Federal Reserve Bank
of Minneapolis

Spring 1989

Quarterly Review

Federal Reserve Bank of Minneapolis

Quarterly Review

Vol. 13, No. 2 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 Warren E. Weber, Kathleen S. Rolfe, and Inga Velde. Graphic design by Barbara Birr, Public Affairs Department.

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

Federal Reserve Bank of Minneapolis
Quarterly Review Spring 1989

Are Economic Forecasts Rational?

Michael P. Keane
Visitor
Research Department
Federal Reserve Bank of Minneapolis
and Assistant Professor
University of Minnesota

David E. Runkle*
Economist
Research Department
Federal Reserve Bank of Minneapolis
and Associate Director
Institute for Empirical Macroeconomics

For at least sixty years, economists have been concerned with how people forecast economic events. Many economists believe that people make the best economic forecasts they can, given the information available to them at the time. This assumption, called the *rational expectations hypothesis,* is of more than academic interest because it may imply that some government policies have little effect on people's behavior and therefore on the economy.

In this paper, we test whether price forecasts are rational. We do so because many of the controversial claims of rational expectations models depend on the rationality of those forecasts. For example, Sargent and Wallace's (1976) conclusion that predictable monetary policies would have no effect on output or employment is based partially on the assumption that people rationally forecast the price level.

Since the rationality of price forecasts is such an important issue in economics, our study is certainly not the first to test whether survey price forecasts are rational. In general, previous studies have found that price forecasts are not rational. (See Lovell 1986 for a review of the literature.) However, we believe that those studies used either inappropriate data or incorrect statistical methods. Most of the studies compared forecasters' price predictions with revised data on actual prices. And the ones that looked at data on individual forecasters assumed that the errors made by any one forecaster were not related to the errors made by any other forecaster.

Our data and statistical methods differ from those of the earlier studies: We use unrevised data on actual prices, and we assume that errors made by individual forecasters are correlated with the errors of other forecasters. As a result of these differences, our findings differ as well. We find strong evidence that price forecasts are indeed rational. Our study thus provides fresh support for the rational expectations hypothesis.

Testing for Rationality

Before we can test whether or not price forecasts are rational, we must first define what we mean by a *rational* forecast. To do this, we need to introduce some simple mathematical notation.

Suppose a forecaster is trying to predict the price level one period from now. We call the current period time t and the next period time $t+1$. The forecaster makes a prediction at t about the price level at $t+1$. We call the actual price level in the next period P_{t+1} and the forecast of that price level ${}_tP^f_{t+1}$. Such a forecast is called a *one-step-ahead* forecast because it predicts what will happen in the very next period.

If the forecast does not equal the actual price level in the next period, then it contains *forecast error.* That error, which we call ϵ_{t+1}, is simply the difference between the actual price and the forecast:

*Also Adjunct Associate Professor of Finance, University of Minnesota.

$$(1) \qquad \epsilon_{t+1} = P_{t+1} - {}_tP^f_{t+1}.$$

We say a forecast is *rational* if its forecast error is unpredictable, given what the forecaster knew when making the forecast.[1]

Having defined what we mean by a rational forecast, we can conduct a relatively simple test of forecast rationality. Suppose, for example, that we have a group of one-step-ahead price forecasts made by an economist. To test for rationality, we need to see whether any part of the economist's forecast error is predictable.

One way to test for predictable error is to show what happens if we try to predict the actual future price (P_{t+1}) using a constant, the forecaster's prediction (${}_tP^f_{t+1}$), and any other variable, say X_t, that the forecaster knew when making the forecast. Expressed mathematically, this test takes the form of the following linear regression equation:

$$(2) \qquad P_{t+1} = \alpha_0 + \alpha_1\, {}_tP^f_{t+1} + \alpha_2\, X_t + u_{t+1}.$$

For equation (2), we want to choose values for the coefficients α_0, α_1, and α_2 to predict P_{t+1} as accurately as possible. With given choices for those coefficients, u_{t+1} is the *residual*—the part of P_{t+1} that the equation does not explain. A commonly used criterion for choosing coefficients in an equation like this one is to choose values that minimize the sum of the squared residuals. We do this by using a statistical method called an *ordinary least squares* (*OLS*) regression.

