey previous estimates of what has been called, a little misleadingly, the 'NAIRU' in the United Kingdom; Table 3 summarises the results. This NAIRU, more informatively called a Non-Increasing-Inflation-Rate-of-Unemployment, is not strictly comparable with our u_t^{nat} series. In particular, the NAIRU will differ from u_t^{nat} if factors generating nominal inertia are present.(39) But we persist with this comparison because it is the most complete one available (the Table 3 figures are based upon a large number of studies). And it is apparent that the movements in, if not always the level of, our u_t^{nat} are broadly consistent with the previous NAIRU estimates presented in Table 3.(40)

Table 3: Summary of UK NAIRU estimates

	Period within which estimate falls:			
	1969-73	1974-80	1981-87	1988-90
Range of NAIRU estimates	1.6%-5.6%	4.5%-7.3%	5.2%-9.9%	3.5%-8.1%
Average of NAIRU estimates	2.9%	5.7%	7.0%	6.1%
Actual unemployment rate	2.5%	3.8%	10.1%	6.8%

Source: Cromb (1993), Coulton and Cromb (1994)

A potentially puzzling aspect of our results is that they suggest that the natural rate of unemployment increased in 1991/92. The plausible aspect of this finding is that our

---

(38) Strictly speaking, of course, the decomposition of measured unemployment into the unemployment gap and the natural rate of unemployment does not make sense if the labour market is characterised by full hysteresis — the limiting case of persistence. In that case, any change in unemployment represents a change in the natural rate. We will proceed on the assumption that IS, LM and technology shocks do have persistent effects (consistent with our impulse responses) but that this persistence falls short of full hysteresis (consistent with our FEVD results).

(39) Factors generating real inertia underlie movements in both the NAIRU and u_t^{nat}.

(40) In particular, our finding that the natural rate rises gradually, if somewhat erratically, from the mid 1970s to the mid to late 1980s is also apparent in Table 3. Similarly, our finding that the natural rate falls — again somewhat erratically — from the late 1980s onwards is apparent in Table 3, at least up until 1990 when the estimates end. Saleheen (1998) and Saleheen and Westaway (1997) provide more up-to-date NAIRU estimates (using a Kalman filter and a 'structural' Layard-Nickell approach respectively). Both studies show that the NAIRU fell from the late 1980s onwards, as does our natural rate (as well as again broadly supporting our findings for earlier in the sample period).

results suggest that part of this increase reflected the delayed(41) impact of the positive (gulf war induced) oil shock that we uncover in 1990. Less plausible, however, is our finding that a large part of the increase reflected (positive) natural rate of unemployment shocks; the factors that such natural rate shocks represent (eg union militancy) seem unlikely to rise during a recession. We are not alone, however, in uncovering an increase in the natural rate of unemployment around this period — Saleheen (*op cit*) uncovers a similar result.(42)

Chart 10: SVAR generated unemployment gap

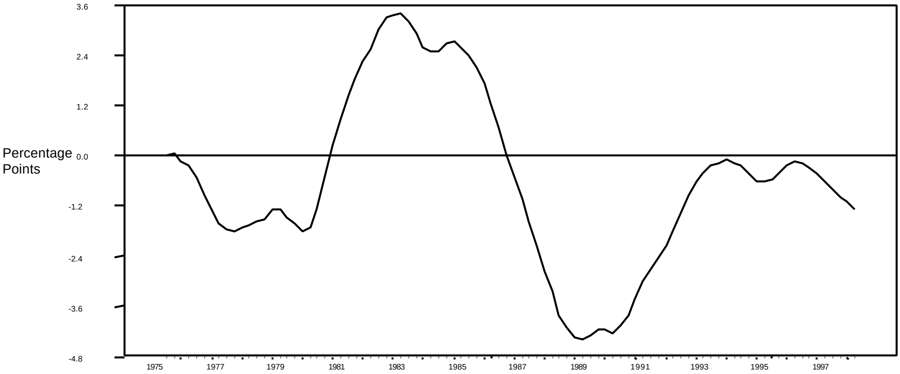

Chart 11: Actual unemployment vs SVAR generated natural rate

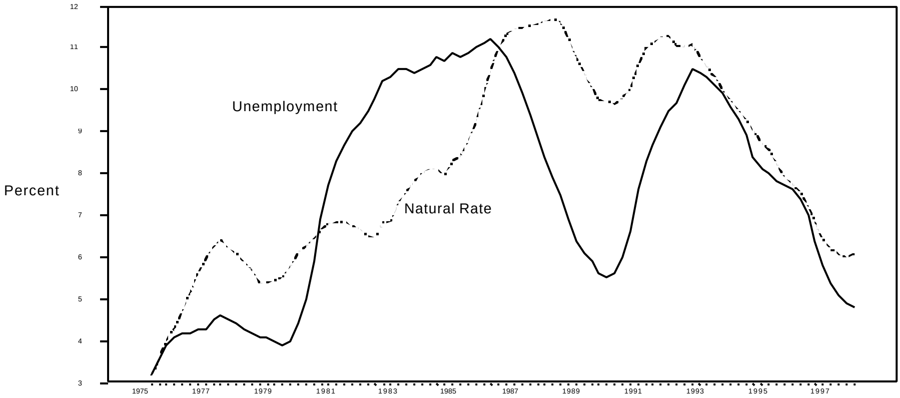

---

(41) Section 3.3.2 showed that it takes over a year for a positive oil shocks to induce a rise in unemployment.

