Federal Reserve Bank
of Minneapolis

Fall 1989

# Quarterly Review

Federal Reserve Bank of Minneapolis

# Quarterly Review

Vol. 13, No. 4 ISSN 0271-5287

This publication primarily presents economic research aimed at improving policymaking by the Federal Reserve System and other governmental authorities.

Produced in the Research Department. Edited by Preston J. Miller, Kathleen S. Rolfe, and Inga Velde. Graphic design by Barbara Birr and Phil Swenson, Public Affairs Department.

Address questions to the Research Department, Federal Reserve Bank, Minneapolis, Minnesota 55480 (telephone 612-340-2341).

# $P^*$: Not the Inflation Forecaster's Holy Grail†

Lawrence J. Christiano
Research Officer
Research Department
Federal Reserve Bank of Minneapolis

Last summer, the Board of Governors of the Federal Reserve System unveiled a new, experimental way to forecast trends in inflation, which it calls *P-Star* ($P^*$). The announcement received an enormous amount of publicity throughout—and even outside—the United States. One journalist captured the excitement and optimism generated by the new forecasting method particularly well: "Economists have long been searching for the holy grail—an accurate thermometer with which to forecast inflation. . . . Some think they have found it."[1]

They haven't. At first glance, $P^*$ may look like an especially good way to forecast inflation. But a closer look raises doubts about that. And those doubts are confirmed by some simple tests of its forecasting ability. Had $P^*$ been used to forecast inflation in the 1970s and 1980s, its track record would not have been much better than those of other forecasting methods. While $P^*$ may not be a bad way to forecast inflation, it is certainly not an exceptionally good way either.

## What's So Appealing About $P^*$?

This new inflation forecasting method is appealing to many people because it is fairly simple, it seems to make sense, and it is consistent with a widely respected theory of what causes inflation in the long run.

### *It's Simple . . .*

As an inflation forecasting method, $P^*$ can be described with just a few simple equations.

One equation states an obvious fact: at any time in an economy, the number of dollars spent or received equals the number of dollars changing hands. Since early in the century (Fisher 1911), economists have expressed this fact as the *equation of exchange:*

$$(1) \qquad P \times Q = M \times V$$

where

$P =$ the current price level
(here, I'll use the implicit price deflator of gross national product, or GNP)

$Q =$ the current level of output, adjusted for inflation
(here, real GNP)

$M =$ the current money supply
(here, the Fed's M2 definition of money)[2]

$V =$ the velocity of money
(the number of times each dollar of the money supply is spent each year).

---

†Revision of a speech to the Board of Directors of the Federal Reserve Bank of Minneapolis on November 16, 1989. The speech updates the author's report of December 12, 1988. The author has benefited from comments by Jeffrey Hallman, Richard Porter, and David Small.

[1]The Board's new forecasting method is described in Hallman, Porter, and Small 1989. The above quotation is from the *Economist* (Business/Economics focus, 1989). Other $P^*$ publicity includes a front-page story in the *New York Times* (Kilborn 1989), two stories in the business section of a Sunday *New York Times* (Hunt 1989 and Lee 1989), and stories in *Business Week* (McNamee 1989) and the *American Banker* (Heinemann 1989).

[2]M2 includes (1) the components of the Fed's M1 definition of money [currency held by the public, travelers' checks of nonbank issuers, demand deposits at banks and thrifts, negotiable order of withdrawal (NOW and Super-

Those interested in using this equation to study inflation can rearrange its parts to solve for the current price level:

$$(2) \qquad P = M \times V/Q.$$

Economists at the Federal Reserve Board have transformed this equation into a potentially useful forecasting tool by making two assumptions about two of its parts: $V$ and $Q$. First, they assume that no matter where $V$ and $Q$ happen to be at any time, these variables always tend reasonably quickly (say, within two or three years) toward equilibrium (or long-run) values, which the Board economists call $V^*$ and $Q^*$. Second, the economists assume that at any time these equilibrium values are easily computed from the available historical data. In particular, $V^*$ is just the annual average of $V$ since World War II, which is around 1.65. (See Chart 1.) And $Q^*$ is the full capacity level of output—potential output—in the economy, which is assumed to be the smooth trend path of real GNP. The Board economists measure $Q^*$ as growing at a constant 2.5 percent annual rate since about 1981.[3] (See Chart 2.)

Charts 1 and 2

Two Parts of $P^*$

Quarterly, 1959:1–1989:3

Chart 1 The Velocity of Money

Number of Times Each Dollar of M2 Is Spent Each Year

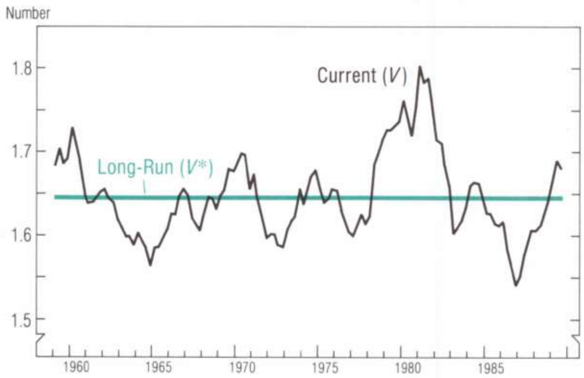


Chart 2 Real Output

Level of Gross National Product, Adjusted for Inflation

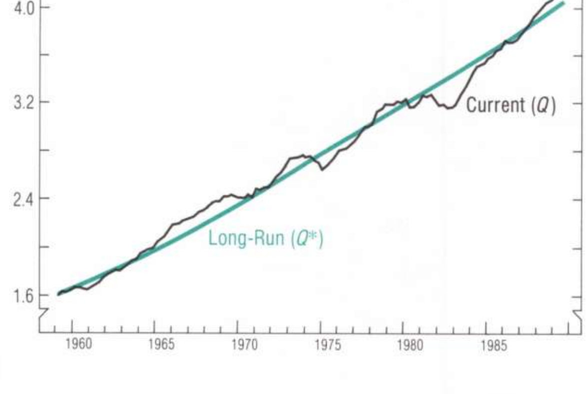


Sources: Federal Reserve Board of Governors; U.S. Department of Commerce; Hallman, Porter, and Small 1989

The Board economists combine $V^*$ and $Q^*$ with the rearranged equation of exchange (2) to produce one more equation, this one for what they call $P^*$: the price level that would occur with the current stock of money if $V$ and $Q$ were at their equilibrium values,

$$(3) \qquad P^* = M \times V^*/Q^*.$$

The significance of this equation for forecasting inflation over the next few years can be seen by comparing it with equation (2) and making use of the Board economists' assumptions. Obviously, the only way $P$ and $P^*$ can differ is by $V$ differing from $V^*$ or $Q$ differing from $Q^*$ or both. But, the assumptions say, if these variables ever differ, then over a few years, their equality will be restored as $V$ and $Q$ drift toward their long-run values of $V^*$ and $Q^*$. What this means, of course, is that over a few years $P$ will tend to drift toward $P^*$. Therefore, if at any time $P$ exceeds $P^*$, then the rate of growth in $P$ (inflation) can be expected to fall for the next few years as $P$ moves down toward $P^*$; and if $P$ is ever less than $P^*$, inflation can be expected to rise. This temporarily lower or higher rate of inflation can be expected to continue until $P = P^*$, at which time the inflation rate will settle

NOW) accounts, automatic transfer service (ATS) accounts, and credit union share draft accounts], (2) savings and small-denomination (less than \$100,000) time deposits, (3) money market deposit accounts (MMDAs), (4) shares in noninstitutional money market mutual funds, (5) overnight repurchase agreements (RPs), and (6) overnight Eurodollar deposits issued to U.S. residents by foreign branches of U.S. banks. For more details on the Fed's money stock definitions, see Walter 1989.

