

DP2000/05

Is the output gap a useful indicator of inflation?

By Iris Claus

March 2000

JEL classification: C32, E31

Discussion Paper Series

Abstract¹

One of the main indicators of inflationary pressures used by the Reserve Bank of New Zealand is the output gap. The output gap is not directly observable and estimates have to be inferred from the data. This paper evaluates whether the output gap, however measured, is a good indicator of inflationary pressures in New Zealand. The results suggest that the output gap provides a useful signal to the monetary authority. When the output gap is positive (negative) two times out of three inflation will increase (decrease) in the next quarter and three times out of five it will increase (decrease) the following year.

¹ I would like to thank Ben Hunt, John McDermott and Christie Smith and other colleagues for valuable comments. The views expressed are those of the author and do not necessarily represent those of the Reserve Bank of New Zealand.

1 Introduction

The Phillips curve has been an important tool for policy analysis over the past decades. It is conventionally specified as a relationship between unemployment and inflation though other measures of economic activity are used besides unemployment. In the Reserve Bank of New Zealand's Forecasting and Policy System (FPS), domestic inflation is determined within a Phillips curve framework using the output gap.² The output gap is defined as deviations of actual from potential output, where potential output is the level of output that is consistent with a stable rate of inflation given the productive stock of capital. A sustained positive output gap is indicative of demand pressures and a signal that inflationary pressures are increasing and that policy may need to be tightened. A level of real output below potential, ie a negative output gap, has the opposite implication. Potential output and the output gap are not directly observable and estimates have to be inferred from the data. Various methods for estimating potential output and the output gap have been developed. However, substantial uncertainty surrounds these estimates.

This paper evaluates whether the output gap, given the uncertainty surrounding its measurement, is still a useful indicator of inflationary pressures in New Zealand.

Four different measures of the gap are used. Measures of the output gap are derived from i) the Hodrick and Prescott filter, ii) the Reserve Bank of New Zealand's multivariate filter, iii) a structural vector autoregression and iv) an unobserved components model. The four measures of the gap are evaluated on how well they explain inflation pressures, as measured by the consumer price index.

Two simple versions of the gap model are estimated as a single equation and as a system of equations. First, the change in inflation is related to the *level* of the output gap. In this specification, inflation will tend to rise if the level of real output is above productive capacity through time, ie inflation will rise if the output gap is positive. Inflation will tend to fall if the gap is negative and remain stable if the gap is zero. Second, the change in inflation is related to the *change* in the output gap: inflation will tend to remain stable as long as the level of the gap is unchanged. The change in inflation, rather than the level, is included for two reasons. The null hypothesis of a unit root in the level of inflation could not be rejected at conventional levels of significance over the sample period 1971q1 to 1999q3. Moreover, first differencing removes mean shifts in inflation that are likely to have occurred over the sample period when, for example, moving from a high inflation regime to price stability. Under an inflation targeting regime what matters is not the level of inflation, but whether inflation is above or below the target. Under a credible inflation regime, the mean of inflation will coincide with the target and first differencing will approximate deviations of inflation from its target.

To foreshadow the results of the paper, the output gap appears to provide a useful signal to the monetary authority. When the output gap is positive (negative) two times

² For a discussion of the core model in the Forecasting and Policy System see Black, Cassino, Drew, Hansen, Hunt, Rose and Scott (1997).

out of three inflation will increase (decrease) in the next quarter and three times out of five it will increase (decrease) the following year.

The remainder of this paper proceeds in four further sections. Section 2 briefly summarises the historical development of the Phillips curve and reviews the current literature on different specifications. Section 3 discusses the methodologies used to estimate potential output and the output gap. In section 4 the data are described and analysed, and the empirical results are presented. Section 5 concludes with some suggestions for further research.

2 The Phillips curve

In his seminal 1958 article, A W H Phillips fitted a statistical equation between the change in nominal wages and the unemployment rate in the United Kingdom. In the original curve, the inverse of the unemployment rate was used as a proxy for excess demand for labour. Low unemployment implied high excess demand and upward pressure on wages. The Phillips curve quickly gained acceptance among policymakers and economists. One factor that contributed to the success of the Phillips curve was its ability to accommodate a wide variety of inflation theories, such as demand-pull and cost-push theories. Moreover, the Phillips curve was appealing because it provided a convincing rationale for policymakers' apparent inability, at the time, to achieve zero unemployment with price stability.

The original version of the Phillips curve has undergone several modifications since 1958. To make it more useful to policymakers, the Phillips curve was transformed from a wage-change equation to a price-change equation. This was achieved by assuming that prices are set by applying a constant mark-up to unit-labour cost. From the slope of the price-change Phillips curve, policymakers could then determine how much unemployment would be associated with any target rate of inflation. These early versions of the Phillips curve reflected the prevailing economic thinking in the 1960s, namely that the supply side of the economy was deterministic and that changes in demand were the prime factor in economic fluctuations. It was not until the early 1970s that price expectations and a stochastic supply side were introduced into Phillips curve analysis. This resulted in the expectations-augmented equation. The demand variable was re-specified more explicitly in terms of excess demand. Originally defined as an inverse function of the unemployment rate, the demand variable was redefined as the gap between the "natural" and actual rates of unemployment. The adoption of the unemployment gap in the Phillips curve finally incorporated the view that economic fluctuations are the result of both demand and supply shocks. The inclusion of price expectations recognised the importance of expectations in the inflation process and was able to account for shifts in the Phillips curve.

The natural rate of unemployment is defined as the rate that prevails when expectations are fully realised and incorporated into all wages and prices, and inflation is neither accelerating nor decelerating. This rate is also called the non-accelerating inflation rate of unemployment or NAIRU. The natural rate of unemployment is unobservable and subject to substantial uncertainty. Conventional estimates for the United States suggested a typical value of the NAIRU of 6.2 percent in 1990 with a 95

percent confidence interval of 5.1 to 7.7 percent (see Staiger, Stock and Watson 1996). Although there are no comparable studies for New Zealand, it is unlikely that the results would be more accurate and more reliable than those obtained for the United States. Nevertheless, Gordon (1996) rejects the argument that the band of statistical uncertainty surrounding the NAIRU is so broad as to render the concept useless for the conduct of policy. The NAIRU is determined by the microeconomic structure and behaviour of the economy and should thus shift only slowly. As a result, Gordon (1996) proposes evaluating alternative NAIRU estimates for any given measure of inflation based on smoothness rather than a statistical criterion.

