

**Inside and Outside Bounds:
Threshold Estimates of the Phillips Curve**

Michelle L. Barnes and Giovanni P. Olivei*

Over the past 30 years, debates about the usefulness of the Phillips curve for explaining inflation behavior have been ongoing. At the end of the 1970s, the Phillips relationship was heavily criticized for its apparent inability to characterize inflation dynamics in the face of oil shocks. The relationship was also deemed inappropriate as a policy tool because of its potential instability in the face of changes in the economic environment. Subsequent empirical work in the early 1990s showed that such criticism was largely unjustified. By controlling for supply shocks as shifters in the Phillips relationship, these studies provided evidence of a stable and significant tradeoff between inflation and unemployment.¹ Indeed, the Phillips relationship continues to feature prominently in several macroeconometric models used for policy analysis.²

More recently, though, as a result of the striking combination of declining unemployment with no attendant inflationary pressures during the second half of the 1990s, criticism has mounted once again. This development has led many researchers to question again the viability of the Phillips curve and its usefulness as a macroeconomic policy tool. While the debate continues, the fall in inflation coupled with significant slack in the economy since late 2002 seems consistent with a standard Phillips curve tradeoff.

¹See, for example, Fuhrer, 1995 and Tootell, 1994.

One of the reasons for the recurring debates about the existence of an inflation and unemployment tradeoff is that there have been several instances when large movements in the unemployment rate have elicited little response in the inflation rate. Figure 1 shows the behavior of inflation, measured by the change in the core personal consumption expenditures (PCE) deflator, and the unemployment rate for each of the four decades, beginning in the 1960s. In the 1960s (Chart A), the unemployment rate fluctuated between 4 percent and 7 percent for many years, with little change in the inflation rate. The tradeoff was evident only at the end of the decade, when the unemployment rate dropped below 4 percent.

After the oil price shock of the early 1970s and the associated stagflation (Chart B), the years 1976 to 1979 saw the unemployment rate decrease noticeably, with only a modest change in inflation.

Then, the relationship seemed to re-establish itself; the recession of the early 1980s (Chart C) saw the expected tradeoff, with increasing rates of unemployment associated with a sharp decline in inflation. However, this tradeoff was not apparent from 1985 to 1989, even though the unemployment rate dropped appreciably.

Similarly, the recession of the early 1990s (Chart D) witnessed a large fall in inflation, but since 1994 and through late 2002, large movements in the unemployment rate, at first decreasing and then increasing, have been associated with relatively small changes in inflation.

² Blinder (1998), for example, notes that "a linear Phillips curve fits the data extremely well." Not everyone would agree with such a statement. See, for example, Atkeson and Ohanian (2001).

In principle, these episodes of horizontal movement -- that is, episodes of inflation stability coupled with large changes in the unemployment rate -- are consistent with a Phillips curve relationship. They just require the curve to shift in the same direction as the movement in the unemployment rate. For example, the experience of the second half of the 1990s can be reconciled, at least partly, with a Phillips curve tradeoff, by arguing that over this period the natural rate of unemployment declined.

Econometric representations of the Phillips relationship usually incorporate factors that can cause the Phillips curve to shift over time. Typical control variables are food and energy shocks, inflation expectations, and, in some instances, a changing natural rate of unemployment.³ So far, however, the literature has not provided a test of whether such controls are sufficient to explain the episodes of horizontal movement.

In this paper, we test the explanatory power of a piecewise linear specification of the Phillips relationship against a simple linear specification. The piecewise linear specification allows the inflation and unemployment tradeoff to vary with the level of the unemployment rate. Such a specification maintains that the tradeoff may vary, depending on whether the unemployment rate lies within or outside some range. If the usual shifters are sufficient to characterize periods of horizontal movement, then a piecewise linear specification should not improve upon the standard linear Phillips curve.⁴

³ Some would argue that the acceleration of productivity in the second half of the 1990s was an important contributing factor to low inflation. There is little compelling evidence, though, supporting the inclusion of productivity as a control variable in the Phillips relationship.

⁴ This is not the first attempt to characterize a changing tradeoff over different levels of the unemployment rate, but it has the advantage of parsimony. In 1958, A. W. Phillips originally drew a nonlinear curve

Instead, our findings strongly support a piecewise linear Phillips relationship. The evidence indicates that the traditional shifters in the relationship are insufficient to characterize the episodes of horizontal movement. Apparently, the gap between the unemployment rate and the natural rate of unemployment must be larger or smaller than some threshold values before triggering a response in inflation. When the unemployment gap lies within the range defined by the thresholds, there is no evidence of a significantly and economically relevant tradeoff between inflation and unemployment.

Granted that some features of the Phillips curve remain theoretically puzzling,⁵ what factors might explain a piecewise linear specification? Bargaining between firms and workers implies that the negotiated wage lies within a range defined by the reservation wage for the firm and the reservation wage for the worker. This range will shift when economic conditions change. Typically, the bargaining process has the feature that at the time of renegotiating the wage, the new wage changes by the smallest possible amount necessary to bring it within the range defined by the reservation wages of the worker and the firm (Thomas and Worrall 1988). As a result, changes in economic conditions usually lead to small changes in wages and prices. It is only when changes become very large that a discrete jump in the wage is necessary to bring it within the new range. This, in turn, will lead to a discrete jump in prices (Hall 2003).

Other explanations for a piecewise linear Phillips curve relate to the shape of the demand curve faced by firms. Suppose

such that low unemployment raises wage inflation more than high unemployment lowers it. Other types of nonlinearities in the Phillips relationship have been explored by Eisner, 1997.

⁵ See Blanchard and Fischer, 1989.

firms face kinked demand curves, with the demand for their product decreasing sharply if firms increase their price and increasing little if firms decrease their price. Then it is possible to show that shifts in demand can be accompanied by little or no movement in prices, unless such shifts become very large (Woglom 1982).

The rest of the paper continues with a description of the linear and piecewise linear Phillips curve representations, along with an explanation of the testing and estimation methodology. The empirical results then follow, along with a conclusion that relates the findings to the most recent inflation experience.

I. Linear vs. Piecewise Linear Phillips Curves

The standard approach to estimating the Phillips curve posits a linear short-run tradeoff between inflation and the level of an indicator of the strength of demand in the economy, such as the unemployment rate. In its linear form, a general Phillips curve is given by the following specification:

$$\mathbf{p}_t = \mathbf{m} + \sum_{j=1}^s \mathbf{a}_j \mathbf{p}_{t-j} + \mathbf{b}u_t + \sum_{j=0}^k \mathbf{d}_j \Delta u_{t-j} + \sum_{j=0}^l \mathbf{x}_j Z_{t-j} + \mathbf{e}_t . \quad (1)$$

The dependent variable \mathbf{p}_t is the rate of inflation. Lagged inflation captures the inertia in the way inflation expectations are formed. Inflation expectations enter the Phillips relationship because workers, concerned with real wages, take into account expected changes in prices when contracting changes in nominal wages. Equation (1) then posits that expected inflation, \mathbf{p}_t^e , is formed as a weighted average of past inflation, with the weights \mathbf{a}_j estimated on actual data.⁶ The variable u_t is the indicator of the intensity of demand in the economy, which in the present analysis takes the form of the unemployment rate. The unemployment rate enters equation (1) both in levels and in first differences (with the first difference $u_t - u_{t-1}$ denoted by Δu_t). Thus, the coefficient β measures the effect of the level of the unemployment rate on inflation, while the sum of coefficients \mathbf{d}_j measures the effect of current and lagged changes in unemployment on inflation (often called the "speed limit" effect). The crucial parameter of interest in the Phillips relationship is, of course, β . In the empirical section that follows, however, we will also

⁶ The sum of the coefficients on lagged inflation is usually constrained to unity, so that in the long run actual inflation equals expected inflation.

mention some results concerning the "speed limit" effect. Finally, the Phillips curve relationship incorporates current and lagged supply shock variables, here summarized by the vector Z_t , (normalized so that $Z_t=0$ indicates absence of supply shocks), while e_t is a serially uncorrelated error term.

The *prima facie* evidence on the relationship between inflation and unemployment discussed in the previous section suggests much horizontal movement in the Phillips curve. In terms of equation (1), this horizontal movement could be accounted for by the supply shocks Z_t , or by changes in inflation expectations. Figure 2 shows a scatterplot of the partial correlation between inflation and the level of the unemployment rate that arises from estimating equation (1) over the period from the third quarter of 1961 through the fourth quarter of 2002.⁷ The estimation uses observations for which the unemployment rate u_t ranges only between 4.0 percent and 7.5 percent over the period considered. The figure shows that within the specified range of the unemployment rate, little inverse relationship between inflation and unemployment is apparent even after controlling for supply shocks and inflation expectations.

