COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 774

FORECASTING EFFICIENCY:

CONCEPTS AND APPLICATIONS

by

W. NORDHAUS

SEPTEMBER, 1985

W. Nordhaus
September 1985

FORECASTING EFFICIENCY:
CONCEPTS AND APPLICATIONS*

In numerous areas of economics, expectations and forecasts play a central role. The importance of forecasting is obvious. Corporations make investment decisions based on future output forecasts. Governments plan budgets, along with macroeconomic policies, based on forecasts of future economic conditions. Individuals make savings or educational decisions based on forecasts of future incomes and needs or labor market conditions.

Expectations have also been of great significance in economics, and, under the influence of the rational-expectations revolution, theories involving expectation have recently become enormously controversial as well. This recent significance arises because, if expectations are rational and if markets clear quickly, then certain kinds of anticipated macroeconomic policies will have no real impact upon the economy.

Given the heightened importance attached to forecasts and expectations, it is natural to ask how accurate and adequately constructed are forecasts of future events? Are forecasts wildly biased, erratic, and unreliable? Or do they tend to use existing information in a reasonable efficient manner?

---

*This is a revised version of a paper presented at the Third International Energy Workshop, Laxenburg, Austria, June 1985. This work was supported by grants from the National Science Foundation and the Sloan Foundation.

The question of forecasting accuracy is, of course, one that has been the subject of numerous investigations over the last two decades. The present study contributes to this line of research in two ways. First, we introduce a new concept, called ''forecast efficiency,'' that measures the extent to which information is incorporated into forecasts. This concept is closely related to concepts of efficiency used in the analysis of stock and other financial markets. The paper proves two readily testable propositions about efficient forecasts. Second, the empirical part of the study examines forecast efficiency by looking at *forecast revisions* (''fixed-horizon forecasts''), rather than a series of forecasts of different events (''rolling-horizon forecasts'') as is the case for most studies of forecasting. This new approach to estimation in certain circumstances will provide a more powerful test of forecast efficiency. A number of fixed-horizon forecasts are collected and these are tested for forecast efficiency.

A. Efficient Forecasts: Theory

1. Efficient Forecasts

Assume that a forecaster is asked to provide a forecast that is reliable. Operationally, I assume that ''reliability'' is measured by the *ex post* root mean square error of the forecast.

For example, say an energy analyst is asked to forecast oil prices in the year $T = 1985$; call this q_T. The forecaster will make a forecast at date t, call this forecast ${}_tq_T$. The forecast error, ${}_tu_T$, is

(1) $${}_tu_T = {}_tq_T - q_T$$

Under the concept of reliability used here, the analyst desires to minimize a loss function which is the squared ex-post error:

$$(2) \qquad L({}_t u_T) = ({}_t u_T)^2 \; .$$

Next introduce the concept of <u>forecast efficiency</u>. By analogy with the theory of finance (see Fama [1976]), call a forecast <u>efficient</u> if it minimizes the loss function in (2) subject to available information. It will be useful to distinguish strong from weak efficiency in the empirical tests that follow. The definition of a <u>strongly efficient forecast</u> is one that minimizes the loss function in (2) when all information available at time t (I_t being the set of such information) is used. That is:

(3) <u>Strong efficiency</u>. A forecast is strongly efficient if $E\{({}_t u_T)^2 | I_t\}$ is minimized, where I_t is all information available at time t .

Strong efficiency is closely related to the concept of rational expectations (RE). RE is usually taken to mean that expectations are formed by taking into account all relevant available information, including appropriate knowledge about the structure of the economy. Because the information content in both RE and strongly efficient forecasts is so high, it has proven very difficult to test strong efficiency in practice, for tests involve complete knowledge about the structure of the economy and access to private data that is not available to most econometricians. Put differently, to test strong forecast efficiency requires knowing the exact relation between the information set I_t and the forecast ${}_t y_T$; otherwise it is not possible to determine whether the information set has been efficiently used.

Because of the practical limitations on testing strong efficiency, we are forced to turn to other tests. One of the most attractive is one that I label <u>weak efficiency</u>. Weak efficiency is defined as a forecast that efficiently incorporates information about past forecasts. Put differently, weak efficiency hold where the information set J_t is all past forecasts, i.e., $I_t \supset J_t = \{{}_tq_T, {}_{t-1}q_T, \ldots\}$ = set of all past forecasts. Formally:

(4) <u>Weak efficiency</u>. A forecast is weakly efficient if it minimizes $E_t\{({}_tu_T)^2 | J_T\}$, where J_T is the set of of all past forecasts.

Weak efficiency is an attractive concept for two reasons. First, it is likely that past forecasts play a very important role is determining current forecasts. Forecasters tend to have a certain consistency (stickiness?) in their views of the world, so that recent forecasts will go far in explaining current forecasts. Second, of all variables that seem plausible candidates for including in a forecaster's information set, surely the forecaster's own views must rate quite high. Hence we can be reasonably sure that we are not including a variable in the forecaster's data set that he or she was not aware of! Finally, as we will next show, there is a very simple and powerful set of tests that can be run to investigate whether forecasts are weakly efficient.

2. <u>Tests for Weak Efficiency</u>

Weak inefficiency of forecasts is relatively easy to detect. The following procedure provides the major conditions on weakly efficient forecasts.

