

Are Phillips Curves Useful for Forecasting Inflation?

Andrew Atkeson
Lee E. Ohanian
Visiting Scholars
Research Department
Federal Reserve Bank of Minneapolis
and Professors
Department of Economics
University of California, Los Angeles

Abstract

This study evaluates the conventional wisdom that modern Phillips curve-based models are useful tools for forecasting inflation. These models are based on the non-accelerating inflation rate of unemployment (the NAIRU). The study compares the accuracy, over the last 15 years, of three sets of inflation forecasts from NAIRU models to the naive forecast that at any date inflation will be the same over the next year as it has been over the last year. The conventional wisdom is wrong; none of the NAIRU forecasts is more accurate than the naive forecast. The likelihood of accurately predicting a change in the inflation rate from these three forecasts is no better than the likelihood of accurately predicting a change based on a coin flip. The forecasts include those from a textbook NAIRU model, those from two models similar to Stock and Watson's, and those produced by the Federal Reserve Board.

The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

A *Phillips curve* is an equation that relates the unemployment rate, or some other measure of aggregate economic activity, to a measure of the inflation rate. Modern specifications of Phillips curve equations relate the current rate of unemployment to future changes in the rate of inflation. These specifications are based on the idea that there is a baseline rate of unemployment at which inflation tends to remain constant. The idea is that when unemployment is below this baseline rate, inflation tends to rise over time, and when unemployment is above this rate, inflation tends to fall. The baseline unemployment rate is known as the *non-accelerating inflation rate of unemployment* (the *NAIRU*), and modern specifications based on it are known as *NAIRU Phillips curves*.

NAIRU Phillips curves are widely used to produce inflation forecasts, both in the academic literature on inflation forecasting and in policymaking institutions.¹ This wide use is based on the view that inflation forecasts made with these equations are more accurate than forecasts made with other methods. For example, Blinder (1997, p. 241), the former Vice Chairman at the Board of Governors of the Federal Reserve System, argues that “the empirical Phillips curve has worked amazingly well for decades” and concludes, on the basis of this empirical success, that a Phillips curve should have “a prominent place in the core model” used for macroeconomic policymaking purposes.

This study critically evaluates the conventional wisdom that NAIRU Phillips curve–based models are useful tools for forecasting inflation. We examine the accuracy of three sets of NAIRU Phillips curve–based inflation forecasts. One set of inflation forecasts is obtained from a simple textbook model of the NAIRU Phillips curve. This textbook model is presented by Stock and Watson (1999b) and others as evidence that the historical data contain a stable negative relationship between the current rate of unemployment and subsequent changes in the rate of inflation which might be exploited to forecast inflation.

Another set of inflation forecasts comes from two NAIRU Phillips curve–based inflation forecasting models similar to those proposed by Stock and Watson (1999a). Their work is a comprehensive study of the accuracy of inflation forecasts from NAIRU Phillips curve–based models and has attracted a great deal of attention, both in the academic literature and in the Federal Reserve System. (See Mankiw 2000 and J. Fisher 2000.) These NAIRU Phillips curve–based models represent the state of the art in the academic inflation forecasting literature.

A third set of forecasts is those produced by the research staff at the Federal Reserve Board of Governors and reported in the Greenbook, the internal collection of materials prepared routinely for meetings of the Federal Open Market Committee. The staff at the Board of Governors uses a large econometric model to help produce the Greenbook forecast. A NAIRU Phillips curve plays a significant role in this model.² In particular, the most recent version of the model predicts that, “all else being equal, if the unemployment rate is held 1 percentage point below its equilibrium level on a sustained basis, inflation should climb steadily about 0.4 percentage point a year” (Reifschneider, Tetlow, and Williams 1999, p. 7). The Greenbook forecasts are a key ingredient in the monetary policy debate at the Federal Reserve.

To evaluate the usefulness of Phillips curves for forecasting inflation, we compare the accuracy of these three sets of inflation forecasts at a one-year forecast horizon to that of a naive model that makes a simple prediction: at any date, the inflation rate over the coming year is expected to be the same as the inflation rate over the past year.³ We establish this naive forecast as our benchmark not because we think that it is the best forecast of inflation available, but rather because we think that any inflation forecasting model based on some hypothesized economic relationship cannot be considered a useful guide for policy if its forecasts are no more accurate than such a simple atheoretical forecast.

Our result contrasts sharply with the conventional wisdom. We find that over the last 15 years, all three sets of NAIRU Phillips curve–based inflation forecasts have been no more accurate than the forecast from our naive model, that inflation over the next year will be equal to inflation over the previous year. We conclude that NAIRU Phillips curves are not useful for forecasting inflation.

A Short History of Phillips Curves

Useful forecasting models exploit stable relationships among variables. Forecasting models that are not based on such stable relationships are not useful because they lead to inaccurate forecasts when the relationships among the variables in the forecasting model change. Do Phillips curves capture stable relationships between unemployment or other measures of economic activity and future inflation? Our review of the evidence indicates that they do not.