Forecast rationality implies restrictions on the values of α_0, α_1, and α_2 in equation (2) because rationality means that forecast errors are not predictable. This can easily be seen by subtracting ${}_tP^f_{t+1}$ from both sides of equation (2). The result is

$$(3) \qquad P_{t+1} - {}_tP^f_{t+1} = \alpha_0 + (\alpha_1 - 1)\, {}_tP^f_{t+1} + \alpha_2\, X_t + u_{t+1}.$$

Note that the left side of equation (3) is simply the forecast error, ϵ_{t+1}. But rationality implies that nothing the forecaster knew at time t can be used to predict the forecast error. Therefore, in equation (3), the values of α_0, $\alpha_1 - 1$, and α_2 should not be significantly different from zero. Otherwise, some portion of the forecast error could be predicted, and the forecast would not be rational.

In testing for rationality, we have only one thing left to decide: what information the forecaster should have used in constructing a forecast—that is, what additional variables we could include as an X_t in equation (2) to improve the forecast. Certainly, in making a prediction, a forecaster should consider any relevant publicly available information. By relevant *public* information we mean any data that can be used to improve a forecast's accuracy. For example, if you were trying to rationally forecast which team will win a particular baseball game, you should take into account the teams' overall hitting and fielding strength, their recent win–loss records, the records of their pitchers, the effects of injured players, and the game's location. But you only need to know those data that could help you make a better forecast. You don't need to know recent home-game attendance or the salaries of the general managers, as long as those factors don't affect the game's outcome.

But publicly available information isn't the only data a forecaster should consider when making predictions. If a forecaster knows something that is not publicly known but is relevant to the outcome, this *private* information should be used to improve the forecast. If, to continue our example, you see the starting pitcher for one team become ill at a restaurant four hours before the game, you know this information will certainly affect the forecast, even if no one else knows about the illness.

So we can require forecasters to use both public and private information that is relevant to their forecasts. But there are two kinds of information about time t that we can't require forecasters to incorporate into their predictions of events in time $t+1$.

First, we can't expect forecasters to consider public announcements about the economy at time t if they are released *after* the forecast is made. Such delayed announcements are particularly relevant for economic forecasts because many leading indicators are reported late. For example, the trade deficit for February is announced almost 45 days after the month has ended, in mid-April. So we can't assume that a forecaster knows the February trade deficit when predicting the March deficit. In fact, at the time of the forecast, the forecaster only knows the January trade deficit.

The second kind of information we can't expect forecasters to know is other people's private information. For example, a forecaster won't know someone

[1]In a strict sense, *rationality* requires that the forecaster's subjective probability distribution be the same as the objective probability distribution of the forecasted variable. Thus, for a forecast to be rational, it is necessary but not sufficient that forecast errors be unpredictable. Note, too, that a forecast can still be rational even though it is incorrect.

else's forecast unless it has been made public. Or, a forecaster can't know with certainty the Federal Reserve's current monetary policy, which is a closely guarded secret.

So far we've set up the general framework of our tests of rationality by constructing a simple regression equation, by determining the restrictions that rationality imposes on the coefficients in that equation, and by assuming what types of information forecasters should incorporate into their predictions. Now we must find appropriate price forecast data to analyze, decide on the proper statistical methods to apply, and make accurate assumptions about what information forecasters knew when they made their predictions.

What Data Should We Use?

At first it seems fairly easy to determine what data we should use to test whether or not price forecasts are rational. We only need two different kinds of data: survey data on price forecasts and data on actual price levels. At least three surveys of price predictions exist, and apparently any one of them could provide us with our measure of the price forecast ${}_tP^f_{t+1}$. And the data for the actual price level P_{t+1} could be obtained from a commercial database. On closer analysis, however, choosing the appropriate data is more complicated than it initially seems.