[3]The Board economists estimated $Q^*$ for the period from the first quarter of 1952 (1952:1) to the third quarter of 1989 (1989:3). To see how their measure of $Q^*$ is constructed, it is revealing to look not directly at $Q^*$, but rather at $400 \times \log(Q^*)$. This is simply four straight lines, with different slopes, joined end to end. Because the slopes refer to $400 \times \log(Q^*)$, they are the annualized percentage changes in $Q^*$. The beginning and ending dates and the slopes of the four linear segments of $400 \times \log(Q^*)$ are 1952:1–1965:2, 3.373; 1965:3–1973:4, 3.466; 1974:1–1980:1, 2.795; and 1980:2–1989:3, 2.497.

down to its long-run average, the growth rate of $P^*$.

Since the Board economists assume that $V^*$ is constant and that $Q^*$ will continue to grow at 2.5 percent, the growth of $P^*$ depends on the growth of $M$, how much greater that is than the growth of $Q^*$. In other words, the long-run inflation rate implied by the $P^*$ analysis equals the difference between the long-run growth rates of the money supply and potential output.

The $P^*$ inflation prediction is easy to illustrate graphically. Consider Chart 3. There, the $P^*$ curve assumes that the long-run annual growth rates of money and output are 7 and 2.5 percent, so that the annual growth rate of $P^*$—that is, the long-run inflation rate—is 4.5 percent. The curve traces the path $P^*$ takes over six years, assuming (for simplicity) it starts at 1. The chart depicts two imaginary situations with regard to the current price level, $P$. In one, $P$ starts above $P^*$, with a value of 1.2; in the other, $P$ starts below $P^*$, with a value of 0.8. Note how, from either side, $P$ moves substantially closer to $P^*$ within three years and how this requires that for a time $P$ grow less or more rapidly than $P^*$. The relative speed of growth shows in the relative slopes of the curves. This is dramatized in Chart 4, which plots the year-over-year average growth rates of all variables in Chart 3. Note how $P^*$ grows 4.5 percent throughout and how inflation is temporarily higher than that when $P$ starts low and temporarily lower when $P$ starts high.

### *. . . And Plausible and Well-Connected*

In addition to this simplicity, the $P^*$ analysis has at least two other characteristics that some people find appealing.

One is that its assumptions seem plausible. Recall that the $P^*$ analysis assumes that the long-run rate money changes hands ($V^*$) and the rate potential output ($Q^*$) grows are constant over the few years that $P$ moves toward $P^*$. Many people find those assumptions reasonable because a few years seems too short a time for significant changes to occur in long-run variables such as these which are dependent on slowly changing things like habits and technology.

The other appealing characteristic is $P^*$'s consistency with the widely accepted quantity theory of money (particularly as defined by Lucas 1986). Loosely, this theory says that over a period of many years, where $M$ goes, only $P$ follows. That is, in terms of long-run average growth rates, a one percentage point change in money growth shows up only as a one percentage point change in inflation, without changing output growth at all.[4] This idea can be seen in the $P^*$ analysis,

Charts 3 and 4

## How $P^*$ Works

Imaginary Situations in Which the Price Level ($P$) Is Higher or Lower Than Its Long-Run Average ($P^*$)†

Chart 3 The Price Level

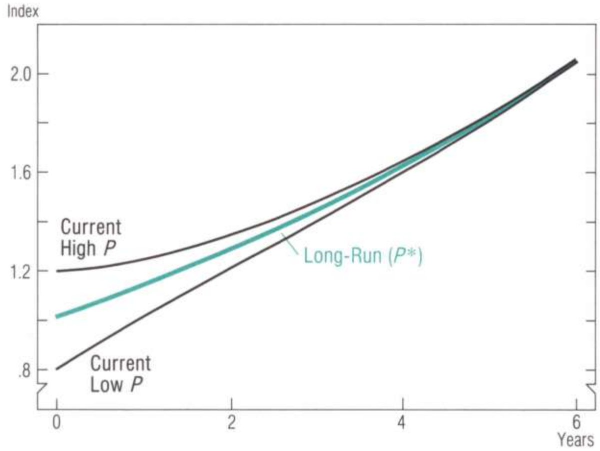


Chart 4 The Inflation Rate

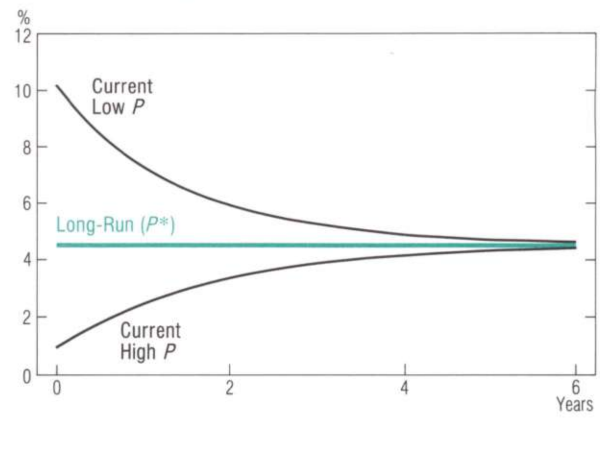


†The $P^*$ curves assume that money and output grow at annual rates of 7% and 2.5%.

[4] For a discussion of the empirical plausibility of the proposition that, in the long run, an increase in the money growth rate ultimately shows up one-for-one in inflation, not in output, see Lucas 1980, 1986; Barro 1987, chap. 7; and Dwyer and Hafer 1988.

too: its assumption that $Q^*$ grows smoothly is consistent with the idea that $Q^*$'s growth is not affected by changes in $M$. Many people seem to think that consistency with the quantity theory adds credence to $P^*$.

### A Closer Look

It doesn't. In fact, neither of the last two reasons people seem to like $P^*$ holds up under scrutiny. And at least one of these failings has potentially serious implications for the usefulness of $P^*$ as a monetary policy tool.

True, $P^*$ is consistent with the quantity theory. But that consistency is irrelevant to $P^*$'s forecasting ability. For the consistency here is between ideas about inflation over the long run—a period of perhaps 10–20 years. As we have seen, $P^*$'s forecasts are for inflation over a much shorter period, only 2–3 years. The quantity theory says nothing about that period. So the agreement between $P^*$ and the quantity theory says nothing about $P^*$'s ability to predict inflation in the shorter run.

Also questionable is the plausibility of $P^*$'s assumptions. Recall that these assumptions are, in short, that velocity will eventually return to its historical average ($V^*$) and that the level of output will return to its historical trend ($Q^*$). If either of these assumptions is wrong, then the $P^*$ analysis could seriously mislead anyone using it to make policy decisions. And today there is ample evidence that both of these assumptions may be wrong.