Define ${}_t v_T$ as the <u>forecast</u> <u>revision</u> from period $t-1$ to period t. Then

(5) $${}_t q_T = {}_0 q_T + {}_1 v_T + \dots + {}_t v_T .$$

We assume that the truth is known on date T, so ${}_T q_T = q_T$. The (*ex-post*) forecast error is

(6) $${}_t u_T = {}_t q_T - q_T = {}_{t+1} v_T + \dots + {}_T v_T .$$

Under the conditions given above, techniques that minimize the squared error of forecasts will be unbiased (as, for example, by the Gauss-Markov theorem) and forecast errors will therefore have zero expected value. Hence

(7) $$E({}_t u_T \mid J_t) = 0$$

which implies

(8) $$E({}_t u_T \mid {}_t v_T, {}_{t-1} v_T, \dots, {}_0 v_T) = 0 , \quad \text{all } t .$$

From (6), (8) implies a series of expectations:

(9.T) $$E({}_T v_T \mid {}_{T-1} v_T, \dots, {}_0 v_T) = 0$$

(9.T-1) $$E({}_{T-1} v_T + {}_T v_T \mid {}_{T-2} v_T, \dots, {}_0 v_T) = 0$$

(9.t) $$E({}_t v_T + \dots + {}_T v_T \mid {}_{t-1} v_T, \dots, {}_0 v_T) = 0$$

Working recursively down this list, (9.t-1) implies the following for (9.t):

(10.T) $E({}_T v_T \mid {}_{T-1}v_T, \ldots, {}_0 v_T) = 0$

(10.T-1) $E({}_{T-1}v_T \mid {}_{T-2}v_T, \ldots, {}_0 v_T) = 0$

$\vdots$

(10.t) $E({}_t v_T \mid {}_{t-1}v_T, \ldots, {}_0 v_T) = 0$.

In words:

Proposition 1. Equations (9.t) show that the forecast error at date t is independent of all forecast revisions up to time $(t-1)$.

Proposition 2. Equations (10.t) show that the forecast revision at date t is independent of all forecast revisions up to time $(t-1)$. [Nordhaus and Durlauf [1984] derived and tested equation sets (10.t).]

Another way of understanding equations (9) and (10) is to say that forecasts should look like a random walk (or, technically, a martingale). If forecast revisions are unpredictable, then in any period, the forecast of oil prices or GNP or whatever should wander around trendlessly. If, instead, the forecast of, say, 1995 oil prices (for forecasts made in 1980, 1981, 1982, 1983,...) moves steadily up or down, then (subject to sampling error) we can say that forecasters have not efficiently incorporated past information. Figure 1 shows a simple example of a jagged efficient forecast and a smooth but inefficient forecast.

The tests outlined in Propositions 1 and 2 are particularly useful for evaluating forecasts because they involve observable phenomena, forecast errors and forecast revisions; they need no knowledge of the structure of the problem being investigated and can apply to energy, GNP forecasting, political polling, and a wide variety of phenomena.

FIGURE 1

EFFICIENT AND INEFFICIENT FORECASTS

Efficient forecasts appear jagged because they incorporate all news quickly. Inefficient forecasts appear smoother and more consistent, for they left the news seep in slowly. In examples shown, the inefficient forecast allows news to seep in at a rate of 10 percent per period. Shaded area represents the conditional absolute forecast errors at every point of forecast period.

(In this simulation, ε_t were drawn from a uniform distribution over the range [-1/2, +1/2].)

