

CENTRE FOR ECONOMIC RESEARCH

WORKING PAPER SERIES

2000

The Behaviour of Inflation and Unemployment in the United States

Vincent Hogan, University College Dublin

Working Paper
WP00/16

June 2000

**DEPARTMENT OF ECONOMICS
UNIVERSITY COLLEGE DUBLIN
BELFIELD DUBLIN 4**

The Behaviour of Inflation and Unemployment in the United States

By

Vincent Hogan¹

June 2000

Abstract

Low rates of inflation have been recorded in the United States in recent years despite a decline in the unemployment rate. This phenomenon could be the result of a series of transitory shocks or of a permanent change in the structure of the economy leading to a lower NAIRU. The paper suggests that while the NAIRU may have fallen slightly, it has not fallen by an amount sufficient to explain the recent behaviour of inflation. The main explanation for recent inflation performance appears to be that there have been favorable price shocks; in particular, the cost of imports has fallen sharply as the dollar has appreciated.

JEL Classification Numbers: E31, E37

Keywords: Inflation, Philips Curve

¹ This research was started while the author was an economist with the International Monetary Fund. The views expressed are those of the author alone and do not necessarily represent those of the IMF.

Author's E-Mail Address: Vincent.hogan@ucd.ie

Summary

Low rates of inflation have been recorded in the United States in recent years despite a decline in the unemployment rate to levels that previously would have been associated with rising inflation. Indeed, over the period since 1990, there has been a positive correlation between the inflation rate and the unemployment rate. This indicates that more than the traditional Philips curve relationship is at work. This phenomenon could be the result of a series of fortuitous, transitory shocks or of a permanent change in the structure of the economy leading to a lower natural rate of unemployment (NAIRU).

The results presented here suggest that the NAIRU may have fallen slightly, but that this, on its own, is not enough to explain the recent history of inflation. The main explanation for recent inflation in the U.S. appears to be that there have been unexpected declines in some cost variables. In particular the cost of imports has fallen in recent years due to the appreciation of the dollar, a decline in commodity prices and the Asian crisis.

This conclusion has important implications for policy makers. It implies that the structure of the economy has not altered fundamentally. Thus, if and when the supply shocks disappear, unemployment will have to increase to the level of the NAIRU.

A. Introduction

Historically, U.S. inflation has followed a fairly predictable course in relation to the business cycle. Typically, inflation rises during an economic expansion. It reaches its maximum slightly after the end of the boom and declines thereafter during the recession until the economy has been in recovery for a year or two. It appears, however, that recent history has departed from this pattern. Over the period since the end of the last recession there has been a prolonged expansion, unemployment has fallen, but inflation has not increased.

The Philips curve, the traditional model used by economists to study the link between inflation and unemployment, worked well in the past. But it has difficulty accommodating the recent behaviour of unemployment and inflation. The basic Philips curve posits a negative relationship between inflation and unemployment in excess of the natural rate. But, over the period since 1990, there has been a positive correlation between the inflation rate and the unemployment rate. This indicates that more than the basic Philips curve relationship is at work.

This phenomenon has two important implications. Firstly, the traditional Philips curve has systematically over-predicted inflation since 1994. This makes the task of the Federal Reserve more difficult. Secondly, the apparent breakdown of the traditional relationship between inflation and unemployment may be a symptom of more fundamental changes in the economy, which themselves could have profound implications for policy.

The purpose of this paper is to attempt to explain this recent history. In particular, we wish to determine what modifications need to be made to the traditional Philips curve model in order that it can accommodate recent history. Two possible alternative explanations are examined in detail. The first, and perhaps most obvious, potential explanation is that the NAIRU has fallen sufficiently to enable unemployment to remain at an historical low even as inflation is falling. The alternative explanation is that the NAIRU has not really changed, but the US has been fortunate in recent years to have been hit by some positive supply shocks that have overwhelmed inflationary pressures caused by unemployment being below its natural level.