Instead of the unemployment gap, the output gap is often included as a measure of excess demand in Phillips curve analysis.³ In the Reserve Bank of New Zealand's macroeconomic model, FPS, domestic inflation is determined within a Phillips curve framework using the output gap. The output gap is defined as the gap between actual and potential output, where potential output is the level of output that is consistent with a stable rate of inflation given the productive stock of capital. The natural rate of unemployment and potential output are conceptually similar. Potential output suffers from the same problems as the NAIRU; it is unobservable and also subject to substantial uncertainty. To date, it is unknown whether or not empirical measures of potential output are subject to more or less measurement error than measures of the natural rate of unemployment. Nevertheless, the output gap has one significant advantage over the unemployment gap in New Zealand. The advantage of the output gap framework is that developments in the goods market tend to provide more leading information than developments in the labour market in New Zealand. Preparative research in constructing FPS showed that the goods market adjusts faster than the labour market following a positive demand shock (see Razzak 1996). This suggests that the output gap rather than the unemployment gap should be used in the Phillips curve to ensure that the Reserve Bank can react as early as possible to signs of inflation pressures arising from excess demand. The prompt reaction to excess demand shocks becomes even more important if the Phillips is asymmetric (discussed below).

The Phillips curve framework has proved to be a useful tool for policy analysis and considerable evidence exists that the output gap is an important determinant of inflation for a number of countries. The most common formulation implies that inflation will be stable when the level of output equals potential. A sustained positive output gap is indicative of demand pressures and a signal to the monetary authority that inflationary pressures are increasing and that a policy tightening may be required. A level of real output below potential, ie a negative output gap, has the opposite implication. Recent empirical evidence, however, suggests that the linear and symmetric specification of the relationship between inflation and the output gap may be incorrect for some countries.

One form of asymmetry implies that excess demand conditions are more inflationary than excess supply conditions are deflationary. Ball and Mankiw (1994) provide some microeconomic basis for asymmetric inflation effects. They describe a menu-cost

³ The link between the unemployment gap and the output gap is through Okun's (1962) law, which relates the change in unemployment to the change in output.

model of pricing behaviour where positive shocks to firms' prices generate faster adjustment than negative shocks, since in the latter case the adjustment of relative prices can be, at least partially, achieved by the presence of trend inflation.⁴ The important policy implication of this asymmetry is that it may be very costly if the economy overheats because this will require a severe tightening in monetary conditions to re-establish inflation control.

Razzak (1995) tests the proposition that the Phillips curve is linear and symmetric for New Zealand over the 1983q3 to 1994q4 period and rejects both hypotheses. Clark, Laxton and Rose (1995) present evidence that there is a significant asymmetry in the output-inflation process for the United States using quarterly data from 1964q1 to 1990q4. Laxton, Rose and Tetlow (1993) also find evidence of asymmetry for Canada when estimating a reduced-form inflation equation on annual data over the period 1973 to 1991. Moreover, the inflationary effects of positive output gaps are found to be four times larger than the deflationary effects of a negative gap for the United States, Japan and Canada using annual data starting from the early 1960s (see Turner 1995).

An alternative version of the Phillips curve focuses on the change in the output gap rather than the level of the output gap – the speed limit effect. A speed limit effect occurs if a rise in inflation can be attributed to a reduction in the output gap despite output not rising above the level of potential output. Price acceleration during the upper phase of the cycle may be more pronounced than the deceleration observed when the output gap is negative. This asymmetry is based on the Keynesian idea of an inflected supply curve that is almost vertical beyond the level of potential output.

Empirical evidence of speed limit effects is mixed. Lown and Rich (1997) find a statistically significant rate of change effect for the output gap in the estimation of a price-inflation Phillips curve for the United States over the period 1965q1 to 1996q3. Estimating a price-inflation Phillips curve for Canada over the period 1961q2 to 1996q3, Gerlach and Smets (1997) provide evidence of a substantial acceleration effect. Coe and McDermott (1997) estimate a price-inflation Phillips curve for a group of thirteen developing and newly industrialised countries and industrial Asian economies using annual data.⁵ The authors find that for New Zealand, the change in inflation is more closely related to the change in the output gap than to the level of the output gap. For Japan and the Philippines, both the level of the output gap and the change in the output gap appear to be important.

In contrast, using quarterly data from 1981q1 to 1996q3 for Japan, Watanabe (1997) finds that the level of the output gap affects inflation significantly, but identifies no speed limit effects – that is, the change in the output gap does not affect inflation.

⁴ Note that the extent of the asymmetry is related to the level of general inflation, that is, at low levels of inflation the asymmetry will decline and at zero inflation the asymmetry would disappear.

⁵ The data typically start in the early 1960s and end in 1994.

Baude and Cette (1997) are also unable to find speed limit effects for France over a truncated sample period that includes 1973q1 to 1982q1 and 1987q2 to 1994q4.⁶

3 Output gaps

Potential output and the output gap are not directly observable and estimates have to be inferred from the data. Various methods for estimating potential output and the output gap have been developed. Broadly speaking, these techniques can be divided into two classes: structural and non-structural (statistical) methods.

An example of a structural method is the production function approach. This approach is probably the most desirable on theoretical grounds, but there are two considerable disadvantages: i) it is not clear what the appropriate production function for the macroeconomy is, and ii) the data for the inputs (typically capital, labour, a measure of productivity, and sometimes intermediate inputs) are often of poor quality, are infrequently measured, or may even be non-existent.

Statistical methods for inferring the level of potential output and the output gap do not condition their estimates on a particular economic model. Instead, they use statistical criteria to decompose the trend and cyclical components of output. The typical approach is to apply some sort of “mechanical” filter to output data. Traditional methods of estimating potential output include drawing a straight line through the peaks or fitting a linear or polynomial trend through actual output. Statistical techniques that have been developed more recently include the Hodrick and Prescott (1997) filter. The advantage of statistical approaches is that they are simple. The disadvantage is that decomposing output based on purely statistical criteria can be arbitrary from an economic perspective. For instance, fitting a linear deterministic time trend through output assumes that the supply side of the economy is non-stochastic and that changes in demand are the prime factor in economic fluctuations. The assumption of a deterministic supply-side is difficult to reconcile with the idea that the permanent component in output is, at least in part, driven by technology shocks. Although the Hoderick and Prescott (HP) filter incorporates a stochastic supply side, some priors are imposed on the relative incidence of temporary and permanent shocks to output.

To overcome the limitations of purely statistical approaches, methods have been developed that add conditioning information derived from economic theory to statistical techniques. This semi-structural approach is followed in the estimation of the output gap in FPS. Conway and Hunt (1997) augment the stochastic-trend estimation of the HP filter with information from a Phillips curve relationship and an Okun’s law relationship, as proposed by Laxton and Tetlow (1992), and also introduce a survey measure of capacity utilisation.

⁶ The first period ends in 1982q1 before the French wage and price freeze was put into effect. The beginning of the second period, 1987q2, marks the last devaluation of the French franc over the estimation period.