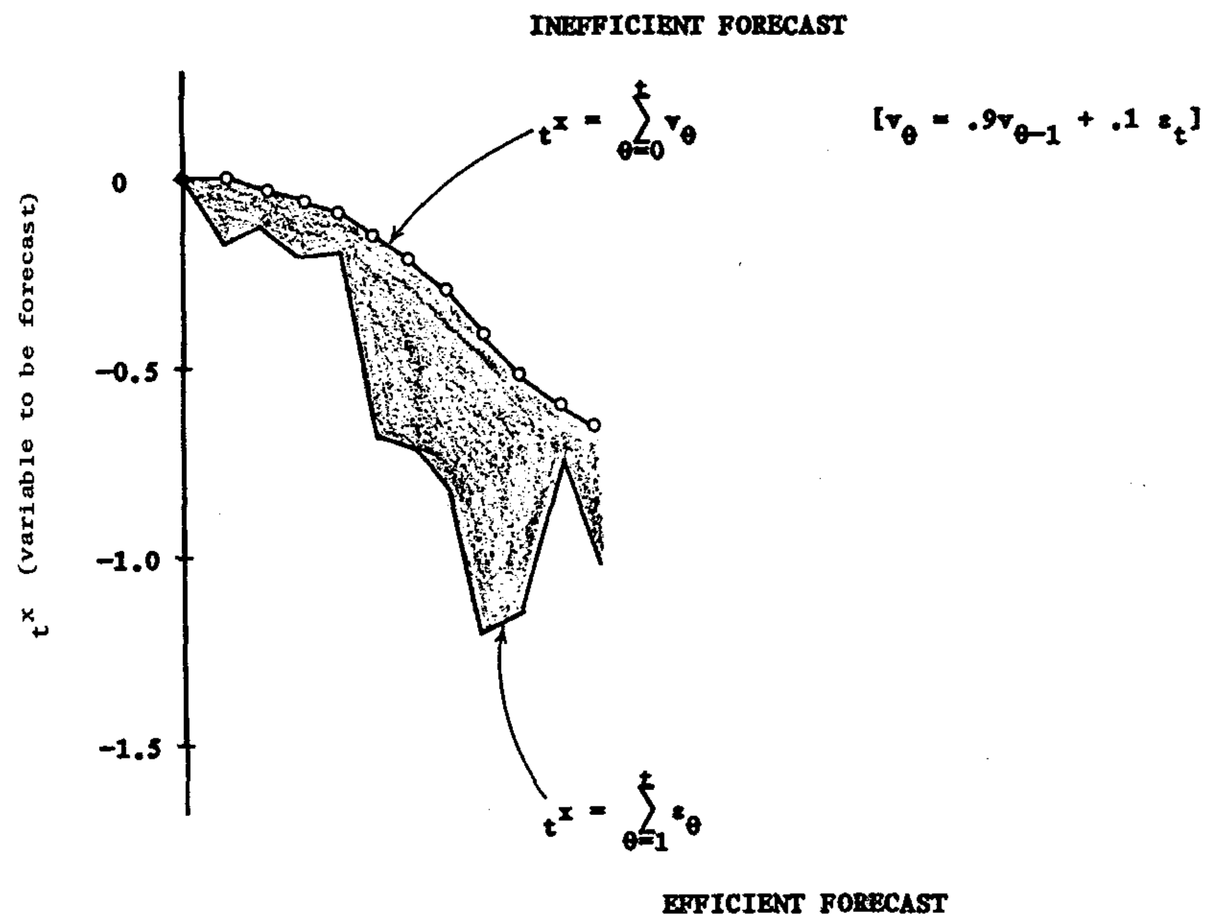

In order to make the tests operational, we need to specify the nature of the process generating forecast errors (or forecast revisions). Equations (10) show that forecast revisions have expected values of zero (conditional on past forecast revisions). In addition, in our statistical tests, we assume:

(11) $\{{}_t v_T\}$ are normally distributed.

The assumption of normality is critical for testing weak efficiency, although many studies indicate that modest degrees of non-normality do not distort test statistics markedly.

B. Empirical Tests

1. DRI Forecasts of Oil Prices

Data Resources, Inc. (DRI) has for a number of years made forecasts of future oil prices. Figure 2 shows the forecasts used in the present tests. See DRI Review [various years]. Each line represents the forecast of a single event, where the forecast is revised from month to month.

We can test for weak efficiency of these forecasts by asking whether the current forecast revision can be predicted using past forecast revisions. In an earlier study with Durlauf, we uncovered a tendency for first-order serial correlation of revision: a positive forecast revision tends to be followed by a positive forecast revision, and vice versa (see Nordhaus and Durlauf [1984]). A first test, then, is to examine whether such a pattern shows up in oil price forecasts.

Table 1 reports on the first-order forecast revisions. By Proposition 1, these coefficients should be 0. All are positive, although only 1 is

FIGURE 2
Forecasts of Future Oil Prices
by DRI for year 1982 through 1986
(Note that right-hand element for each price series is the actual price)