If we instead estimate equation (1) using all the observations -- and not just the observations for which the unemployment rate lies between 4.0 percent and 7.5 percent -- the estimated tradeoff between inflation and unemployment becomes statistically significant and economically relevant. One potential explanation for this phenomenon is that the tradeoff

⁷ Details about the estimation are given in Section III. The measure of inflation used in deriving Figure 2 is the change in the core PCE deflator.

between inflation and unemployment differs at different levels of the unemployment rate. Broadly speaking, the tradeoff could be more pronounced for high and low values of the unemployment rate than for "normal" values. To assess whether this is indeed the case, we consider a more general specification of equation (1) that takes the following form:

$$p_t = \underline{\mu} + \sum_{j=1}^s \underline{a}_j p_{t-j} + \underline{\beta} u_t + \sum_{j=0}^k \underline{d}_j \Delta u_{t-j} + \sum_{j=0}^l \underline{x}_j Z_{t-j} + e_t, \quad (1')$$

where

$$\underline{\mu} = \mathbf{m}_1 I(u_t) + \mathbf{m}_0 (1 - I(u_t)), \quad \underline{\beta} = \mathbf{b}_1 I(u_t) + \mathbf{b}_0 (1 - I(u_t)), \quad \text{and}$$

$$\underline{d} = \mathbf{d}_1 I(u_t) + \mathbf{d}_0 (1 - I(u_t)).$$

The indicator variable $I(u_t)$ takes the value of 1 when the unemployment rate u_t lies within a specified interval $[\mathbf{g}_L, \mathbf{g}_H]$, and the value of 0 when u_t lies outside that interval. Simply put, equation (1') allows the intercept \mathbf{m} and the coefficients on the level and first difference of the unemployment rate, \mathbf{b} and \mathbf{d}_j , to take different values according to whether u_t lies inside or outside a specified range.

The piecewise linear form of equation (1') is illustrated in Figure 3, which depicts the partial correlation between inflation and the level of unemployment. The tradeoff between inflation and unemployment takes the value of \mathbf{b}_1 when the unemployment rate is inside the interval between \mathbf{g}_L and \mathbf{g}_H , and it takes the value of \mathbf{b}_0 when the unemployment rate is outside this interval. In principle, there is no reason to constrain the tradeoff between inflation and unemployment to be the same for "low" and "high" values of the unemployment rate. In equation (1'), we do

constrain the coefficients m , b , and d_j to be the same when the unemployment rate is below threshold g_L or above threshold g_H only to preserve degrees of freedom at the estimation stage. Since we are mainly interested in assessing the strength of the tradeoff between inflation and unemployment when the unemployment rate is not particularly high or low, this simplification is not crucial. Moreover, to the extent that the piecewise linear relationship in (1') is statistically better than the linear relationship (1), then, *a fortiori*, a piecewise linear relationship that is even more flexible than (1') will perform better than the linear relationship (1).

Another potential explanation for the weak tradeoff between inflation and unemployment depicted in Figure 2 is that the linear relationship in equation (1) omits some important explanatory variable that acts as a shifter to the inflation-unemployment tradeoff. An obvious candidate is a time-varying natural rate of unemployment.⁸ As shown in the previous section, during the years 1994 to 1999, the unemployment rate fell from 6.5 percent to 4.1 percent, while the rate of price inflation actually declined slightly. Instead of reflecting a lack of a significant tradeoff over this range of the unemployment rate, the failure of inflation to increase over this period could be due to a decline in the natural rate of unemployment. Such a decline would lead to a "shifting in" of the Phillips curve, with potentially little effect on inflation. If this is indeed the case, then the suggested piecewise linear representation (1') would grossly misrepresent the nature of the inflation-

⁸ Equation (1) assumes that, absent supply shocks, there is a constant natural level of the unemployment rate that is consistent with a

unemployment tradeoff, even if it were to provide a better approximation to the data than the linear relationship (1).

For this reason, we consider an alternative formulation of the linear Phillips curve relationship that explicitly allows for a changing natural rate of unemployment:

$$\mathbf{p}_t = \sum_{j=1}^s \mathbf{a}_j \mathbf{p}_{t-j} + \mathbf{b}(u_t - u_t^N) + \sum_{j=0}^k \mathbf{d}_j \Delta u_{t-j} + \sum_{j=0}^l \mathbf{x}_j Z_{t-j} + \mathbf{e}_t \quad (2)$$

where u_t^N denotes the time-varying natural rate of unemployment. Factors related to labor demand and supply, such as demographics and productivity, could in principle lead to changes in the natural rate of unemployment. Unfortunately, the natural rate of unemployment is not known, and the series u_t^N needs to be estimated. None of the methods commonly used for estimating a time-varying natural rate of unemployment is foolproof, and the uncertainty surrounding the estimated u_t^N is usually large. Thus, while equation (2) is potentially more informative than (1), estimates of a Phillips relationship of the form of (2) should still be regarded with caution.

The piecewise linear counterpart to the linear specification (2) is given by:

$$\mathbf{p}_t = \underline{\mu} + \sum_{j=1}^s \mathbf{a}_j \mathbf{p}_{t-j} + \underline{\beta}(u_t - u_t^N) + \sum_{j=0}^k \underline{\mathbf{d}}_j \Delta u_{t-j} + \sum_{j=0}^l \mathbf{x}_j Z_{t-j} + \mathbf{e}_t, \quad (2')$$

where

$$\underline{\mu} = \mathbf{m}_1 I(u_t - u_t^N) + \mathbf{m}_0 (1 - I(u_t - u_t^N)), \quad \underline{\beta} = \mathbf{b}_1 I(u_t - u_t^N) + \mathbf{b}_0 (1 - I(u_t - u_t^N)), \quad \text{and}$$

$$\underline{\mathbf{d}} = \mathbf{d}_1 I(u_t - u_t^N) + \mathbf{d}_0 (1 - I(u_t - u_t^N)).$$

The interpretation of the coefficients in equation (2') is the

constant level of inflation. This constant natural level of the unemployment rate is estimated as $-\mu/\beta$.

same as in equation (1'), but now $I(u_t - u_t^N)$ takes the value of 1 when the unemployment gap -- i.e., the deviation of the unemployment rate from the time-varying natural rate of unemployment -- lies within a specified interval $[g_L, g_H]$, and it takes the value of 0 when the unemployment gap lies outside that interval. If the presence of a changing tradeoff between inflation and the level of unemployment over different ranges of the unemployment rate is largely the result of a time-varying natural rate of unemployment, then equation (2') should not provide a better representation of the data than the linear specification (2).

In the next section, we proceed to estimate the piecewise linear, or threshold, relationships (1') and (2') and their nested linear versions, equations (1) and (2), respectively. In so doing, we assess whether the difference between a linear and a threshold representation of the Phillips relationship is meaningful, from both a statistical and an economic standpoint.

II. Estimation Method and Data Description

In order to evaluate empirically a threshold relationship such as (1') or (2'), it is necessary to estimate, in addition to all other parameters in the relationship, the threshold parameters \mathbf{g}_L and \mathbf{g}_H . Estimation is greatly simplified by noting that for given values of \mathbf{g}_L and \mathbf{g}_H , it is possible to estimate the threshold relationship by ordinary least squares (OLS). As a result, the method for estimating a threshold relationship consists of a sequential process that performs OLS on (1') or (2') over different values of the pair $(\mathbf{g}_L, \mathbf{g}_H)$. The threshold estimates are then given by the pair $(\hat{\mathbf{g}}_L, \hat{\mathbf{g}}_H)$ for which the OLS regression of (1') or (2') yields the minimum sum of squared residuals. The other parameters' estimates in relationship (1') or (2') result from the OLS regression that uses $(\hat{\mathbf{g}}_L, \hat{\mathbf{g}}_H)$ as the pair of threshold estimates.

The linear relationships (1) and (2) of the Phillips curve can be thought of as restricted versions of (1') and (2'), respectively.⁹ In particular, in the linear case the parameters are restricted to be equal inside and outside the interval $[\mathbf{g}_L, \mathbf{g}_H]$ for any value of \mathbf{g}_L and \mathbf{g}_H .¹⁰ Since the threshold specification nests its linear version, the threshold model will always provide at least as good a fit as the linear model. The question is whether the fit is significantly better from a statistical

⁹ Of course, relationships (1) and (2) can be estimated directly by OLS.

¹⁰ If the parameters are restricted to be equal inside and outside a given interval $[\mathbf{g}_L, \mathbf{g}_H]$, then the restriction will apply to any interval $[\mathbf{g}_L, \mathbf{g}_H]$, as one can easily infer from Figure 3. As a result, when the linear restriction holds, the threshold values \mathbf{g}_L and \mathbf{g}_H are not identified.

standpoint. To address this issue, we perform an F -test that uses the sum of squared errors from the (restricted) linear and the (unrestricted) threshold model, respectively. Because under the null hypothesis of linearity the values of the upper and lower thresholds \mathbf{g}_L and \mathbf{g}_H from the piecewise linear alternative are not identified, the F -test has a non-standard distribution. The distribution theory for such a test, however, is now well developed (Hansen 1996 and 2000), and we rely on the extant literature to construct p -values for the test. Box 1 provides additional information about estimation and inference of threshold models -- including information on how to obtain confidence intervals for the estimates $\hat{\mathbf{g}}_L$ and $\hat{\mathbf{g}}_H$.

The data used in the estimation are quarterly and cover the period from the third quarter of 1961 through the fourth quarter of 2002. We consider three basic measures of inflation: the change in the core component of the chain-weighted deflator for personal consumption expenditures, the change in the core component of the Consumer Price Index (CPI), and the change in the chain-weighted GDP deflator.¹¹ Our measure for u_t is given by the civilian unemployment rate. When evaluating relationships (2) and (2'), we also need to form an estimate of the time-varying natural rate of unemployment, u_t^N . We use the Congressional Budget Office (CBO) measure of u_t^N , and we discuss alternative measures later in the section. The supply shock variables included in Z_t are the change in the relative price of food and energy and Gordon's (1982) series for wage and price controls.