Both the HP and MV filters impose prior beliefs about how potential output evolves. An alternative method that is somewhat less restrictive than the use of HP and MV filters is the structural vector autoregression (SVAR) methodology with long-run restrictions proposed by Blanchard and Quah (1989). With this approach, the assumption that movements in measured output are the result of cyclical shocks arising from demand-side developments, and permanent shocks arising from supply-side developments provides a set of long-run restrictions to identify potential output and the output gap. Claus (2000) obtains a measure of potential output for New Zealand by means of a three-variable SVAR model that includes real output, full-time employment and a survey measure of capacity utilisation.

Another alternative, also less restrictive than filters, is the unobserved components approach, which, like the SVAR, estimates potential output as part of a system. Scott (2000) uses a multivariate unobserved components (UC) model that provides an estimate of cyclical activity for New Zealand. Estimated via the Kalman filter and maximum likelihood, the model identifies a common, trend-reverting component to real output, unemployment and capacity utilisation. The output gap is measured as actual output less the underlying trend.

4 Testing the gap model

Two simple versions of the gap model are estimated as a single equation and as a system of equations. The change in inflation is related to the *level* of the output gap. In this specification, inflation will tend to rise if the level of real output is above productive capacity through time, ie inflation will rise if the output gap is positive. Inflation will tend to fall if the gap is negative and remain stable if the gap is zero. In the alternative version, the change in inflation is related to the *change* in the output gap: inflation will tend to remain stable as long as the level of the gap is unchanged. The relationship between four different measures of the output gap and inflation is evaluated over the period 1971q1 to 1999q3.⁷ Measures of the output gap are obtained using the techniques discussed in the previous section: i) the HP filter, ii) the Reserve Bank's MV filter, iii) the structural vector autoregression methodology, and iv) the unobserved components approach. All measures of the gap are obtained from real seasonally adjusted production gross domestic product (GDP). Inflation is measured by the consumer price index excluding credit services adjusted for the introduction of and increase in the good and services tax (GST).

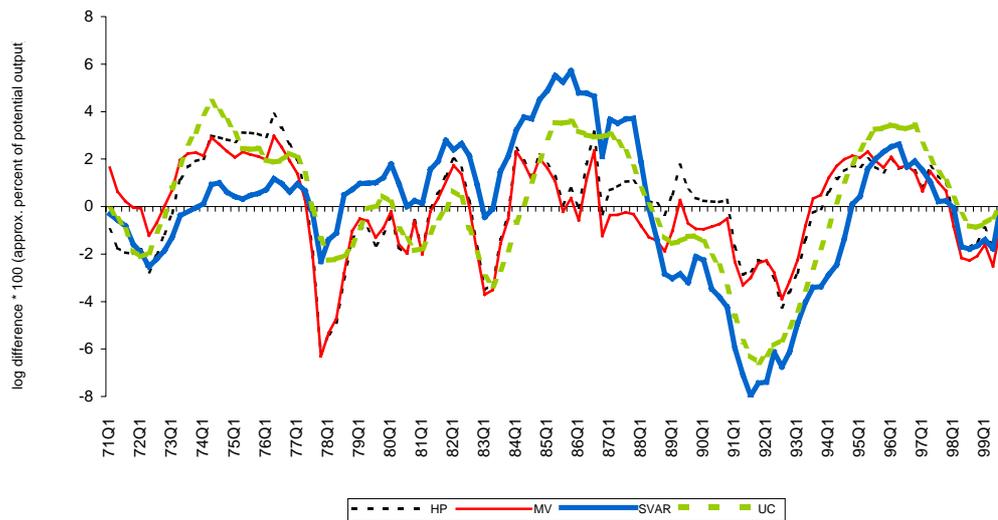
Figure 1 plots the four measures of the output gap from the HP and MV filters and the SVAR and UC models.⁸ The estimates of the output gap show that the New Zealand economy was in excess demand during the 1970s. This was largely because of stimulatory fiscal and monetary policy although large declines in the measures of the output gap were registered during the second oil price shock in 1977. Monetary and fiscal policies were stimulatory during the 1970s for several reasons. As a consequence of the oil price shocks in 1975 and 1977, the world economy went into

⁷ The choice of the sample period is dictated by the availability of the SVAR gap.

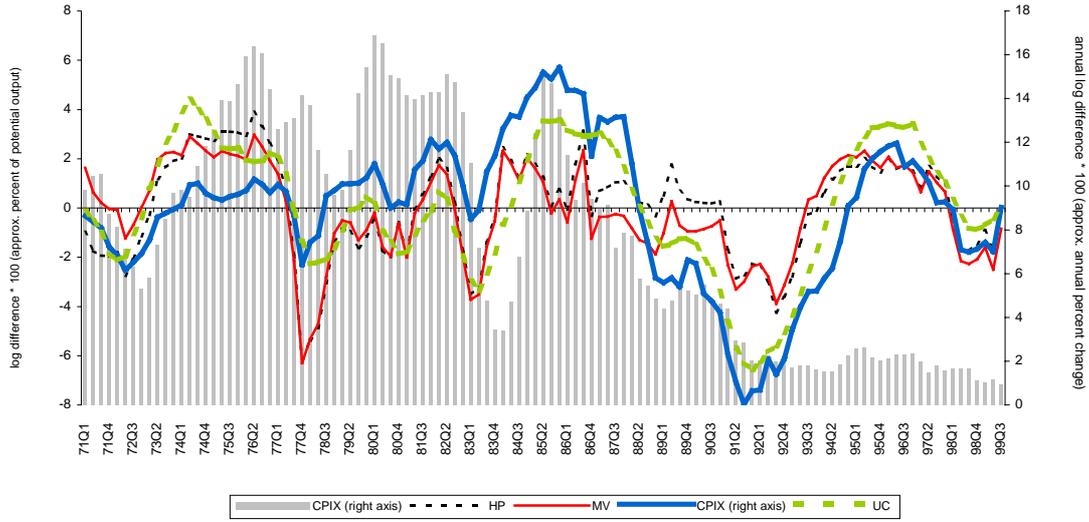
⁸ The smoothness parameters, λ , in the HP and MV filters is set to 1600.

recession lowering the demand for New Zealand's exports. Moreover, access to a key market was increasingly restricted upon Britain's entry to the European Economic Community (EEC) in 1972 and many commodity producers in other countries (Europe in particular) increased their effective rates of protection for agricultural products. In response the government of the day ran stimulatory fiscal and monetary policy in an attempt to protect living standards. Lawrence and Diewert (1999) also note generally poor productivity growth during the 1970s.