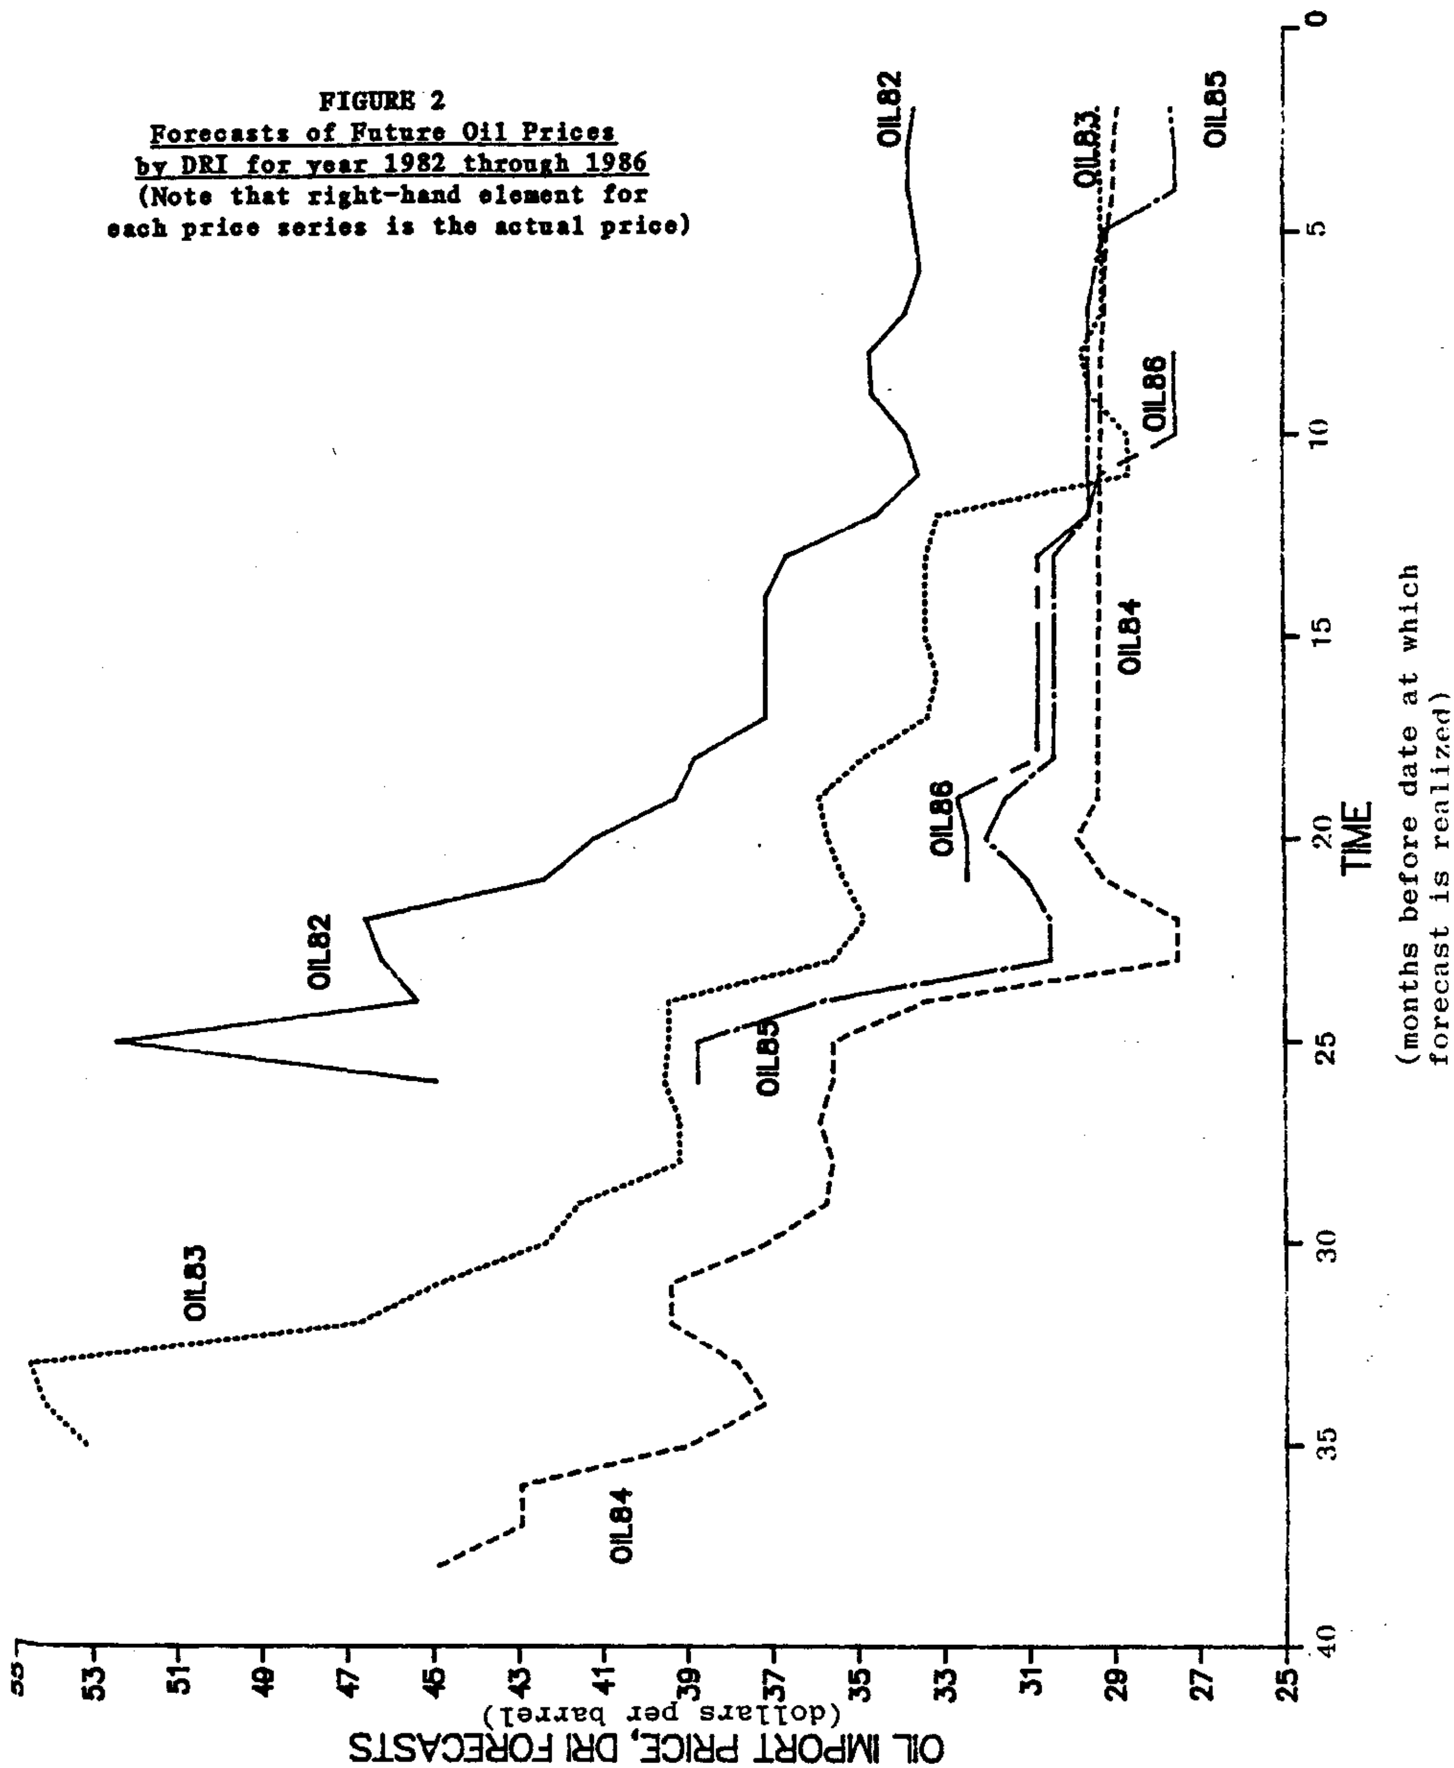

TABLE 1

First-order Association of Revisions*

Target Year of Forecast	Coefficient	t-Statistic	p**	Number of Observations
1982	.31	1.4	< .10	22
1983	.20	1.1	< .20	32
1984	1.28	7.6	< .005	37
1985	.14	0.6	< .30	21
1986	.16	0.4		18

*Reported coefficient is that in a regression

$$_{t}v_{T} = b[_{t-1}v_{T}]$$

where $_{t}v_{T}$ = DRI forecast of U.S. imported oil price for year T, where T = 1982 through 1986.