¹¹ Denoting the price index by p_t , inflation is defined as $p_t = 100 * ((p_t / p_{t-1})^4 - 1)$.

When the measure of inflation is given by the change in the GDP deflator, Z , also includes the change in the trade-weighted dollar. Data sources for all series are given in the data appendix.

III. Empirical Findings

We first discuss estimation results for the linear and the piecewise linear forms of the Phillips curve with a constant natural rate of unemployment, equations (1) and (1'), respectively. In the piecewise linear case, the threshold parameters g_L and g_H are estimated over a wide range of values taken by the unemployment rate during the sample period. Table 1 summarizes the main findings for the three measures of inflation considered. The table reports estimates for the main parameter of interest in the Phillips relationship, the coefficient b on the level of the unemployment rate, for both the linear and the piecewise linear specifications. In the piecewise linear specification, the coefficient takes the value b_I when the unemployment rate lies inside the interval $[g_L, g_H]$, and the value b_O when the unemployment rate is either below g_L or above g_H -- outside the interval. The table also reports estimates of the thresholds g_L and g_H , and, in square brackets, the 95 percent confidence interval associated with these estimates. The last column of the table shows the p -value of an F -test of the null hypothesis of a linear model against the estimated threshold alternative.

As usual, the estimated coefficient on the level of the unemployment rate in the linear model is highly significant and economically relevant for all three measures of inflation. The piecewise linear representation of the Phillips curve shows that the tradeoff between inflation and unemployment is only significant for values of the unemployment rate below \hat{g}_L or above \hat{g}_H . For all three measures of inflation, g_L and g_H are estimated

at about 4 percent and 7.5 percent, respectively. There is no evidence of a tradeoff between inflation and the level of the unemployment rate when the unemployment rate lies between 4 percent and 7.5 percent. In addition, it is possible to show that "speed limit" effects are estimated to be insignificantly different from zero when the unemployment rate lies inside the interval $[\hat{g}_L \hat{g}_H]$. However, for values of the unemployment rate below \hat{g}_L or above \hat{g}_H , these effects tend to be significant and larger than the ones estimated by means of a linear specification.¹²

The last column in the table shows that, when inflation is measured by either the core PCE deflator or the core CPI index, the null hypothesis of a linear model is rejected in favor of a threshold specification. The small p -values indicate that the piecewise linear representation of the Phillips curve is, from a statistical standpoint, highly significant. When inflation is instead measured by the GDP deflator, the null hypothesis is rejected in favor of a threshold specification only marginally at the asymptotic 5 percent level. This, however, has to do with the presence of two influential observations.¹³ As shown in the last row of the table, when these two observations are excluded from the sample, the piecewise linear representation of the Phillips curve again becomes highly significant.

¹² This is true when inflation is measured either by the core CPI index or by the GDP deflator. However, when inflation is measured by the core PCE deflator, speed limit effects are estimated to be insignificantly different from zero in both the linear and the piecewise linear specifications.

¹³ These observations are for the second quarter of 1981 and the first quarter of 1996.

Confidence intervals for the estimated \hat{g}_L and \hat{g}_H in the table are often large and run into the lower and upper bounds for the unemployment rate, 3.9 percent and 7.6 percent, respectively, over which we are searching for threshold effects. Still, it is interesting to note that, when inflation is measured by either the core PCE deflator or the core CPI index, the t -statistic associated with the estimated coefficient b_i is never greater than 2 for any pair (g_L, g_H) within the 3.9 percent to 7.6 percent range. When inflation is measured by the GDP deflator and we drop the two outliers from the sample, the maximum t -statistic for b_i is 2.1, with the t -statistic above 2 in only 4 of the approximately 800 different pairs (g_L, g_H) we consider over the mentioned range. Overall, the findings in the table lend support to a piecewise linear version of the Phillips curve, with a statistically significant and economically relevant tradeoff between inflation and unemployment only for particularly high or low values of the unemployment rate.

We next ask whether these findings continue to hold over a shorter sample period. Table 2 reports estimation results for the same specifications of the Phillips relationship as in Table 1, but over the period from the third quarter of 1961 through the fourth quarter of 1986. There are two reasons to be interested in this particular sample period. First, some have argued that the natural rate of unemployment was quite stable over the 1960s, 1970s, and 1980s.¹⁴ If so, then estimating relationships (1) and (1') over the period 1961 to 1986 should alleviate the criticism that the relationships are misspecified by not allowing for a

¹⁴ See, for example, Fuhrer (1995), Gordon (1998), and Tootell (1994).

time-varying natural rate of unemployment. Second, over the period 1987 to 2002, the unemployment rate has been outside the range of 4 percent to 7.5 percent only in three quarters. It seems natural, then, to ask to what extent the findings in Table 1 are driven by the experience of the past 15 years.

Overall, the results in Table 2 confirm our previous findings. For the piecewise linear specification, we find evidence of a statistically significant tradeoff between inflation and the level of the unemployment rate only outside the estimated interval $[\hat{g}_L, \hat{g}_H]$. The point estimates for g_L and g_H are essentially the same as before for all three measures of inflation. The null hypothesis of a linear model is rejected in favor of the threshold alternative when inflation is measured by the core CPI index or the GDP deflator. One cannot reject the hypothesis of a linear specification when inflation is measured by the core PCE deflator, but the t -statistic associated with the estimated coefficient b_i is never greater than 2 for any pair (g_L, g_H) within the 3.9 percent to 7.6 percent range.

It is possible to compare the performance of the threshold specification (1') vis-à-vis the linear specification (1) by means of a "dynamic" simulation, in which predicted values of inflation for the current period are fed into subsequent periods as lagged values. In essence, the simulation is a multi-period in-sample forecast of inflation that does not refer to an actual inflation rate over the simulation horizon. Figure 4 depicts the result of this exercise for core CPI inflation over the period 1961 to 1986, using the estimates reported in Table 2. Considering that the simulation runs over 25 years, both the linear and the piecewise linear model perform well. It is

evident, though, that the threshold model tracks actual inflation much better than the linear model.

Unfortunately, the reported success of the piecewise linear specification is mirrored by its failure to track inflation in a dynamic simulation over the period 1987 to 2002, using the estimates in either Table 1 or Table 2. The reason is that the threshold g_H is usually estimated at about 7.5 percent. The threshold model largely fails to capture the decline in inflation associated with the recession and the slow initial recovery of the early 1990s, when the unemployment rate was above the 7.5 percent mark in only two quarters.¹⁵

This observation, per se, does not mean that a piecewise linear representation of the Phillips curve cannot capture recent inflation dynamics. Indeed, it is possible to show that the piecewise linear relationship (1') estimated over the period 1987 to 2002 is highly significant, with the null hypothesis of a linear relationship (1) always rejected in favor of the threshold alternative. But the estimated interval over which there is little or no tradeoff between inflation and the unemployment rate now ranges from about 4 percent to 6.5 percent. With such an estimate of g_H , it is not surprising that the threshold relationship performs very well in a dynamic simulation of inflation over the past 15 years. Yet, we find it somewhat unappealing to resort to a change in \hat{g}_H in order to explain recent inflation behavior. Parameter instability could in fact be the result of having specified equation (1') incorrectly. For this reason, in the rest of this section we explore whether a

¹⁵ As a result, in a dynamic simulation over the period 1987 to 2002, the piecewise linear model would predict core CPI inflation in the range of 5 percent to 6 percent over the past decade.

Phillips curve relationship that allows for a time-varying natural rate of unemployment provides a more stable guidance for inflation dynamics.

There are several approaches to estimating the path of a time-varying natural rate of unemployment. One approach estimates a changing natural rate of unemployment from a linear Phillips relationship in which the intercept is allowed to vary over time. Another approach takes a *constant* natural rate of unemployment estimate from a linear Phillips curve such as equation (1), but then adjusts the estimate to account for demographic factors. Changes in these factors cause the "demographically adjusted" estimate of the natural rate of unemployment to change over time. An example of the latter approach is the Congressional Budget Office (CBO) measure of the natural rate of unemployment, which we take as our estimate of the time-varying natural rate of unemployment in what follows. In essence, such a measure varies over time because the shares of different demographic groups (broken down by age, sex, and race) in the labor force vary.¹⁶

Figure 5 depicts the CBO estimate of the natural rate of unemployment, the actual value of the civilian unemployment rate, and the "unemployment gap," which is the difference between the unemployment rate and the natural rate of unemployment. The figure shows an increase in the estimate of the natural rate of unemployment from about 5.5 percent at the beginning of the 1960s, to about 6.25 percent in the late 1970s. Since then, the

¹⁶ The CBO estimates, by means of a linear Phillips curve such as equation (1), a natural rate of unemployment for married males. This natural rate of unemployment in turn is used to estimate a natural rate of unemployment for different demographic groups. The overall natural rate of unemployment is then computed as a weighted average of the natural rate of unemployment for the different demographic groups, with the weights set equal to each group's share of the labor force.

natural rate of unemployment has declined, most notably in the past decade, to close to 5 percent. The increase in young workers -- who have higher unemployment rates than older workers -- accounts for much of the rise in the natural rate of unemployment before 1980, while the recent decrease in the number of young workers explains much of the recent fall in the natural rate of unemployment. Because changes in the estimated natural rate of unemployment occur very gradually over time, the unemployment gap depicted in the figure tracks the actual unemployment rate closely.