Figure 1: Measures of the output gap



The SVAR gap suggests excess demand until 1988, while the UC gap indicates excess demand from about 1984 until 1988. Productive capacity was insufficient to meet demand, which in turn led to increasing inflationary pressures. Indeed, until the mid-1980s, consumer price inflation was at double digit rates except for the price and wage freezes in 1983-84 (see figure 2). Dissatisfaction with the systematic under-performance of the New Zealand economy compared to other countries during the 1970s and 1980s prompted the government of the time to embark on a period of comprehensive economic reform. Entrenching macroeconomic stability through fiscal restraint and monetary policy focused on reducing inflation was one of the reforms' key themes. During this period of reform, the measures of the gaps are quite volatile (in particular the HP and MV gaps), but tend to co-move again over the 1990s. Following the reforms over the latter part of the 1980s and the early 1990s, the four gaps are negative and indicate large excess supply as the rigours of structural and fiscal adjustment, along with disinflation and a collapse in the New Zealand sharemarket, culminated in a very poor rate of economic growth. The SVAR and UC gaps suggest a more prolonged and deeper recession in the early 1990s than the HP and MV filters. This level difference in the estimated output gaps is in part because of a "starting point" problem inherent in the SVAR and UC models and due to a definitional difference of what constitutes a demand shock and what constitutes a supply shock. This issue is discussed in detail in Conway and Frame (2000).

Figure 2: Measures of the output gap and inflation

The time series properties of the data are examined to determine the order of integration of the variables included. A series is said to be integrated of order d , denoted $I(d)$, if the series becomes stationary or $I(0)$ after being differenced d times. The augmented Dickey and Fuller (Said and Dickey 1984) test and the Phillips and Perron (1988) Z_α test are performed. Both the augmented Dickey-Fuller (ADF) and the Phillips-Perron Z_α statistics allow one to test the null hypothesis that a series is $I(1)$ against the alternative that it is $I(0)$. The results from both tests can be found in appendix 1. The results suggest that the null hypothesis of a unit root in the level of the output gap can be rejected for all four measures and that inflation is a difference-stationary process at conventional levels of significance.⁹

To assess the relationship between the estimated output gaps and inflation, a simple version of the gap model is tested following Coe and McDermott (1997).

$$\text{Model 1: } \Delta\pi_t = \alpha_1 + \sum_{k=0}^p \beta_{1k} GAP_{t-k} + \varepsilon_{1t} \quad (1)$$

where π_t is the annual logarithmic difference of the consumer price index, GAP_t is the logarithmic difference between actual and potential output, ε_{1t} is a stochastic disturbance term, and Δ is the first difference operator. A constant is included to avoid imposing the constraint that the non-inflationary level of the output gap is exactly zero.¹⁰ The model is derived from an inflation-expectations-augmented Phillips curve with adaptive expectations

⁹ The results of the ADF test also provide evidence that inflation is trend stationary at the ten percent level of significance; that is, the series becomes stationary after subtracting a deterministic time trend.

¹⁰ Clark, Laxton and Rose (1995) show that asymmetry in the inflation-output process implies that the measure of excess demand, which is appropriate in estimating a Phillips curve, cannot have a zero mean.

$$\pi_t = \alpha_1 + \sum_{k=0}^p \beta_{1k} GAP_{t-k} + \pi_t^e + \varepsilon_{1t} \quad (2)$$

with $\pi_t^e = \pi_{t-1}^e$, where π_t^e is inflation expectations at time t . First differencing removes mean shifts in inflation that are likely to have occurred over the sample period when, for example, moving from a high inflation regime to price stability. Under an inflation targeting regime what matters is not the level of inflation, but whether inflation is above or below the target. For example, a forecast inflation rate of 1.5 percent would have had different policy implications in 1996 than in 1999. Under a credible inflation regime, the mean of inflation will coincide with the target and first differencing will approximate deviations of inflation from its target. Under a credible inflation regime, inflation expectations will also coincide with the target.

Model 2 relates the change in inflation to the change in the gap and is a special case of model 1. It constraints the coefficients of the level of the gap to alternate in sign and the pairs β_{10} and β_{11} , β_{12} and β_{13} , and so on, to sum to zero. Model 1 also allows for the level of the gap as well as the change in the gap to affect inflation. In this case some of the β_{1i} s would be negative but their sum would be positive.

$$\text{Model 2:} \quad \Delta\pi_t = \alpha_2 + \sum_{k=0}^p \beta_{2k} \Delta GAP_{t-k} + \varepsilon_{2t} \quad (3)$$

Both models (including the gap in levels and in first differences) are estimated with four lags. The lag-length was determined using an F-test that tests the included lagged terms for significance at the ten percent level. To evaluate the different models, F-tests that the β_i s, for $i = 1, 2$, are jointly significantly different from zero, ie the output gap is a significant determinant of inflation, are performed. Moreover, diagnostic tests on the regression residuals are used. A test for normality is based on a Lagrange Multiplier (LM) test for joint skewness and kurtosis proposed by Jarque and Bera (1980). The null hypothesis of a linear model with homoscedastic errors is tested using the modified Breusch and Pagan (1979) test proposed by Koenker (1981). The results can be found in appendix 2.

F-tests indicate that the level of all four output gaps (model 1) is a significant determinant of the change in inflation and explains between sixteen and thirty percent of the total variance in inflationary pressures. The SVAR gap has a significant impact on the change in inflation, but the sum of the estimated coefficients of the gap variables is negative (although virtually zero). The signs of the estimated coefficients on the HP and UC gaps alternate and their sum is close to zero suggesting that the change in the HP and the UC gaps may be a better indicator of inflationary pressures. Including the change in the gap rather than the level (model 2) results in an increase of the R^2 for the HP gap from 0.26 to 0.30, but remains unchanged for the UC gap model. Estimation results with the constraint implied by model 2 suggest that the change specification may indeed work slightly better than model 1. This finding is consistent with Coe and McDermott (1997). The sums of the estimated coefficients on the output gap variables are all positive. Moreover, the R^2 s are higher in all cases except for the UC gap.

Residual diagnostic tests suggest that the gap model only including inflation and the output gap is an over-simplification and detect some mis-specification. Lagrange Multiplier tests indicate that the null hypothesis of no serial correlation is rejected at conventional levels of significance for all eight equations.¹¹ Also, the Jarque and Bera (1980) tests indicate that the residuals are not normally distributed and the null hypotheses that the models are linear with homoscedastic errors are rejected for some of the equations.