Consider the assumption about velocity. Suppose for some reason the average value of $V$ were to drop permanently, but the $P^*$ analysis were used with the historical average. Then $P$ in (2) would drop below $P^*$ as $V$ fell and $V^*$ in (3) continued to be held fixed, and the $P^*$ analysis would falsely signal an imminent danger of inflation. Actually, this may have happened after late 1982 when financial innovations such as money market deposit accounts and Super-NOW accounts were introduced. Since then, M2 velocity has almost always been below $V^*$ (Chart 1). Those who think this reflects a permanent drop in velocity say it reflects the reduced opportunity cost of holding M2 due to the fact that these new types of accounts pay explicit interest.

Similarly, consider the assumption about the economy's underlying trend level of output. Suppose that its growth rate were to accelerate, but that this acceleration were not taken into account in the $P^*$ analysis.[5] Then $P$ would be driven below $P^*$, and once again the analysis would falsely warn of inflation. The notion that the growth rate of potential or trend output should shift abruptly is actually pretty reasonable. Even the Board's historical record of $Q^*$ displays periodic, abrupt shifts in its growth rate. Before the first oil price shock, in the early 1970s, for example, $Q^*$'s growth rate was around 3.4 percent. Just after the oil shock, its growth rate dropped to 2.8 percent, and with the second oil shock, in the late 1970s, it fell again, to 2.5 percent. (See footnote 3 for more details.)

In either of the above two situations, if the Fed were paying close attention to the $P^*$ analysis, it might be tempted into a needless and potentially harmful credit contraction.

### How Well Does It Forecast?

Still, the fact that the $P^*$ method requires some assumptions that may be false does not necessarily mean it will not work well. All forecasting methods require assumptions that are false at some level. This does not necessarily imply that they will forecast poorly. The only reasonable basis for evaluating a forecasting method is how well it does what it's supposed to do: forecast. I test $P^*$ here, using two different approaches. And while $P^*$ doesn't do badly, it doesn't do especially well either.

#### *Qualitatively, Not Clear*

My first approach is qualitative. How well would $P^*$ have predicted the direction of the major changes in U.S. inflation since World War II? The answer, unfortunately, is not clear.

The postwar price experience is dominated by two events: the dramatic rise of inflation in the 1970s and its fall in the 1980s. I do a graphical analysis to investigate whether $P^*$ would have helped policymakers anticipate these events.

Chart 5 shows measures of inflation, the price level, and $P^*$ over the last 30 years. Look first at inflation's unprecedented rise to double digits in the late 1970s. Would the $P^*$ analysis have predicted that rise? Maybe, according to this chart. From 1963 until 1970, $P^*$ exceeded $P$, apparently successfully forecasting the accelerating pace in inflation. But twice $P^*$ stopped signaling the alarm: during 1970 and from the second half of 1974 until the second half of 1976. It did this even though the underlying inflation rate clearly had not yet fallen. Thus, the evidence on whether $P^*$ would have signaled the big rise is somewhat mixed, though on balance it seems to be reasonably positive. What about

[5]The work of Nelson and Plosser (1982) suggests another sense in which this assumption may be wrong: output may not revert to any underlying, stable trend. More recently, the academic literature appears to be moving toward a consensus that little can reliably be said about the underlying trend properties of real GNP. See Cochrane 1988, Christiano and Eichenbaum 1989, and Sims 1989.

the big fall? Here $P^*$ gives a sharper, less ambiguous signal. The actual price level, $P$, was consistently above $P^*$ from mid-1978 to mid-1985, clearly predicting lower inflation ahead.

What about since the big fall? Since 1985 until very recently, the $P^*$ analysis predicted a rise in inflation, though inflation has remained fairly stable at a low level. What this means is hard to say. It could have at least two interpretations. One is that $P^*$ was right. It was correctly signaling the threat of rising inflation, and the Fed's gradual credit-tightening moves since 1987 eliminated the threat. This idea is consistent with the facts that starting in 1987 there was some upward pressure on inflation and now $P^*$ coincides with the actual price level. The other interpretation of the recent data, of course, is that $P^*$ was wrong. This could be true, as noted before, if its assumptions are wrong. In particular, the fall of $P$ below $P^*$ could just reflect some combination of a rise in the growth rate of potential output and a fall in velocity.

Obviously, although a qualitative analysis like this is a good, simple way to start evaluating a new forecasting method, it is not good enough for most people who want to know how well that method can really forecast. As we have seen, a qualitative approach doesn't give clear answers about $P^*$'s ability to predict even dramatic changes in inflation. Partly that is because it doesn't give quantitative forecasts of inflation, which is what most forecasters want anyway: how much will inflation change and when? A quantitative analysis should be less ambiguous, too, since it can weigh the pluses against the minuses. Also, the evaluation has so far not been comparative: where does $P^*$ rank relative to other methods of forecasting inflation? Those who call $P^*$ the holy grail of inflation forecasting seem to think it is far superior to all competitors. Is it?

### *Quantitatively, Not Clearly Better*

To answer such questions, we need more than just the simple separate $P$ and $P^*$ equations. A comparison of $P$

Chart 5

**How Well $P^*$ Works: Qualitatively . . .**

Level and Annualized % Change in the GNP Deflator and Level of $P^*$†

Quarterly, 1959:1–1989:3

Shaded areas indicate periods when $P^*$ signals an imminent rise in inflation ($P^* > P$).

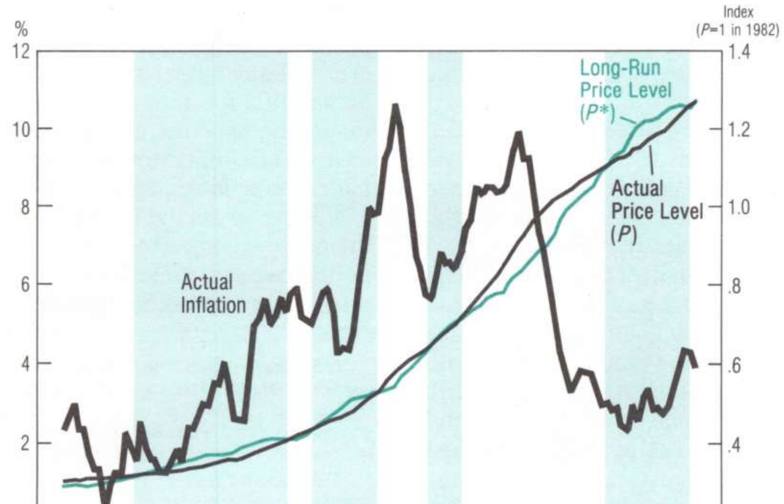


†$P^*$ is constructed using equation (3).

Sources: Federal Reserve Board of Governors; U.S. Department of Commerce; Hallman, Porter, and Small 1989 and $P^*$ can suggest the direction in which inflation will be moving, but not how much it will move or when. To determine things like that, we must build $P$ and $P^*$ into an explicit statement of what determines future inflation and how, that is, into a mathematical model. Economists at the Federal Reserve Board have done this by starting with a simple benchmark model that forecasts future inflation by just extrapolating, or projecting, past inflation. The Board economists incorporate $P^*/P$ into that model. According to the resulting model, if $P^*/P$ is large (meaning $P^*$ is much greater than $P$), then the forecasted level of inflation is higher than would be indicated by extrapolating past inflation alone. This is the $P^*$ model I will test.

