FEDERAL RESERVE BANK OF NEW YORK

CURRENT ISSUES IN ECONOMICS AND FINANCE

April 2000 Volume 6 Number 4

The Unreliability of Inflation Indicators

Stephen G. Cecchetti, Rita S. Chu, and Charles Steindel

Analysts seeking evidence of rising inflation often focus on the movements of a single indicator—an increase in the price of gold, for example, or a decline in the unemployment rate. But simple statistical tests reveal that such indicators, used in isolation, have very limited predictive power.

Controlling inflation—perhaps the most vital responsibility of the Federal Reserve—requires a high degree of foresight. Because policy actions to curb inflation typically take effect only after a long lag, the Federal Reserve needs to know in advance when inflation is likely to rise. Consequently, to understand where prices are headed and what policy steps are appropriate, policymakers turn to forecasts of inflation.

In this edition of *Current Issues*, we consider the usefulness of certain "indicator variables" in forecasting inflation. These variables—which include commodity prices, financial market measures, and measures of real economic activity—are widely thought to be linked to movements in the consumer price index (CPI) and other gauges of inflation.

To determine whether indicator variables can in fact predict inflation, we construct a simple forecasting model that incorporates the variables. Our tests of the model suggest that the indicators, used individually, fail to provide accurate signals of inflation. Indeed, it appears that analysts would achieve better results by basing their predictions on the past behavior of inflation. The forecasts based on indicators also prove to be less reliable than those produced by using a well-known econometric model or by averaging the predictions of a panel of business economists. Although some researchers have argued that the indicator variables *in combination* produce reasonably good forecasts,[1] simple models based on single indicators fail to yield consistently useful information.

Inflation Indicators

Movements in certain economic and financial market variables are thought to presage changes in the CPI and other broad inflation measures. Nineteen variables often proposed as inflation indicators are listed in Table 1. For ease of understanding, we can group all but one of these variables into three broad classes:

- *Commodity prices, such as specific prices for oil and precious metals, or indexes of a group of such goods.* Increases in these measures—either in their level or in their rate of growth—are frequently linked to higher inflation.

- *Financial indicators, such as exchange rates, monetary aggregates, or term premia (the difference between long-term and short-term interest rates)*. A decline in the exchange rate, faster growth of the monetary aggregates, and a widening of the term premium are all supposed to signal increasing inflation.

- *Indicators of the status of the real economy, such as the capacity utilization and unemployment rates.* Higher capacity utilization and lower unemployment (both presumably exceeding some threshold) are regarded as signs that inflation is on the rise.

The one measure listed in the table that does not fit any of these classes is average hourly earnings. Faster growth in this measure is often assumed to be closely related to rising inflation.

Table 1
Candidate Inflation Indicators

Indicators
Journal of Commerce price index for industrial materials, growth
Journal of Commerce price index for industrial materials, level
National Association of Purchasing Management price diffusion index
Price of gold, London fixed
Price of oil, West Texas Intermediate
Exchange rate, trade-weighted U.S. dollar against basket of Group of Ten currencies, growth
Exchange rate, trade-weighted U.S. dollar against basket of Group of Ten currencies, level
Monetary base, growth
M1, growth
M2, growth
Federal funds rate
Spread between the interest rate on the ten-year Treasury bond and the federal funds rate
Spread between the prime commercial paper rate and the federal funds rate
Index of weekly hours worked in private nonfarm business, growth
National Association of Purchasing Management composite index of manufacturing activity
Capacity utilization rate
Unemployment rate
Employment-to-population ratio
Average hourly earnings in private nonfarm business, growth

Developing a Test of the Indicator Variables

To test the predictive power of these variables, we build a simple statistical model that approximates how an analyst might use indicators to forecast a change in the CPI.[2] We begin by regressing CPI growth on its own past values. This "autoregression" essentially uses the recent behavior of inflation to predict inflation's future course. We then incorporate each of the indicators in the model to determine whether the addition of particular variables improves the accuracy of the resulting inflation forecast.