Table 3 summarizes estimation results for the unemployment gap version of the linear and piecewise linear Phillips curves, equations (2) and (2'), respectively. In the piecewise linear specification, the threshold parameters g_L and g_H are now estimated over a wide range of values taken by the unemployment gap during the sample period. Thus, the tradeoff between inflation and the unemployment gap is given by b_I when the unemployment gap lies inside the interval $[g_L, g_H]$, and by b_O when the unemployment gap is either below g_L or above g_H .

In the linear specification, estimates of the tradeoff, \hat{b} , are significant and economically relevant for all three measures of inflation. In the piecewise linear specification, we find evidence of a significant tradeoff only for values of the unemployment gap below \hat{g}_L or above \hat{g}_H . In order to observe this tradeoff, the unemployment gap must be below -1.4 percent or above 1.4 percent, approximately. It is interesting to note that over the period 1961 to 2002, values of the gap below -1.4 percent usually map into values of the unemployment rate at or

below 4 percent, and values of the gap above 1.4 percent usually map into values of the unemployment rate at or above 7.5 percent. Such a mapping is precise, but not perfect. In particular, in the early 1990s, the gap was above 1.5 percent for six consecutive quarters, whereas the unemployment rate was above 7.5 percent in just two quarters. Still, there is a close similarity between the findings in Table 1 for the constant natural rate of unemployment specification (1') and the findings in Table 3 for the gap specification (2').

The last column in Table 3 shows that the null hypothesis of a linear model is strongly rejected in favor of a threshold specification when inflation is measured by the core CPI index or by the GDP deflator. The linear model is rejected marginally when inflation is measured by the core PCE deflator. In this case, however, the t -statistic associated with the estimated coefficient b_i is always considerably smaller than 2 for any of the different pairs (g_L, g_H) we consider within the -1.4 percent to 1.4 percent range. Overall, we conclude that the piecewise linear model continues to be significant from a statistical standpoint even when we allow for a time-varying natural rate of unemployment.

One could argue that such a finding is an artifact of our chosen measure of the time-varying natural rate of unemployment. For this reason, we experimented with other estimates of the natural rate of unemployment -- specifically with estimates that exhibit more variability than the CBO measure -- but we reached similar conclusions.¹⁷ In general, the more variable the estimated

¹⁷ For example, we have estimated specifications (2) and (2') using a long-run, two-sided moving average of the actual unemployment rate as the estimate for the time-varying natural rate of unemployment (see

natural rate of unemployment, the smaller the estimated interval $[g_L g_H]$ over which there is an insignificant tradeoff between inflation and unemployment. A more variable natural rate of unemployment tracks the actual unemployment rate more closely, and this leads to a decline in the average size of the unemployment gap. But then the estimated interval $[g_L g_H]$ over which there is no significant tradeoff, while smaller, still contains almost two-thirds of the observations, as is the case with the findings in Table 3.

Before concluding, let us briefly revisit the issue of whether the piecewise linear model can track inflation over the years 1987 to 2002. In order to do so, we now estimate equations (2) and (2') for the period 1961 to 1986, and we use these estimates to forecast inflation over the next 16 years. The results of this dynamic simulation for core CPI inflation are displayed in Figure 6. In the first half of the simulation, the linear and threshold specifications perform equally well. Since the unemployment gap is above the estimated threshold value g_H of 1.4 percent for six quarters during the early 1990s, the threshold model is able to capture the downturn in inflation that occurred at that time. Not surprisingly, the threshold model does significantly better than the linear model in the latter half of the simulation. The natural rate of unemployment, while declining over this period, is not declining fast enough to close the negative unemployment gap. As a result, the linear model predicts a surge in inflation that never materialized. However, since the

Staiger, Stock and Watson 2001). We have also derived an estimate of the time-varying natural rate of unemployment from a linear Phillips curve specification such as (1) that allows for a time-varying intercept (see Gordon 1998).

negative gap is never below the estimated g_L for extended periods of time, the threshold model performs reasonably well.

IV. Concluding Remarks

Overall, we find that a piecewise linear specification of the Phillips curve provides a good characterization of inflation dynamics over the past 40 years. The estimated relationship implies a range of values for the unemployment rate, or the unemployment gap, over which we find no significant tradeoff between inflation and unemployment. Outside of this range, a significant tradeoff obtains.

How should we interpret our findings in light of the most recent inflation and unemployment experience? As is well known, core inflation measures have dropped significantly over the period from the fourth quarter of 2002 through the second quarter of 2003, with the unemployment rate rising to 6.4 percent by the end of the first half of 2003. Unless the decline in inflation is temporary, it is difficult to reconcile this recent experience with the estimates obtained from a piecewise linear Phillips relationship using a constant natural rate of unemployment. In such a specification, the estimated tradeoff between inflation and unemployment is insignificant when the unemployment rate ranges from about 4.0 percent to 7.5 percent.

It is likely, though, that while this range is well suited to characterize the path of inflation over the years 1960 to 1989, it is not appropriate for the most recent period. The experience of the second half of the 1990s appears, in fact, to be consistent with a value of the natural rate of unemployment that is lower than the average natural rate of unemployment of the previous three decades. For this reason, the estimates that arise

from an "unemployment gap" specification of the piecewise linear Phillips relationship may have more bearing on the current situation. If the natural rate of unemployment is now near 5 percent, then, according to our upper threshold estimate of 1.4 percent, we should expect inflation to start dropping when the unemployment rate is close to 6.5 percent. This is broadly consistent with recent experience, in that inflation has been relatively stable over most of the period from 2000 to 2002, when the unemployment gap was small. By contrast, it started to decline noticeably only when the unemployment gap became relatively large.

From a monetary policy standpoint, the finding of a range of values for the unemployment rate or the unemployment gap over which there is no significant tradeoff between inflation and unemployment means that there is scope for targeting a level of the unemployment rate at which inflation is stable. Clearly, the lower the targeted level of the unemployment rate associated with no inflationary pressures, the higher society's welfare.¹⁸ This is an argument for monetary policy to try cautiously to drive the unemployment rate lower until the lower limit of the stable inflation range is reached. Such a strategy, however, is complicated by the fact that the lower level of the unemployment rate or unemployment gap associated with no inflationary pressures is estimated with uncertainty. As a result,

¹⁸ The Federal Reserve Act specifies that the Federal Reserve System and the Federal Open Market Committee should seek "to promote effectively the goals of maximum employment, stable prices, and moderate long-term interest rates."

policymakers would have to determine when to stop by being alert to incipient signs of accelerating inflation.

Thus, it remains an open question whether a trigger strategy in which monetary policy aims to drive the unemployment rate or unemployment gap to the lower threshold and then responds vigorously would deliver a satisfactory stabilization performance. The answer depends on one's judgment concerning policymakers' ability to note the emergence of incipient inflation-acceleration indicators, their ability to respond rapidly with appropriate measures, and the lead-time for such measures to take effect.

* *Economist and Senior Economist, respectively, Federal Reserve Bank of Boston. We thank several colleagues at the Boston Fed for helpful comments and suggestions.*

References

Atkeson, Andrew, and Ohanian, Lee E. 2001. "Are Phillips Curves Useful for Forecasting Inflation?" *Federal Reserve Bank of Minneapolis Quarterly Review* 25 (1) Winter: 2-11.

Blanchard, Olivier Jean, and Fischer, Stanley. 1989. *Lectures on Macroeconomics*. Cambridge, MA: MIT Press.

Blinder, Alan S. 1998. *Central Banking in Theory and Practice*. Cambridge, MA: MIT Press.

Eisner, Robert. 1997. "A New View of the NAIRU." In *Improving the Global Economy; Keynesianism and the Growth in Output and Employment*, edited by Paul Davidson and Jan A. Kregel. Lynne, NH: Edward Elgar.

Fuhrer, Jeffrey C. 1995. "The Phillips Curve Is Alive and Well." *New England Economic Review* March/April: 41-56.

Gordon, Robert J. 1982. "Price Inertia and Ineffectiveness in the United States, 1840-1980." *Journal of Political Economy* 90(6) December: 1087-1117.

Gordon, Robert J. 1998 "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU." *Brookings Papers on Economic Activity* 0 (2): 297-333.

Hall, Robert E. 2003. "Wage Determination and Employment Fluctuations." NBER Working Paper No. w9967, September.

Hansen, Bruce E. 1996. "Inference When A Nuisance Parameter Is Not Identified Under The Null Hypothesis." *Econometrica* 64(2): 413-430.

Hansen, Bruce E. 2000, "Sample Splitting and Threshold Estimation." *Econometrica* 68(3): 575-603.

Phillips, A. W. 1958. "The Relation between Unemployment and the Rate of Change of the Money Wage Rates in the United Kingdom." *Economica* 25 November: 283-299.

Staiger, Douglas, Stock, James H. and Watson, Mark W. 2001. "Prices, Wages and the U.S. NAIRU in the 1990s." In *The Roaring Nineties: Can Full Employment Be Sustained?*, edited by Alan B. Krueger and Robert M. Solow, pp. 3-60, New York : Russell Sage Foundation, Century Foundation Press.

Thomas, Jonathan and Worrall, Tim. 1988. "Self-Enforcing Wage Contracts." *Review of Economic Studies* 55(4): 541-553.

Tootell, Geoffrey M. B. 1994. "Restructuring, the NAIRU, and the Phillips Curve." *New England Economic Review* September/October: 31-44.

Woglom, Geoffrey. 1982. "Underemployment Equilibrium with Rational Expectations." *Quarterly Journal of Economics* 97(1): 89-107.