To keep this test simple, I will start by comparing the performance of the $P^*$ model to that of just one other forecasting method. Experience in the Research Department of the Minneapolis Federal Reserve Bank suggests that the quarterly change in the yield of 90-day Treasury bills (T-bills) is helpful in forecasting inflation. So, to construct an alternative model, I inserted this change into the Board's benchmark model. In the resulting model, when the yield on 90-day T-bills increases, the inflation forecast is higher than what a simple extrapolation of past inflation rates would indicate.

Now I have two competing models: a $P^*$ model and a T-bill model. (For technical descriptions of these, see Appendix A.) How shall I compare them? The natural way to do that is to find out what would have happened if two forecasters had started out long ago—one with each model—and had used the models to periodically forecast inflation, using at the time of each forecast only the information real forecasters would have had.

Thus, I started by giving each forecasting model all the available quarterly historical data it needs up to and including the fourth quarter of 1969. I used these data to estimate the historical relationships between each model's variables. (In each model, the first period used for this estimation is the third quarter of 1960.[6]) Then I had each model compute a series of one-year-ahead average inflation forecasts, each forecast using one more quarter of actual data. For example, I began with a forecast of the average inflation rate from the fourth quarter of 1969 to the fourth quarter of 1970 (which I dated 1970:4). Then I added the actual data for the first quarter of 1970 to both models' historical data, used the updated data to reestimate the variables' relationships, and computed a forecast of average inflation from the first quarter of 1970 to the first quarter of 1971. I proceeded in this way until I had a series of one-year-ahead forecasts from the fourth quarter of 1970 through the third quarter of 1989. To evaluate the models' longer-run performance, I also computed sets of two- and three-year-ahead forecasts in the same way. (For technical details of my procedures, see Appendix B.)[7]

The results of this test are shown in Charts 6–9.

First consider Chart 6. Plotted there are the one-year-ahead inflation forecasts of the $P^*$ and T-bill models and the corresponding actual inflation rates. Both the forecast and actual inflation rates are measured at an annual rate and in percentage terms. Perhaps the most noticeable thing on this chart is the way the forecasts seem to shadow, or lag, the actual inflation changes: instead of showing what inflation will be, these models seem to show what it has been. In that sense, neither model seems much of a forecasting tool.

Overall in Chart 6, the two models perform about equally well. Note that the chart has, roughly, four episodes. In the first, which extends until the late 1970s, the two models' forecasts nearly coincide. In the second, around the turn of the decade, the T-bill model forecasts better than the $P^*$ model. But the reverse is true in the third episode: In the early 1980s, the $P^*$ model does better. In the final episode, during recent years, the two models' forecasts are close together again, though not as close as they were at first. Considering all four episodes, we must conclude that neither of these models is superior to the other.

These observations are much clearer when we chart just the models' forecast errors. In Chart 6, the forecast error for any model is the vertical distance between the actual inflation rate and the model's forecast. Mathematically, an error equals the actual minus the forecasted inflation rates, so positive numbers represent underpredictions; negative numbers, overpredictions.

Chart 7 plots the forecast errors that correspond to the one-year-ahead forecasts in Chart 6. The episodic pattern we saw there is easier to see here: first the two models do about equally well, then the T-bill model

---

[6]My data set starts in the first quarter of 1959 because that is the first quarter for which M2 data are available. The estimation period starts in the third quarter of 1960 because the forecasting equation requires six initial observations on the price level: four reflecting the number of lags in the model and two more reflecting second-differencing of the price data.

[7]Each forecast of the $P^*$ model requires a forecast of $P^*$ itself. According to equation (3), the forecast of $P^*$ requires forecasts of $M$, $V^*$, and $Q^*$. Appendix B explains how forecasts of these variables were computed using only data that would have been available at the time of each forecast.

Note that my calculations actually only approximate how well real forecasters would have been able to do with these models during the forecasting period. This is because I use revised data here, while real forecasters would have had to use preliminary data.

Charts 6–9

## . . . And Quantitatively

A Comparison of Predicted and Actual Annualized % Changes in the GNP Deflator
Quarterly, 1970:4–1989:3†

Predicted: *P** Model — T-Bill Model — Actual

### One-Year-Ahead Forecasts

Chart 6 **Inflation Rates**

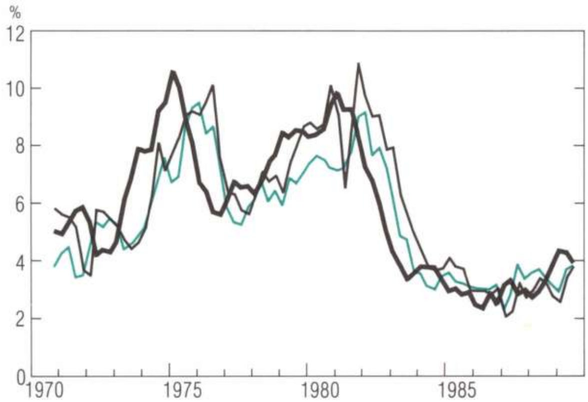


Chart 7 **Inflation Forecast Errors**
(Actual – Predicted)

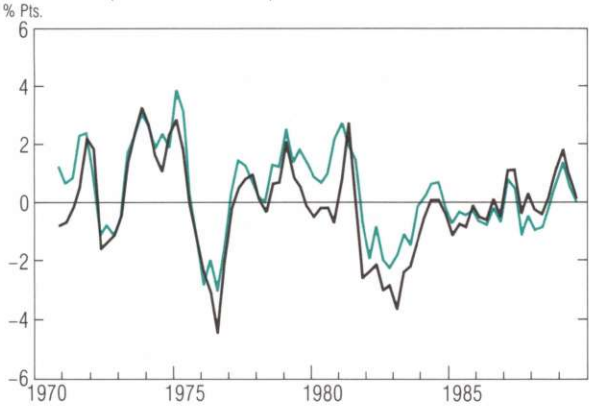


### Forecast Errors (Actual – Predicted)

Chart 8 **Two Years Ahead**

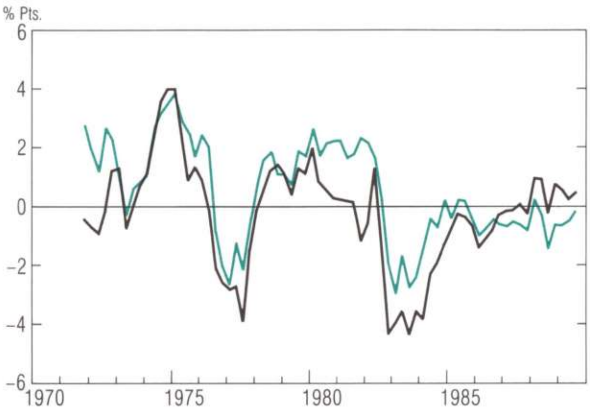


Chart 9 **Three Years Ahead**

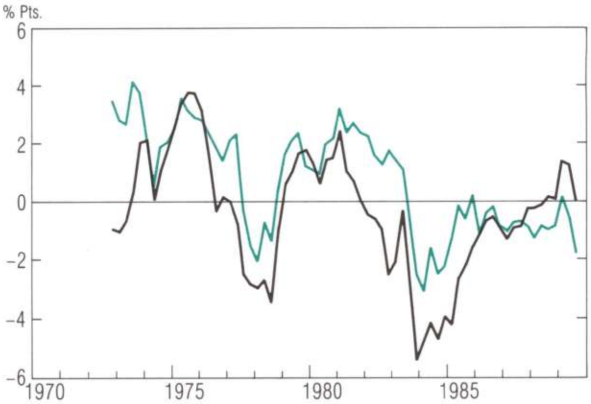


†All the forecasts begin in 1969:4, but the curves in Charts 8 and 9 start later than the others because these curves are for longer forecast periods (two and three years, not just one). Throughout, each forecast is plotted on the last quarter of its period. For example, for the period starting in 1969:4, the one-year-ahead forecast is plotted at 1970:4; the two-year-ahead forecast, at 1971:4; the three-year-ahead forecast, at 1972:4. For technical descriptions of the models, the forecasting procedures, and the forecasts, see Appendixes A, B, and C.