Early Specifications

Unemployment has been suggested as an indicator of future inflation on the basis of early empirical work documenting a statistical relationship between these variables. I. Fisher (1926) was the first to document such a relationship using data from the United States. Later studies by Phillips (1958) and Samuelson and Solow (1960) attracted great attention. These studies all document a negative relationship between the unemployment rate (unemployment as a percentage of the labor force) and either the rate of nominal wage growth or the rate of inflation. Equations relating the unemployment rate to the inflation rate were the first called *Phillips curves*.

These empirical studies initiated a long debate on the usefulness of Phillips curves for forecasting inflation. Much of this debate has centered on the question of whether the statistical relationship between unemployment and inflation documented in these early empirical studies should be expected to remain stable over time. As argued by Friedman (1968), Phelps (1969), Lucas (1972), Fischer (1977), and Taylor (1980), among others, economic theory does not predict a stable and systematic relationship between current unemployment and future inflation. Instead, theory predicts that observed relationships between these variables should change with changes in agents’ expectations of inflation. Since theory predicts that agents’ expectations of inflation should vary as the economic environment changes, theory predicts that any relationship between current unemployment and future inflation observed in historical data should be expected to change as the economic environment changes. Thus, there is no theoretical presumption that a statistical relationship observed in one economic environment would be stable enough to be useful for forecasting inflation when that economic environment changes.

The theoretical prediction that historical Phillips curves should change as the economic environment changes is borne out in the data. Charts 1 and 2 illustrate the empirical breakdown of Samuelson and Solow's (1960) specification of the Phillips curve, relating the rate of unemployment to the rate of inflation. In Chart 1, the horizontal axis shows the unemployment rate for each quarter from the first quarter of 1959 through the fourth quarter of 1969, while the vertical axis shows the subsequent inflation rate, as measured by the percentage change in the implicit price deflator for the gross domestic product (the GDP deflator) over the next four quarters. The chart also shows a linear regression line through these data. This line can be interpreted as a forecast of the inflation rate one year ahead given the current level of the unemployment rate. The line is clearly downward-sloping, which represents a definite negative relationship between the two variables during the 1960s.⁴

After 1970, however, many aspects of the economic environment changed. For example, inflation was both higher and more volatile in the 1970s than it had been in the 1960s. As the economic environment changed, the negative relationship between unemployment and future inflation observed in data from the 1960s, as illustrated in Chart 1, disappeared. Chart 2 documents the disappearance of this negative relationship after the 1960s. This chart displays quarterly data on the unemployment rate for the first quarter of 1970 through the first quarter of 1999 and the inflation rate over the next four quarters. The chart also shows two regression lines: the original regression line from Chart 1, computed from the 1960s data, and a second regression line through the 1970–99 data. In contrast to the downward-sloping regression line from the 1960s, the regression line from the more recent data shows virtually no relationship between unemployment and subsequent inflation. Moreover, any inflation forecasts for post-1970 data based on the 1960s regression line clearly would be inaccurate. Lucas and Sargent (1979, p. 6) argue that the breakdown of this simple Phillips curve relationship, as well as that of the more sophisticated econometric models based on it, represents an “econometric failure on a grand scale.” Thus, both theory and data seem to be telling economists not to use Phillips curves to forecast inflation.

The NAIRU Specification

Yet some still do. Economists have persisted in arguing that there is an empirical relationship of some kind between unemployment and future inflation that can be used to forecast inflation. These economists have focused on versions of the NAIRU Phillips curve, which differs from the early specification presented in Charts 1 and 2. In a NAIRU Phillips curve, unemployment or some other measure of economic activity is used to forecast future changes in the inflation rate rather than the inflation rate itself.⁵

Chart 3 illustrates a textbook specification of a NAIRU Phillips curve.⁶ In this chart we show on the horizontal axis quarterly data for the unemployment rate from the first quarter of 1960 through the fourth quarter of 1983 and on the vertical axis the change in the inflation rate (as measured by the GDP deflator) over the subsequent four quarters relative to the inflation rate over the previous four quarters. The line in the chart is the regression line through these data. This regression line can be interpreted as a forecast of the change in inflation over the next four quarters

relative to inflation over the previous four quarters given the current unemployment rate. The regression line shows a negative relationship between unemployment and subsequent changes in inflation. Specifically, the line shows that, during this time period, when the unemployment rate was low, there was a tendency for the inflation rate to rise, and when the unemployment rate was high, there was a tendency for the inflation rate to fall. The regression line identifies a NAIRU of about 6 percent: This, again, is the rate of unemployment at or near which, according to this regression, the inflation rate has no tendency to either rise or fall.