The crucial issue in choosing a set of survey data to test for rational expectations is finding survey data that reflect the respondents' *true* expectations. If the forecasters' responses don't reflect their true expectations, then our tests of rationality will not be valid, so we won't learn much from them.

A frequent criticism of using survey forecasts to test the rationality of expectations is that many surveys collect data from academic economists or consumers—people with little to lose if their predictions are wrong or if they report them inaccurately.

Academic economists who are not professional forecasters have little incentive to make accurate forecasts. They get paid for their research and teaching—not their forecasting—so we might expect that the forecasts they report would not necessarily represent their true expectations. An example will clarify how this disparity could occur.

Suppose you're an economics professor who's busy writing a paper to be presented in three days. Someone phones you for your forecast of next quarter's interest rate on three-month Treasury bills. Because you're paid for writing papers, not making forecasts (and because you're very busy), you quickly tell the caller the first number that pops into your head—8 percent. Later in the day, while reading the *Wall Street Journal*, you see that the forward rate on three-month T-bills is 9 percent. Since bond prices rise when interest rates fall, it seems logical that you'd rush out and buy bonds. But you don't because, when you think about it, 9 percent seems reasonable. Thus, 8 percent is an *erroneous* measure of your true expectation because you don't act in the market as if it really were your expectation.[2]

The same argument applies to data from surveys of consumers' expectations, for consumers also have little to lose if their predictions are wrong (as long as they don't *act* on those predictions) or if they report them inaccurately.

The best way to ensure that survey data on forecasts reflect the true expectations of their forecasters is to use only the forecasts of people with an economic incentive to report their expectations accurately. Professional forecasters have such an incentive because they sell their reported expectations on the market and will lose business if their predictions are consistently less accurate than those of other professional forecasters.

When we looked for a survey whose respondents were limited to professional forecasters, we found only one—a survey of economic forecasters conducted since 1968 by the American Statistical Association (ASA) and the National Bureau of Economic Research (NBER). The survey's respondents produce quarterly forecasts of gross national product (GNP), its major components, and other economic indicators for each of several quarters ahead. Because these professionals report to the survey the same forecasts they sell on the market, their responses provide an accurate measure of their expectations. Thus, their forecasts should be free of the type of errors that arise when respondents have no incentive to report their true expectations.[3] We chose to examine the respondents' predictions of the price level, as measured by a price index known as the GNP deflator.

To the best of our knowledge, only studies by Zarnowitz (1984,1985) have previously used the ASA–NBER survey data to test for the rationality of price forecasts. All other previous studies that test the rationality of survey price forecasts have used data

[2]See Cukierman 1986 for a discussion of this problem.

[3]Interestingly, the ASA–NBER also conducts an annual survey of a much larger group of economists of whom three-fourths are occasional rather than professional forecasters. In that survey, Zarnowitz (1969, p. 15) found that "a number of the occasional forecasters submitted extreme and . . . rather unreasonable predictions."

from people who were not professional forecasters. They have used either the Livingston survey of economists' predictions about the consumer price index or the University of Michigan's survey of consumers' predictions about inflation.[4] Since neither of these survey groups is composed primarily of professional forecasters, we think the results of studies based on those two surveys are suspect.

The second data decision we had to make was which data to use for the actual price series—that is, for the values of P_{t+1}. This choice proved more difficult than it first seemed because all price series are continually revised. Every previous study of price forecast rationality has run tests using revised data on actual prices. Thus, those studies implicitly test whether the forecasters rationally predicted a revised price level using the revised price level for previous periods.[5] The forecasters, however, couldn't have known these data at the time they made their predictions. Therefore, using these data to test forecast rationality could lead to incorrect conclusions.

We have corrected for this second problem. Instead of using revised data, we compare the forecasters' one-quarter-ahead predictions of the GNP deflator with the actual deflator data initially released by the Commerce Department 45 days after the quarter's end. Our choice of unrevised data best reflects what the forecasters knew at the time they made their predictions.