Sources: Federal Reserve Board of Governors; U.S. Department of Commerce; Hallman, Porter, and Small 1989 does better, then the $P^*$ model does better, and finally the two come together again.

This same pattern appears in the two- and three-year-ahead forecast errors, which are plotted in Charts 8 and 9. Overall in these charts, the two models perform similarly, but there are subperiods in which one dominates the other.

In sum, neither the $P^*$ nor the T-bill model is clearly much better than the other at forecasting inflation from one to three years out.[8] A comparison of $P^*$ and seven other models (in Appendix C) produces the same result. If $P^*$ is, indeed, the inflation forecaster's holy grail, then there is little hint of it in the historical data.

[8] For a formal quantitative analysis of the results in Charts 6–9, see Appendix C.

# Appendix A
# Describing the $P^*$ and T-Bill Models

Here I describe the two competing models in the preceding paper: the $P^*$ model and the T-bill model. The $P^*$ model is

$$\Delta\pi_t = \alpha\log(P^*_{t-1}/P_{t-1}) + \beta_1\Delta\pi_{t-1} + \beta_2\Delta\pi_{t-2} + \beta_3\Delta\pi_{t-3} + \beta_4\Delta\pi_{t-4} + u_t$$

where $P_t$ is the quarterly price level; $\pi_t$ is the quarterly inflation rate, defined as $\pi_t = \log(P_t) - \log(P_{t-1})$; $\Delta\pi_t = \pi_t - \pi_{t-1}$; and $u_t$ is the regression error term. The T-bill model is the same as the above except that $\log(P^*_{t-1}/P_{t-1})$ is replaced by $R_{t-1} - R_{t-2}$, where $R_t$ is the average of annualized returns for 90-day Treasury bills purchased in quarter $t$. Both models' parameters are estimated by the method of ordinary least squares.

The accompanying table displays estimation results, based on data covering the period 1960:3–1989:3, for the T-bill model and two versions of the $P^*$ model. The two versions of the $P^*$ model use two different measures of $P^*$: mine and the Board's. For reasons given in Appendix B, I could not use the Board's measure of $P^*$ for the calculations reported in the paper. In the table, the coefficients of the $P^*$ model seem relatively insensitive to which measure of $P^*$ is used. This is consistent with the results reported in Appendix B, according to which forecasts from the $P^*$ model do not deteriorate when my measure of $P^*$ is used instead of the Board's.

Estimation of Inflation Models

1960:3–1989:3

| Parameters and Statistics | Values (and Standard Errors) for Each Model | | |
|---|---|---|---|
| | Board $P^*$ | My $P^*$ | T-Bill |
| $\alpha$ | .032 | .021 | .0012 |
| | (.0085) | (.0077) | (.00042) |
| $\beta_1$ | −.61 | −.57 | −.59 |
| | (.090) | (.089) | (.092) |
| $\beta_2$ | −.44 | −.39 | −.40 |
| | (.10) | (.10) | (.10) |
| $\beta_3$ | −.27 | −.23 | −.27 |
| | (.10) | (.10) | (.10) |
| $\beta_4$ | −.13 | −.094 | −.13 |
| | (.087) | (.087) | (.090) |
| Residual S.E. | .003960 | .003984 | .004064 |
| $R^2$ | .3172 | .3087 | .2808 |

# Appendix B
# Calculating the Forecasts

Here I describe the calculations underlying Charts 6–9 in the preceding paper. Throughout, forecasts only use data available at the date of the forecast. Forecasts computed in this way are called *real-time* forecasts.

Let $P_{69:4}$ and $P_{70:4}$ denote the price levels in the fourth quarters of 1969 and 1970. The actual average inflation rate between these quarters is $\log(P_{70:4}) - \log(P_{69:4})$. To compute the forecast of this as of 1969:4, the $P^*$ model requires forecasts of $P^*_{70:1}$, $P^*_{70:2}$, and $P^*_{70:3}$, and the T-bill model requires forecasts of $\Delta R_t$ for those periods. Here, $\Delta R_{70:1} = R_{70:1} - R_{69:4}$.

Forecasts of $P^*$ require forecasts of $V^*$, the Federal Reserve Board's M2 measure of money, and the full capacity level of output, $Q^*$. Real-time forecasts of $V^*$ were computed using the sample mean of $V$ from 1959:1 until the date of the forecast. Real-time forecasts of M2 were computed using this second-order autoregressive [AR(2)] model:

$$\text{(B1)} \quad \log(M2_t) = \alpha_0 + \alpha_1 \log(M2_{t-1}) + \alpha_2 \log(M2_{t-2}) + \epsilon_t.$$

When estimated over the period 1960:3–1989:3, the parameters of this model are as shown in the M2 column of Table B1. The Box-Pierce $Q$-statistic and associated significance level indicate no significant serial correlation in the fitted residuals.

Unfortunately, my attempts to construct a real-time version of the Federal Reserve Board economists' measure of $Q^*$ were not successful. However, it turns out that the Board measure of $Q^*$ resembles an exponential trend for real GNP. This is fortunate, since an exponential trend model of $Q^*$ is easy to implement in real time. Thus, I modeled $\log(Q^*)$ at date $t$ as $\beta_0 + \beta_1 t$, where $\beta_0$ and $\beta_1$ were obtained by least squares regressions of log(real GNP) on a constant and on time.

The results of the following experiment suggest that the $P^*$ model is not placed at a disadvantage by using my measure of $Q^*$ rather than the Board's. I constructed a Board measure of $P^*$ for the period 1959:1–1989:3 using actual M2 and the Board's reported measure of $Q^*$ (Hallman, Porter, and Small 1989). I constructed my version of $P^*$ for the same period using actual M2 and my measure of $Q^*$, computed using estimates of $\beta_0$ and $\beta_1$ obtained from real GNP data covering the period 1959:1–1989:3. In both versions of $P^*$, the measure of $V^*$ used was the sample average of $V$ over that period. I then computed out-of-sample one-, two-, and three-year-ahead inflation forecast errors as I describe in the paper, except here I assumed that the whole time series on $P^*$ is known. My results are as shown in Table B2. Note that, in terms of both root mean squared error and bias, my measure of $P^*$ produces uniformly better forecasts than the Board's does.