**Probability that value of one-tailed t-statistic could arise if true value were zero.

TABLE 2

Future Forecast Revisions as Related to Cumulative Revision*

Target Year of Forecast	Coefficient	t-Statistic	p**
1982	1.59	2.7	$< .01$
1983	1.68	2.0	$< .05$
1984	.04	0.1	
1985	.02	0.0	
1986	.44	0.4	

*Estimated equation is

$$\left(\sum_{\theta=t}^{T} {}_{\theta}v_T\right) = b\left[\sum_{\theta=0}^{t-1} \left({}_{\theta}v_T\right)\right]$$

where ${}_{\theta}v_T$ is defined in footnote to Table 1.

**See Table 1.

significant at high levels of confidence.

Table 2 reports the cumulative revision from date t on as a function of the total forecast revisions up to date t. These show a pattern of positive signs, indicating that a past history of positive revisions tends to be followed by further positive revisions. The wide variation in coefficients suggests, however, considerable instability in the structure of forecast revisions.

A Measure of Forecast Inefficiency

Evidence presented here suggests that forecasters consistently display inefficiency in that they fail to incorporate all the information from their own past forecasts. An important question is, how large is the degree of inefficiency? Table 3 shows the ratio of the actual forecast error to an *ex post* weakly efficient forecast error. An ''*ex post* weakly efficient forecast error'' is one that, for each year, corrects the forecast for the estimated inefficiency. It is therefore the absolute best that could be done; indeed, this correction is far more than we could hope for given the instability in the coefficients apparent from Tables 1 and 2.

In general, the possible forecast improvements are modest. The average is a 27 percent improvement for the total period forecasting error — enough to worry about but probably not enough to write home about.

2. Forecasts of Nuclear Capacity

While the DRI forecasts of oil prices show but a modest degree of forecasting inefficiency, other forecasts, particularly those of quantities made by non-econometrically-based groups, show much high levels of inefficiency. As an example, we show forecasts made by the OECD of the

TABLE 3

Forecast Error Relative to Weakly Efficient Forecast

Oil Price for Target Year	Ratio of forecast errors: One-period Forecast Revision	Total Forecast
1982	1.04	1.31
1983	1.28	1.12
1984	1.89	1.03
1985	1.28	1.23
1986	1.05	1.43

*This table shows the ratio of actual forecast error to the ideal weakly efficient forecast error. It is calculated as $(1 - \bar{R}^2)^{-1/2}$. Due to sampling variation, this will generally be greater than unity.

TABLE 4

First-order Coefficient of Forecast Inefficiency, Nuclear Power

Forecasted Variable:

Year-end Nuclear Generating Capacity, 1985 for:	Coefficient	Standard Error	p*
OECD	.88	.28*	$<$ .005
US	.50	.35	$<$ .10
OECD except US	.74	.33*	$<$ .025

Equation estimated is ${}_{1985}v_t = a[{}_{1985}v_{t-1}]$, where ${}_{1985}v_t$ is the revision of the forecast series shown in Figure 3.

*See Table 1.