What Statistical Methods Should We Use?

We still have to decide exactly how to test whether or not the price forecasts made by members of the ASA–NBER survey are rational. To test for rationality, we need to compare the forecast made by each individual forecaster with the actual price level initially reported in the next period.

We can do that by estimating the following equation:

$$(4) \qquad P_{t+1} = \alpha_0 + \alpha_1\, {}_tP^f_{i,t+1} + \alpha_2\, X_{i,t} + u_{i,t+1}$$

where ${}_tP^f_{i,t+1}$ is the one-period-ahead forecast of the price level made by forecaster i in period t and where $X_{i,t}$ is any other variable known to forecaster i at time t. For example, the variable $X_{i,t}$ could be forecaster i's own price level prediction made in the previous period—that is, ${}_{t-1}P^f_{i,t}$.

The simplest way to estimate equation (4) would be to use an OLS regression. That method is the one preferred by previous studies that use individual survey data on price forecasts. (See Figlewski and Wachtel 1981 and Zarnowitz 1984, 1985.) Statistical tests based on OLS regressions, however, assume that errors are uncorrelated across forecasters—an assumption that is unrealistic. If prices in a particular period rise faster than the average forecaster predicted, then the forecasters' prediction errors will, on average, be positive in that period. If prices rise slower than the average forecaster predicted, then the forecasters' prediction errors will, on average, be negative. In either case, the errors are correlated; that is, for any two forecasters, i and j, $\mathrm{cov}(u_{i,t+1}, u_{j,t+1}) > 0$. If the prediction errors are correlated across forecasters, then an additional observation on a forecaster does not really provide an additional degree of freedom. But the OLS method assumes that each additional observation *does* provide an additional degree of freedom. Thus, OLS incorrectly suggests there is less uncertainty than there actually is about the regression coefficients.[6]

To correct this problem, we developed a new statistical method that allows for the correlation in the forecasters' prediction errors. We treated the data as a panel and used a generalized method-of-moments estimator to estimate the coefficients, taking into account the correlation of errors among forecasters within a given time period. (For a technical explanation of our procedures and additional results, see Keane and Runkle 1988.) This procedure let us properly test the rational expectations hypothesis using the ASA–NBER data.

What Did the Forecasters Know?

Having decided what data and statistical methods to use, we still must determine what to assume the forecasters knew when they made their predictions. This decision is important because it helps us choose which economic variables to include in our tests of forecast rationality.

Time series models of economic data have shown that the most helpful information in predicting the future values of most economic variables is the current and past values of that variable. But the timing of data collection in the ASA–NBER survey makes it impossible for an individual forecaster to know the current

---

[4]Previous studies using the Livingston survey include Brown and Maital 1981, Carlson 1977, Figlewski and Wachtel 1981, Mullineaux 1978, Pearce 1979, and Pesando 1975. Studies using the Michigan survey data include Gramlich 1983 and Rich 1987.

[5]Zarnowitz (1985) is the only author who recognized that the use of revised data might affect his inference, but he believed that the effect would not be large.

[6]Zarnowitz (1985) recognized this problem with using the OLS method on panel survey data.

price level P_t at the time of the forecast ${}_tP^f_{i,t+1}$. We compare each forecaster's prediction, reported 20 days after the end of each quarter, with the GNP deflator announcement made 45 days after the end of the quarter. The relative timing of these events is shown in the accompanying chart. The forecaster predicts ${}_tP^f_{i,t+1}$ knowing P_{t-1} but not P_t, even though P_t would be very helpful for predicting P_{t+1}. Therefore, we can't use P_t to test the rationality of the forecaster's prediction because it was not available when the forecast was made. But we certainly can use P_{t-1} to test forecast rationality because that information was released two months before the forecaster's prediction ${}_tP^f_{i,t+1}$.