Data Appendix

Price Series (p):

- Core CPI: CPI-U, All Items Less Food and Energy (SA, 1982-84=100), Bureau of Labor Statistics, Quarterly Percentage Change, annualized.
- GDP Deflator: Gross Domestic Product: Chain-type Price Index (SA, 1996=100), Bureau of Economic Analysis, Quarterly Percentage Change, annualized.
- Core PCE Deflator: PCE less Food and Energy: Chain Price Index (SA, 1996=100), Bureau of Economic Analysis, Quarterly Percentage Change, annualized.

Real Activity Variables (u):

- Civilian Unemployment Rate: 16 yr + (SA, %), Bureau of Labor Statistics.
- Non-accelerating Inflation Rate of Unemployment {CBO} (%), Congressional Budget Office.

Other Variables (z):

- Relative Price of Food and Energy: $(CPI_t / \text{Core } CPI_t)$, where CPI is CPI-U: All Items (SA, 1982-84=100), Bureau of Labor Statistics, Quarterly Percentage Change.
- U.S.: Nominal Effective Exchange Rate (1995=100), International Financial Statistics (IMF), Quarterly Percentage Change.
- Gordon (1982) Wage and Price Control Variable: equal to 0.8 from 1971:Q3 through 1972:Q3, -0.4 for 1974:Q2, -1.6 from 1974:Q3 through 1974:Q4, -0.4 for 1975:Q1, and 0 for all other dates.

Box 1

A general and compact way of representing the piecewise linear or threshold model in the text is:

$$y_t = \mathbf{b} x_t(\mathbf{g}) + \mathbf{e}_t$$

with $\mathbf{b} = (\mathbf{b}_I, \mathbf{b}_O)$, $\mathbf{g} = (\mathbf{g}_L, \mathbf{g}_H)$, and

$$x_t(\mathbf{g}) = \begin{pmatrix} x_t I(\mathbf{g}_L \leq q_t \leq \mathbf{g}_H) \\ x_t (1 - I(\mathbf{g}_L \leq q_t \leq \mathbf{g}_H)) \end{pmatrix},$$

where x_t is the vector of right-hand side variables, q_t is the threshold variable, \mathbf{g}_L is the lower threshold value of the threshold variable, \mathbf{g}_H is the upper threshold value of the threshold variable, and I is the indicator function which takes on a value of one when the expression in parenthesis is true.

Let S_0 represent the sum of squared residuals under the null hypothesis of a linear model, and $S_1(\mathbf{g})$ the sum of squared residuals of the piecewise linear model as a function of \mathbf{g} . In this case an F -test of the null hypothesis of a linear Phillips curve is based on:

$$F_1 = n \frac{S_0 - S_1(\hat{\mathbf{g}})}{S_1(\hat{\mathbf{g}})},$$

where $S_1(\hat{\mathbf{g}})$ is the minimum sum of squared residuals for the piecewise linear model obtained by searching over a grid of possible upper and lower threshold values. Hansen (1996) shows that the asymptotic statistical distribution for this test statistic is nonstandard and strictly dominates the χ^2 distribution. This nonstandard distribution can be approximated to the first order by a bootstrap procedure, and p -values constructed from the bootstrap for the test of the linear null

against the threshold alternative hypothesis are asymptotically valid.

In order to test the hypothesis $H_0:\mathbf{g}=\mathbf{g}_0$, the likelihood ratio test is to reject for large values of $LR_1(\mathbf{g}_0)$, where:

$$LR_1(\mathbf{g}) = n \frac{S_1(\mathbf{g}) - S_1(\hat{\mathbf{g}})}{S_1(\hat{\mathbf{g}})}.$$

In addition, the asymptotic distribution of $LR_1(\mathbf{g}_0)$ can be used to form valid asymptotic confidence intervals about the estimated threshold values. These confidence intervals are the set of values of \mathbf{g} such that the likelihood ratio lies below the critical values tabulated by Hansen (2000) for the desired level of confidence.

Finally, the estimator $\mathbf{b} = \mathbf{b}(\mathbf{g})$ depends on the threshold estimate. Since the dependence on the threshold estimate is not of first-order asymptotic importance, inference on \mathbf{b} can proceed as if the threshold estimate were the true value.

Table 1

Phillips Curve Estimates with Constant Natural Rate of Unemployment: 1961 to 2002

$$\pi_t = \underline{\mu} + \sum_{j=1}^s \alpha_j \pi_{t-j} + \underline{\beta} u_t + \sum_{j=0}^k \underline{\delta}_j \Delta u_{t-j} + \sum_{j=0}^l \xi_j Z_{t-j} + \varepsilon_t$$

Inflation Measure	n	Linear Model	Threshold Model				p-value of F-test
		$\hat{\beta}$	$\hat{\beta}_I$	$\hat{\beta}_O$	$\hat{\gamma}_L$	$\hat{\gamma}_H$	
Core PCE Deflator	166	-.1734 ** (.0484)	.1008 (.0792)	-.2768 ** (.0607)	3.95 [3.90 , 5.30]	7.40 [5.80 , 7.60]	.016
Core CPI	166	-.2410 ** (.0745)	.0926 (.1053)	-.2868 ** (.0974)	3.95 [3.90 , 4.80]	7.45 [7.30 , 7.60]	.001
GDP Deflator	166	-.2569 ** (.0587)	-.0372 (.0953)	-.3261 ** (.0739)	3.95 [3.90 , 5.50]	7.35 [6.00 , 7.60]	.048
GDP Deflator excluding outliers	164	-.2419 ** (.0573)	-.0135 (.0823)	-.2990 ** (.0783)	3.95 [3.90 , 5.50]	7.45 [7.30 , 7.60]	.001

Note: Estimation period is 1961:Q3 to 2002:Q4. For the core PCE deflator and the Core CPI index, the estimated equations include seven lags of inflation, the contemporaneous unemployment rate, the contemporaneous and two lags of the first difference of the unemployment rate, two lags of the change in the relative price of food and energy, and the Gordon variable. For the GDP deflator, the estimated equations include eight lags of inflation, one lag of the unemployment rate, four lags of the first difference of the unemployment rate, one lag of the change in the relative price of food and energy, two lags of the change in the exchange rate, and the Gordon variable. For the core PCE deflator and the Core CPI index, the threshold variable is the contemporaneous unemployment rate, while for the GDP deflator it is the first lag of the unemployment rate.

** indicates significance at the 5 percent level.

Table 2

Phillips Curve Estimates with Constant Natural Rate of Unemployment: 1961 to 1986

$$\pi_t = \underline{\mu} + \sum_{j=1}^s \alpha_j \pi_{t-j} + \underline{\beta} u_t + \sum_{j=0}^k \underline{\delta}_j \Delta u_{t-j} + \sum_{j=0}^l \underline{\xi}_j Z_{t-j} + \varepsilon_t$$

Inflation Measure	n	Linear Model	Threshold Model				p-value of F-test
		$\hat{\beta}$	$\hat{\beta}_I$	$\hat{\beta}_O$	$\hat{\gamma}_L$	$\hat{\gamma}_H$	
Core PCE Deflator	102	-.2003 ** (.0589)	.2687 (.1434)	-.2782 ** (.0664)	3.95 [3.90 , 5.50]	7.25 [5.80 , 7.60]	.205
Core CPI	102	-.3037 ** (.1045)	.1513 (.1888)	-.3262 ** (.1223)	3.95 [3.90 , 4.80]	7.45 [7.00 , 7.60]	.004
GDP Deflator excluding outliers	100	-.2808 ** (.0782)	-.0737 (.1354)	-.3092 ** (.0957)	3.95 [3.90 , 5.50]	7.45 [7.40 , 7.60]	.048

Note: Estimation period is 1961:Q3 to 1986:Q4. For the core PCE deflator and the Core CPI index, the estimated equations include seven lags of inflation, the contemporaneous unemployment rate, the contemporaneous and two lags of the first difference of the unemployment rate, two lags of the change in the relative price of food and energy, and the Gordon variable. For the GDP deflator, the estimated equation includes eight lags of inflation, one lag of the unemployment rate, four lags of the first difference of the unemployment rate, one lag of the change in the relative price of food and energy, two lags of the change in the exchange rate, and the Gordon variable. For the core PCE deflator and the Core CPI index, the threshold variable is the contemporaneous unemployment rate, while for the GDP deflator it is the first lag of the unemployment rate.