Note that the inflation forecast produced by this textbook NAIRU Phillips curve is quite similar to that produced by the large econometric model used by the staff at the Federal Reserve Board. Recall that this model predicts that if unemployment is one percentage point below the NAIRU “on a sustained basis,” then inflation is forecast to rise about 0.4 of a percentage point per year (Reifschneider, Tetlow, and Williams 1999, p. 7). In Chart 3, we see that when unemployment is one percentage point below the NAIRU, at 5 percent, inflation is forecast to rise at 0.6 of a percentage point over the next year.

Of course, there is no theoretical presumption that this NAIRU Phillips curve should be any less susceptible to instability with changes in the economic environment than was the early Phillips curve. In fact, there are good reasons to expect the NAIRU Phillips curve to be unstable since many aspects of the U.S. economy have changed since the 1980s: the business cycle, monetary policy,⁷ and inflation have all been less volatile since 1984 than they were in the previous 15 years.

Did these changes in the economic environment affect the NAIRU Phillips curve observed in the data? We address this question by extending the plot of the textbook NAIRU Phillips curve past 1983. Chart 4 illustrates the same textbook specification of the NAIRU Phillips curve shown in Chart 3, except that it uses data on unemployment and changes in inflation starting in 1984. The regression line through the 1984–99 data is shown and, for comparison, so is the regression line through the 1960–83 data which we saw in Chart 3.

Chart 4 shows that the relationship between unemployment and future changes in inflation has shifted. In particular, the regression line through the 1984–99 data is much flatter than the regression line through the 1960–83 data. According to that earlier regression line, the current U.S. unemployment rate of about 4 percent is associated with about a one percentage point increase in inflation from one year to the next; according to the 1984–99 regression line, 4 percent unemployment is associated with an increase in inflation of only about one-quarter of one percentage point.⁸

While the breakdown of the NAIRU Phillips curve shown in Chart 4 is not as severe as the breakdown of the early version of the Phillips curve shown in Chart 2, this instability of a textbook NAIRU Phillips curve raises an obvious question: Are NAIRU Phillips curves stable enough to produce accurate inflation forecasts in the economic environment of low and stable inflation that the United States has experienced since the early 1980s?

Simple Tests of Forecasting Accuracy

As we now show, they are not.

We test the accuracy of these Phillips curve–based forecasts in two ways. To assess the usefulness of the textbook NAIRU Phillips curve model and two new ones developed by Stock and Watson (1999a), we consider what are called *simulated forecasting exercises*. In this sort of exercise, a simulated series is constructed of the forecasts of inflation that a model would have produced had it been used historically to generate forecasts of inflation. To assess the accuracy of the inflation forecasts produced by the staff at the Federal Reserve Board, we need not simulate forecasts. Instead, we consider the historical record of actual inflation forecasts reported in the various Greenbooks prepared for the regular meetings of the Federal Open Market Committee.

A Textbook NAIRU Model

In our first simulated forecasting exercise, we construct the simulated record of inflation forecasts produced by our naive model and by a textbook NAIRU model starting with the first quarter of 1984 and ending with the third quarter of 1999. For this exercise, we use the GDP deflator as the measure of inflation.

Our naive model predicts that inflation over the next four quarters is expected to be equal to inflation over the previous four quarters:

$$(1) \quad E_t(\pi_{t+4} - \pi_t) = 0.$$

Here π_t is the percentage change in the inflation rate between quarters $t - 4$ and t .

The forecasts from the textbook NAIRU model specify that the expected change in the inflation rate over the next four quarters is proportional to the unemployment rate, u_t , minus the NAIRU, \bar{u} :

$$(2) \quad E_t(\pi_{t+4} - \pi_t) = \beta(u_t - \bar{u}).$$

Here u_t is the unemployment rate in quarter t , \bar{u} is the model's NAIRU (where the change in inflation will be zero), and β is the slope of the Phillips curve. To construct the forecast for each quarter from this textbook NAIRU Phillips curve, we estimate the parameters β and \bar{u} with ordinary least squares, using the data for the unemployment rate and changes in the inflation rate from the first quarter of 1970 up to the specific forecast quarter. Note that our naive inflation forecast is the same as the forecast from the textbook NAIRU Phillips curve model with the restriction that unemployment is unrelated to future inflation: $\beta = 0$.

We compare the accuracy of the inflation forecasts from this textbook NAIRU Phillips curve model to our naive forecast by comparing the root mean squared error (RMSE) of these two sets of forecasts. The RMSE for any forecast is the square root of the arithmetic average of the squared differences between the actual inflation rate and the predicted inflation rate over the time period for which simulated forecasts are constructed:

$$(3) \quad \text{RMSE} = \left((1/T) \sum_{t=1}^T \{ [\pi_{t+4} - E_t(\pi_{t+4})]^2 \} \right)^{1/2}.$$

We compare the two forecasts by forming the ratio of the RMSE for the NAIRU model to the RMSE for the naive model. A ratio greater than 1 thus indicates that the NAIRU model's forecast is less accurate than the naive

model's. Subtracting 1 from the ratio and multiplying the result by 100 gives the percentage difference in RMSE between the two models.