The chart shows other data the forecaster knew at the time of the prediction. For example, the forecaster's own past prediction of the price level, ${}_{t-1}P^f_{i,t}$, was certainly known, so we can use that previous forecast to test for rationality. The forecaster also knew $\epsilon_{i,t-1}$, that forecaster's own error made in predicting P_{t-1} at time $t-2$. If the forecaster learns from that mistake, the forecaster's past error should be uncorrelated with any of the forecaster's future errors.

The chart also shows that the forecaster's own current forecast error $\epsilon_{i,t}$ was unknown when the forecaster made the new prediction. That forecast error (equal to $P_t - {}_{t-1}P^f_{i,t}$) wasn't known because P_t hadn't yet

What Each Forecaster Knew When Predicting
Next Period's Price Level

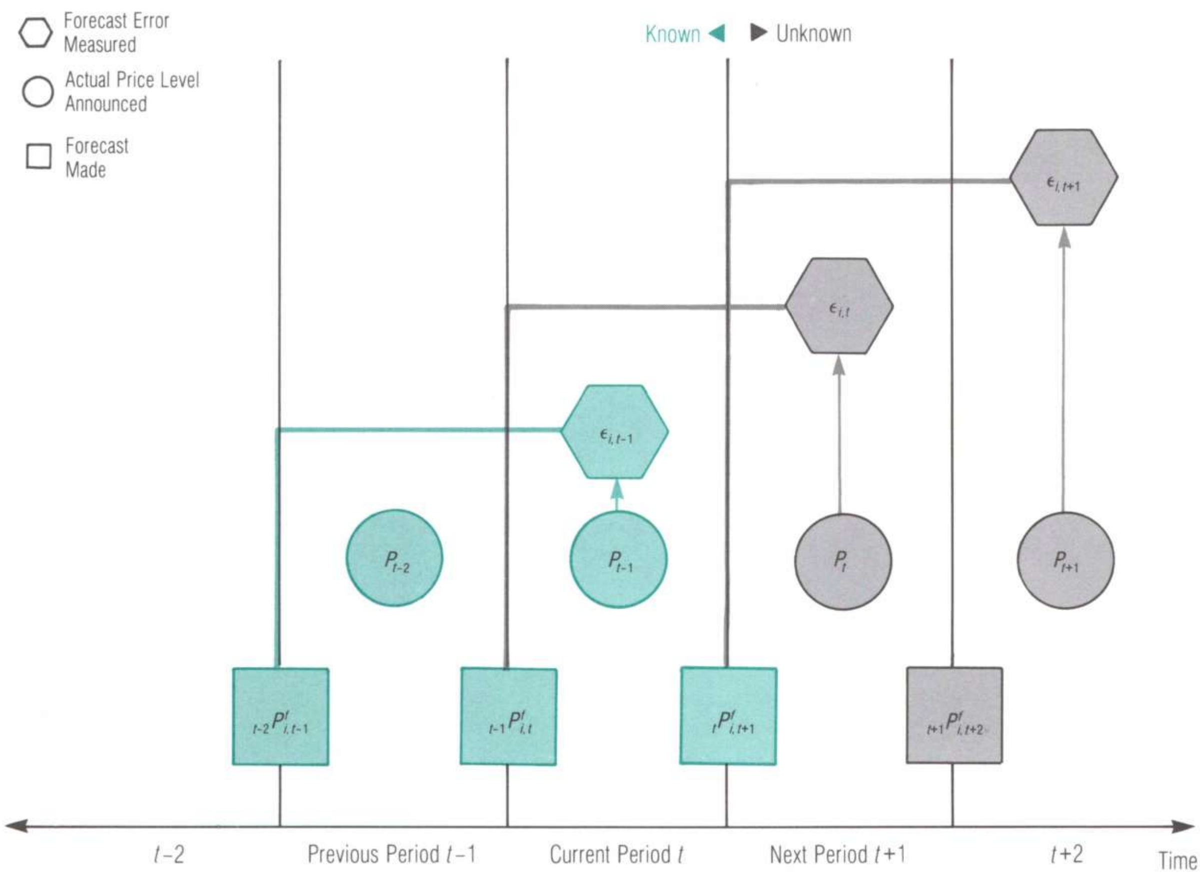

been released at the time of the forecaster's prediction. Since the forecaster was unaware of that error $\epsilon_{i,t}$, it will be correlated with $\epsilon_{i,t+1}$—the forecaster's error in the next period.[7] So, for any individual forecaster, $\text{cov}(u_{i,t+1}, u_{i,t}) \neq 0$.

But past prices, price forecasts, and forecast errors aren't the only variables that forecasters should take into account in making their predictions. We should expect that any other relevant publicly available information will also affect their forecasts. During the 1970s and 1980s, for example, changes in oil prices and the money supply influenced price growth. Economic forecasters were accused of systematically underestimating the effects of those variables on price behavior. If price forecasts are rational, then neither past values of the money supply nor past values of oil prices should help predict future prices once the forecasters' price predictions have been taken into account.

Test Results

Our tests of rationality were made using a sample of 1,613 forecasts made by 60 professional forecasters from the fourth quarter of 1968 to the third quarter of 1986. The results of our tests are shown in the accompanying table.

In the table, some of the reported regressions include fewer observations because data were missing for other variables. In each regression, the variable being forecasted is P_{t+1}. Each row of test results lists the variables used in a given regression. We use a statistic called the *chi-square* (χ^2) *test* to measure whether the coefficients in a regression differ in any significant way from what they are expected to be.[8] Specifically, we use this test to see whether α_0 and α_2 differ from zero and whether α_1 differs from one.