We first estimate the model using data from the first quarter of 1975 through the end of 1984 to produce inflation forecasts for the eight quarters of the 1985-86 period. We then reestimate the model using data through the end of 1985 to predict inflation for the 1986-87 period, data through the end of 1986 to predict inflation for the 1987-88 period, and so forth. This procedure is designed to reproduce the reality that forecasters have data only up to the starting point of their forecasts. Moreover, tracking the performance of the indicators in different years allows us to assess the robustness of their forecasting power.

Model Specification

The model we estimate relates quarterly CPI growth to "lagged," or past, values of CPI growth and lagged values of individual indicators. The lag periods are the same in all regressions: we use the values of CPI growth in the preceding four quarters and the value of the indicator in the preceding quarter.[3]

To be sure, this specification of the model is somewhat arbitrary. In principle, the model could include more than one indicator. Furthermore, the optimal lag periods for CPI growth and the indicator could vary from period to period, as well as from indicator to indicator. However, we believe that our simple model is sufficiently representative to reveal whether one or more of the indicators are consistently useful for inflation forecasts.

Producing Forecasts

To generate inflation forecasts using this model, we enter the lagged values for inflation growth and the individual indicators in our estimated regressions. This poses little problem when we are forecasting the first of the eight quarterly CPI values in each projection period. For example, to forecast inflation in the first quarter of 1985, we enter separately the values for CPI growth in each of the four quarters of 1984 and the value of the indicator for the fourth quarter of 1984. The difficulty comes when we put ourselves in the position of an analyst who is seeking to forecast CPI growth for, say, the second quarter of 1985 but who has hard data on inflation and the indicator only through the end of 1984. The model specification requires us to enter values for inflation and the indicator for first-quarter 1985. A straightforward method of estimating the value for inflation—and the method we use in our analysis—is to adopt the inflation forecast for first-quarter 1985 produced by the model. Estimating the lagged value of the indicator, however, is more complicated, because the model does not generate predictions of the future value of this variable.

One way to circumvent this problem would be to produce inflation forecasts only for those periods in which the analyst had available the data on the lagged value of the indicator. However, since the value of the indicator for the quarter immediately preceding the projection period is included in the model, this strategy would limit our inflation forecast to only one period.

Perhaps a better solution to the problem is to make the extreme assumption that an analyst with expertise in, say, oil prices could predict the future values of this indicator with complete accuracy.[4] This assumption gives us a clean and simple way to test the value of the indicator. Accordingly, in our model, we capture the idea of a perfect forecast by using the actual values of the indicator for the lagged values in the forecast period. For example, to arrive at an inflation forecast for the second quarter of 1985, we use the actual price of oil in first-quarter 1985.

In sum, the steps of our analysis are as follows: We estimate our model for every indicator over sample periods starting with the first quarter of 1975 and ending with the fourth quarter of every year from 1984 to 1996. Using the regression estimates for the sample periods ending in one year, we then produce forecasts of inflation. Our forecasting horizon is eight quarters, a period often used for policy discussions. We find little qualitative change in our results when we adopt horizons of four or twelve quarters.[5]

A Measure of Forecast Performance

To assess the accuracy of our inflation forecasts, we use the root-mean-squared-error (RMSE) statistic. This statistic measures the degree to which the predicted change in the CPI deviates from the actual change over the forecast period.[6] The conventional procedure for testing a variable's ability to forecast inflation involves determining whether the variable, when added to a model, lowers the RMSE.