These alternative explanations have dramatically different policy implications. If the NAIRU has truly fallen, it is very good news for the U.S. economy. A fall in the NAIRU implies that there has been change in the structure of the economy that would tend to keep unemployment low in the long run without the danger of rising inflation. If, however, it is true that recent inflation performance is the result of temporary shocks, then there has been no permanent change in the structure of the economy. Inflationary pressures will re-assert themselves once the temporary supply shocks have ceased. At some point in the near future, policy makers will have to allow unemployment to rise if inflation is not to accelerate.

We investigate the question as follows: In section one we show that the traditional Philips curve has systematically over-predicted inflation since 1994. In section two we explore the

possibility that the natural rate of unemployment may have changed. We present estimates that suggest that the natural rate may have fallen to below 5.5% in recent years. But we also show that this change, on its own, is not sufficient to explain recent inflation. In section three we examine the possibility that low inflation is the result of a series of fortuitous shocks. We find that recent inflation performance of the U.S. can be best explained by changes in the dollar price of imported goods. Forecasts of inflation based on estimates of a Philips Curve that includes import prices as a regressor, do not over predict inflation. Interestingly, we also find that changes in the cost of labour do not contribute much to the explanation of recent inflation performance. It appears that rising productivity and/or falling (non-wage) labour costs have not been critical factors in explaining recent inflation.

B. The Problem

We begin the analysis by estimating a standard Philips curve model. This model worked well in the past, but we will show that it has over-estimated inflation in recent years. The basic Philips curve is given in equation (1), where U is the unemployment rate, U^N is the NAIRU,

$$p_t = \mathbf{a}(L)p_{t-1} + \mathbf{b}(L)(U_t - U_t^N) + e_t$$

δ is the rate of inflation and $\hat{a}(L)$ and $\hat{b}(L)$ are polynomials in the lag operator.

Equation (1) posits a negative relationship between unemployment and inflation in the short run. More precisely, it predicts that inflation will accelerate if the economic activity is above potential i.e. if the unemployment rate is below the NAIRU. The NAIRU is a similar concept to Friedman's natural rate of unemployment. Both concepts are intended to capture the notion that there is a rate of unemployment that is determined by the fundamental micro economic structure of the country. The unemployment gap (defined as the difference between actual unemployment and the NAIRU) captures the effect of excess demand on inflation. Thus if the economy is operating above potential, labour (and other resources) will be in short supply. This scarcity will cause prices to be bid up leading to inflation. The inclusion of lagged variables captures the inertia in this process.

Economic theory imposes two restrictions on the model. Firstly, there should be no constant term in the regression and secondly, the sum of the \hat{a} coefficients should be unity. When these two conditions are met inflation will be stable if the actual unemployment rate is equal to the NAIRU.

We estimate equation (1) using quarterly U.S. data over the period 1960:1 to 1993:4. The inflation variable is the annualized percentage change in the CPI for urban consumers excluding food and energy costs (This is often referred to as "core" inflation). The NAIRU is assumed constant over the period at 6 per cent, which until recently, was widely agreed to be approximately its true level. Four lags of each variable and the contemporaneous unemployment gap are included in the regression. These four lags were sufficient to eliminate

serial correlation in the estimated residuals, thus ensuring the consistency of the estimators. The inclusion of the contemporaneous unemployment gap variable could, in principle, cause a degree of simultaneity bias. However, Saiger et. al. (1997) reported that, in practice, this bias was not significant. The results presented here are qualitatively unchanged if the contemporaneous value of the variable is excluded.

The first column of table 1 shows the results. Each cell in the table shows the sum of the coefficients, and in parentheses, the p-value on a Wald test that the sum of the coefficients is equal to zero.² The estimated coefficients take on plausible values. The coefficients on the lagged inflation variables sum to unity as expected and the sum of the coefficients on the employment gap are significantly different from zero.