Our first test examines whether the forecasters' average prediction error is significantly different from zero. If the forecasts are rational, there should be no important difference. In test (1), the difference between the coefficients' values and what we expected them to be would be significant at the 5 percent level only if the χ^2-statistic were larger than 5.99. Since the value of that statistic is only 0.900, we cannot reject the rationality hypothesis.

We conducted single-variable tests to see whether any variable X_t that the forecasters knew when making their predictions could be used to improve their forecasts. If an improvement showed up as a result of the added variable, forecast rationality would be refuted. For each of these tests (2)–(7), we show a χ^2-statistic that tests the restrictions of rationality: $\alpha_0 = 0$, $\alpha_1 = 1$, and $\alpha_2 = 0$. The estimated coefficients would be significantly different from those values at the 5 percent level if the χ^2-statistic were to exceed 7.81.

We concluded earlier that the forecasters should have known P_{t-1} when they made their forecasts. Test (2) tries to find out whether the forecasters correctly used that information in predicting P_{t+1}. The coefficient for P_{t-1} is not significantly different from zero, and the value of the χ^2-statistic is low enough to indicate that knowledge of P_{t-1} cannot be used to predict the price forecast errors.

Test (3) tries to find whether each forecaster's one-step-ahead prediction from the previous period, ${}_{t-1}P^f_{i,t}$, was correctly incorporated into that forecaster's current prediction. Neither α_2 nor the χ^2-statistic is significant, so we cannot reject rationality.

Test (4) tries to see whether the forecasters used the information in their past prediction errors to improve their predictions—that is, whether they learned from past mistakes. Again, α_2 is not significantly different from zero, and the χ^2-statistic is not significant. Thus, this test also supports forecast rationality.

So far, we have not considered whether forecasters correctly used information about other variables—such as oil prices and the money supply—in making their price forecasts. We include these variables in tests (5) and (6). Test (5) shows the results for the previous quarter's nominal crude oil price, $P_{O,t-1}$, and test (6) gives the results for the previous quarter's money supply, $M1_{t-1}$. In each case, the coefficient for the additional variable is not significantly different from zero. The χ^2-statistics are also insignificant. These test results suggest that professional forecasters in this survey couldn't improve their forecasts by making better use of data on oil prices or the money supply.