(42) Interestingly, however, Saleheen and Westaway (*op cit*) do not.

How does our unemployment gap compare with the one derived from an HP filter(43) (presented in Chart 12 overleaf)? It is apparent that the two series share the same broad cyclical pattern. In particular, both methods point to: negative unemployment gaps (measured unemployment below its natural rate) at the end of the 1970s; positive unemployment gaps in the early to mid 1980s; negative unemployment gaps in the mid to late 1980s (ie around the 1980s boom); and a negative unemployment gap from the start of 1997 until the end of the sample period (1998 Q2). But there are also some differences. First, the SVAR-derived unemployment gaps shows larger and more persistent swings than the HP filter series. Second, the SVAR-derived series does not pick up the positive early 1990s unemployment gap apparent in the HP filter derived series.

Given these differences it again seems sensible to take a closer look at our results with a view to determining their economic content. In particular, as with our output gap results, we examine which of the three shocks (IS, technology and LM) were the prime source of our estimated movements in the unemployment gap (in conjunction with their persistent effects) and consider whether we can relate those shocks to observed (‘real world’ or ‘off model’) economic developments. We are encouraged that our results suggest that:

- The positive unemployment gaps in the early to mid 1980s reflect the persistent effect of negative IS and LM shocks that often occurred during, and prior to, that period. Such shocks seem reasonably congruent with economic history — corresponding to the fiscal retrenchment of the first Thatcher government, and the attendant monetary tightening associated with the Medium Term Financial Strategy.

- The switch from a positive unemployment gap in the mid 1980s to a negative unemployment gap in the late 1980s was due to a series of positive technology shocks (the ‘productivity miracle’?) and, to a lesser extent, positive IS shocks(44) that reversed the effects of the earlier negative IS and LM shocks.(45) The positive IS shocks — which coincided with fiscal loosening — are familiar from our discussion of the positive mid/late 1980s output gap — representing the Lawson fiscal loosening.

---

(43) Calculated using the same λ values used in the equivalent output decomposition.

(44) Though we also uncover positive LM shocks in the mid 1980s — probably reflecting high money growth around that period — our results suggest that these shocks played a small role in reducing the unemployment gap (but contributed to the positive output gap).

(45) We think of the series of positive technology shocks that we uncover between 1983 and 1988 as representing the improvement in UK productivity that Holland and Scott (1997) found evidence for over a similar period.

- The gradual elimination of the negative unemployment gap in the early 1990s was due to the impact of negative technology shocks and negative IS shocks. The former can be linked to the fall in UK productivity that Holland and Scott (*op cit*) uncovered over this period. The latter probably reflect the downwards revision of income expectations associated with the early 1990's recession.

Chart 12: Unemployment gap generated by HP filter

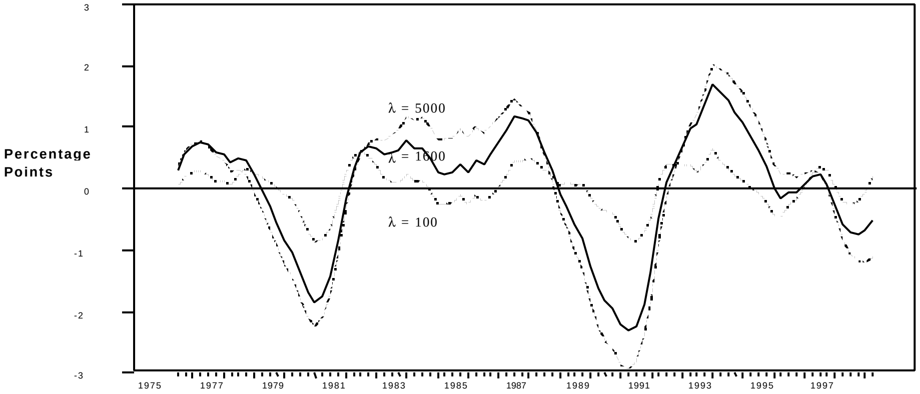

As previously mentioned, our SVAR estimates suggest — in accordance with the HP filter results — that a negative unemployment gap existed in 1998 Q2 (our most recent data point). Our results point to two factors here. First, a series of small positive technology shocks that occurred from 1996 onwards. Second, the positive IS shocks uncovered in 1996/97— which Section 3.5.1 argued reflected the positive wealth effects of tax cuts, equity price increases and building society demutualisations and contributed to the positive end-sample output gap.