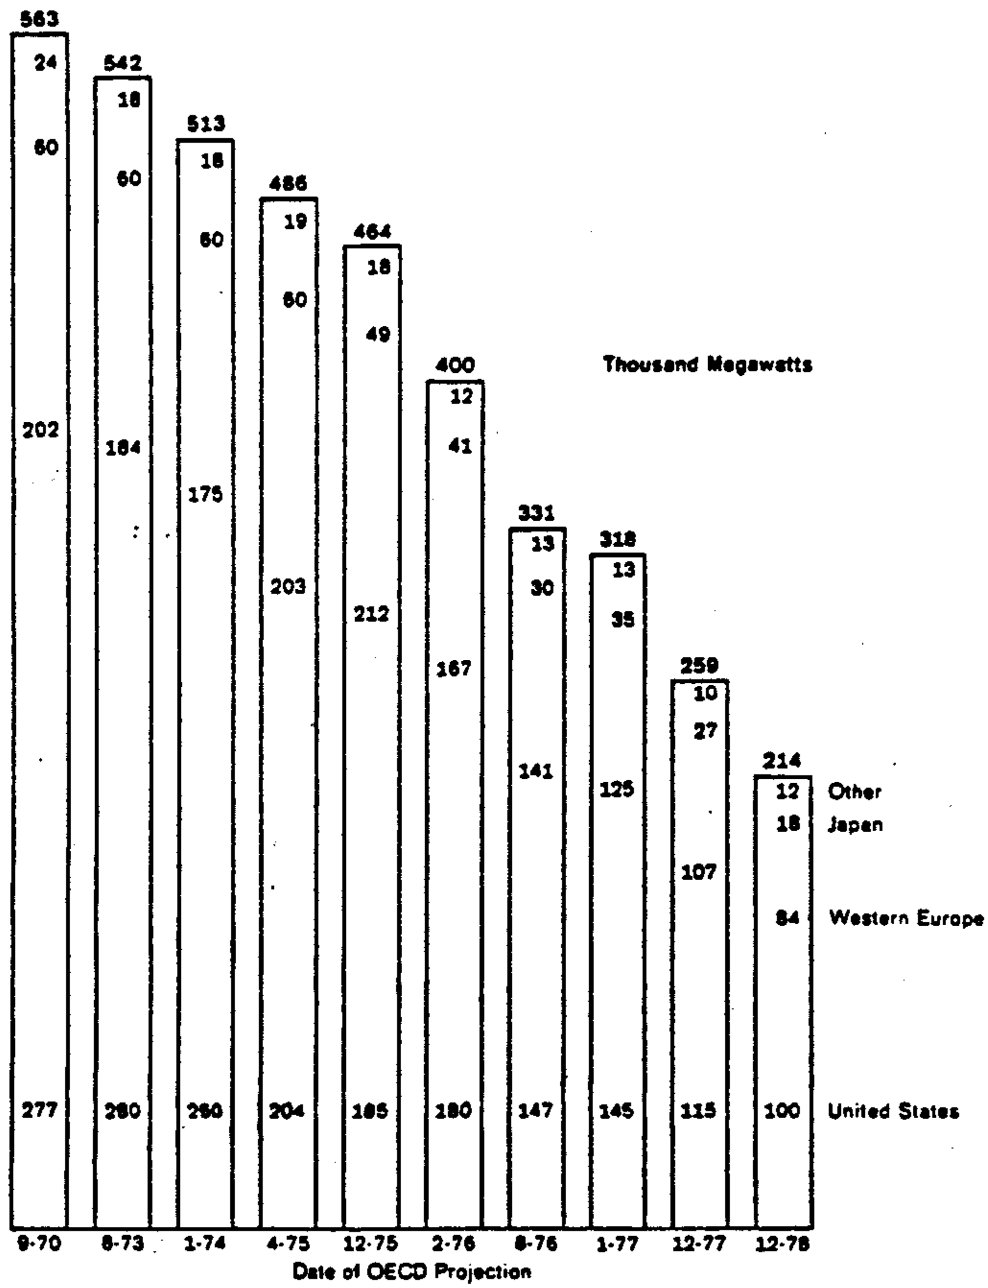

FIGURE 3.

Past Projections of Year-end 1985 Nuclear Generating Capacity

Made by OECD

Source: Central Intelligence Agency [1980].

installed capacity of nuclear power for advanced industrial countries. Figure 3 shows a collection of forecasts, while Table 4 presents estimates of the autoregressive structure of the forecast revisions. Although the size of the sample is quite small, the large and significant coefficient indicates that the nuclear-power forecasts were very inefficiently constructed.

3. Energy Forecasts by European Community

The European Community has made energy forecasts for the European region since 1973, where the objects of forecast were energy consumption for the years 1985, 1990, and 2030. Table 5 shows a series of tests of forecast efficiency. These show extremely high degrees of first-order association, much in the way that the nuclear-power forecasts of the OECD do.

4. Real GNP Forecasts

A final example, based on a study by Nordhaus and Durlauf, examines monthly forecasts for real GNP growth for 5 years (1978 to 1982) for four major U.S. forecasting services and for the ''Eggert consensus,'' which is an average of 30 individual forecasters. Figure 4 shows the time series of the consensus forecast for 1982 (which surely does not look like a random walk).

Table 6 presents the statistics on a single regression of the following form:

$$(9) \qquad {}_t v_T = \alpha({}_{t-1}v_T) + {}_t \varepsilon_T ,$$

where the estimates in Table 6 are the estimates of α , with standard errors in parentheses. The coefficient in the upper left [.81 with a

TABLE 5

First-order Coefficient of Forecast Inefficiency
European Energy Consumption

Terminal Year of Forecast	Coefficient	Standard Error	p*	Number of Observations
1985	.36	.21	< .04	8
1990	.28	.36	< .30	7
2030	1.65	.30	< .025	2

Equation estimated is ${}_{T}v_{t} = a[{}_{T}v_{t-1}]$, where ${}_{T}v_{t}$ is the logarithm of the forecast of European energy consumption made by the European Community.

Data Source: Haerter [1985], pp. 207, 209, 212.

*See Table 1.

TABLE 6

Coefficient of Current with One-period Lagged Forecast Revisions for Real GNP Growth*

Forecaster	1982	1981	1980	1979	1978
Eggers Consensus	.81 (.13)	.52 (.18)	.57 (.19)	.49 (.20)	.11 (.26)
DRI	.30 (.20)	.003 (.20)	.56 (.18)	.30 (.20)	.27 (.23)
UCLA	.15 (.21)	.16 (.20)	.24 (.23)	-.06 (.22)	1.90 (.25)
Wharton	.30 (.20)	.12 (.27)	.07 (.22)	.26 (.22)	.24 (.23)
Chase	.57 (.18)	.38 (.19)	.33 (.18)	.17 (.22)	.17 (.23)

*Table shows results for a regression of forecast adjustments on lagged forecast adjustments; i.e., ${}_t v_T = \alpha\{{}_{t-1}v_T\}$, where ${}_t v_T$ is the change in the forecast of real GNP growth for year T over year T-1 , where the forecast was made at date t . Table reports the coefficient and standard error on α with the estimated standard error on α in parentheses.