Figure 3: Actual and predicted changes in inflation
(percent)

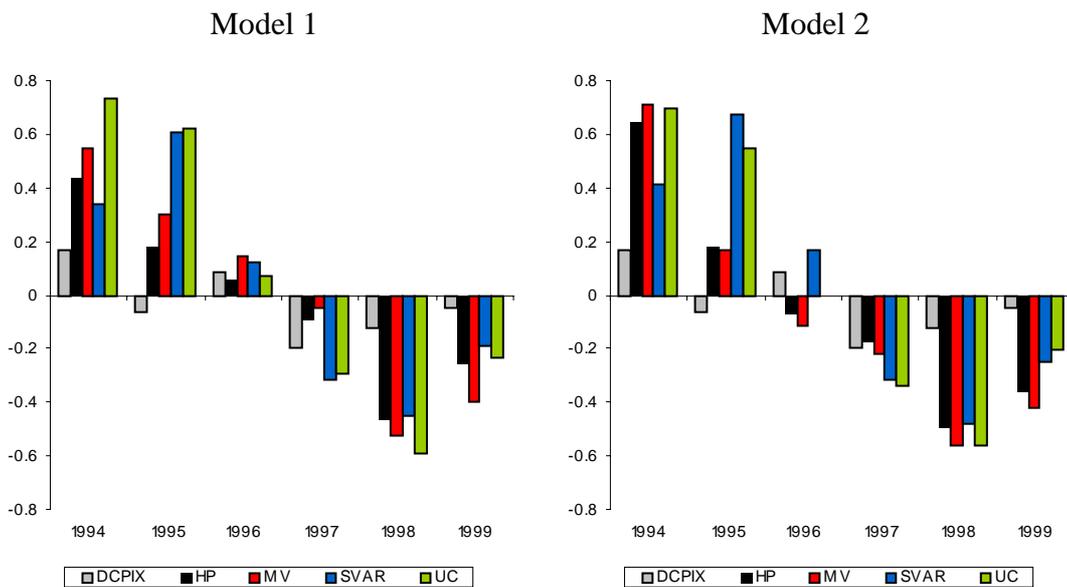


Figure 3 plots actual and predicted changes in inflation for the six years 1994 to 1999. Although the estimation is performed using quarterly data, annual averages are plotted for ease of illustration.¹² The left-hand graph shows the results from model 1 including the gap in levels and the right-hand graph from model 2 including the change in the gap. The figures show the change in actual inflation (first bar) and the in-sample predictions from the HP, MV, SVAR and UC gaps (the other four bars). The plots confirm the finding that the output gap overall is a good indicator of inflationary pressures as actual and predicted inflation tend to move in the same direction. In 1995 both models predicted an increase in inflationary pressures, while actual inflation dropped. Model 2 also failed to identify the increasing inflation pressure in 1996, except when using the SVAR estimates of the output gap.

Table 1 summarises the number of times (in percent) that the output gap correctly predicts, in-sample, the direction of inflationary pressures. The level of the output gap appears to be a slightly better indicator than the change in the gap of whether inflation

¹¹ Test statistics are not reported, but are available upon request.

¹² Because output gaps tend to be quite persistent and therefore conceptually are a good indicator of general inflation pressures, the appropriate frequency may indeed be annual. However, quarterly data were chosen here to be able to incorporate new information as it becomes available.

is accelerating or decelerating except when using the HP gap. The HP filter performs better in model 2 than it does in model 1.

Table 1: In-sample predictions of inflation pressures
(number of times correctly predicted in percent)

	Model 1	Model 2
HP	71.2	71.8
MV	71.2	67.3
SVAR	69.4	68.2
UC	72.1	70.0

The fact that the output gap is useful for explaining movement in inflation within sample does not necessarily indicate that it will be useful for forecasting in real time. To evaluate the *out-of-sample* forecast performance a simple vector autoregression (VAR) model is estimated including the change in inflation and the output gap. A vector autoregression model is a system of equations, in which each variable is regressed on lagged values of itself and the other variables in the system. Hence, the change in inflation depends on its past realisations and past realisations of the output gap. The output gap also depends on its past realisations and past realisations of inflation. The reduced form vector autoregressive representation is given by

$$X_{it} = \Phi(L)X_{it-1} + \varepsilon_{it} \quad (4)$$

for $i = 1, 2$, where X_{it} is a covariance stationary vector process, ie $X_{1t} = [\Delta\pi_t, GAP_t]'$ and $X_{2t} = [\Delta\pi_t, \Delta GAP_t]'$, ε_{it} is a vector of random disturbances and $\Phi(L)$ is a polynomial lag operator. As both equations in the system share the same matrix of regressors, estimation of the reduced form VAR amounts to applying ordinary least squares (OLS) separately to each equation in (4) given the optimal number of lags to eliminate serial correlation from the residuals. Testing down the lag structure of the model using likelihood ratio tests at the ten percent level chose a lag length of four.

The information content of the output gap is first assessed by examining the in-sample results of the VAR model.¹³ The results from model 1 and model 2 are qualitatively similar although the R^2 's are higher for the level specification except when using the SVAR gap. The SVAR gap performs better in model 2 than it does in model 1. F-tests that the coefficients on the output gap are jointly significantly different from zero indicate that the output gap is a significant determinant of inflation. The R^2 's for the inflation equation in the VAR, for each of the four output gap measures, is higher than for the single equation case, suggesting that past changes in inflation are also an important determinant of current inflation. However, even with the inclusion of lags of the change in inflation diagnostic tests still detect some mis-specification. This implies that past inflation and the output gap alone cannot explain inflationary pressures in the economy.

¹³ The results are not reported, but are available upon request.

The estimated VAR equations can be inverted so that each variable appears as a function of lags in the disturbances or innovations to all variables in the system. The variance decomposition of the change in inflation is the percentage of the within-sample forecast error that is accounted for by the output gap at various forecast horizons.¹⁴ If the forecast horizon is large and the time series vector is stationary, then the forecast error variance approximates the actual variance of the series. If a series is truly exogenous then all of its forecast variance will be explained by its own innovations. If a large fraction of the forecast variance of the change in inflation is explained by the innovations in the output gap then the output gap is important in explaining inflationary pressures. The variance decomposition of the change in inflation can therefore add important information about the relative influence of the output gap.

Variance decomposition results are shown in table 2 as a percent of the variance in the forecast error of the change in inflation that is attributable to an innovation in each of the four output gap measures. The decompositions are shown for a forecast horizon of 36 quarters and indicate that roughly 20-30 percent of the variance in the change of inflation can be attributed to the output gap. The level of the HP and MV gaps explain a larger fraction of the forecast variance of the change in inflation than their change. When using the SVAR and UC gaps the change in the gap performs better than the level of the gap in terms of contribution to forecast error variance in inflationary pressures.