Test (7) examines whether P_t could have been used to improve the predictions of the forecasters. The regression coefficient for that variable is significantly different from zero and the χ^2-statistic shows that we must reject the rationality restrictions $\alpha_0 = 0$, $\alpha_1 = 1$, and $\alpha_2 = 0$. But these results are not a rejection of

---

[7] In testing forecast rationality, one common fallacy is to assume that if forecast errors are correlated over time, then the forecast is not rational. That would be true if the forecaster knew the current forecast error when making the new prediction. In our data, that is not true, and a forecaster's predictions could be rational even if $\epsilon_{i,t+1}$ is correlated with $\epsilon_{i,t}$. However, if $\epsilon_{i,t+1}$ is correlated with $\epsilon_{i,t-1}$, then the forecast would not be rational because $\epsilon_{i,t-1}$ was known when the prediction ${}_tP^f_{i,t+1}$ was made. We assume that the forecast errors are moving average of order 1 for all our reported results.

[8] We use χ^2-tests rather than F-tests because the forecast errors would have to be normally distributed for the F-tests to be valid. The χ^2-tests are asymptotically valid.

Results of Tests for the Rationality of Price Forecasts

Test Using	Constant α_0	Coefficient on Forecast α_1	Coefficient on Other Variable α_2	χ^2-Statistic* [Significance Level]	Regressors	Number of Observations**
(1) Forecast Only	5.853	.997	—	.900	$1, {}_tP^f_{i,t+1}$	1,613
	(6.227)	(.005)	—	[.638]		
Forecast Plus						
(2) Previous Price Level	7.243	.878	.120	4.012	$1, {}_tP^f_{i,t+1}, P_{t-1}$	1,079
	(7.646)	(.068)	(.066)	[.260]		
(3) Previous Forecast	7.471	.911	.086	2.331	$1, {}_tP^f_{i,t+1}, {}_{t-1}P^f_{i,t}$	1,119
	(7.132)	(.068)	(.068)	[.507]		
(4) Past Forecast Error	7.078	.996	-.002	.875	$1, {}_tP^f_{i,t+1}, \epsilon_{i,t-1}$	728
	(7.692)	(.004)	(.111)	[.832]		
(5) Previous Oil Price	3.092	.999	-.001	1.057	$1, {}_tP^f_{i,t+1}, P_{0,t-1}$	1,613
	(7.549)	(.006)	(.002)	[.787]		
(6) Previous Money Supply	4.481	1.000	-.001	1.964	$1, {}_tP^f_{i,t+1}, M1_{t-1}$	1,613
	(6.011)	(.004)	(.001)	[.580]		
(7) Current Price Level	8.602	.709	.291	10.760	$1, {}_tP^f_{i,t+1}, P_t$	1,129
	(7.272)	(.097)	(.098)	[.013]		
(8) Variables of Tests (2)–(6)	21.613	.761	—	10.960	$1, {}_tP^f_{i,t+1}, P_{t-1}, {}_{t-1}P^f_{i,t}, \epsilon_{i,t-1}, P_{0,t-1}, M1_{t-1}$	580
	(11.374)	(.076)	—	[.140]		

Note: The sample includes 1,613 quarterly forecasts of the GNP deflator made by 60 professional forecasters from 1968:4 to 1986:3. Standard errors are shown in parentheses under the coefficients.

*The χ^2-statistic tests the null hypothesis that $\alpha_1 = 1$ and all other coefficients equal zero.

**The number of observations varies because not all observations include the relevant variables. Single-variable tests for rationality are included because they contain more observations.

Sources of basic data: ASA–NBER, U.S. Department of Commerce.

rationality because the forecasters didn't know P_t when they made their forecasts. In fact, predictions are often inaccurate because forecasters don't know the current values of the variables they are forecasting. Thus, understanding *when* forecasters knew P_t is crucial for correctly testing forecast rationality.