Assuming that the coefficients are distributed as the t-statistic and are independently generated, the expected and actual distribution of coefficients is as follows:

t-statistic	Probability level (p)	Expected	Actual
Less than 1.36	Greater than 20%	20 cases	9 cases
Greater than 1.36	Less than 20	5	16
Greater than 1.78	Less than 10	2.5	12
Greater than 2.68	Less than 1	.25	9
Greater than 3.06	Less than .5	.125	4

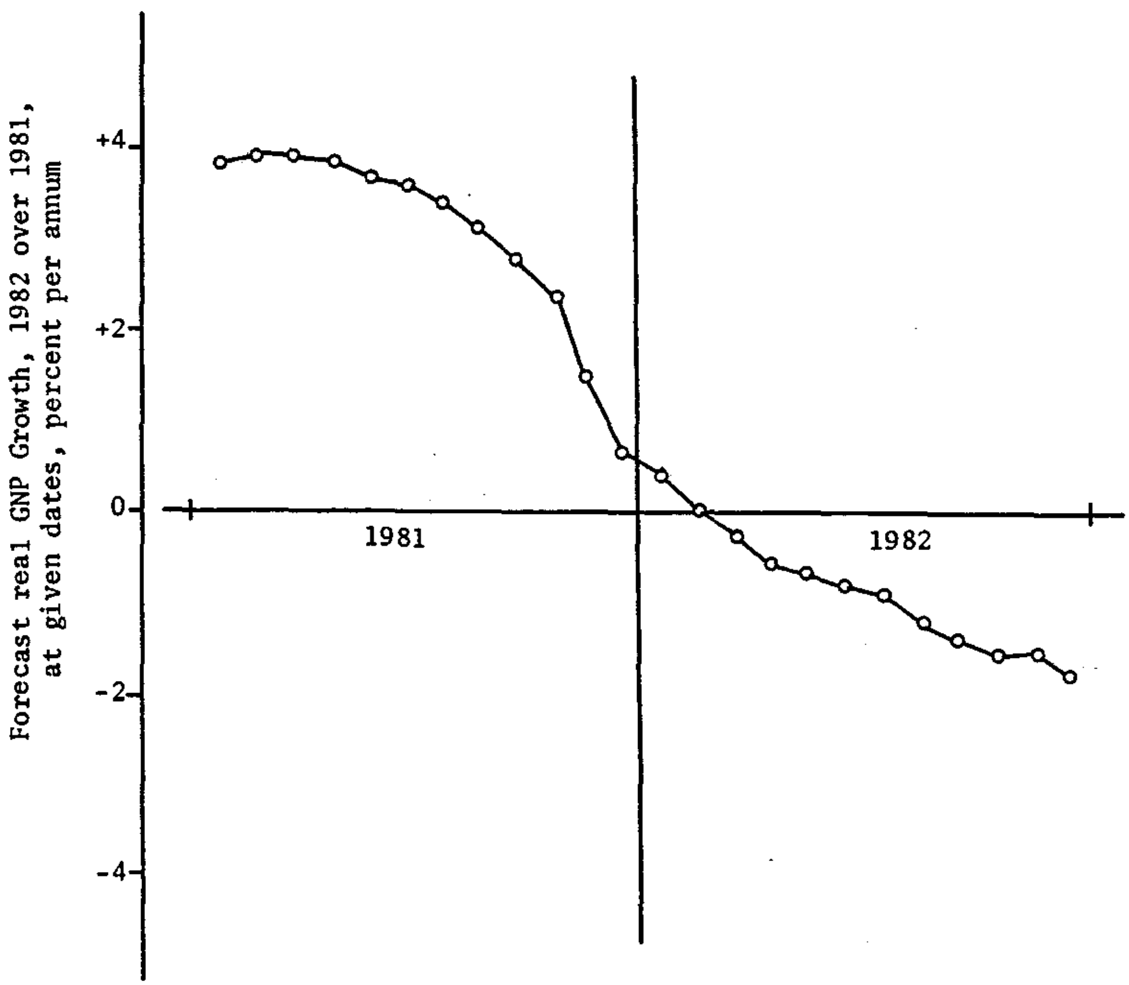

FIGURE 4

Forecast of Real GNP Growth for 1982 over 1981 for Eggers Consensus forecast, for periods from January 1981 to December 1982

standard error of .13] confirms the impression from Figure 4 that forecast revisions are highly correlated for the consensus.

The pattern of results in Table 6 shows in a striking way that forecast revisions tend to be smoothed. If equation (9) is a stable structural equation, then the mean reaction time, M, is given by

$$M = \frac{\alpha}{1-\alpha} .$$

Hence, for the consensus the real amount of time required to incorporate information ranged from 4.3 months in 1982 to .12 months for 1978. For most forecasters, the figures in Table 5 suggest a pattern of incorporation of information in the weeks rather than in the months.

5. A Digression on Bubbles

Recent behavior in financial and commodity markets — the rise and fall of oil prices, the boom and bust in copper and many non-fuel minerals, the surging dollar — has led some analysts to ask whether speculative forces lie behind such price movements (see particularly Shiller [1984]). Some would say, for example, that the dollar is riding the crest of a speculative bubble reminiscent of the Dutch tulipmania or the stock-market boom of the 1920's.

The approach suggested here can be used to test for expectational bubbles. An expectational bubble is one where prices are driven by price expectations of market participants. In such conditions, the price of a tulip, of the dollar, or of oil rises above its intrinsic value, only to be followed by a bursting of the bubble when price falls back toward or even below its intrinsic value. Figure 5 pictures the movement of actual price (p_t) and of intrinsic value (p_t^*) during a complete bubble cycle.