Note, however, that the RMSE measure does not provide a definitive test of a variable's ability to predict inflation. People can react to an economic forecast in such a way as to change the path of the economy from that initially predicted. Most important, if the Federal Reserve and other policymakers take seriously a forecast of increased inflation and then adopt stabilizing policies, the increase in inflation may not materialize. Thus, demonstrating that a particular indicator accurately forecasts inflation could invite offsetting policy actions and lead to a weakening of the indicator's predictive performance in the future. Likewise, a particular inflation forecast might miss the mark quite dramatically because policymakers heeded its prediction and reacted appropriately.[7]

One can, of course, make forecasts of inflation that take into account the reactions of policymakers. Private forecasters presumably factor into their predictions the likely responses of policymakers. Nevertheless, representing the reactions of policymakers in the framework of a simple model is extremely difficult.[8]

These considerations limit the conclusions we can draw from statistical tests of the accuracy of the inflation forecasts generated by our model. However, given that policy moves to control inflation take effect over a relatively long horizon, our two-year forecast period may be brief enough to preclude much of the impact of such moves. Consequently, we may be justified in viewing the RMSE statistic as a reasonably reliable measure of forecasting performance.

Table 2

Performance of Indicators in Forecasting Inflation Eight Quarters Ahead

Model: $\Delta \text{CPI}_t = \alpha + \Sigma^4_{i=1} \beta_i \Delta \text{CPI}_{t-i} + \delta \text{IND}_{t-1} + \varepsilon_t$

		Number of Indicators That Performed		Root-Mean-Squared Error		
Regression Period	Projection Period	Better Than Autoregression	Worse Than Autoregression	Autoregression	Best Indicator	Worst Indicator
1975:1 1984:4	1985:1 1986:4	12	7	3.19	1.90 (NAPM composite index)	6.59 (Employment/pop. ratio)
1975:1 1985:4	1986:1 1987:4	10	9	2.77	1.94 (*JOC* index, growth)	5.22 (M1)
1975:1 1986:4	1987:1 1988:4	9	10	2.37	1.94 (Unemployment rate)	4.44 (Price of gold)
1975:1 1987:4	1988:1 1989:4	2	17	0.99	0.91 (Federal funds rate [r^{ff}])	5.55 (M1)
1975:1 1988:4	1989:1 1990:4	2	17	1.44	1.36 (Exchange rate, growth)	5.68 (MI)
1975:1 1989:4	1990:1 1991:4	8	11	2.02	1.74 (Price of oil)	3.51 (Exchange rate, level)
1975:1 1990:4	1991:1 1992:4	12	7	3.21	1.44 (NAPM diffusion index)	4.89 (Exchange rate, level)
1975:1 1991:4	1992:1 1993:4	12	7	1.24	0.51 (10-year bond rate - r^{ff})	3.36 (M1)
1975:1 1992:4	1993:1 1994:4	9	10	1.59	0.84 (M2)	2.91 (Monetary base)
1975:1 1993:4	1994:1 1995:4	7	12	1.31	0.90 (M2)	2.95 (*JOC* index, growth)
1975:1 1994:4	1995:1 1996:4	6	13	0.74	0.53 (Price of gold)	3.38 (M1)
1975:1 1995:4	1996:1 1997:4	9	10	1.25	0.78 (Price of oil)	2.83 (M1)
1975:1 1996:4	1997:1 1998:4	8	11	2.12	0.34 (M1)	2.85 (Capacity utilization rate)

Model Results: How the Indicators Performed

What does our analysis reveal about the predictive power of the indicators? First, we learn whether, in each period, forecasters could have improved their predictions by choosing one of the indicators at random and incorporating it in the forecasting model along with past values of inflation. Table 2 shows for every period the RMSE of the forecast produced by the autoregression (that is, the forecast generated solely from past values of inflation), the number of indicators that produced a better forecast, the number that produced a worse one, and the RMSEs of the forecasts produced by the best and worst indicator for that period.

We see that, in most periods, incorporating an indicator selected at random would have produced a worse inflation forecast than that produced by a model based simply on inflation's past values. The total number of indicators that produced forecasts better than the autoregression fluctuated between twelve and two. Typically, the majority of indicator-based forecasts were less accurate than the autoregression benchmark. And in some periods, the best indicator outperformed the autoregression by only the most modest of margins.

Second, we learn whether some indicators provided more reliable forecasts than others. Although forecasters would usually have been better off ignoring all the indicators than choosing an indicator at random, they may still have found some indicators to be superior to others. Table 3 reports how well each indicator performed relative to the autoregression over all periods and how frequently each produced the best or the worst forecast for the full set of indicators examined.