We can use these estimated coefficients to make out-of-sample forecasts of inflation.³ If the estimated model is a correct representation of the true behaviour of inflation, we would expect that any difference between the predicted and actual inflation rates would turn out to be random. The first panel of figure 1 shows the actual and predicted inflation for the period 1994-1997. It is clear that the Philips curve has consistently over-predicted the rate of inflation. The last two rows of table 1 show the mean and standard deviation of the forecast error. The mean forecast is greater than zero although only by two thirds of a standard deviation. A formal Chow test of the forecast error cannot reject the hypothesis that the forecasts error is zero. However, a Chow test focuses on the average size of the forecast error and not on their sign. The phenomenon of positive shocks has been too pervasive to have been the result of random shocks or measurement error. If the observed forecast errors are truly random draws from a symmetric distribution, the probability that 10 of them are positive is 0.5^{10} which is less than 1 out 1,000. We must conclude therefore, that the estimated model does not accurately account for the behaviour of inflation and unemployment during the late 1990s.

C. Has the NAIRU Fallen?

The most obvious, potential explanation for recent inflation performance is that the NAIRU has fallen to less than the 6 percent that was assumed previously. This would explain why

²If, for example, the p-value is 0.04, the null hypothesis of zero effect cannot be rejected at the 1 percent significance level, but it can be rejected at the 5 percent significance level. Thus, the lower the p-value, the lower the probability that the variable has no effect on inflation.

³The forecasts shown in this paper are static as opposed to dynamic forecast. In other words, when forecasting in to the future, the lagged actual values of the inflation variable are used rather than the values previously predicted by the model.

inflation is falling even as unemployment is much less than 6 percent. A decrease in the NAIRU is very good news for the economy. It implies that there has been structural change in the US that will allow unemployment to be kept low, even in the long term, without the danger of increased inflation.

Broadly speaking, two methods have been employed in the literature to estimate the NAIRU. One method, exemplified by Adams and Coe (1990), involves estimating a structural model of the economy, with the NAIRU being a function of one or more of the estimated parameters. The model includes variables that might be expected to influence the structure of the labour market as well as the standard macro variables. In practice there are three serious problems with this procedure. Firstly, it is very difficult to account for all the changes in the structure of the labour market. Secondly, as the choice of variables and functional form are not clear a priori and as these models involve consideration of expectations (at least implicitly), they are highly vulnerable to the Lucas critique. Thirdly, the models often make ad-hoc assumptions regarding the time path of the NAIRU or equivalently potential GDP.

The alternative method is to estimate the NAIRU from within the Philips curve itself. This is the method adopted by Gordon (1997) and Staiger, Stock and Watson (1997). In a sense, this method reverses the logic of the Philips curve. The econometrician finds that NAIRU which is implied by the behaviour of inflation and unemployment over a given time period i.e. find the U^N that would solve equation (1). The procedure can be repeated for a different period and thus a time series for the NAIRU can be calculated. Of course, if that was all that was done, the exercise would be trivial and worthless. By definition, the estimated NAIRU would explain recent behaviour and the problem discussed in the previous section could never be observed. Furthermore, one would doubt that changes in the NAIRU measured in this way would correspond to underlying changes in the structure of labour markets, which is what most economists have in mind when they talk about a change in the NAIRU. In order to accommodate this criticism, the authors restrict the degree to which the estimated NAIRU can move from year to year.⁴ Nevertheless, these models are fundamentally unidentified in the sense that the imposition of the restriction on the variability of the NAIRU is essentially arbitrary. It is also the case that these models not attempt to identify the causes in any change in the NAIRU.

For reasons of simplicity, we estimated a semi-structural model of the NAIRU. The model is semi-structural because observed unemployment is regressed on variables that may be supposed to affect it, but without attempting to identify behavioural relationships. In terminology of simultaneous equations models, the estimated equation is reduced form. The model is structural in the sense that it does not follow the procedure of Gordon (1997) and Staiger, Stock and Watson (1997) discussed above. The actual rate of unemployment was

⁴Gordon (1997) forces the NAIRU to follow a stochastic trend with a restriction on the variance of its first difference. Staiger, Stock and Watson (1997) force it to follow a cubic spline.