3.5.3 The SVAR capacity utilisation gap

Our capacity utilisation gap (CU_t^{gap}) is given by:

$$cu_t^{gap} = \sum_j C_{51}(j)e^{OIL}{}_{(t-j)} + \sum_j C_{54}(j)e^{UNEM}{}_{(t-j)} + \sum_j C_{55}(j)e^{LM}{}_{(t-j)}$$

Chart 13 plots our estimate of the time series profile of the capacity utilisation gap. This shows that negative CU gaps (capacity utilisation below its 'natural' rate) occurred around the recessions of the early 1980s and early 1990s. Our results indicate that the former reflected a combination of positive unemployment shocks, negative LM shocks and positive oil shocks. But we find that negative LM shocks played a smaller role in the early 1990s negative CU gap — which was due more to positive unemployment shocks and positive oil shocks (perhaps related to the gulf war).

Chart 13: SVAR derived capacity utilisation gap

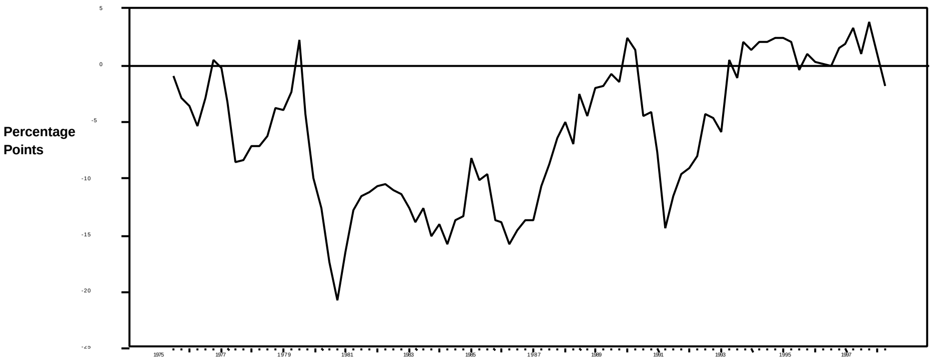

The CU gap was only mildly positive during what conventional wisdom might call the ‘late 1980s boom’. Our results show that this was mainly due to the negative unemployment shocks around that period. Interestingly, the positive CU gap apparent from late 1993 is only manifest from early 1997 in the output gap series. (The CU gap was mainly due to the series of negative unemployment shocks that our system uncovers from early 1993 onwards, reinforced by positive LM shocks in 1996/97.) Note also that our results suggest that end-period positive CU gap was *larger* than its late 1980s equivalent, which contrasts with our estimates of the output and unemployment gap — both of which were smaller at the end of the sample than in the late 1980s.(46)

How do our estimates of the CU gap correspond to those derived from the HP filter? Chart 14 plots the HP filter derived. Like our SVAR-derived series, these point to negative CU gaps in the early 1980s and early 1990s, and positive CU gaps in the late 1980s and at the end of our sample (notably again from late 1993). But there are several differences.(47)

---

(46) This may reflect the previously noted fact that our capacity utilisation data relate to the manufacturing sector whereas our output and unemployment data relate to the whole economy. Or it could reflect the fact that the gaps contain different shocks.

(47) For the sake of completeness, the main differences are that the HP filter series: (i) pick up a positive CU gap in the late 1970s that is far less apparent in the SVAR series; (ii) point to a less prolonged negative CU gap in the early 1980s than the SVAR does; (iii) indicate a larger CU gap in the late 1980s than the SVAR does; and (iv) points to a smaller CU gap at the end of the sample than does the SVAR. Indeed, notably, the HP filter series suggest that the end-sample positive CU gaps was smaller than their late-1980s equivalent.

Chart 14: Capacity utilisation gaps generated by HP filter

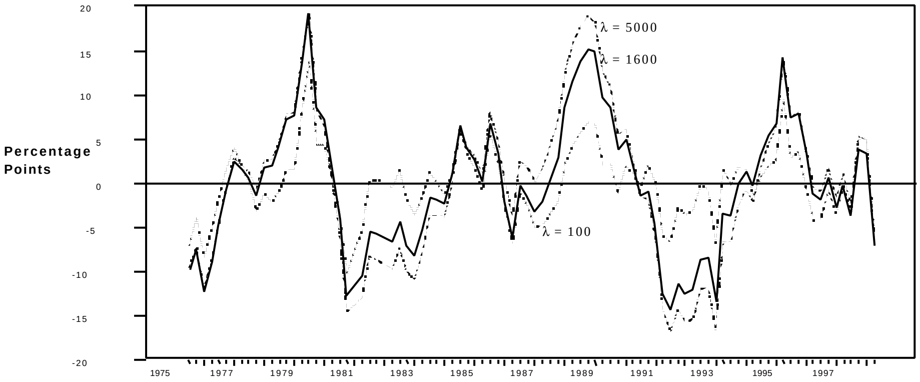

3.5.4 Comparing the three SVAR-generated gaps

Chart 15 summarises what should already be apparent from our discussion, that the size and dynamics of our three gaps are quite different. This is hardly surprising: as we pointed out when we discussed the impulse responses, it was clear that output, unemployment and capacity utilisation respond quite differently to the shocks that are common to these gaps.(48) And those shocks that are not common to each gap clearly do not work to offset the differences.