Ten of the nineteen indicators were consistently worse at forecasting inflation than the autoregression—that is, they underperformed the autoregression more than half the time. In particular, incorporating the future path of the exchange rate level and the growth of the monetary aggregate M1 almost always weakened the performance of the model and produced RMSEs that were often larger than those of other indicators. The interest rate variables, the unemployment rate, the monetary base, the employment-to-population ratio, the capacity utilization rate, and the composite index produced by the National Association of Purchasing Management (NAPM) also repeatedly reduced the accuracy of the forecasts generated.

A number of indicators, however, seemed to work well—among them, growth in the *Journal of Commerce* (*JOC*) price index for industrial materials, M2 growth, and growth in average hourly earnings and weekly hours worked. Adding these indicators to the model consistently increased the precision of the inflation forecasts. Nevertheless, as aids in forecasting, these indicators may not warrant much confidence. Although growth in the *Journal of Commerce* index proved superior to the autoregression nine times out of thirteen, on one occasion it yielded the single worst forecast.

M2 growth and growth in the labor market measures—average hourly earnings and weekly hours worked— are problematic in a different way. In our formulation, forecasters must predict the future values of the indicators in order to produce an inflation forecast for more than one quarter ahead. But the future values of M2 and the labor market measures are closely tied to inflation itself: movements in M2 are highly dependent on movements in nominal interest rates and output, and growth in average hourly earnings and in weekly hours worked is easily influenced by changes in inflation and real output. This interrelationship with the variable being predicted undercuts the usefulness of these indicators; forecasters need an indicator whose future values can be predicted independently of inflation.

Three of the other variables tested—the price of gold, the price of oil, and the level of the *Journal of Commerce* price index for industrial materials—also usually increased the accuracy of the inflation forecasts. And unlike movements in M2 and the labor market measures, moderate movements in the price of all these commodities might be predicted with little reference to future movements in overall price inflation.

Table 3
Ranking the Inflation Indicators

	Number of Times the Indicator			
Indicator	Outperforms Autoregression	Underperforms Autoregression	Produces Best Forecast	Produces Worst Forecast
JOC index, growth	9	4	1	1
JOC index, level	10	3	—	—
NAPM diffusion index	7	6	1	—
Price of gold	9	4	1	1
Price of oil	8	5	2	—
Exchange rate, growth	8	5	1	—
Exchange rate, level	0	13	—	2
Monetary base	4	9	—	1
M1	1	12	1	6
M2	8	5	2	—
Federal funds rate (r^{ff})	2	11	1	—
10-year bond rate – r^{ff}	5	8	1	—
Comm. paper rate – r^{ff}	2	11	—	—
Weekly hours	8	5	—	—
NAPM composite index	5	8	1	—
Capacity utilization rate	5	8	—	1
Unemployment rate	2	11	1	—
Employment/pop. ratio	4	9	—	1
Average hourly earnings	9	4	—	—

Table 4
Effect of Unit Increase in Indicators on Inflation, Estimated from 1975–97 Sample
Percentage Points

Indicator	Four Quarters	Eight Quarters	Twelve Quarters
JOC index, level	-0.04	-0.06	-0.08
Price of gold	-0.01	-0.01	-0.02
Price of oil	-0.22	-0.41	-0.60

Note: The long-run effects of the unit increases are -0.12 for the *JOC* index, -0.03 for the price of gold, and -3.42 for the price of oil.

Nevertheless, a closer look at the evidence suggests that we should be wary of interpreting movements in these variables as forerunners of changes in inflation trends. Table 4 reports the effect on inflation of a unit increase in the *Journal of Commerce* index and a dollar increase in the price of gold and the price of oil over horizons of four, eight, and twelve quarters.[9] One would expect increases in the price of industrial materials, gold, and oil to be associated with a *rise* in inflation. What the table shows, however, is that increases in these commodity prices precede future *declines* in inflation. According to our statistical results, an increase of one dollar in the price of oil is associated with a drop in the annual rate of inflation of 0.22 percentage points after four quarters.