regressed on quadratic time trend, several structural and several cyclical variables. The structural variables included the dependency ratio, an index of minimum wages and an index of unionization. The cyclical variables included a dummy variable reflecting the NBER's definition of recession, an index of co-incident indicators and a variable representing the deviation of capacity utilization from its trend. The model was estimated using the Cochrane-Orcutt method to correct for serial correlation in the residual. The results are shown in Table 2. Of the structural variables only the index of unionization turned out to be significant, but all the cyclical variables are significant. Also note that the coefficient on the autoregressive residual is highly significant. This reflects the tendency of the natural rate not to stray too far from its previous value. As was noted earlier Gordon (1997) and Staiger et. al. (1997) achieved this by restricting the NAIRU to follow a stochastic or deterministic trend.

Setting the cyclical variables equal to zero and simulating the estimated equation can calculate an estimate of the NAIRU. The estimated NAIRU and the two standard deviation error bands are shown in the second panel of figure 2. Figure 2 also shows the actual level of unemployment and a Holdrick-Prescott filtered unemployment series. The Holdrick-Prescott filter is equivalent to regressing actual unemployment on a series of time trends and then taking the fitted value. The idea is that short run changes are "filtered" out and what we are left with is the long run trend rate of unemployment. This can be thought of as an alternative measure of the NAIRU.

As can be seen from the second panel of figure 2 the point estimate of the NAIRU has fallen from 6.1 percent in 1994:1 to 5.3 percent in 1997:4. However, the two standard error bands around the point estimate is approximately 1.3 percent. This implies that a 95 percent confidence interval would place the NAIRU anywhere between just over four and just under seven percent during 1997. For example, the point estimate for the third quarter of 1997 was 5.3 percent with a standard error of 0.65 percent. This gives 95 percent confidence interval for the NAIRU of 4.1 percent to 6.6 percent of the labour force. These point and interval estimates are similar to those found in Staiger, Stock and Watson (1997) and also to those in Gordon (1997). The lack of precision, while disappointing, is typical of estimates of the NAIRU in general. It explains why there is so much uncertainty regarding the level of the NAIRU. Standard statistical tests have almost no power to distinguish between the hypothesis that the NAIRU is 6 percent versus the alternative hypothesis that it is 4 percent. Yet whether the NAIRU is truly 4 percent or 6 percent is of great interest to policy makers.

The point estimate of the NAIRU can be used to re-calculate the unemployment gap and then re-estimate the Philips curve. The results are shown in the second column of table 1. As before, these estimated coefficients, can be used to make out-of-sample forecasts of inflation. It is clear that the Philips curve still consistently over predicted the level of inflation, but by less than the previous model. The average forecast error is greater than zero, but smaller than for the previous model. This can be seen from the second panel of figure 1 which shows the actual and predicted inflation for the period 1994-1997. As before, that it

is not so much the size of the forecast errors that gives cause for concern, as the fact that they are systematically positive.⁵

In summary, although the measured NAIRU has fallen since 1994, it has not fallen sufficiently far in order to explain inflation. The NAIRU has still been above actual unemployment in every quarter since 1994:1. This should have caused a sustained acceleration in inflation. Instead inflation was stable or falling over the entire period. Hence the Philips curve with the new measure of the NAIRU still over-predicts inflation.

D. Have Costs Fallen?

If changes in the NAIRU do not explain the recent behaviour, what does? As mentioned previously, the alternative is that the U.S. has been fortunate enough to have been on the receiving end of a series of beneficial supply shocks that have helped keep inflation under control. For example, suppose oil prices were to fall unexpectedly. This would reduce the cost of energy to U.S. firms, which would enable them to keep their prices relatively low, even if unemployment were so low (i.e. below the NAIRU) that wage costs were rising. If the change in the cost of energy is transitory then the NAIRU itself would be unaffected. Thus, we would observe inflation falling even as unemployment is below the NAIRU.