** indicates significance at the 5 percent level.

Table 3

Phillips Curve Estimates with a Time-Varying Natural Rate of Unemployment: 1961 to 2002

$$\pi_t = \underline{\mu} + \sum_{j=1}^s \alpha_j \pi_{t-j} + \underline{\beta}(u_t - u_t^N) + \sum_{j=0}^k \underline{\delta}_j \Delta u_{t-j} + \sum_{j=0}^l \xi_j Z_{t-j} + \varepsilon_t$$

Inflation Measure	<i>n</i>	Linear Model	Threshold Model				<i>p</i> -value of <i>F</i> -test
		$\hat{\beta}$	$\hat{\beta}_I$	$\hat{\beta}_O$	$\hat{\gamma}_L$	$\hat{\gamma}_H$	
Core PCE Deflator	166	-.2328 ** (.0551)	-.0640 (.0933)	-.2954 ** (.0612)	-1.40 [-1.60 , -1.10]	1.30 [1.20 , 1.60]	0.06
Core CPI	166	-.3277 ** (.0849)	-.0138 (.1366)	-.2823 ** (.0978)	-1.34 [-1.40 , -1.30]	1.40 [1.30 , 1.45]	0.001
GDP Deflator excluding outliers	164	-.3193 ** (.0626)	-.1314 (.1034)	-.3080 ** (.0719)	-1.04 [-1.60 , -0.80]	1.56 [1.20 , 1.60]	0.001

Note: Estimation period is 1961:Q3 to 2002:Q4. For the core PCE deflator and the Core CPI index, the estimated equations include seven lags of inflation, the contemporaneous unemployment gap, the contemporaneous and two lags of the first difference of the unemployment rate, two lags of the change in the relative price of food and energy, and the Gordon variable. For the GDP deflator, the estimated equation includes eight lags of inflation, one lag of the unemployment gap, four lags of the first difference of the unemployment rate, one lag of the change in the relative price of food and energy, two lags of the change in the exchange rate, and the Gordon variable. For the core PCE deflator and the Core CPI index, the threshold variable is the contemporaneous unemployment gap, while for the GDP deflator it is the first lag of the unemployment gap.

** indicates significance at the 5 percent level.

Figure 1

Chart A
Inflation and Unemployment: 1960 to 1969

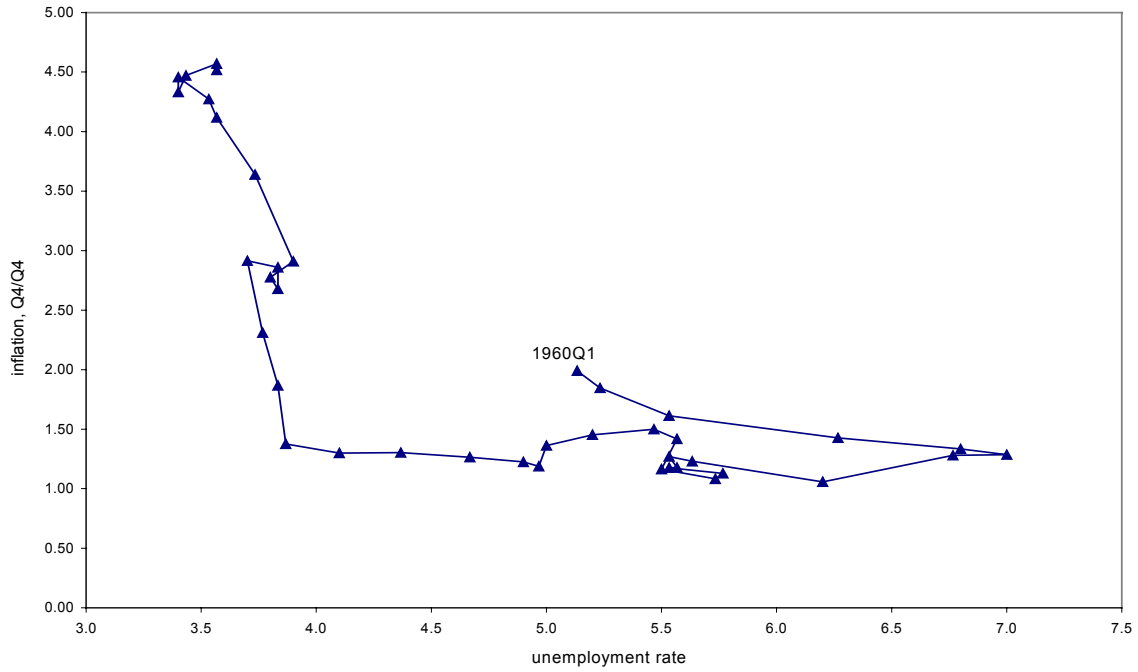


Chart B
Inflation and Unemployment: 1970 to 1979

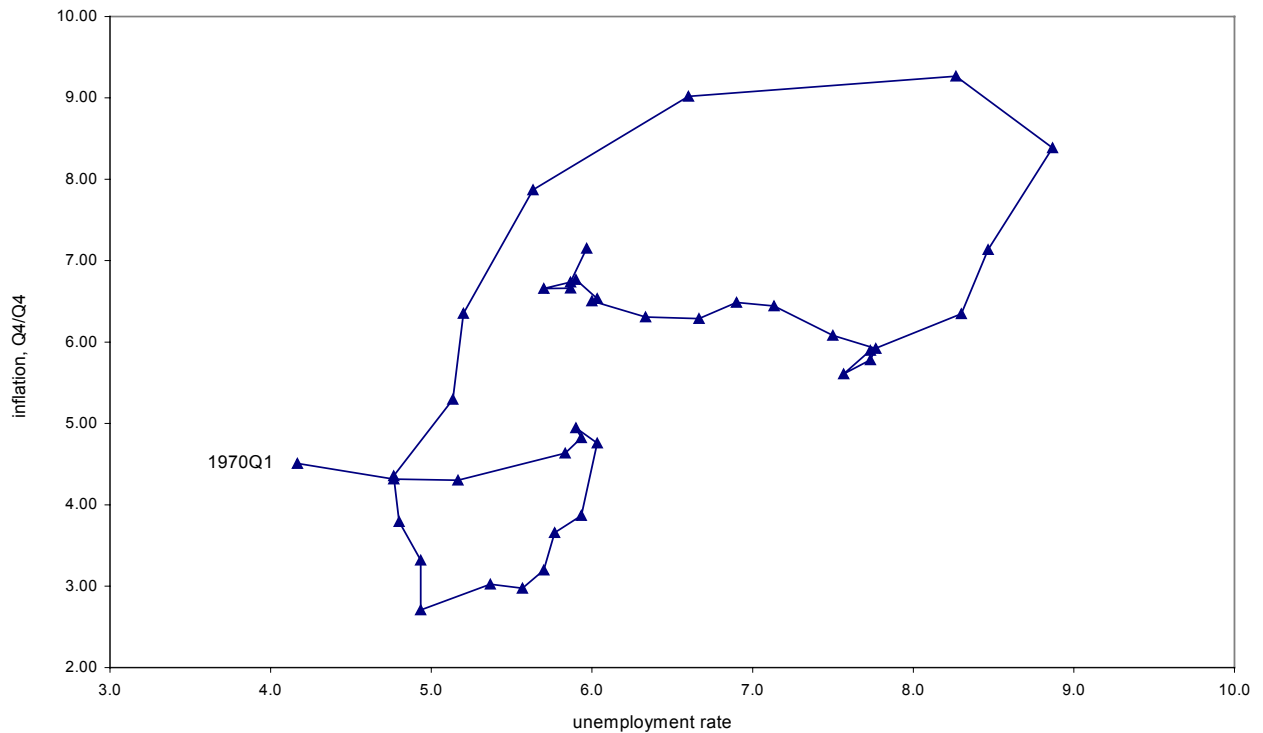


Figure 1 (continued)