We find that the forecasts from the textbook NAIRU Phillips curve model are considerably less accurate than those from the naive model. The ratio of the NAIRU RMSE to the naive model RMSE is 1.88. This indicates that the forecast error is 88 percent higher for the NAIRU model than for the naive model. We conclude from this evidence that the textbook NAIRU Phillips curve model has not been a useful inflation forecasting tool over the last 15 years.

Stock and Watson's NAIRU Models

Now we turn to new versions of the NAIRU Phillips curve model developed by Stock and Watson (1999a).

These researchers conduct simulated forecasting exercises to evaluate the performance of a wide array of inflation forecasting models using monthly data on inflation as measured by the implicit price deflator for personal consumption expenditures (the PCE deflator) and the consumer price index (CPI). They focus on the performance of two NAIRU Phillips curve–based models to forecast inflation over a 12-month horizon. One of these models uses the unemployment rate to forecast future changes in the inflation rate. The other uses a broader measure of economic activity to forecast inflation, which Stock and Watson call an *activity index*. Both of these NAIRU Phillips curve models differ from the textbook NAIRU Phillips curve model in that they include some lagged values of the unemployment rate or the activity index and the inflation rate, rather than just the current unemployment rate, to forecast inflation.

In their (1999a) study, Stock and Watson do not compare the forecasts from either of their NAIRU Phillips curve models with the forecast from a naive model that predicts that inflation over the next 12 months will be equal to inflation over the previous 12 months. We do that here.

□ Notation

To address the question of whether NAIRU Phillips curve models have been stable enough to be useful for forecasting inflation in the current U.S. economic environment, we present the results of a simulated forecasting exercise for the 1984–99 time period. We use the notation p_t to denote the level of the price index in month t , π_t to denote monthly inflation as measured by $100[\log(p_t) - \log(p_{t-1})]$, and π_t^{12} to denote inflation over a 12-month period as measured by $100[\log(p_t) - \log(p_{t-12})]$. Our naive forecast of inflation is then given by

$$(4) \quad E_t(\pi_{t+12}^{12} - \pi_t^{12}) = 0.$$

To construct simulated forecasts from Stock and Watson's NAIRU Phillips curve models that can be compared directly to this forecast from the naive model, we consider a slight modification of their forecasting regressions. They construct inflation forecasts using regressions of the form

$$(5) \quad \pi_{t+12}^{12} - \pi_t = \alpha + \beta(L)x_t + \gamma(L)(\pi_t - \pi_{t-1}) + \varepsilon_{t+12}$$

where x_t is a candidate inflation indicator such as the unemployment rate or the activity index and $\beta(L)$ and $\gamma(L)$

are polynomials in the lag operator L that specify the number of lagged values included in the regression. The term on the left side of equation (5) is the difference between inflation over the next 12 months and inflation in the current month. To facilitate the comparison of Stock and Watson's NAIRU forecasts with the naive forecast of inflation, we construct the NAIRU forecasts using regressions of the form

$$(6) \quad \pi_{t+12}^{12} - \pi_t^{12} = \alpha + \beta(L)x_t + \gamma(L)(\pi_t - \pi_{t-1}) + \varepsilon_{t+12}$$

where the left side term is the difference between inflation over the next 12 months and inflation over the previous 12 months. Note that this regression produces the naive forecast when the parameters $\alpha = \beta(L) = \gamma(L) = 0$.

We consider three measures of inflation in our simulated inflation forecasting exercises: the PCE deflator, the CPI for all items, and the CPI for all items except food and energy, which is often referred to as the *core CPI*. We conduct two simulated forecasting exercises. One uses the unemployment rate, and the other uses a version of Stock and Watson's activity index.

□ *With the Unemployment Rate*

First consider the results when we use the unemployment rate as the inflation indicator x_t . We use monthly data from January 1959 through November 2000. For each month t from January 1984 through November 1999, we construct simulated forecasts of inflation over the next 12 months (π_{t+12}^{12}) by estimating the regression (6) using all of the data from January 1959 up through the month t . We consider specifications of $\beta(L)$ running from 1 through 12 lags of x_t and specifications of $\gamma(L)$ running from 1 through 11 lags of $(\pi_t - \pi_{t-1})$. Altogether, we thus consider 132 specifications of this regression. For each specification of the lags in the regression (6), we compute the ratio of the RMSE of the forecast from this regression with the RMSE of our naive model's forecast. Again, values of this ratio that are greater than 1 indicate that the given specification of the NAIRU Phillips curve-based forecast is less accurate than the naive model's forecast.