---

(48) LM shocks for all three gaps, IS and LM shocks for output and unemployment gaps.

Chart 15: SVAR estimates of real disequilibria

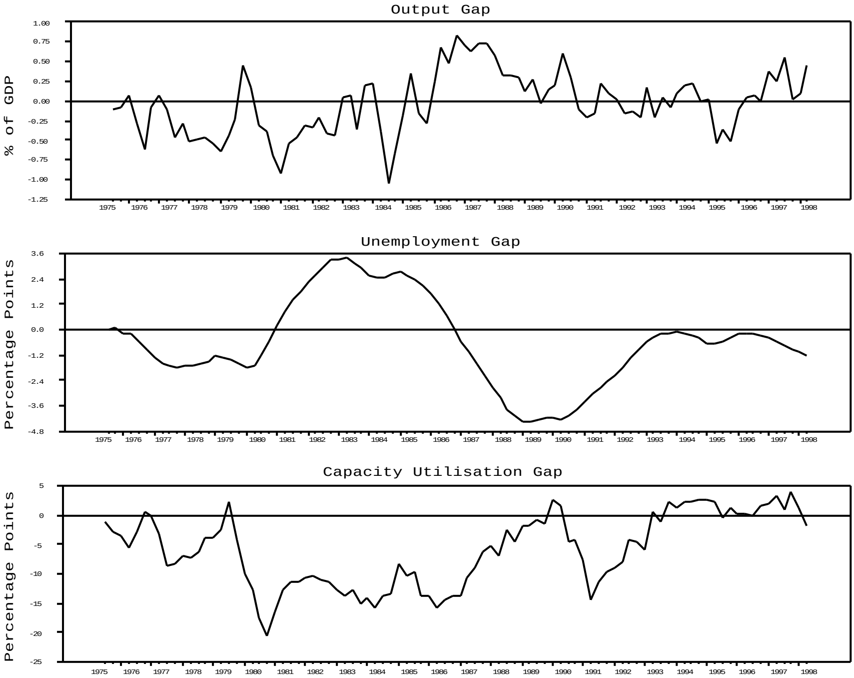

Table 4 below makes this point a little more formally, by showing how the gaps are correlated at different leads and lags. The table does not give us any new information, but it does bring out what is embedded in the comparisons of the different underlying shocks and impulse responses.

Table 4: Correlation between SVAR output, unemployment and capacity utilisation gaps

		Correlation of (i) at t with (ii) at t - n						
(i)	(ii)	n=-3	n=-2	n=-1	n=0	n=1	n=2	n=3
YGAP	UGAP	-0.54**	-0.51**	-0.45**	-0.37**	-0.26*	-0.14	-0.03
YGAP	CUGAP	-0.12	-0.24*	-0.27**	-0.29*	-0.24*	-0.23*	-0.18
UGAP	CUGAP	-0.56**	-0.51**	-0.46**	-0.42**	-0.39**	-0.37**	-0.34**

Key: *=significant at 5%; **=significant at 1%.

Several things are clear. First, as would be expected, the output gap is significantly negatively correlated with the unemployment gap. Second, the output gap also tends to lead the unemployment gap. Third, contrary to expectations, the output gap is negatively correlated with the CU gap, with no tendency for one to lead the other. Fourth, the unemployment gap is significantly negatively correlated with the capacity utilisation gap (although there is no clear pattern of the unemployment gap leading the capacity utilisation gap or *vice versa*).

In some sense these results suggest that output provides an ‘earlier’ signal of real disequilibria than does unemployment or capacity utilisation. Chart 15 suggests that this is primarily due to the output gap being the first of our three gaps to pick up the mid/late 1980s boom. The difference has been less apparent in the most recent expansionary phase.

Taken at face value, these results emphasise the concern that motivated our paper: that care needs to be taken in interpreting the signals about demand conditions coming from temporary movements in output, unemployment and capacity utilisation.