Chart C Inflation and Unemployment: 1980 to 1989

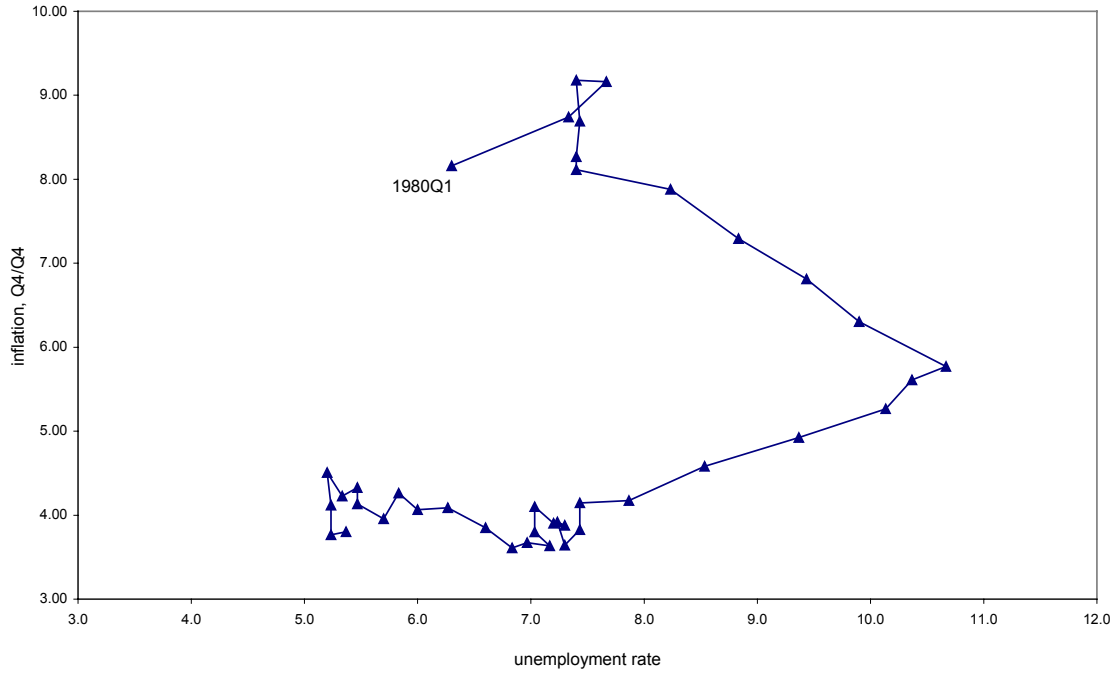


Chart D Inflation and Unemployment: 1990 to 2002

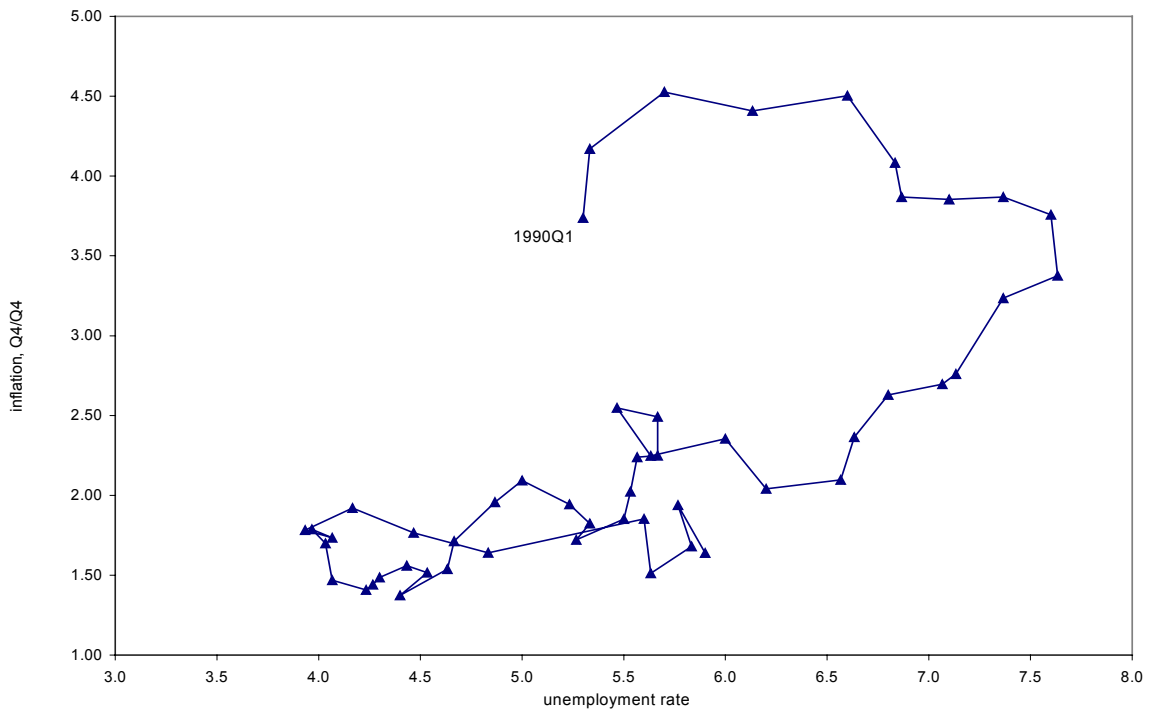


Figure 2

Partial Correlation Between Inflation and the Unemployment Rate when the Unemployment Rate is between 4 and 7.5 percent: 1960 to 2002