The accompanying table reports the best and worst results of the exercise with the NAIRU specifications. The table shows that none of the 132 specifications of these unemployment-based NAIRU Phillips curve forecasts has been substantially more accurate than our naive inflation forecast for predicting inflation over the past 15 years. In particular, the RMSE of the best specification of the unemployment-based NAIRU Phillips curve is slightly higher than that of our naive forecast for two measures of inflation and slightly lower for one measure.⁹ We include the maximum as well as the minimum ratio in the table in order to demonstrate how much the RMSEs of these NAIRU Phillips curve forecasts vary across specifications.

□ *With the Activity Index*

Next consider the inflation forecasting results when we replace the unemployment rate with a Stock and Watson-style activity index. The particular activity index we use for x_t is an implementation of the Stock and Watson procedure developed at the Federal Reserve Bank of Chicago. This index is intended to capture the information in 85 monthly indicators of national economic activity.¹⁰ We

perform simulated forecasting exercises using all of the available data.

The table also shows the results of these exercises. The data, again, are the minimum and maximum ratios, across the same 132 specifications considered earlier, of the RMSE of the NAIRU Phillips curve forecast to the RMSE of the naive forecast for our three measures of inflation. As is clear from these data, the RMSEs vary quite a bit across specifications. And even at their best, the activity index-based NAIRU Phillips curve forecasts have not been more accurate than our naive inflation forecast over the past 15 years.¹¹

In sum, we find that since 1984 neither the unemployment-based nor the activity index-based NAIRU Phillips curve inflation forecasts studied by Stock and Watson (1999a) have been substantially more accurate than our naive forecast. This finding indicates that neither of these models would have been useful for forecasting inflation over the past 15 years.

The Federal Reserve's Greenbook

We now examine the accuracy of inflation forecasts produced by the staff at the Board of Governors of the Federal Reserve System. Again, we obtain these forecasts from past issues of the Greenbook. As we did with the other inflation forecasts, we compare the Greenbook forecasts to our naive forecast that inflation over the next year will be equal to inflation over the previous year. The Fed researchers have used two measures of inflation over the period we are examining—the gross national product (GNP) deflator through 1991 and the GDP deflator after that—so we use these here as well. Because the Fed treats Greenbook forecasts as confidential within the Federal Reserve System for five years after they are produced, we can only evaluate forecasts now available to the public. Today that includes forecasts made through 1995.

To construct our naive forecast so that it is comparable to the Greenbook forecasts, we use only the data that would have been available historically at the time that the Greenbook forecasts were made. These historical inflation data are obtained from the Real-Time Data Set compiled by researchers at the Federal Reserve Bank of Philadelphia.¹² This data set is intended to be a record of the major macroeconomic data that were available to the public in the middle of each quarter, starting with the fourth quarter of 1965.

Our naive forecast is constructed as follows. Let $P(t)$ denote the level of the GNP or GDP deflator in quarter t . Then the forecasts that we construct are forecasts of $100[P(t+4)/P(t) - 1]$. Thus, for example, the forecasts that we construct for the fourth quarter of 1988 are forecasts of inflation over the four quarters of 1989. The naive forecast that we use is inflation over the previous four quarters measured in the historically available data as $100[P(t-1)/P(t-5) - 1]$. This choice of timing in our construction of the naive forecast differs from the timing used in the simulated forecasting exercises. The difference arises from the fact that the price level in quarter t is not actually known until the next quarter.

We compile a series of quarterly forecasts of inflation over the subsequent four quarters from back issues of the Greenbook.¹³ Specifically, we select Greenbook forecasts prepared for FOMC meetings that occurred on or after November 13 for the fourth quarter of each year from 1983

through 1995. These forecasts cover inflation over the years 1984 through 1996. (Again, we cannot use forecasts from the Greenbook in more recent years because these forecasts are confidential.) Note that our choice of timing in selecting forecasts implies that the Greenbook forecasts were compiled no more than a few days earlier and often as much as six weeks later than the date at which the historical data used for the naive forecast were published. This timing suggests that the Greenbook forecasts should be more accurate, on average, than the naive forecast, if for no other reason than that more historical data are available when the Greenbook forecasts are made.

We compare both the Greenbook and the naive model inflation forecasts against the data on realized inflation computed using current data on the quarterly GDP deflator. We follow the design of our previous forecasting experiments by comparing the relative RMSE of the Greenbook forecasts to the RMSE of the naive forecast starting in 1984. We find that the RMSEs for the Greenbook and the naive forecasts are basically the same; the ratio of their RMSEs is 1.01. In other words, the Greenbook's forecast has on average been no better than the naive model's. Given the particularly poor performance of NAIRU-based inflation forecasts in recent years (as reported by Gordon 1998 and Brayton, Roberts, and Williams 1999, for example), we strongly suspect that this finding would hold up if data from more recent years were included in our analysis.

We conclude from this historical record that the Phillips curve-based model which helps the staff at the Federal Reserve Board forecast—just like other Phillips curve-based models—has not proved to be useful for forecasting inflation for the past 15 years.