3.5.5 Can we forecast inflation with real disequilibria?

One objective of our paper was to comment on how previous work has modelled the relation between real disequilibria and inflation, as reflected in equations resembling **(1)-(6)**. Section 2.3 argued that such equations were flawed — because they assumed that causality ran from real to nominal variables, rather than treating both as endogenous variables and influenced by more fundamental demand and supply shocks. So a short answer to the question ‘Can we forecast inflation with real disequilibria?’ might be: ‘no’. Instead we would reformulate the question as ‘how does inflation respond to shocks that also cause real disequilibria to open up?’ Our answer to this question is contained in the discussion of the estimated impulse responses (of inflation). But note that in general we are unlikely to observe much of a correlation between inflation and real disequilibria. The main reason for this is that the shocks that dominate inflation movements (LM shocks) are different from those that dominate movements in our real variables (mainly technology and unemployment shocks). On top of this, we are unlikely to observe our measures of real disequilibria leading inflation movements. This is because our impulse responses reveal little evidence of real variables responding more quickly than inflation to structural shocks.(49)

It is not surprising therefore that Table 5, which shows the correlation of real disequilibria with inflation at different leads and lags, tells us (i) that the correlations are pretty weak and (ii) that there is no lead of real variables over nominal ones: if

---

(49) In particular: unemployment appear to respond as quickly as inflation to IS shocks but slower than inflation following technology and LM shocks; output responds at least as quickly as inflation to IS shocks, but more slowly following LM shocks; capacity utilisation responds as quickly as inflation to LM shocks and unemployment shocks, but slower following oil shocks.

anything, the reverse is true.(50) In other words, these results confirm our theoretical arguments against using *reduced-form* Phillips curve equations to forecast inflation.

Table 5: Correlations between the gaps and lead/lagged RPIX inflation (1981Q1-1998Q1)

		Correlation of (i) at t with (ii) at t - n						
(i)	(ii)	n=-3	n=-2	n=-1	n=0	n=1	n=2	n=3
YGAP	Inflation	-0.01	-0.06	-0.17	-0.22	-0.20	-0.31**	-0.28*
UGAP	Inflation	-0.05	-0.01	0.05	0.10	0.15	0.18	0.23*
CUGAP	Inflation	-0.17	-0.17	-0.20	-0.24	-0.24*	-0.26*	-0.34**

Key: *=significant at 5%; **=significant at 1%.

This result leads to the obvious question: 'how should we understand or forecast inflation movements then?'. Our answer is to repeat the reasoning of Section 2.3: a better approach is to identify the structural shocks hitting and translate them, through the estimated impulse responses, into inflation movements. And it turns out that some fairly plausible results are obtained when we do this for inflation movements observed during our sample period. In particular, and as previously stated, our FEVD results show that LM shocks are the main source of inflation movements over the whole sample period, especially at long horizons. But we also uncover other shocks having plausible effects on sub-sample inflation movements. For example, we find that positive oil shocks and positive unemployment shocks — reflecting the second OPEC oil price increase and increased union militancy ('winter of discontent') — played important roles in the late 1970s increase in inflation.

4 Conclusions

This paper has tried to contribute to the debate on how to construct real gaps, or disequilibria and model their relationship with inflation. Rather than use a single equation Phillips curve with causality running from the real gap to inflation, we suggest treating both variables as endogenous. Rather than constructing the real gap separately from the task of modelling its relation to inflation using statistical detrending methods, we suggest using both variables to identify fundamental demand and supply (and real and nominal) shocks and then construct the gap from these shocks. Rather

---

(50) We calculate those correlations over the post-1981 period because most of the gaps will not, by definition, be affected by the positive oil shocks that our results indicate underlie a high proportion of the end-1970s rise in inflation. Table 5 shows that the correlation coefficients between the gaps and inflation are often contrary to the priors of users of Phillips curves. In particular, the capacity utilisation gap is *negatively* correlated with inflation at all leads/lags and the output gap is *negatively* correlated with inflation except at the longest lead. And while the unemployment gap is — in accordance with Phillips curves - negatively correlated with future inflation movements, these correlations are insignificant.

than consider only one measure of the real gap, we consider an output gap, unemployment gap and a capacity utilisation gap all in the same system: in doing so we discover what we would expect from theory, that these real variables move differently in response to common shocks.

We presented a structural vector autoregression (SVAR) for the United Kingdom to act as a vehicle for our comments and to offer some new results. Our SVAR appears to be economically sensible: supply and demand shocks generate, broadly, the expected correlation between real and nominal variables; and, in the long run, our nominal variable, inflation, is primarily caused by nominal shocks rather than real shocks. We construct three real disequilibria in output, unemployment and capacity utilisation. Our output gap is smaller and more volatile than our other gaps, and than gaps constructed in other papers; our unemployment gap is larger and shows more persistent swings. We are less certain of our capacity utilisation gap.

We use the SVAR to make some tentative inferences about the economy. For a start, we can put aside the hypothetical world of no factor substitution, one source of nominal rigidity and capital adjustment costs, a world where the movement in all three variables in response to demand shocks is equivalent. In our system this isn’t the case. We also find that changes in inflation are costly in terms of output, which suggests that either we should prefer a Fuhrer and Moore (1995) style model over Taylor (1980), or that expectations are sluggish.