I obtained real-time forecasts of $\Delta R_t$ for the T-bill model using this AR(2) model:

$$\text{(B2)} \quad \Delta R_t = \beta_0 + \beta_1 \Delta R_{t-1} + \beta_2 \Delta R_{t-2} + \nu_t.$$

When this model's parameters are estimated over the period 1960:3–1989:3, the results are as shown in the $\Delta R$ column of

Table B1

**Estimation of Money and Interest Rate Models**

1960:3–1989:3

| Lags on Explanatory Variable and Statistics | Parameter Values (and Standard Errors) for Each Model† | |
|---|---|---|
| | M2 | $\Delta R$ |
| Zero | .012<br>(.0058) | .0435<br>(.082) |
| One | 1.59<br>(.074) | .20<br>(.089) |
| Two | -.59<br>(.074) | -.31<br>(.089) |
| Residual S.E. | .0064 | .88 |
| $R^2$ | .99992 | .12 |
| $Q$-Stat. at Lag 30 | 26.82 | 36.34 |
| Significance Level | .63 | .20 |

†These models correspond to equations (B1) and (B2).

Table B2

## Comparison of Two Versions of the P* Inflation Forecasting Model

| Forecast Horizons (Years) | Root Mean Squared Errors (and Means) | |
|---|---|---|
| | Board P* | My P* |
| One | 1.41 | 1.35 |
| | (.19) | (.020) |
| Two | 1.61 | 1.46 |
| | (.28) | (-.0060) |
| Three | 1.81 | 1.55 |
| | (.45) | (.030) |

Note: The forecast errors underlying these results are out of sample: for each forecast, the coefficients of the P* model are estimated using only data available at the time of the forecast. These forecast errors are not real time, however, since the forecasts use actual—not forecasted—values of P*.

Table B1. Here, too, the *Q*-statistic and associated significance level indicate no significant serial correlation in the residuals.

# Appendix C
# Analyzing the Forecasts

Here I further analyze the results in Charts 6–9 in the preceding paper, while also describing the results of other tests not detailed there. Besides the two models tested in the paper, I also examined the real-time forecasting performance of seven other models. The results of all the tests are the same: $P^*$ is not far superior to other simple methods of forecasting inflation.

### 7 More Models and 3 Statistics

Among the seven extra models are two *money growth* models and one *term structure* model. In the money growth models, I replaced $R_{t-1} - R_{t-2}$ in the T-bill model (described in Appendix A) with $\log(M2_{t-1}) - \log(M2_{t-2})$ and with $\log(MB_{t-1}) - \log(MB_{t-2})$, where $MB$ is the Federal Reserve Board's measure of the monetary base. I used an AR(2) model [equation (B1) in Appendix B] to forecast M2 in the M2 model and an AR(2) model to forecast MB in the MB model. In the term structure model, I replaced $R_{t-1} - R_{t-2}$ in the T-bill model with $R_{10t-1} - R_{t-1}$, where $R_{10}$ is the return on 10-year Treasury bonds. To forecast $R_{10} - R$ in this model, I used an AR(2).

Besides these three models, I also considered four more. One is the *benchmark* model mentioned in the paper; it simply extrapolates inflation from its own past because it sets $\alpha = 0$ in the $P^*$ model. Another I call *combination;* it combines the $P^*$ and T-bill models in the obvious way, by adding $\delta(R_{t-1} - R_{t-2})$ as an explanatory variable to the $P^*$ model, where $\delta$ is an unknown parameter to be estimated. I also used what I call a *level T-bill* model; it is the same as the T-bill model, except that it replaces $\log(P^*_{t-1}/P_{t-1})$ in the $P^*$ model with $R_{t-1}$ instead of $R_{t-1} - R_{t-2}$ and it obtains forecasts of $R_t$ using an AR(3) instead of an AR(2) model. The final model I call *premium;* it uses $R_{ct-1} - R_{t-1}$ instead of $\log(P^*_{t-1}/P_{t-1})$ in the $P^*$ model, where $R_{ct}$ is the yield on commercial paper. In the premium model, forecasts of $R_{ct} - R_t$ come from an AR(2) model.

For each of the above seven models, I computed one-, two-, and three-year-ahead forecast errors using the same real-time procedure applied to the $P^*$ and T-bill models (as described in Appendix B). Then I used the forecast errors of all the models to compute three statistics by which to evaluate and compare their inflation forecasting performance. One of these statistics, the simple average of the forecast errors (or the *mean*), is designed to assess whether a model's forecasts are biased. The two others measure the typical absolute size of a model's forecast errors: the mean of the absolute values of the errors (the *MAVE*) and the square root of the mean squared error (the *RMSE*). The basic difference between these two is that the RMSE weighs large forecast errors relatively more heavily than the MAVE does. All three statistics are measured in percentage terms at an annual rate.

### The Results

The results for each model at each forecast horizon are displayed in Table C1.

*Eyeballing the 3 Statistics*

Consider first the results for the $P^*$ and T-bill models, shown in the table's first two rows. They are based on the forecast errors in the paper's Charts 7–9, and they confirm the paper's conclusions: though on some dimensions one model seems to outperform the other, overall the two models perform about equally well. In terms of bias and average absolute error, the T-bill model outperforms the $P^*$ model at all three forecast horizons, but in terms of the RMSE, the $P^*$ model outperforms the T-bill model at all horizons. In either case, these differences seem too small to be economically or statistically significant. For example, in RMSE terms, the largest improvement from using the $P^*$ model rather than the T-bill model occurs at the three-year horizon. That improvement is a mere 0.15 of a percentage point, an economically negligible amount in view of the roughly 6 percent inflation we have averaged over the postwar period. Also, a glance back at Charts 7–9 shows that the magnitude of fluctuations in the forecast errors is much larger than 0.15 of a percentage point. This suggests that differences in RMSE of such magnitude are not statistically significant, or that the superior RMSE performance of the $P^*$ model can't be counted on to persist. The formal statistical analysis reported in the next subsection supports this conclusion.

Bringing the results of the combination and benchmark models into the analysis offers additional ways to assess how much information about inflation is contained in $P^*$. Thus, comparing the combination and T-bill model results shows that there is little information in $P^*$ that is not already in T-bill yield changes. Incorporating $P^*$ into the T-bill model (to get the combination model) reduces the bias and magnitude of the T-bill model's forecast errors by a trivial amount. Comparing the benchmark and $P^*$ model results suggests a sense in which $P^*$ actually contains disinformation about inflation. According to the mean and MAVE, introducing $P^*$

Table C1

Comparison of Real-Time Performance of Nine Inflation Forecasting Models

Measured by Three Statistics†

(The shaded numbers identify the best model in each column.)

| | Value of Each Statistic at Each Forecast Horizon (Years) | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | One | | | Two | | | Three | | |
| Models | Mean | MAVE | RMSE | Mean | MAVE | RMSE | Mean | MAVE | RMSE |
| $P^*$‡ | .38 | 1.26 | 1.53 | .55 | 1.46 | 1.74 | .79 | 1.71 | 1.98 |
| T-Bill | -.12 | 1.22 | 1.58 | -.26 | 1.37 | 1.84 | -.39 | 1.65 | 2.13 |
| Combination | .15 | 1.18 | 1.54 | .17 | 1.37 | 1.73 | .24 | 1.64 | 1.99 |
| Benchmark | -.07 | 1.25 | 1.60 | -.16 | 1.38 | 1.86 | -.24 | 1.64 | 2.09 |
| M2 | -.34 | 1.30 | 1.70 | -.58 | 1.44 | 2.01 | -.83 | 1.72 | 2.29 |
| MB | -.62 | 1.31 | 1.72 | -1.0 | 1.60 | 2.14 | -1.5 | 2.02 | 2.58 |
| Term Structure | .05 | 1.32 | 1.67 | -.025 | 1.48 | 1.93 | -.10 | 1.69 | 2.14 |
| Level T-Bill | -.44 | 1.30 | 1.76 | -.74 | 1.55 | 2.22 | -1.1 | 2.00 | 2.71 |
| Premium | -.32 | 1.31 | 1.72 | -.54 | 1.47 | 2.10 | -.76 | 1.84 | 2.45 |

†The *mean* is the simple average of the model's forecast errors, the *MAVE* is the mean of the absolute values of those errors, and the *RMSE* is the square root of the mean squared error.