How should we interpret this surprising finding? We observed earlier that a weak statistical relationship between an indicator and inflation could arise if policymakers heeded the inflation signals of an indicator and took some offsetting action. In that case, inflation would be unchanged, and the statistical relationship we expected to find would diminish or disappear. However, it is quite unlikely that a policy response would result in a statistical relationship that was the *opposite* of what we expected.[10] Indeed, the illogical nature of this relationship prevents us from putting much stock in the seemingly good showing of gold prices, oil prices, and the *Journal of Commerce* index as inflation indicators.

Alternative Forecasting Tools
The indicators do poorly both in an absolute sense and relative to some obvious alternatives. Consider the performance of the two-year inflation forecasts of the Blue Chip Consensus (Table 5). These forecasts, a simple average of the predictions of several dozen business economists who use a variety of techniques, almost always proved more accurate than the autoregression. Forecasts based on the complex, multi-equation econometric model of Data Resources Incorporated (DRI) outperformed the autoregression eleven times out of thirteen, and on two of these occasions—most notably in 1991-92—also beat the forecasts of both the Blue Chip Consensus and the best of the indicators (compare Tables 2 and 5).

Table 5
Forecasting Performance of Blue Chip, DRI, and Autoregressive Models

Projection Period	Root-Mean-Squared Error: Autoregression	Blue Chip	DRI
1985:1 1986:4	3.19	2.96	2.50
1986:1 1987:4	2.77	2.14	1.97
1987:1 1988:4	2.37	0.78	0.51
1988:1 1989:4	0.99	0.99	0.94
1989:1 1990:4	1.44	1.73	1.86
1990:1 1991:4	2.02	1.90	2.03
1991:1 1992:4	3.21	0.91	0.50
1992:1 1993:4	1.24	0.78	0.78
1993:1 1994:4	1.59	0.84	0.79
1994:1 1995:4	1.31	0.77	0.90
1995:1 1996:4	0.74	0.79	0.60
1996:1 1997:4	1.25	0.83	0.71
1997:1 1998:4	2.12	1.32	1.12

Even the Blue Chip and DRI forecasts have their limitations, however. The RMSE of the forecasts is often more than 1 percentage point—a large miss when inflation is in the area of 2 or 3 percent a year. More important, perhaps, the information in these forecasts is cloudy from a policy perspective. The Blue Chip panel and the DRI modelers surely take into account likely future changes in policy and adjust their forecasts accordingly. Consequently, policymakers who might consult these forecasts would need to filter out such adjustments to obtain a true reading of the likely course of inflation in the absence of policy changes.

Conclusion
No single indicator in our simple statistical framework clearly and consistently improved autoregressive projections. The indicators we found to be reasonably well correlated with overall price inflation either are inherently difficult to forecast independently of inflation or bear an inverse relationship to inflation that seems to defy all logic.

Where does our analysis leave the general topic of inflation indicators? It is conceivable that more complex statistical methods will show the superiority of one of these indicators. Researchers might also want to pursue the possibility that a combination of these indicators could produce a consistently useful measure.[11] But clearly there is good reason to be wary of placing too much confidence in forecasts of inflation that rest on the signals from a single indicator.

Notes

1. See Stock and Watson (1999).

2. We have no specific knowledge of how analysts actually use these indicators in making inflation forecasts.

3. Akaike information criterion tests on every sequential combination from one-quarter up to eight-quarter lags on CPI growth and the indicator showed that the combination of a four-quarter lag on CPI growth and a one-quarter lag on the indicator was almost always better than any other. Adjusting lag lengths for the five top-performing indicators had a marginal effect on their performance.

4. An alternative solution might be to develop an independent or recursive forecasting methodology for the indicator, such as a vector autoregression (VAR) model of overall inflation and the indicator.