The empirical model presented above was intended as much to embody our comments on previous work as to generate new results. The model itself could and should be extended in many ways. By explicitly positing a monetary policy rule, and identifying monetary policy shocks, rather than embodying all nominal shocks in a single nominal variable; by extending the analysis to include the exchange rate as an additional endogenous variable. The model we have presented would also allow us to construct a measure of core inflation akin to that in Quah and Vahey (1995). But we leave those tasks for future research.(51)

---

(51) We have already constructed a core inflation from our five-variable system, and compared it to the results obtained from an update of the Quah and Vahey bivariate system. We do not present these results here as they are tangential to the main focus of this paper.

Appendix A: An explanation of some Phillips curves

Phillips curve **(2)** comes out of the Rotemberg (1982) model, where firms face costs of adjusting prices that are proportional to the square of the price change. In **(2)**, *b* is a measure of 'real rigidity', the elasticity of firms' desired prices with respect to a change in demand; *c* measures the cost of changing prices relative to the cost of being away from the optimal price; and *e* is a random process moving the optimal price over time (in other words, it is a supply shock). Equation **(2)** looks like our traditional Phillips curve **(1)**: inflation increases if actual output exceeds potential output (because of a demand shock), the more so the lower are menu costs or the higher is the elasticity of the optimal price with respect to demand.

Phillips curve **(3)** emerges from Calvo's (1983) model of staggered price-setting. Here *g* is a parameter that measures the proportion of firms that change prices in any period (strictly, the proportion of firms who receive a random signal that they can change prices). All the other parameters have the same interpretation as in **(2)**. Phillips curve **(3)** also tells us that demand shocks will increase inflation. And this inflation response is larger the higher is the elasticity of the optimal price with respect to changes in demand or the greater the proportion of firms that change prices each period.

Phillips curve **(4)** is based upon the Taylor (1979, 1980) model of overlapping nominal wages, rather than nominal rigidity in prices. In that model employees bid for wages that compensate them for past errors in predicting prices, and for future predicted inflation. Here *s* is simply the constant in the labour supply curve; *b* is the slope of the labour supply curve (how much each extra unit of unemployment reduces the wage at which labour offers itself); and *u* is the unemployment rate. The ε's now represent labour supply shocks, rather than the output supply shocks that they represented in **(2)** and **(3)**. Phillips curve **(4)** says that inflation is a function of current expectations of inflation, a moving average of current and lagged unemployment, current and lagged *expected* unemployment, a moving average of shocks to labour supply and a term in last period's price prediction error.

Phillips curve **(5)** is due to Fuhrer and Moore (1995), who pointed out that **(2)-(4)** do not deliver the persistence or stickiness in inflation observed in the data. This model comes out of a labour supply schedule which posits that workers bargain over *real* wages relative to those set for competing groups in the past and relative to those expected to be set for competing groups in the next period. Roberts (1997) calls it a 'slipped derivative' version of equation **(4)** — because it gives us a Phillips curve that is expressed in terms of rates of change of the variables in **(4)**. Importantly, however, the unemployment expectations terms present in **(4)** are absent from **(5)**.

Phillips curve **(6)** is derived from the Layard, Nickell and Jackman (1991) model. Prices *and* wages are set in advance, and monopolistically competitive firms bargain over their producer surplus with unions. Both firms and workers have to form expectations about future prices, and seek to compensate themselves for past prediction errors (by λ_p and λ_w for price-setters (firms) and wage-setters (workers) respectively). Here β'' is the elasticity of the real wage bid with respect to unemployment, and so is akin to the *b*'s in equations **(2)-(5)** (the slope of the labour supply curve). Finally u_t^* is the rate of unemployment that would prevail in the absence of demand shocks. It is a function of real variables describing the structure of the labour market, and we can think of it as being shocked around by ε's.

Annex B: Data definitions and sources

Oil prices — Average of dollar spot oil prices from the West Texas, Brent and Dubai-Fateh oil fields, which were taken from the Bloomberg database (codes USCRWTIC, EUCRBRDT and PGCRDUBA respectively).

RPIX inflation — Quarterly change in index of retail prices excluding mortgage interest payments. The ONS code for RPIX price index is CHMK. The series can be found in *Economic Trends* (Table 3.1).

GDP — Gross domestic product at constant factor cost (1995 prices). The ONS code is YBHH. The series can be found in the *UK National Accounts* (Table AA1).

Unemployment Rate — Claimant count unemployment rate. The ONS code is BCJE. The series can be the *Economic Trends* (Table 4.4).

Capacity Utilisation — Percentage of manufacturing firms not operating at full capacity. Specifically, the percentage of ‘no’ responses to question 4 of the Confederation of British Industry’s (CBI) quarterly Industrial Trends survey.

References

Adams, C and Coe, D (1990), 'A Systems Approach to Estimating the Natural Rate of Unemployment and Potential Output for the United States', *IMF Staff Papers*, Vol 37, pages 232-93.