‡The results in this row are based on the version of the $P^*$ model that uses my measure of $P^*$, not the Board's. For details on the two versions, see Appendix B.

Table C2

**A Closer Look at the Models' Root Mean Squared Errors**

Difference Between the RMSEs of the Indicated Model and the *P** Model and *t*-Statistic Associated With That Difference†

| | Values at Each Forecast Horizon (Years) | | | | | |
|---|---|---|---|---|---|---|
| | One | | Two | | Three | |
| Models | Diff. | *t*-Stat. | Diff. | *t*-Stat. | Diff. | *t*-Stat. |
| T-Bill | .048 | .30 | .099 | .37 | .15 | .40 |
| Combination | .0070 | .07 | −.0059 | −.04 | .0075 | .04 |
| Benchmark | .065 | .44 | .12 | .51 | .12 | .32 |
| M2 | .16 | .77 | .27 | .80 | .32 | .65 |
| MB | .19 | .79 | .40 | 1.1 | .60 | 1.2 |
| Term Structure | .13 | .95 | .19 | .84 | .16 | .44 |
| Level T-Bill | .22 | .92 | .48 | 1.2 | .73 | 1.3 |
| Premium | .18 | .88 | .36 | 1.1 | .47 | .99 |

†Each *t*-statistic is the ratio of the difference in RMSEs (the indicated model's minus the *P** model's) to the standard error of that difference. For an argument that this ratio has, approximately, a standard normal distribution under the null hypothesis that the underlying RMSE difference is actually zero, see Appendix D.

into the benchmark model (to get the *P** model) reduces forecast accuracy at all horizons.

What evidence there is that *P** improves inflation forecasts appears to be greater when the *P** model is compared with the remaining five models. This improvement in forecast performance seems most pronounced with the RMSEs. For example, at the three-year horizon, the level T-bill model's RMSE is 0.73 of a percentage point higher than the *P** model's. But even RMSE differences of this magnitude are not statistically significant.

### *Formally Analyzing 1 Statistic*

To determine that, I computed *t*-statistics for testing the null hypothesis that the true RMSE of the *P** model is identical to that of each of the other eight models. For any model and forecast horizon, this *t*-statistic is just the ratio of the difference between the model's RMSE and *P**'s RMSE to the standard error of that difference. In Appendix D, I argue that these *t*-statistics have, approximately, a standard normal distribution under the null hypothesis that the true underlying RMSE differences are in fact zero.

The *t*-statistics are displayed in Table C2. For convenience, the RMSE differences—the numerators in the *t*-statistic ratios—are also displayed there. (Apart from discrepancies due to rounding, the RMSE differences can be obtained by subtracting the appropriate elements in Table C1.) The fact that all but one of the differences are positive reflects the fact that all but one of the *P** model's RMSEs are lower than those of the other models.

Despite that, the *t*-statistics in Table C2 are all consistent with the hypothesis that each model's true RMSE performance is actually identical to that of the *P** model. All of the *t*-statistics are very close to zero, the central tendency of the standard normal distribution. For example, the *t*-statistic for the T-bill model's one-year-ahead forecasts is only 0.30. The probability that a standard normal random variable will exceed 0.30, in absolute value, is 76 percent. Even the largest *t*-statistics, those in the last four rows of the table, are quite small. For example, the probability of a standard normal random variable exceeding (in absolute value) 1.3, the largest *t*-statistic in the table, is 19 percent.†

## The Conclusion

Thus, my analysis has failed to turn up any convincing

†Some readers of early drafts of this paper were concerned that the simplicity of the AR(2) representation I used to forecast money for the *P** model may have placed that model at an unfair disadvantage. They suggested

evidence that $P^*$ far outperforms other models as an inflation forecaster. In my tests, $P^*$ does worse than its competitors on some dimensions and better on others. The dimension on which $P^*$ looks best is the estimated RMSE of its forecast errors. But the superiority of $P^*$ on this dimension is so small that it could just reflect dumb luck.

an alternative procedure which avoids the need to forecast other variables when forecasting inflation: use annual or biennial rather than quarterly observations to forecast inflation one or two years out. Note that, effectively, this procedure throws away a large number of observations. Other things the same, one expects that to result in a deterioration in forecast performance. Still, to investigate this suggestion, I used annual data to compute a set of real-time, one-year-ahead forecast errors using annual versions of my $P^*$, T-bill, combination, and benchmark models. I found that the forecast performance of all these models is generally inferior to that of their quarterly counterparts. Moreover, the annual $P^*$ model forecasts less well than all of my quarterly models: the quarterly T-bill and benchmark models beat it in terms of all three of my statistics, and the other quarterly models beat it in terms of two of the three, the mean and the MAVE.

# Appendix D Rationalizing the Use of the Standard Normal Distribution

Here I explain why the $t$-statistics in Appendix C's Table C2 can be interpreted as though they were drawn from the standard normal distribution.

### An Approximation

Let $\hat{\gamma}$ denote the estimator underlying the numbers in any of the differences columns in that table. Thus, $\hat{\gamma}$ is the estimated RMSE for the forecasts of one of the models in the table minus the corresponding RMSE for the forecasts of the $P^*$ model. Let $\tilde{\gamma}$ denote the underlying true RMSE difference. This is what $\hat{\gamma}$ would be if it were computed using an unlimited number of observations instead of the 70-odd data points that are actually available. I use Hansen's (1982) generalized method of moments formula to estimate the standard error of $\hat{\gamma}$. To a first approximation, the ratio of $\hat{\gamma}$ to this standard error (the $t$-statistic discussed in Appendix C) has an asymptotic standard normal distribution under the null hypothesis that $\tilde{\gamma}$ is zero.

To use Hansen's standard error formula, I must express $\hat{\gamma}$ as the solution to the sample analog of some first-moment condition. To keep the notation simple, I describe here just the calculations for the T-bill model's one-year-ahead forecast errors. The calculations for all the other models and forecast horizons are analogous.

To start, define this $2 \times 1$ random variable:

$$\text{(D1)} \quad h_t(\gamma,\sigma) \equiv \{\gamma[\gamma + 2\sigma] - e(\Delta R,t)^2 + e(P^*,t)^2, \\ \sigma^2 - e(P^*,t)^2\}'$$

where $\sigma$ is the RMSE for the $P^*$ model's forecasts and $e(\Delta R,t)$ and $e(P^*,t)$ are the date $t$ one-year-ahead forecast errors from the T-bill and $P^*$ models. Also, $\gamma = \sigma_R - \sigma$, where $\sigma_R$ is the RMSE for the T-bill model's forecasts, so that $\gamma[\gamma + 2\sigma] = \sigma_R^2 - \sigma^2$. Therefore, if $\tilde{\sigma}$ is the (unknown) true value of $\sigma$, then $Eh_t(\tilde{\gamma}, \tilde{\sigma}) = 0$ for each $t$. This is the first-moment condition that underlies my estimator of $\gamma$ (and $\sigma$).