5. Our statistical analysis compares quarterly inflation forecasts with actual quarterly inflation. It is possible that our results could change if we examined actual and forecast changes in the price level over the full eight-quarter projection period. This issue may be worth examining, although quarter-to-quarter changes in inflation tend to be fairly mild, lessening the likelihood that this distinction would be material. Moreover, if one wanted to look at the multiperiod inflation forecast, the regression should be specified to capture the multiperiod inflation process. However, the horizon that forecasters and policymakers focus on can change; for that reason, the flexibility of single-quarter inflation forecasting is advantageous.

6. More precisely, the RMSE is the square root of the average squared deviation of the differences between the forecast and actual values of inflation. Suppose inflation is 2 percent over two periods, but is forecast to be 1 percent in the first period and 5 percent in the second. The errors in the inflation forecast are 1 percentage point in the first period and –3 percentage points in the second. The squares of these errors are 1 in the first period and 9 in the second period; the sum of the squared errors is 10; the mean squared error is then 5, and its square root is 2.24. Thus, the RMSE of this inflation forecast is 2.24 percentage points. In contrast, the average error of the forecast is –1 percentage point (1 plus –3 divided by 2), and the average absolute error is 2 percentage points (1 plus 3 divided by 2).

7. The paradox that a good forecast of inflation may not hold up while a bad forecast may provide useful information is an example of what has been referred to as either the "Lucas critique" or "Goodhart's law." The critique is named for Robert Lucas, and the law for Charles Goodhart—economists who explained in the 1970s how such ironies plagued the statistical evaluation of economic stabilization policies and the reliability of monetary policy targets and indicators.

8. Policymakers' reactions are often modeled by a stable linear function relating changes in operating instruments (such as the federal funds rate) to deviations of outcomes from goals (such as inflation targets). The assumption of a stable linear response may be rather difficult to justify.

9. The estimates were derived from the analysis of regressions of the type used in Tables 2 and 3 and estimated over a 1975-97 sample period. Blomberg and Harris (1995) provide a more thorough review of commodity prices as inflation indicators. They find that commodity price movements weaken as signals of inflation after the mid-1980s.

10. One might interpret this result as showing that policymakers consistently overreact to the indicator, but this explanation seems quite far-fetched.

11. As Stock and Watson (1999) attempt to do. For a good discussion of their approach to inflation forecasting, see Fisher (2000).

References

Blomberg, S. Brock, and Ethan S. Harris. 1995. "The Commodity-Consumer Price Connection : Fact or Fable?" Federal Reserve Bank of New York *Economic Policy Review* 1, no. 3.

Blue Chip Economic Indicators. February issues, 1985-96.

Bryan, Michael F., and Stephen G. Cecchetti. 1993. "The Consumer Price Index as a Measure of Inflation." Federal Reserve Bank of Cleveland *Economic Review* 29, no. 3.

Cecchetti, Stephen G. 1995. "Inflation Indicators and Inflation Policy." *NBER Macroeconomics Annual*, 1995.

Cecchetti, Stephen G. 1997. "Measuring Short-Run Inflation for Central Bankers." Federal Reserve Bank of St. Louis *Review* 79, no. 3.

Fisher, Jonas D. M. 2000. "Forecasting Inflation with a Lot of Data." *Chicago Fed Letter*, no. 151, March.

Stock, James H. and Mark W. Watson. 1999. "Forecasting Inflation." *Journal of Monetary Economics* 44, no. 2 (October): 293-335.

The U.S. Economy (formerly DRI/McGraw-Hill *Review of the U.S. Economy* and Standard & Poor's DRI *Review of the U.S. Economy*). February issues, 1985-96.

About the Authors

Stephen G. Cecchetti, a professor of economics at Ohio State University, served as director of research at the Bank between 1997 and 1999. Rita S. Chu, formerly an assistant economist in the Domestic Research Function of the Research and Market Analysis Group, is now a graduate student in economics at New York University. Charles Steindel is a senior vice president in the Business Conditions Function.

Current Issues in Economics and Finance is published by the Research and Market Analysis Group of the Federal Reserve Bank of New York. Dorothy Meadow Sobol is the editor.