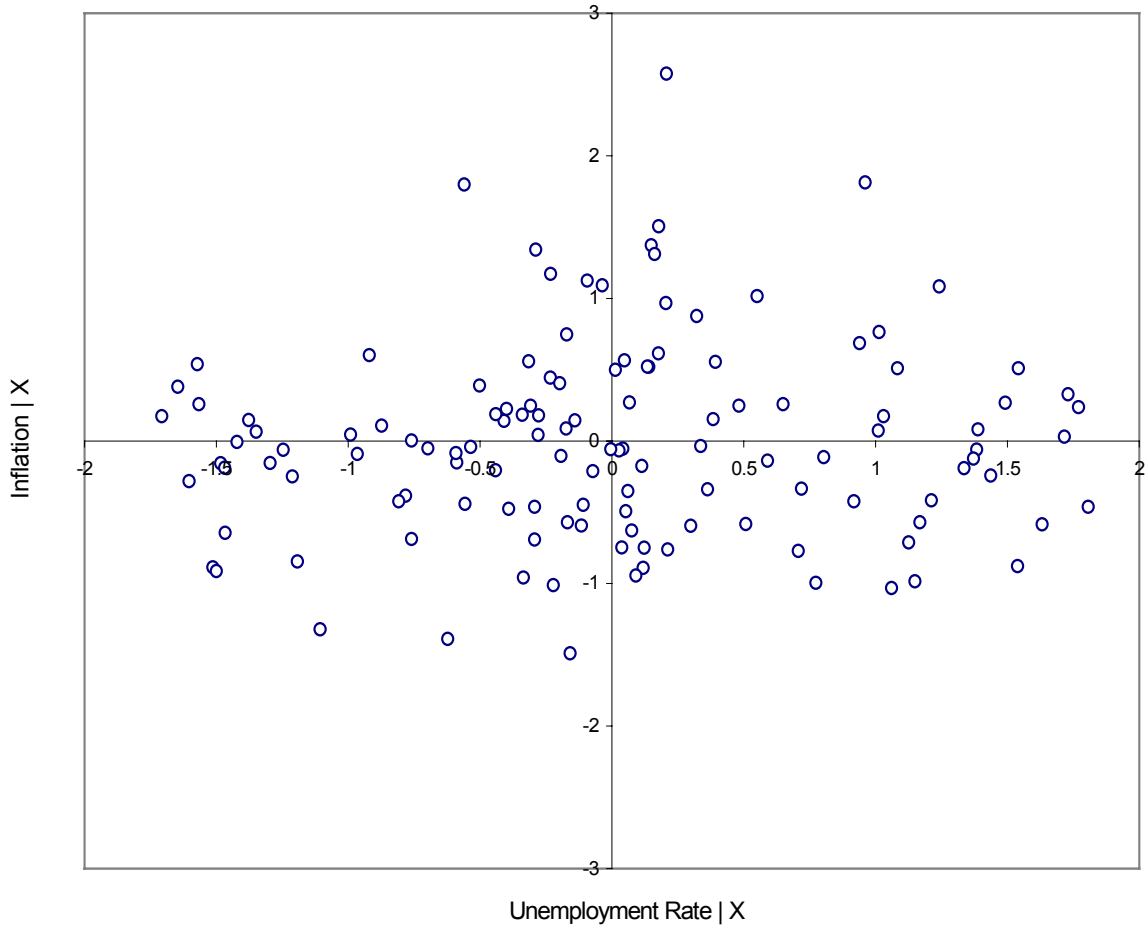


Figure 3
Piecewise Linear Phillips Curve

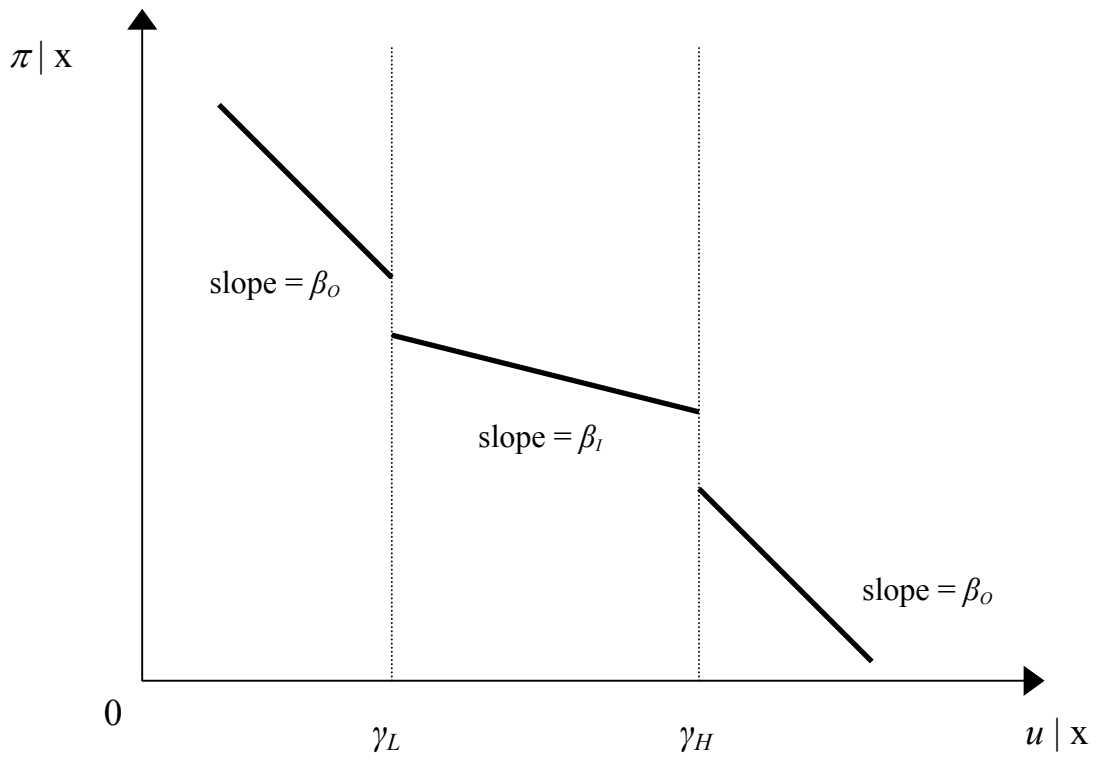


Figure 4
Dynamic Simulation of Core CPI Inflation:
1961 to 1986

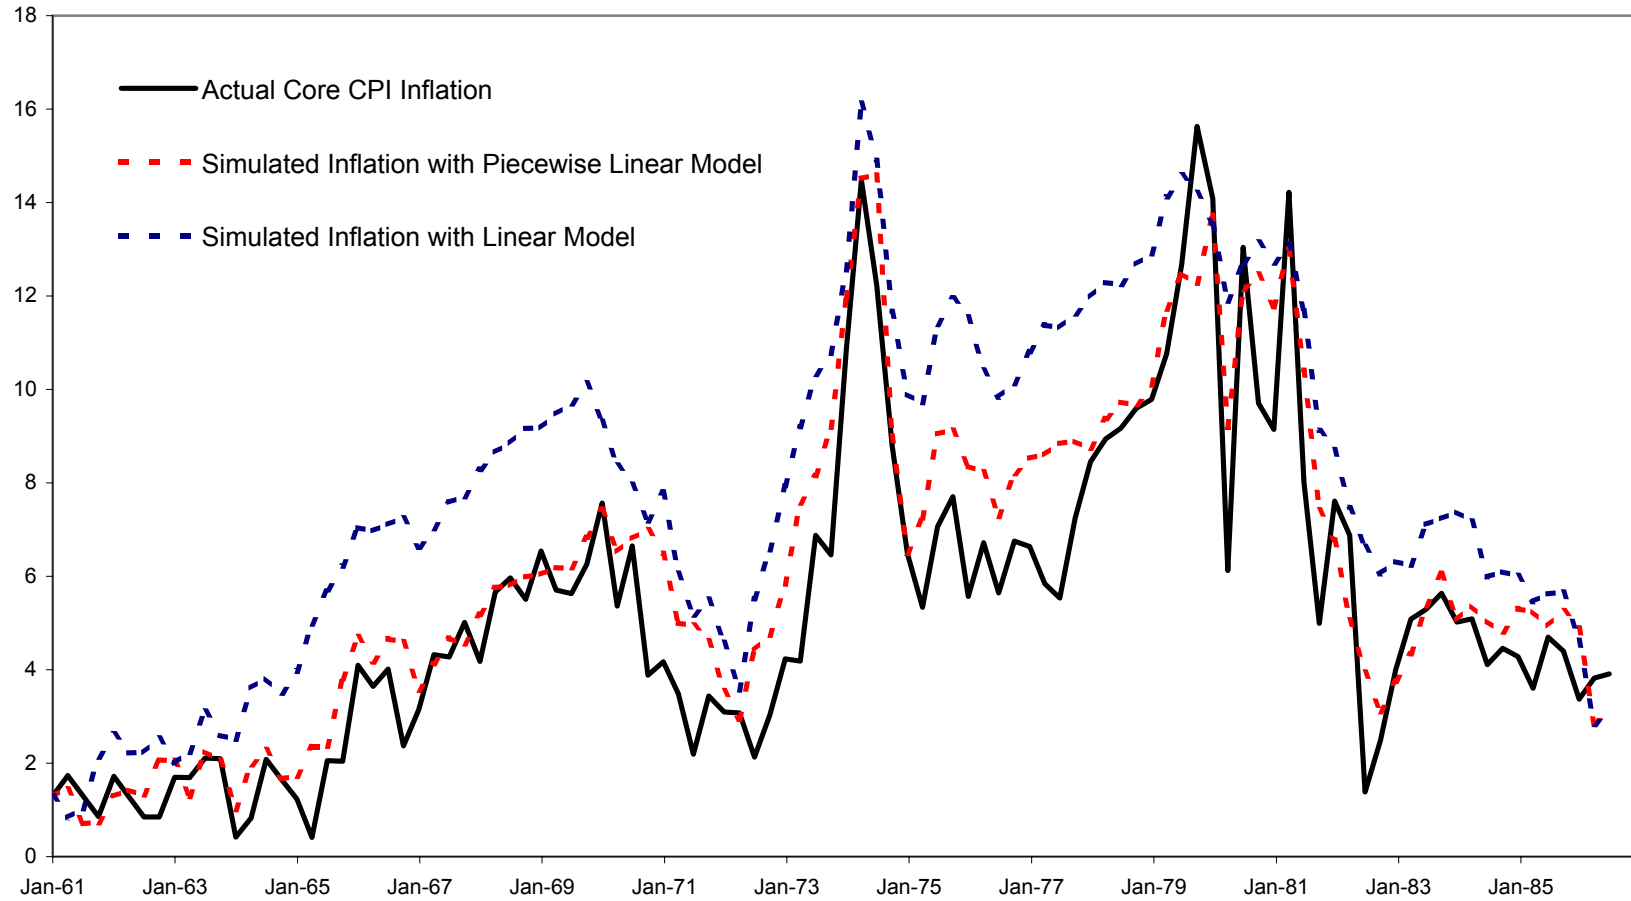


Figure 5
Actual Unemployment Rate, CBO Natural Rate of Unemployment, and Unemployment Gap:
1959 to 2002

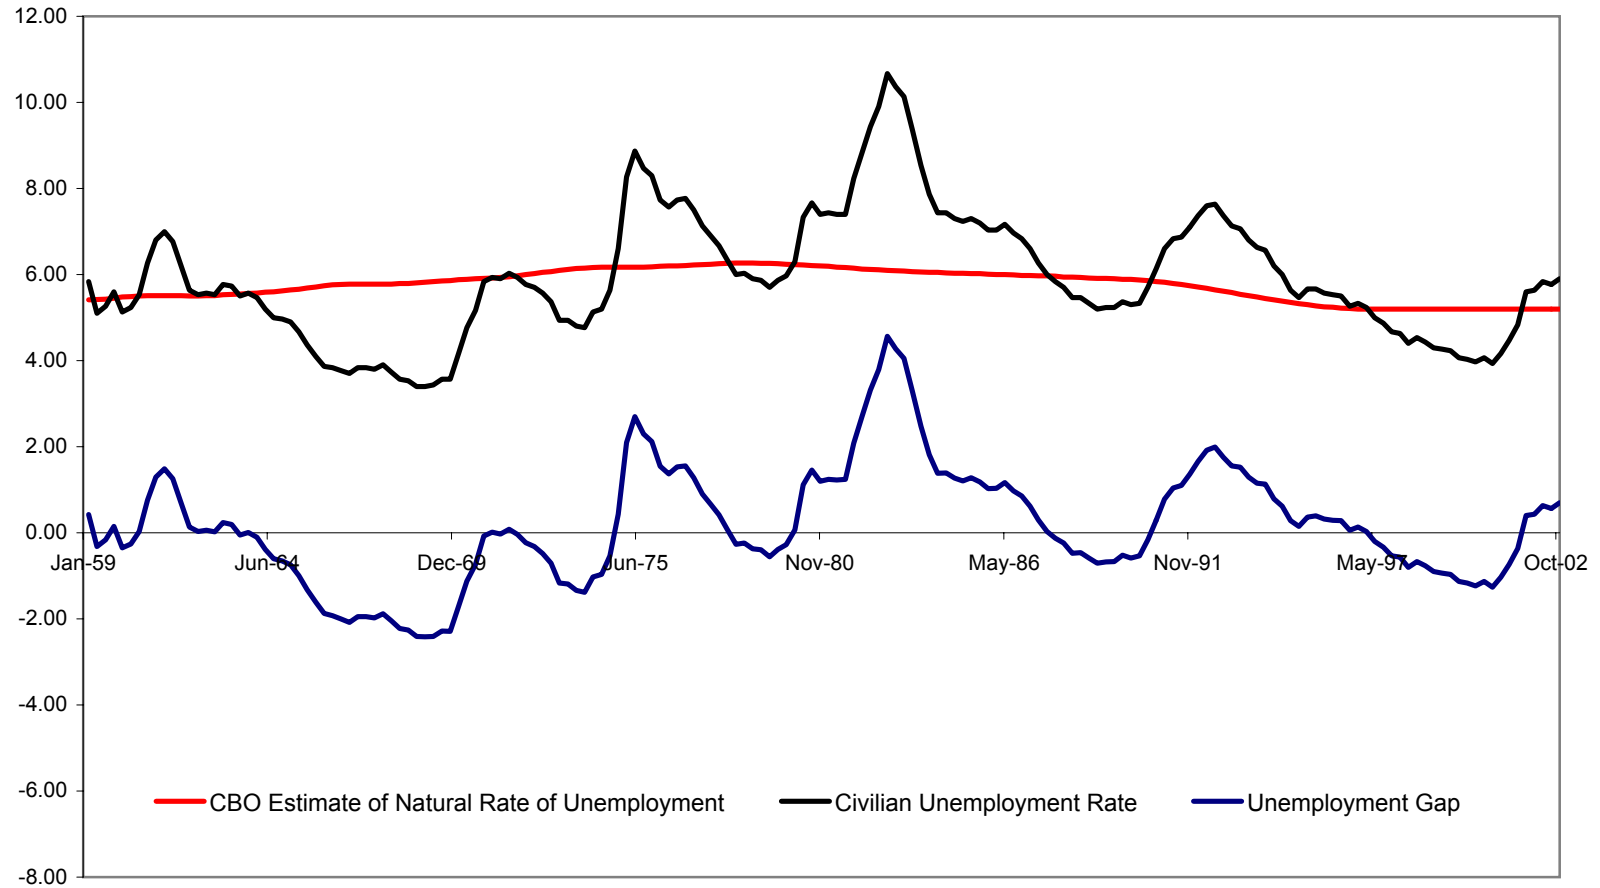


Figure 6
Dynamic Simulation of Core CPI Inflation:
1987 to 2002