Conclusion

Phillips curves of various kinds have been a major component of many macroeconomic models for the past 40 years. Economists such as Blinder (1997) argue that Phillips curves should continue to play such a role because these curves summarize empirical relationships critical for policymaking. Our review of the evidence indicates that this view is mistaken. We find that for the last 15 years, economists have not produced a version of the Phillips curve that makes more accurate inflation forecasts than those from a naive model that presumes inflation over the next four quarters will be equal to inflation over the last four quarters.

Some might conclude from our review that applied economists should renew their search for a stable empirical relationship between unemployment and inflation that might be used to improve inflation forecasts. We conclude otherwise. Given the weak theoretical and empirical underpinnings of the various incarnations of the Phillips curve, we conclude that the search for yet another Phillips curve-based forecasting model should be abandoned.

Over the last 15 years, inflation in the United States has been hard to predict using any method. Stock and Watson (1999a) have evaluated the performance of a wide array of potential inflation indicators, including money and interest rates. None of the candidate indicators is found to perform particularly well. Cecchetti, Chu, and Steindel (2000) conduct a related simulated forecasting exercise, evaluating the performance of many potential inflation indicators, including the unemployment rate, commodity prices, capacity

utilization, the money supply, and interest rates. These researchers also conclude that none of these indicators is particularly useful.

How should policymakers react to this inability to accurately forecast inflation? They should be skeptical of arguments to change policy based on the claim that someone's favorite inflation indicator, whatever it may be, is currently signalling a big change in inflation in the near term. There is no evidence that any such indicator reliably signals short-term changes in inflation.

*The authors thank Art Rolnick, Timothy Kehoe, David Runkle, and Kathy Rolfe for many helpful comments on this work.

¹Stock and Watson (1999a) have done a significant academic study of the accuracy of inflation forecasts made by NAIRU Phillips curves. For a discussion of the use of NAIRU Phillips curves for inflation forecasting in a policymaking institution, see U.S. President 1996, pp. 45–50.

²Brayton and Tinsley (1996), Brayton et al. (1997), and Reifschneider, Stockton, and Wilcox (1997) all describe the historical evolution of the model used at the Federal Reserve Board for forecasting and policy analysis and the role of Phillips curve equations in that model.

Note, however, that the staff at the Board uses this model as only one of several inputs in constructing forecasts for the Greenbook. Ultimately, the staff forecast is judgmental and cannot be said to be tied to any particular forecasting framework.

³We consider forecasts at horizons of either four quarters or twelve months, depending on the frequency of the data used.

⁴Brayton et al. (1997) discuss how a Phillips curve of this kind was built into the early versions of the model developed by the staff at the Federal Reserve Board.

⁵For a discussion of the intellectual history of the NAIRU, see the paper by Gordon (1997).

⁶We follow Stock and Watson 1999b in showing a scatter plot of the unemployment rate against subsequent changes in the inflation rate as a simple presentation of the NAIRU Phillips curve.

⁷Clarida, Gali, and Gertler (2000) discuss how monetary policy changed significantly in the early 1980s.

⁸This finding that a textbook NAIRU Phillips curve is different in data before and after 1984 is consistent with that of Stock and Watson (1999a). They find a break in the parameters of their estimated unemployment-based NAIRU Phillips curve equation in the early 1980s.

⁹Our finding that the unemployment-based NAIRU Phillips curve model does not produce useful inflation forecasts is consistent with the results reported by Stock and Watson (1999a). Specifically, in their Table 2 (p. 303), Stock and Watson report that their unemployment-based NAIRU Phillips curve model has not been significantly more accurate than a statistical model that forecasts inflation on the basis of lagged values of inflation alone between 1984 and 1996. In that table, they report the ratio of the mean squared forecast error (MSE) for a forecasting model that includes lagged values of inflation alone (referred to as the *univariate autoregression* model) to the MSE of the unemployment-based NAIRU Phillips curve model for their two measures of inflation over two time periods, 1970–83 and 1984–96. We see from their table that in the latter time period, these MSE ratios are not significantly greater than 1.

¹⁰The time series is known as the *Chicago Fed National Activity Index* and is available at <http://www.chicagofed.org/economicresearchanddata/national/cfnai.cfm>. At the time we write, the series is available for the months from March 1967 through October 2000.

¹¹A comparison of Stock and Watson's new NAIRU Phillips curve model with a naive forecast was first done by J. Fisher in unpublished work. In particular, Fisher found that the naive inflation forecast was more accurate during 1984–96 than any of the inflation forecasting models, including the new activity index-based forecast reported in the Stock and Watson (1999a) study. Watson replicated Fisher's results and reported on them in an unpublished memo to Fisher (personal communication, 2000).

¹²The Real-Time Data Set is available at <http://www.phil.frb.org/econ/forecast/reaindex.html>.