The sample analog of this condition is

$$\text{(D2)} \quad g_T(\gamma,\sigma) = \{\gamma[\gamma + 2\sigma] \\ - (1/T)\Sigma_{t=1}^{T}[e(\Delta R,t)^2 - e(P^*,t)^2], \\ \sigma^2 - (1/T)\Sigma_{t=1}^{T} e(P^*,t)^2\}'.$$

Here $T$ is the number of observations used in the RMSE calculations (76, 72, and 68 for the one-, two-, and three-year-ahead forecast errors). It is easy to confirm that my estimators of $\gamma$ and $\sigma$ (the $\hat{\gamma}$ defined above and the estimated RMSE of the $P^*$ model's forecast errors) uniquely satisfy $g_T(\hat{\gamma},\hat{\sigma}) = 0$.

Now define

(D3) $\hat{D}_T = \partial g_T(\gamma,\sigma)/\partial(\gamma,\sigma)$

evaluated at $\gamma = \hat{\gamma}$ and $\sigma = \hat{\sigma}$. Hansen (1982) shows that if $h_t(\tilde{\gamma},\tilde{\sigma})$ is strictly stationary and other regularity conditions are satisfied, then for large $T$

(D4) $\begin{bmatrix} \hat{\gamma} \\ \hat{\sigma} \end{bmatrix}$ is normally distributed with mean $\begin{bmatrix} \tilde{\gamma} \\ \tilde{\sigma} \end{bmatrix}$ and variance $\tilde{V}$

where $\tilde{V}$ is a positive definite matrix consistently estimated by

(D5) $\hat{V} = (\hat{D}_T)^{-1}\hat{S}(\hat{D}_T')^{-1}/T.$

Here $\hat{S}$ is a matrix that I will explain in the next paragraph. Meanwhile, we can already see by (D4) and (D5) that the estimated standard error of $\hat{\gamma}$ is the square root of the first diagonal element of $\hat{V}$. (The square root of the second diagonal element is Hansen's estimator of the standard error of $\hat{\sigma}$.) For the T-bill model's one-year-ahead forecast errors, the standard error of $\hat{\gamma}$ is 0.16, thus accounting for the 0.30 $t$-statistic in Table C2. (The standard errors of $\hat{\sigma}$ at the one-, two-, and three-year-ahead horizons are 0.18, 0.20, and 0.40.)

To return to $\hat{S}$ in (D5): It is a consistent estimator of $\tilde{S}$, the spectral density at frequency zero of $h_t(\tilde{\gamma},\tilde{\sigma})$. In particular,

(D6) $\tilde{S} = \Sigma_{k=-\infty}^{\infty} E\{h_t(\tilde{\gamma},\tilde{\sigma})[h_{t-k}(\tilde{\gamma},\tilde{\sigma})]'\}.$

Both $\hat{S}$ and $\tilde{S}$ are $2 \times 2$ matrices, and $\tilde{S}$ is assumed to be positive definite under the regularity conditions. My estimator of $\tilde{S}$, $\hat{S}$, replaces the population second moments on the right side of (D6) by their sample counterparts and truncates the summation for $|k| > 6$. In addition, I linearly damp higher-order covariances according to the formula of Doan (1988, p. 14-143) with $\theta = 1$.

### A Hitch?

An attractive feature of Hansen's estimator of the variance-covariance of $(\hat{\gamma},\hat{\sigma})$ is that it is robust to autocorrelation and conditional heteroscedasticity in $h_t(\tilde{\gamma},\tilde{\sigma})$. However, as Christopher Sims has pointed out to me, Hansen's stationarity requirement may not be satisfied. After all, model parameter estimates are based on less data at the beginning of a sample period than at the end. Therefore, forecast errors may have larger variance at the beginning than at the end.

Charts 7–9 in the paper give (slight) support for this possibility. Thus, it may be worthwhile to investigate whether failure of Hansen's stationarity assumption substantially affects my standard error estimates. Still, I suspect that it does not, for my results are very strong. To overturn my conclusion about the comparable RMSE performance of the $P^*$ and T-bill models, for example, would require showing that my $t$-statistics are biased downward by a factor of at least five. This is because the largest of those $t$-statistics is 0.4 and rejecting a null hypothesis at the conventional 5 percent significance level requires a $t$-statistic roughly equal to 2.

# References


Barro, Robert J. 1987. *Macroeconomics.* 2nd ed. New York: Wiley.

Business/Economics focus. 1989. Star gazing. *Economist* 312 (July 15–21): 65.

Christiano, Lawrence J., and Eichenbaum, Martin. 1989. Unit roots in real GNP: Do we know, and do we care? Working Paper 3130. National Bureau of Economic Research. Also, forthcoming in Carnegie-Rochester Conference Series on Public Policy 32. Amsterdam: North-Holland.

Cochrane, John H. 1988. How big is the random walk in GNP? *Journal of Political Economy* 96 (October): 893–920.

Doan, Thomas A. 1988. User's manual: RATS (Regression analysis of time series). Version 3.00. Evanston, Ill.: VAR Econometrics.

Dwyer, Gerald P., Jr., and Hafer, R. W. 1988. Is money irrelevant? *Review* (May/June): 3–17. Federal Reserve Bank of St. Louis.

Fisher, Irving. 1911. *The purchasing power of money.* New York: Macmillan.

Hallman, Jeffrey J.; Porter, Richard D.; and Small, David H. 1989. M2 per unit of potential GNP as an anchor for the price level. Staff Study 157. Board of Governors of the Federal Reserve System.

Hansen, Lars Peter. 1982. Large sample properties of generalized method of moments estimators. *Econometrica* 50 (July): 1029–54.

Heinemann, H. Erich. 1989. Overview/Commentary: Cure for Greenspan's dilemma is not to be found in formula. *American Banker* 154 (June 19): 4, 8.

Hunt, Lacy H. 1989. Forecasting inflation: The Fed's new tool just doesn't work. *New York Times* 138 (October 22): Business 2.

Kilborn, Peter T. 1989. Federal Reserve sees a way to gauge long-run inflation. *New York Times* 138 (June 13): 1.

Lee, L. Douglas. 1989. 'P-Star' can spot inflationary trends. *New York Times* 138 (October 22): Business 2.

Lucas, Robert E., Jr. 1980. Two illustrations of the quantity theory of money. *American Economic Review* 70 (December): 1005–14.

__________. 1986. Adaptive behavior and economic theory. *Journal of Business* 59 (Part 2, October): S401–26.

McNamee, Mike. 1989. Putting 'Keynes's head on Milton Friedman's body.' *Business Week* (July 31): 66.

Nelson, Charles R., and Plosser, Charles I. 1982. Trends and random walks in macroeconomic time series: Some evidence and implications. *Journal of Monetary Economics* 10 (September): 139–62.

Sims, Christopher A. 1989. Modeling trends. Discussion Paper 22. Institute for Empirical Macroeconomics (Federal Reserve Bank of Minneapolis and University of Minnesota).

Walter, John R. 1989. Monetary aggregates: A user's guide. *Economic Review* 75 (January/February): 20–28. Federal Reserve Bank of Richmond.