In order to account for this possibility, variables representing the potential supply shocks must be included in the Philips curve. Strictly speaking a shock should be measured as the difference between the actual value of the variable and value that was anticipated when agents made the economic decisions that determined the NAIRU. In general, reliable data on expectations are not available, so the shocks are approximated by the rate of change of a variable.⁶

The third column of table 1 shows the results from a Philips curve regression including supply shock variables. These shock variables measure changes in the real price of oil, changes in the dollar price of imported goods and changes in the real unit cost of labour. The changes in the price of imports will pick up the effect of a change in the value of the US dollar as well as changes in foreign currency prices of U.S. imports. The real unit cost of labour will include changes in wages, benefits and also temporary changes to productivity.

All of the point estimates have the expected sign. For example, a fall in the price of oil of 25 percent would lead to a fall in the inflation rate of one quarter of a percentage point. The p-values indicate that the hypothesis that the variable has no impact on inflation cannot not be

⁵The results are almost the same if the Holdrick-Prescott measure of the NAIRU is used instead.

⁶This would be exactly right if the variable followed a random walk.

rejected at conventional significance levels, in the case of all the variables. In the case of the oil price and the import prices, this insignificance is borderline, but in the case of the unit labour cost, the insignificance is overwhelming.

There is some double counting because price of oil is included in the price of imports and this may lead to multicollinearity. The fourth column of table 1 shows the results of estimation of the Philips curve with the oil price variable dropped. As a result, the import price variable becomes highly significant as expected, but unit labour cost variable remains insignificant. These results are similar to those reported by Gordon (1997). He found that import prices had a significant effect on core inflation, but that the deviation of productivity from its trend (which is reflected here in the real unit labour cost variable) did not.

As a forecasting tool, this model performs much better than does the model without supply shocks. The average error of the forecast is close to zero, indicating that this model does not systematically over-predict inflation. This can also be seen from the third panel of figure 1, which shows actual inflation and the inflation rate predicted by the model in column 4 of Table 1, for the period 1994:1 to 1997:4. Note also that errors are not systematically positive.

Comparing the first two panels of figure 1, suggests that the recent U.S. inflation performance has been due to falling import prices and not to fundamental structural change in the economy. Further evidence for this hypothesis can be found in by looking at the price of imports has changed over the last 8 years. Figure 3 shows both the level and the rate of change of the relative price of imports (defined as the implicit import price deflator divided by the GDP deflator). This ratio fell, more or less consistently up to 1994:1. It remained approximately constant until 1995:2, and then it fell dramatically. This latter fall coincides with the period for which the traditional Philips curve over-predicts inflation (see figure 1). On reflection, therefore, it is not surprising, that including import price changes in the Philips curve enables the model to predict inflation more accurately.

It is also worth noting the relative unimportance of the real unit labour cost variable. Some authors⁷ have suggested that the reason inflation was falling faster than appeared warranted by unemployment levels was that the structure of labour costs had changed, with productivity rising faster than normal and non-wage labour costs (such as benefits) being cutback. The results shown here suggest that this was not a critical factor in explaining the recent history of U.S. inflation.

⁷See, for example, Lown and Rich (1997)

E. Conclusions

This paper set out to explain the recent history of U.S. inflation, in particular how inflation could be falling even as unemployment appears to be below the natural rate. Two potential explanations were considered. Firstly, the natural rate of unemployment could have fallen. Secondly, the U.S. could have experienced some temporary disinflationary shocks.

The results presented here suggest that the NAIRU may have fallen slightly, but that this, on its own, is not enough to explain the recent history of inflation. Other researchers have also reached this conclusion.⁸ The main explanation for recent inflation in the U.S. appears to be that there have been unexpected declines in some cost variables. In particular the cost of imports has fallen in recent years due to the appreciation of the dollar, a decline in commodity prices and the Asian crisis. It was also noted that the recent inflation is not well explained by changes in the cost of labour.

The conclusion that the recent good performance of inflation is the result of supply shocks rather than a change in the NAIRU has important implications for policy makers. It implies that the structure of the economy has not altered fundamentally. Thus, if and when the supply shocks disappear, unemployment will have to increase to the level of the NAIRU if inflation is not to rise. In 2000:1 the required increase would be approximately 0.6 percentage points. Conversely, as long as the U.S. continues to be on the receiving end of positive supply shocks, the recent behaviour of inflation and unemployment should be repeated. Given that the Asian crisis has lowered the dollar price of Asian exports and that the dollar has recently appreciated against some major trading partners, the short run prognosis looks good. But in the medium to long run, without a further decline in the NAIRU, the free lunch must eventually end.