Apel, M and Jansson, P (1997), 'System Estimates of Potential Output and the NAIRU', *Sveriges Riksbank Working Paper No. 41*.

Ascari, G (1997), 'Optimising agents, staggered wages and persistence in the real effects of money shocks', *mimeo*, Institute of Economics, University of Copenhagen.

Astley, M and Garratt, T (1998), 'Exchange Rates and Prices: Sources of Sterling Real Exchange Rate Movements 1973-94', *Bank of England Working Paper, No 85*.

Bank of England (1999), *Economic Models at the Bank of England*.

Barrell, R and Sefton, J (1995), 'Output Gaps. Some Evidence from the UK, France and Germany', *National Institute Economic Review*, February, pages 65-73.

Bean, C (1992), 'Identifying the Causes of British Unemployment', *Centre for Economic Performance Discussion Paper No 276*, London School of Economics.

Bean, C (1994), 'European Unemployment: A Survey', *Journal of Economic Literature*, Vol XXXII, pages 573-619.

Beveridge, S and Nelson, C (1981), 'A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the Business Cycle.', *Journal of Monetary Economics*, 7(2), March 1981, pages 151-74.

Bjornland, H (1997), 'Estimating Core Inflation — The Role of Oil Price Shocks and Imported Inflation', *Discussion Paper No 200, Statistics Norway.*

Blanchard, O and Summers, L (1986), 'Hysteresis and the European Unemployment Problem', *NBER Macroeconomics Annual*, pages 15-78.

Blanchard, O, and Quah, D (1989), 'The Dynamic Effects of Aggregate Demand and Aggregate Supply Disturbances', *American Economic Review*, Vol 79, pages 653-58.

Calvo, G (1983), 'Staggered contracts in a utility maximising framework', *Journal of Monetary Economics*, 12 (September), pages 383-98.

Clark, P, Laxton, D and Rose, D (1996), ‘Asymmetry in the US output-inflation nexus’, *IMF Staff Papers* Vol 43, No 1, pages 216-51.

Clarida, R and Gali, J (1994), ‘Sources of Real Exchange-Rate Fluctuations: How Important Are Nominal Shocks?’, *Carnegie Rochester Conference Series on Public Policy*; 41(0), pages 1-56.

Claus, I (1997), ‘A Measure of Underlying Inflation in the United States’, *Bank of Canada Working Paper 97-20*.

Cogley, T and Nason, J (1995), ‘Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research’, *Journal of Economic dynamics and Control*, 19, pages 253-78.

Coulton, B and Cromb, R (1994), ‘The UK NAIRU’, *Government Economic Service Working Paper No. 124*, September.

Cromb, R (1993), ‘A Survey of Recent Econometric Work on the NAIRU’, *Journal of Economic Studies*, Vol 20, No 1/2, pages 27-51.

DeSerres, A, Guay, A and St-Amant, P (1995), ‘Estimating and Projecting Potential Output Using Structural VAR Methodology. The Case of the Mexican Economy ’, *Bank of Canada Working Paper 95-2*.

Dolado, J and Jimeno, J (1995), ‘Why is Spanish Unemployment So High?’, *CEPR Discussion Paper No 1184*.

Dolado, J and Lopez-Salido, J (1996), ‘Hysteresis and Economic Fluctuations (Spain, 1970-94)’, *CEPR Discussion Paper No 1334*.

Donner, A and McCallum, J (1972), ‘The Phillips Curve: An historical note’, *Economica*, pages 322-23.

Evans, G and Reichlin (1994), ‘Information, forecasts, and measurement of the business cycle’, *Journal of Monetary Economics*, 33, pages 233-54.

Faust, J and Leeper, E (1997), ‘When do Long-run Identifying Restrictions give reliable results?’, *Journal of Business and Economic Statistics*, 15 (3), pages 345-53.

Finn, M (1996), ‘A Theory of the Capacity Utilisation / Inflation Relationship’, *Federal Reserve Bank of Richmond Economic Quarterly*, 82/3.

Fisher, I (1926), 'A Statistical Relation Between Unemployment and Price Changes', *International Labour Review*, Vol 13, No. 6, pages 785-92. (Reprinted in Journal of Political Economy, Vol 81, No 2, pages 496-502.)

Fisher, P, Mahadeva, L and Whitley, J (1996), 'The output gap and inflation - experience at the Bank of England', paper for BIS model-builders meeting, Basle, January.

Friedman, M (1968), 'The Role of Monetary Policy', *American Economic Review*, 58, pages 1-17.

Fuhrer, J and Moore, G (1995), 'Inflation persistence' *Quarterly Journal of Economics*, 110, pages 127-160.

Funke, M (1997), 'Supply Potential and output gaps in West German manufacturing', *International Journal of Forecasting*, 13, pages 211-22.