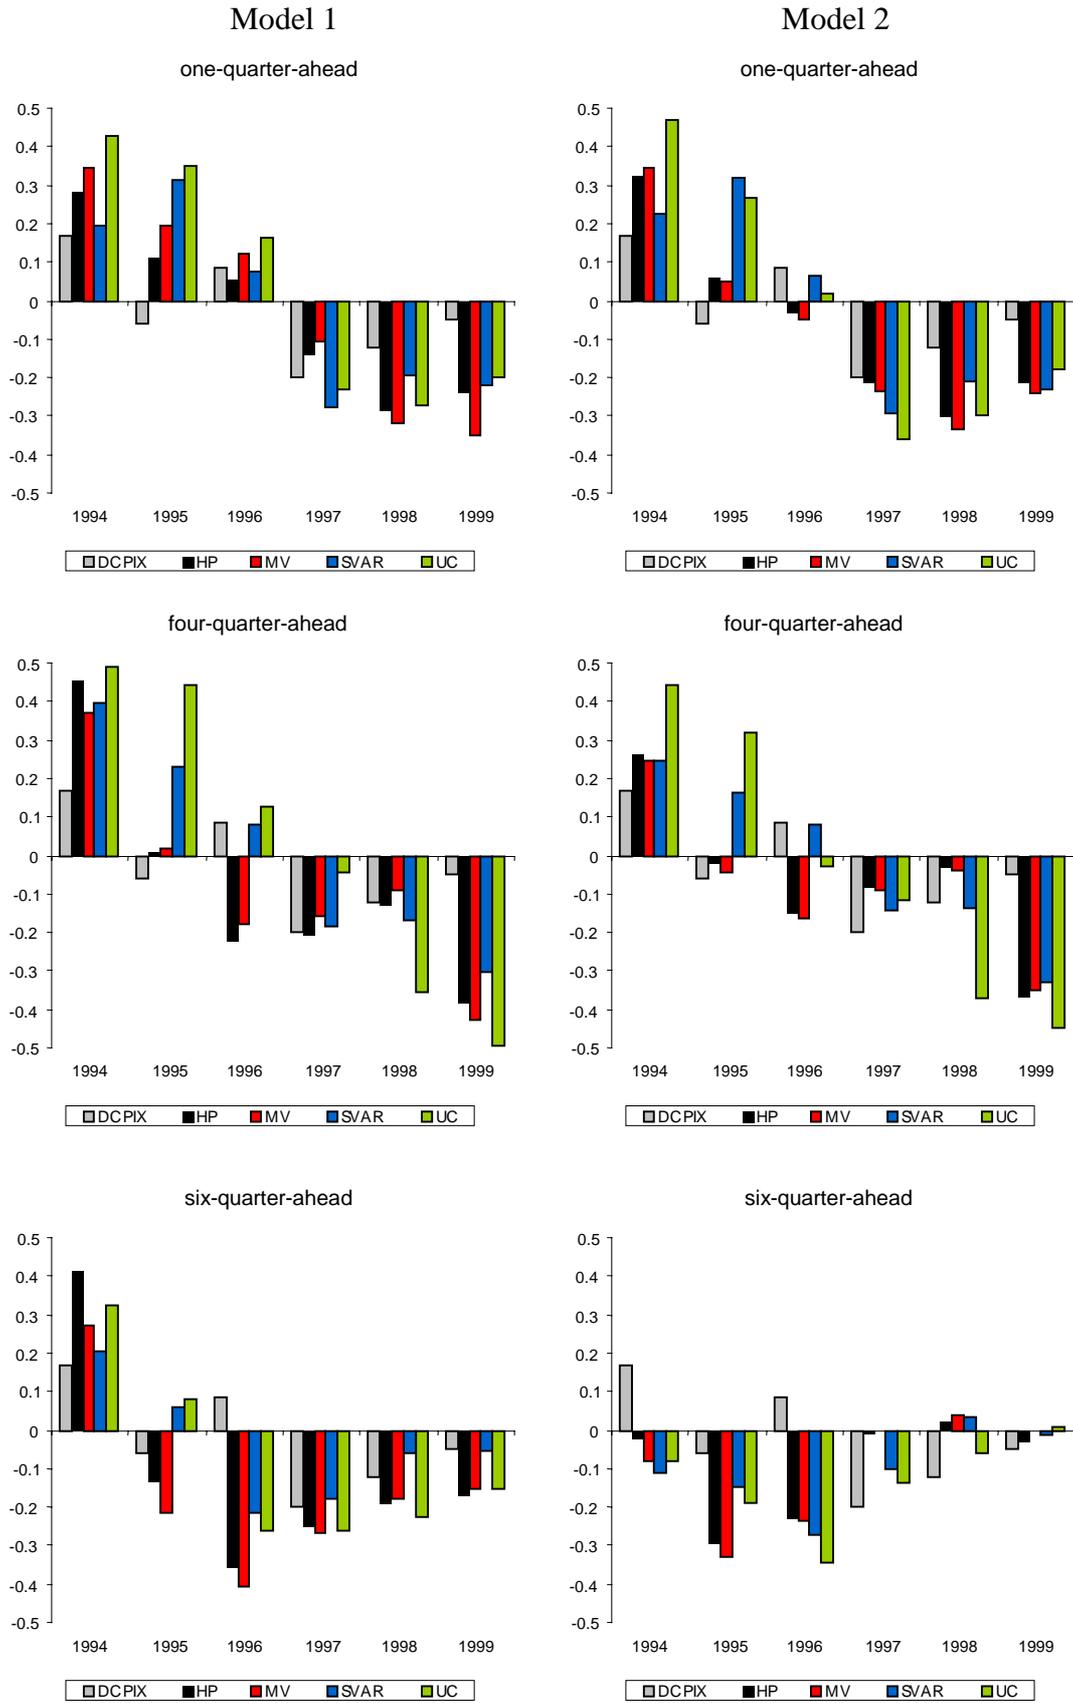
Table 2: Variance decomposition of the change in inflation
(percent of forecast variance explained by innovation in the output gap)