¹³We focus on forecasts of inflation as measured by the GNP and GDP deflators because the historical record of forecasts for these measures of inflation is substantially longer than the record of forecasts of inflation as measured by the CPI. We focus on forecasts over a four-quarter horizon because this is the longest horizon for which there is a consistent quarterly historical record.

References

Blinder, Alan S. 1997. Is there a core of practical macroeconomics that we should all believe? *American Economic Review* 87 (May): 240–43.

- Brayton, Flint; Levin, Andrew; Tryon, Ralph; and Williams, John C. 1997. The evolution of macro models at the Federal Reserve Board. *Carnegie-Rochester Conference Series on Public Policy* 47 (December): 43–81.
- Brayton, Flint; Roberts, John M.; and Williams, John C. 1999. What's happened to the Phillips curve? Finance and Economics Discussion Series 1999-49. Board of Governors of the Federal Reserve System.
- Brayton, F., and Tinsley, P. A., ed. 1996. A guide to FRB/US: A macroeconomic model of the United States. Finance and Economics Discussion Series 1996-42. Board of Governors of the Federal Reserve System.
- Cecchetti, Stephen G.; Chu, Rita S.; and Steindel, Charles. 2000. The unreliability of inflation indicators. *Current Issues in Economics and Finance* 6 (April): 1–6. Federal Reserve Bank of New York.
- Clarida, Richard; Gali, Jordi; and Gertler, Mark. 2000. Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics* 115 (February): 147–80.
- Fischer, Stanley. 1977. Long-term contracts, rational expectations, and the optimal money supply rule. *Journal of Political Economy* 85 (February): 191–205.
- Fisher, Irving. 1926. A statistical relationship between unemployment and price changes. *International Labor Review* 13 (June): 785–92. Reprinted 1973. I discovered the Phillips curve. *Journal of Political Economy* 81 (March–April): 496–502.
- Fisher, Jonas D. M. 2000. Forecasting inflation with a lot of data. *Chicago Fed Letter* 151 (March): 1–3. Federal Reserve Bank of Chicago.
- Friedman, Milton. 1968. The role of monetary policy. *American Economic Review* 58 (March): 1–17.
- Gordon, Robert J. 1997. The time-varying NAIRU and its implications for economic policy. *Journal of Economic Perspectives* 11 (Winter): 11–32.
- _____. 1998. Foundations of the Goldilocks economy: Supply shocks and the time-varying NAIRU. *Brookings Papers on Economic Activity* (No. 2): 297–333.
- Lucas, Robert E., Jr. 1972. Expectations and the neutrality of money. *Journal of Economic Theory* 4 (April): 103–24.
- Lucas, Robert E., Jr., and Sargent, Thomas J. 1979. After Keynesian macroeconomics. *Federal Reserve Bank of Minneapolis Quarterly Review* 3 (Spring): 1–16.
- Mankiw, N. Gregory. 2000. The inexorable and mysterious tradeoff between inflation and unemployment. Working Paper 7884. National Bureau of Economic Research.
- Phelps, Edmund S. 1969. The new microeconomics in inflation and employment theory. *American Economic Review* 59 (May): 147–60.
- Phillips, A. W. 1958. The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861–1957. *Economica* 25 (November): 283–99.
- Reifschneider, David L.; Stockton, David J.; and Wilcox, David W. 1997. Econometric models and the monetary policy process. *Carnegie-Rochester Conference Series on Public Policy* 47 (December): 1–37.
- Reifschneider, David; Tetlow, Robert; and Williams, John. 1999. Aggregate disturbances, monetary policy, and the macroeconomy: The FRB/US perspective. *Federal Reserve Bulletin* 85 (January): 1–19.
- Samuelson, Paul A., and Solow, Robert M. 1960. Analytical aspects of anti-inflation policy. *American Economic Review* 50 (May): 177–94.
- Stock, James H., and Watson, Mark W. 1999a. Forecasting inflation. *Journal of Monetary Economics* 44 (October): 293–335.
- _____. 1999b. Business cycle fluctuations in U.S. macroeconomic time series. In *Handbook of Macroeconomics*, ed. John B. Taylor and Michael Woodford, Vol. 1A, Chap. 1. Amsterdam: Elsevier/North-Holland.
- Taylor, John B. 1980. Aggregate dynamics and staggered contracts. *Journal of Political Economy* 88 (February): 1–23.
- U.S. President. 1996. *Economic Report of the President 1997*. Washington, D.C.: U.S. Government Printing Office.

Charts 1–2

The Breakdown in an Early Phillips Curve

Quarterly Unemployment as a Percentage of the U.S. Labor Force vs. Changes in the Implicit Price Deflator for U.S. GDP Over the Next Four Quarters, 1st Quarter 1959–1st Quarter 1999

Chart 1 A Negative Relationship in 1959–69 . . .