Giorno, C, Richardson, P, Roseveare, D and van de Noord, P (1995), 'Estimating Potential Output , Output Gaps and Structural Budget Balances', *OECD Economic Studies*, No 24.

Gordon, R (1997), 'The time-varying NAIRU and its implications for economic policy', *Journal of Economic Perspectives*, Vol 11, pages 11-32.

Gordon, R (1999), 'Foundations of the Goldilocks Economy: Supply Shocks and the time-varying NAIRU', *Paper Presented at Brookings Panel on Economic Activity*, September.

Greenwood, J, Hercowitz, Z and Huffman, G (1988), 'Investment, Capacity Utilisation and the Real Business Cycle', *American Economic Review*, Vol 78, No 3, pages 402-17.

Holland, A and Scott, A (1997), 'The Determinants of UK Business Cycles', *Bank of England Working Paper, No 58*.

King, R, Plosser, C, Stock, J, and Watson, M (1991), 'Stochastic Trends and Economic Fluctuations', *American Economic Review*, 81, pages 819-40.

King, R, Stock, J and Watson, M (1995), 'Temporal instability of the unemployment-inflation relationship', *Economic Perspectives*, Federal Reserve Bank of Chicago, pages 2-12.

Kuttner, K (1994), ‘Estimating Potential Output as a Latent Variable’, *Journal of Business and Economic Statistics*, pages 361-8.

Layard, R, Nickell, S and Jackman, R (1991), ‘Unemployment: macroeconomic performance and the labour market’, Oxford University Press, Oxford.

Lippi, M and Reichlin, L (1994), ‘Diffusion of Technical Change and the Decomposition of Output into Trend and Cycle’, *Review of Economic Studies*, 61, pages 9-30.

Manning, A (1993), ‘Wage bargaining and the Phillips curve: The Identification and Specification of Aggregate wage Equations’, *Economic Journal*, 103, pages 98-118.

Okun, A (1962), ‘Potential GNP: Its Measurement and Significance’, *Proceedings of the Business and Economics Statistics Section*, American Statistical Association, pages 98-104. (Reprinted in J. Penchman (ed.), *Economics for Policy Making*, MIT Press, 1983.)

Phillips, A W (1958), ‘The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957’, *Economica*, n.s, Vol 25, pages 283-99.

Prachowny, M (1993), ‘Okun’s Law: Theoretical Foundations and Revised Estimates’, *Review of Economics and Statistics*.

Quah, D and Vahey, S (1995), ‘Measuring Core Inflation’, *Economic Journal*, 105, pages 1,130-44.

Roberts, J (1993), ‘The Sources of Business Cycles: A Monetarist Interpretation’, *International Economic Review*, Vol 34, No 4.

Roberts, J (1995), ‘New Keynesian Economics and the Phillips Curve’, *Journal of Money, Credit and Banking*, 27, pages 974-84.

Roberts, J (1997), ‘Is Inflation Sticky?’, *Journal of Monetary Economics*, 39, pages 173-96.

Rotemberg, J (1982), ‘Sticky prices in the United States’, *Journal of Political Economy* 60 (November), pages 1,187-211.

Saleheen, J (1998), ‘Estimating a time-varying NAIRU using the Kalman Filter’, Bank of England, *mimeo*.

Saleheen, J and Westaway, P (1998), 'The Natural Rate of Unemployment: Theory and Practice', Bank of England, *mimeo*.

Sims, C (1980), 'Macroeconomics and Reality', *Econometrica*, Vol 48, No 1, pages 1-48.

Staigar, D, Stock, J and Watson, M (1996), 'How precise are estimates of the natural rate of unemployment?', *NBER Working paper*, No 5477.

St-Amant, P and van Norden, S (1997), 'Measurement of the Output Gap: A Discussion of Recent Research at the Bank of Canada', paper for BIS model-builders meeting, Basle, January.

St-Amant, P and Tessier, D (1998), 'A Discussion of the Reliability of Results Obtained with Long-Run Identifying Restrictions', *Bank of Canada Working Paper*, No 98-4.

Sterne, G and Bayoumi, T (1995), 'Temporary Cycles or Volatile Trends? Economic Fluctuations in 21 OECD Economies', *The Manchester School*, Vol LXIII, No 1, pages 23-51.

Taylor, J (1979), 'Staggered contracts in a macro model', *American Economic Review*, 69 (May), pages 108-13.

Taylor, J (1980), 'Output and price stability: an international comparison' *Journal of Economic Dynamics and Control* 2 (February), pages 109-32.

Thomas, R, Dhar, S and Pain, D (1998), 'A Small Structural Empirical Model of the UK Monetary Transmission Mechanism', Bank of England, *mimeo*, December.

Turner, D (1995), 'Speed limit and asymmetric effects from the output gap in the major seven economies', *OECD Economic Studies*, 24, pages 58-87.

Westaway, P (1997), 'What Determines the Natural Rate of Unemployment? And What Does Not?', Bank of England, *mimeo*.