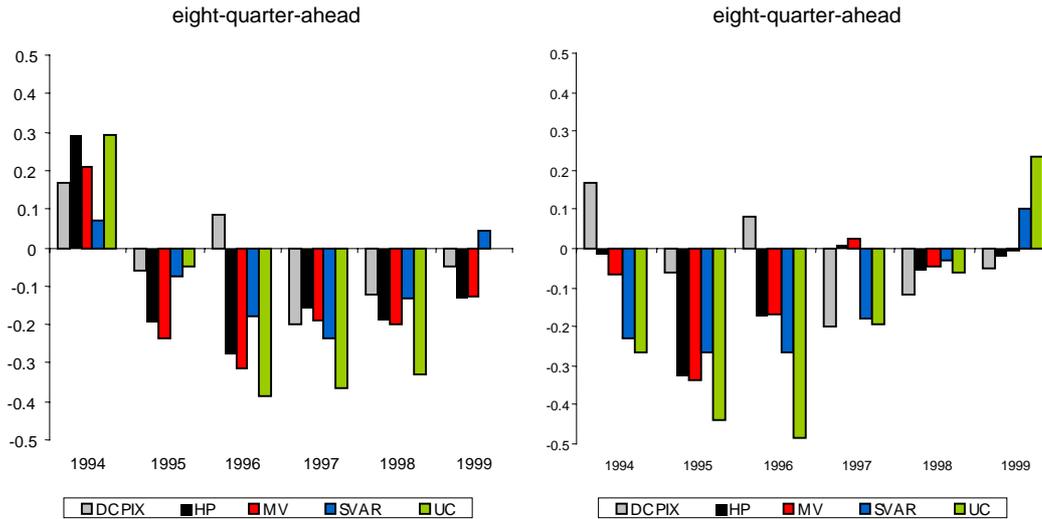
	Model 1	Model 2
thirty-six-quarter-ahead		
HP	33.0	32.1
MV	33.5	31.1
SVAR	18.3	18.7
UC	24.3	25.4

To evaluate the usefulness of the output gap for forecasting in real time, out-of-sample forecasts are obtained from a rolling sample, time varying coefficient approach over a period of 60 quarters or 15 years. With this technique, the four-quarter-ahead forecast model, for instance, is initially estimated for the period 1971q1 to 1984q2. The estimated coefficient are then used to forecast the change in inflation and the output gap one quarter ahead, which are used to produce forecasts for the following quarter, and so on up to four quarters. The model is then rolled forward one quarter and re-estimated. The change in inflation and the output gap are forecast four quarters ahead following the same approach. The process is repeated until the last observation. One-, six- and eight-quarter-ahead forecasts are obtained similarly.

¹⁴ The calculation uses an orthogonalised version of the errors.

Figure 4: Actual and out-of-sample forecast changes in inflation (percent)





Out-of-sample forecast performance is illustrated in figure 4. As before, the estimation uses quarterly data, but annual averages are plotted for ease of illustration. The left-hand column shows the results from model 1 including the gap in levels and the right-hand column from model 2 including the change in the gap. The first bar plots the actual change in inflation and the other four bars are the forecasts from the HP, MV, SVAR and UC gaps respectively. The results from the one-quarter-ahead out-of-sample forecast are similar to the predicted in-sample changes of inflation from the single equation (plotted in figure 3). However, the magnitude of the error is smaller with the VAR model suggesting that past changes in inflation are also an important determinant of current inflation. Overall the results confirm that the output gap is a good indicator of inflationary pressures as actual and forecast inflation tend to move in the same direction. Interestingly, none of the measures of the output gap was able to correctly forecast the direction of inflationary pressures in 1995 and 1996. Moreover, the forecasts from the four gap models in the level specification diverged in 1995 and 1996, although they generally moved together over the other years. This suggests that even if the Reserve Bank had used the output gap as the main indicator of inflation pressures during the mid-1990s, inflation may have still exceeded the target band in 1996. Forecasts using the change in the output gap appear to diverge more than those from the level of the output gap.

Table 3 shows the number of times (in percent) that the output gap correctly predicts the direction of inflationary pressures out-of-sample. This test is quite stringent because it evaluates quarterly outcomes, which can be very volatile.¹⁵ The test is also quite “naïve” as it does not distinguish whether the forecast was only slightly wrong or completely missed. The level of the output gap appears to be a better indicator of whether inflation is accelerating or decelerating. The results for the one-quarter-ahead, out-of-sample forecasts are slightly worse than for the in-sample prediction (reported in table 1). This is possibly the result of small sample bias: the out-of-sample percentages are calculated with fifteen years of observations compared to close to thirty years for the in-sample predictions. The out-of-sample forecast performance

¹⁵ The numbers reported in table 1 and table 3 are about five to ten percent higher when using annual averages.

deteriorates with the forecast horizon. Overall, the results suggest that the output gap provides a useful signal to the monetary authority. When the output gap is positive (negative) two times out of three inflation will increase (decrease) in the next quarter and three times out of five inflation will increase (decrease) the following year.

Table 3: Out-of-sample forecasts of inflationary pressures
(number of times correctly predicted in percent)

	Model 1	Model 2
one-quarter-ahead		
HP	68.9	65.6
MV	67.2	62.3
SVAR	63.9	63.9
UC	63.9	62.3
four-quarter-ahead		
HP	56.9	56.9
MV	60.3	56.9
SVAR	60.3	58.6
UC	63.8	58.6
six-quarter-ahead		
HP	57.1	46.4
MV	51.8	50.0
SVAR	51.8	50.0
UC	57.1	58.9
eight-quarter-ahead		
HP	53.7	50.0
MV	51.9	53.7
SVAR	51.9	51.9
UC	50.0	55.6

5 Concluding remarks

In the Reserve Bank of New Zealand's macroeconomic model, the output gap is a key indicator of domestic inflation. The output gap is unobservable and estimates are subject to substantial uncertainty. This paper asked the question whether the output gap still is a useful indicator of inflationary pressures in New Zealand. The relationship between inflation and four different measures of the output gap was tested over the period 1971q1 to 1999q3. Two simple versions of the gap model were estimated within a single equation framework and as a system of equations. The first version included the change in inflation and the level of the output gap. In an alternative version, the change in inflation was related to the change in the gap. The results showed that the output gap provides a useful signal to the monetary authority. A positive output gap is indicative of demand pressures and a signal that inflationary pressures are increasing and that policy may need to tighten. A negative output gap has the opposite implication.

The research strategy in this paper was to keep the estimation simple. Inflation, of course, does not only depend on the unemployment or output gap, but also on other factors. Changes in the terms of trade or import prices, commodity price shocks, exchange rate changes, indirect taxes, and labour market policies and institutions that affect wage formation can all directly affect inflation and inflation pressures. The inclusion of any of these variables should not alter the results that the gap is a useful indicator of inflationary pressures, but is likely to improve the forecast performance. This area presents an avenue for further research.