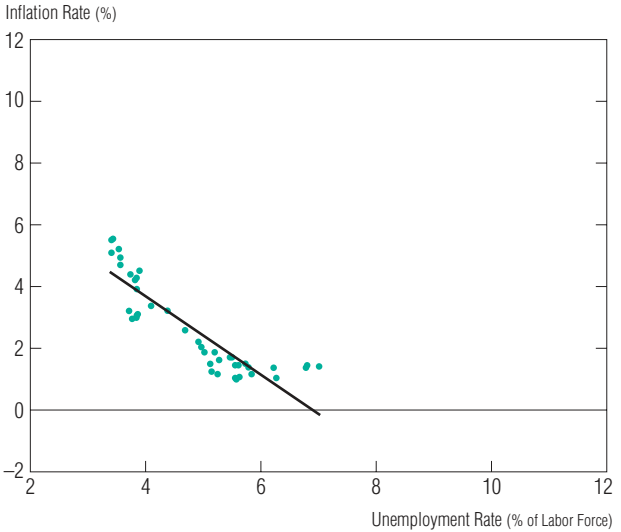


Chart 2 . . . Disappeared in 1970–1999



Charts 3–4

A Shift in the Textbook NAIRU Phillips Curve

Quarterly Unemployment as a Percentage of the U.S. Labor Force vs. Difference Between Change in the Implicit Price Deflator for U.S. GDP Over the Next Four Quarters and Its Change Over the Previous Four Quarters, 1st Quarter 1960–1st Quarter 1999

Chart 3 The Steep Negative Relationship in 1960–83 . . .

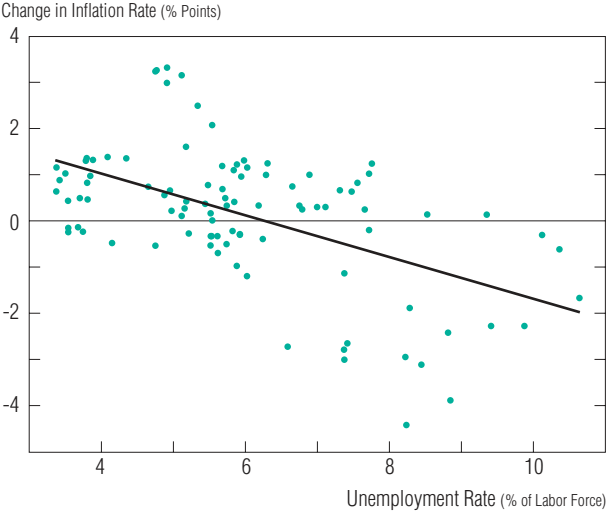
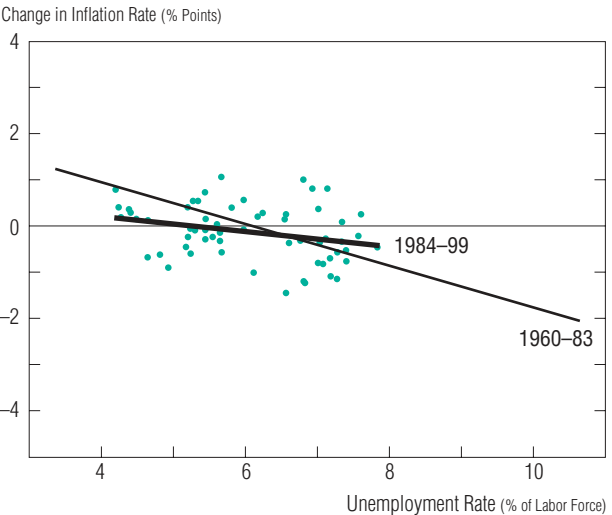


Chart 4 . . . Flattened in 1984–99



Sources: U.S. Departments of Labor and Commerce

Why Use the NAIRU Phillips Curve?

Ratios of Errors of NAIRU and Naive Model* Forecasts of Inflation for 1984–99,
Made With Alternative Indicators and Measures

Inflation Indicator	Inflation Measure†	Range of Ratio of NAIRU/Naive RMSEs**	
		Minimum	Maximum
Unemployment Rate	PCE Deflator	1.02	1.34
	CPI	.99	1.32
	Core CPI	1.06	1.94
Activity Index‡	PCE Deflator	1.04	1.23
	CPI	1.06	1.32
	Core CPI	1.33	1.81

*The NAIRU models are versions of Stock and Watson's (1999a) models. The naive model simply predicts that at any date inflation will be the same as it had been over the past year.

**RMSE=root mean squared error.

†The *PCE deflator* is the implicit price deflator for personal consumption expenditures; the *CPI*, the consumer price index for all items; and the *core CPI*, the consumer price index for all items except food and energy.

‡The *activity index* is the Chicago Fed National Activity Index.

Sources of basic data: U.S. Departments of Labor and Commerce,
Federal Reserve Bank of Chicago