Although we have tested whether the forecasters accounted for the individual effects of each of the previously mentioned variables on future prices, a proper test of rationality must check whether all these variables, together, can be used to make a more accurate price forecast. We perform this multivariate test in test (8). We include the following regressors: a constant, ${}_tP^f_{i,t+1}$, P_{t-1}, ${}_{t-1}P^f_{i,t}$, $\epsilon_{i,t-1}$, $P_{O,t-1}$, and $M1_{t-1}$. The hypothesis of rationality implies that the coefficient on ${}_tP^f_{i,t+1}$ should be equal to one and all the other coefficients should be equal to zero. The χ^2-statistic shows that the hypothesis cannot be rejected because it is smaller than the critical value of 14.07.

In short, the test results reported in the table offer strong support for the hypothesis that the GNP deflator forecasts of these professional forecasters are rational. The average forecast error is not significantly different from zero, and data that forecasters knew at the time they made their forecasts cannot be used to improve their predictions.

Conclusion

The evidence in this paper suggests that the price predictions of professional forecasters are rational because their forecast errors are unpredictable. This result is surprising because previous studies have found that price forecasts are not rational.

Why are our results so different? They differ because we use better data and statistical methods. If we either had used revised data or assumed that the prediction errors of different forecasters were uncorrelated, our tests would have rejected rationality. (See Keane and Runkle 1988 for these and other results.) But we have shown that either of those choices would be mistaken.

Our tests show we cannot reject the hypothesis that forecasters optimally use the information they have when they make their forecasts. Of course, we still do not have evidence about whether or not the expectations of other people besides professional forecasters are rational. Even so, the evidence in this study suggests we should take the rational expectations hypothesis seriously.

References

Brown, Bryan W., and Maital, Shlomo. 1981. What do economists know? An empirical study of experts' expectations. *Econometrica* 49 (March): 491–504.

Carlson, John A. 1977. A study of price forecasts. *Annals of Economic and Social Measurement* 6 (Winter): 27–56.

Cukierman, Alex. 1986. Measuring inflationary expectations: A review essay. *Journal of Monetary Economics* 17 (March): 315–24.

Figlewski, Stephen, and Wachtel, Paul. 1981. The formation of inflationary expectations. *Review of Economics and Statistics* 63 (February): 1–10.

Gramlich, Edward M. 1983. Models of inflation expectations formation: A comparison of household and economist forecasts. *Journal of Money, Credit, and Banking* 15 (May): 135–73.

Keane, Michael P., and Runkle, David E. 1988. Testing the rationality of price forecasts: New evidence from panel data. Unpublished manuscript, Research Department, Federal Reserve Bank of Minneapolis.

Lovell, Michael C. 1986. Tests of the rational expectations hypothesis. *American Economic Review* 76 (March): 110–24.

Mullineaux, Donald J. 1978. On testing for rationality: Another look at the Livingston price expectations data. *Journal of Political Economy* 82, Part 1 (April): 329–36.

Pearce, Douglas K. 1979. Comparing survey and rational measures of expected inflation: Forecast performance and interest rate effects. *Journal of Money, Credit, and Banking* 11 (November): 447–56.

Pesando, James E. 1975. A note on the rationality of the Livingston price expectations. *Journal of Political Economy* 83 (August): 849–58.

Rich, Robert W. 1987. The rationality of price and inflation forecasts: New evidence on the Livingston and the SRC survey series. Chapter 1 of Ph.D. dissertation, Brown University.

Sargent, Thomas J., and Wallace, Neil. 1976. Rational expectations and the theory of economic policy. *Journal of Monetary Economics* 2 (April): 169–83.

Zarnowitz, Victor. 1969. The new ASA–NBER survey of forecasts by economic statisticians. *American Statistician* 23 (February): 12–16.

__________. 1984. Business cycle analysis and expectational survey data. Working Paper 1378. National Bureau of Economic Research.

__________. 1985. Rational expectations and macroeconomic forecasts. *Journal of Business and Economic Statistics* 3 (October): 293–311.