References

- Ball, L and G Mankiw (1994), "Asymmetric price adjustment and economic fluctuations," *The Economic Journal* 104, 247-261.
- Baude, J and G Clette (1997), "An empirical assessment of the link between the output gap and inflation in the French economy," *BIS Conference Papers* 4, 43-65.
- Black, R, V Cassino, A Drew, E Hansen, B Hunt, D Rose and A Scott (1997), "The Forecasting and Policy System: the core model," *Reserve Bank of New Zealand Research Paper* No. 43.
- Blanchard, O and D Quah (1989), "The dynamic effects of aggregate supply and demand disturbances," *American Economic Review* 79, 655-673.
- Breusch T S and A R Pagan (1979), "A simple test for heteroscedasticity and random coefficients variation," *Econometrica* 47, 1287-1294.
- Clark, P, D Laxton and D Rose (1995), "Asymmetry in the US output-inflation nexus: issues and evidence," *International Monetary Fund Working Paper* WP/95/76.
- Claus, I (2000), "Estimating potential output: a structural VAR approach," *Reserve Bank of New Zealand Discussion Paper* DP2000/03.
- Coe, D T and J McDermott (1997), "Does the gap model work in Asia?" *International Monetary Fund Staff Papers* 44, 59-80.
- Conway, P and D Frame (2000), "A spectral analysis of New Zealand output gaps using Fourier and wavelet techniques," *Reserve Bank of New Zealand Discussion Paper* DP2000/06.
- Conway, P and B Hunt (1997), "Estimating potential output: a semi-structural approach," *Reserve Bank of New Zealand Discussion Paper* G97/9.
- Gerlach, S and F Smets (1997), "Output gaps and inflation: unobservable-components estimates for the G-7 countries," *Mimeo*. Bank for International Settlements.
- Gordon, R G (1996), "The time-varying NAIRU and its implications for economic policy," *NBER Working Paper* 5735.
- Hodrick, R J and E C Prescott (1997), "Post-war US business cycles: an empirical investigation," *Journal of Money Credit and Banking* 29, 1-16.
- Jarque, C M and A K Bera (1980), "Efficient test for normality, homoscedasticity and serial independence of regression residuals," *Economic Letters* 6, 255-259.
- Koenker R (1981), "A note on studentizing a test for heteroscedasticity," *Journal of Econometrics* 17, 107-112.

- Lawrence, D and E Diewert (1999), *Measuring New Zealand's productivity*. Report prepared for the Department of Labour, Reserve Bank of New Zealand and The Treasury.
- Laxton, D, D Rose and R Tetlow (1993), "Is the Canadian Phillips Curve non-linear?" *Bank of Canada Working Paper* 93-7.
- Laxton, D and R Tetlow (1992), "A simple multivariate filter for the measurement of potential output," *Bank of Canada Technical Report* No. 59.
- Lown, C S and R W Rich (1997), "A look at the US inflation puzzle," *BIS Conference Papers* 4, 193-219.
- Okun, A (1962), "Potential GNP: its measurement and significance," in *Proceedings of the Business and Economics Section*, American Statistical Association, 98-104.
- Phillips, A W H (1958), "The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861-1957," *Economica* 25, 283-300.
- Phillips, P C B and P Perron (1988), "Testing for a unit root in time series regressions," *Biometrika* 75, 335-346.
- Razzak, W (1995), "The inflation-output trade-off: is the Phillips curve symmetric?" *Reserve Bank of New Zealand Discussion Paper* G95/7.
- Razzak, W (1996), "Do prices lead wages in the demand cycle?" *Reserve Bank of New Zealand Research Memorandum* M96/6.
- Said, S and D A Dickey (1984), "Testing for unit roots in autoregressive-moving average models with unknown order," *Biometrika* 71, 599-608.
- Scott, A (2000), "Estimating common business cycle activity via a multivariate unobserved components model: an application to New Zealand," *Reserve Bank of New Zealand Discussion Paper* DP2000/04.
- Staiger, D, J Stock and M W Watson (1996), "How precise are estimates of the natural rate of unemployment?" *NBER Working Paper* 5477.
- Turner, D (1995), "Speed limit and asymmetric inflation effects from the output gap in the major seven economies," *OECD Economic Studies* No. 24, 57-87.
- Watanabe, T (1997), "Output gap and inflation: the case of Japan," *BIS Conference Papers* 4, 93-112.

Appendix 1: Tests for integration^a

	augmented Dickey-Fuller test (data dependent lag) ^b		Phillips-Perron Z_{α} test ^c	
	no trend	trend	no trend	trend
CPIX	-0.26 (12)	-3.37 (12) *	-5.03	-13.83
Δ CPIX	-4.30 (11) ***	-4.56 (11)***	-47.83 ***	-48.01 ***
HP	-4.22 (11) ***	-4.21 (11) ***	-21.01 ***	-20.97 *
MV	-3.58 (11) ***	-3.61 (11) **	-22.27 **	-22.65 **
SVAR	-3.08 (11) **	-3.27 (11) *	-7.81	-8.11
UC	-4.05 (9) ***	-4.08 (9) ***	-10.84	-11.04
<p>*** H_0 of a unit root is rejected at the 1 percent level. ** H_0 of a unit root is rejected at the 5 percent level. * H_0 of a unit root is rejected at the 10 percent level.</p>				
<p>a. All test regressions include a constant. b. A lag-length selection is used that tests the included lagged terms for significance at the 10 percent level. The initial number of lags is set equal to three times the seasonal frequency, ie, twelve. c. The spectral density is estimated with an AR(4) spectral estimator.</p>				

Appendix 2: Testing the gap model

	lag length ^a	gap coefficients		F-test ^b	R ²	test for normality ^c	test for a linear model with homoscedastic errors ^d
		sum	sign				
Model 1							
HP ¹	4	0.07	- - - -	7.49 ***	0.26	20.04 ***	6.63
MV ¹	4	0.12	+ + + + -	8.11 ***	0.28	16.07 ***	6.42
SVAR ¹	4	-0.00	+ + - - -	4.10 ***	0.16	4.75 *	12.31 **
UC ¹	4	0.02	- - - - -	8.82 ***	0.30	2.56	3.91
Model 2							
Δ HP ¹	4	0.96	- + + + +	9.02 ***	0.30	10.40 ***	11.62 **
Δ MV ¹	4	1.01	+ + + + +	9.19 ***	0.31	9.00 **	9.72 *
Δ SVAR ¹	4	0.87	+ + + + +	4.26 ***	0.17	3.41	4.75
Δ UC ¹	4	0.83	- + + + -	8.84 ***	0.30	3.29	10.55 *
***	H ₀ is rejected at the 1 percent level.						
**	H ₀ is rejected at the 5 percent level.						
*	H ₀ is rejected at the 10 percent level.						
a.	The lag-length was determined using an F-test that tests the included lagged terms for significance at the ten percent level with the maximum number of lags equal to the seasonal frequency, ie four.						
b.	H ₀ : The gap variables jointly have no effect on the change in inflation.						
c.	The test for normality is based on a Lagrange Multiplier (LM) test for joint skewness and kurtosis proposed by Jarque and Bera (1980).						
d.	The null hypothesis of a linear model with homoscedastic errors was tested using the modified Breusch and Pagan (1979) test proposed by Koenker (1981).						
I.	LM tests suggest that the null hypothesis of no serial correlation is rejected at conventional levels of significance. Results are not reported, but available upon request.						