⁸See Gordon (1997), Staiger, Stock and Watson (1997) and also Lown and Rich (1997)

List of References

Adams, Charles and Coe, David, (1990), "A Systems Approach to Estimating the Natural Rate of Unemployment and Potential Output for the United States," *IMF Staff Papers*, vol. 37, no. 2, pp. 232-293.

Gordon, Robert, 1997, "The Time-Varying NAIRU and its Implication for Economic Policy," *Journal of Economic Perspectives*, Vol. 11, No. 1, pp. 11-32.

Lown, Cara and Rich, Robert, 1997, "Is There an Inflation Puzzle?," *Federal Reserve Bank of New York Policy Review*, December.

Staiger, Douglas, James Stock, and Mark Watson, 1997, "The NAIRU, Unemployment and Monetary Policy," *Journal of Economic Perspectives*, Vol. 11, No. 1, pp. 33-49.

Table 1. Estimates of the Philips Curve
(1961 Q1–1993 Q4)

	(1)	(2)	(3)	(4)
Model	Basic Model	Natural Rate Model	Supply Shocks Model I	Supply Shocks Model II
Lagged inflation	1.016 (0.000)	1.006 (0.000)	0.997 (0.000)	1.002 (0.000)
Unemployment gap	-0.246 (0.007)	-0.361 (0.004)	-0.289 (0.018)	-0.255 (0.033)
Change in real unit labour cost	-	-	0.085 (0.37)	0.079 (0.398)
Change in relative price of imports	-	-	0.045 (0.121)	0.072 (0.001)
Change in the relative price of oil	-	-	0.01 (0.156)	-
Adjusted R ²	0.755	0.754	0.804	0.803
Mean forecast error (Standard deviation)	0.38 (0.58)	0.27 (0.52)	0.09 (0.58)	0.002 (0.64)

1/ The table shows the sum of coefficients on four lags of the variables. The statistic in parentheses is the p-value from a Wald Test of the hypothesis that sum of coefficients on the variable is zero. The lower the p-value, the lower the probability that the variable has no effect on inflation. If, for example, the p-value is 0.04, the null hypothesis of zero effect cannot be rejected at the 1 percent significance level, but it can be rejected at the 5 percent significance level.

2/ Rate of change variables are scaled to measure annualized percentage rate of change.

Table 2. NAIRU Equation
(1961 Q2–1997 Q3)

Variable	Coefficient	T-Statistic
Constant	2.6	0.27
Trend	0.13	1.10
Trend ² /1000	-.69	-1.49
Dependency ratio	10.42	0.99
Minimum wage	-0.09	-0.56
Unionization	-25.66	-2.28
Change in co-incident indicators	5.42	1.84
Capacity utilization	-0.23	-14.86
Recession	-0.03	-0.88
Autoregressive error term	0.95	35.59
Adjusted R ²	0.98	

Table 3. Definitions of Variables 1/

Variable	Definition
Capacity utilization	Capacity utilization of all industry.
Coincident indicators	A composite index of coincident indicators.
Dependency Ratio	Population less than 15 years of age , plus the population greater than 65 years of age, all divided by the labour force.
Import prices	Implicit import price deflator from the national accounts, divided by the GDP deflator.
Inflation	Rate of change of CPI for urban consumers excluding the cost of food and energy.
Minimum wage	Minimum wage (dollars per hour) deflated by the CPI for urban consumers.
Unemployment	Civilian unemployment rate.
Unionization	Union membership as a percentage of the civilian labour force.
Recession	Indicator of recession as defined by the National Bureau of Economic Research (NBER)
Real unit labour costs	The unit labour cost of the nonfarm business sector divided by the GDP deflator.

1/ All rates of change are expressed as annual percentage rate of change.