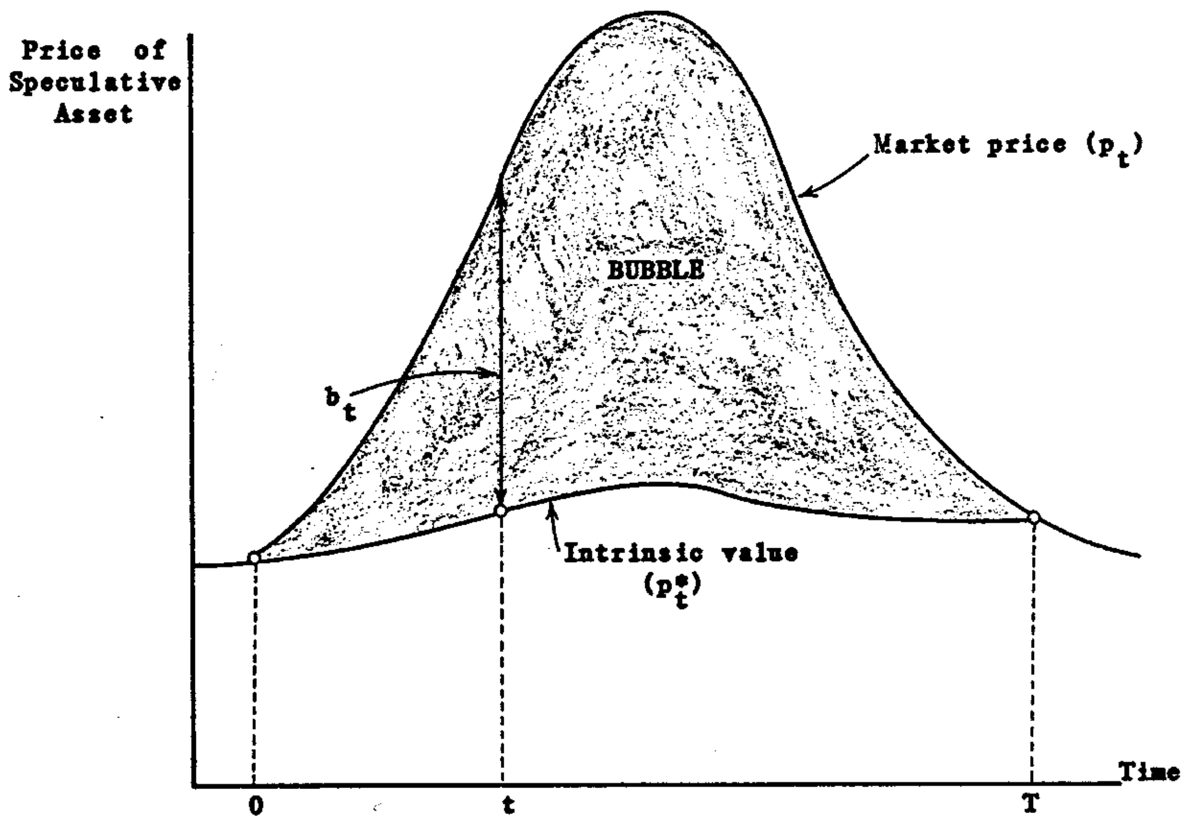

FIGURE 5

During a complete bubble cycle, the price first rises above, then collapses back to, intrinsic value. Empirical results in this paper leave doubts whether any bubble elements arise from dynamics of expectation formation

If we were to track the movement of price expectations during a complete bubble cycle, we would find that price expectations would contain (a) random elements (technically, martingale elements) and (b) bubble elements. As can be seen in Figure 4, the bubble elements are self-reversing. That is, if we compare the future movements of price expectations with past movements, they are of the opposite sign. For example, define the bubble element as $b_t = p_t - p^*_t$ (see Figure 5). If we examine the bubble over a complete cycle (where $b_0 = b_T = 0$), then the movement of the bubble term from t to T must be equal and opposite in sign to the movement of the bubble term from 0 to t.

What do the results of our forecasting survey suggest? They provide little comfort to the bubble view. If expectational bubbles were significant, we should observe forecast revisions reversing earlier forecast revisions; in fact, revisions tend to underreact or grow over time. If markets do indeed show periods when actual prices deviate sharply from intrinsic value, rising and then falling, it is questionable whether such deviations arise purely from expectations.

C. Conclusion

The present paper has investigated the question of the performance of forecasts. It begins by defining the concept of efficiency, which refers to the property that a forecast contains all information available at the time of the forecast. The most useful concept, however, is weak efficiency, under which a forecast minimizes the expected squared error conditional on information about all current and past forecasts. The major result is that forecast revisions should be uncorrelated, or that forecasts of a fixed future event should behave as a random walk (or, more technically, as a martingale).

The test for weak efficiency was applied to a wide variety of circumstances. In 50 of 51 tests, the forecast revisions were found to be positively correlated. The degree of correlation appears to be highest for institutional forecasts (such as those made by international agencies) and lowest for professional forecasters using time-series techniques.

Certain comments on the nature of the results are important to understand. To begin with, the test for weak efficiency is a necessary condition for strong efficiency, or for a good forecast, but it is clearly not sufficient. A babboon could construct a series of weakly efficient forecasts by simply wiring himself to a random-number generator, but such a series of forecasts would be completely useless. Hence we should look at weak efficiency as merely one attribute of well constructed forecasts.

In this respect, one early reader of this paper commented that the idea that forecasts should be a random walk was a ''counsel of despair.'' This reflects a misunderstanding about the central proposition. To say a forecast is a random walk simply means that <u>forecast revisions should not be forecastable</u>. If I could look at your most recent forecasts and accurately say, ''Your next forecast will be 2 percent lower than today's,'' then you can surely improve your forecast. Put differently, if a forecast contains all current information, then forecasts revisions should be unforecastable.

The evidence for positive correlation of forecast revisions in this paper appears quite overwhelming. The reason for such forecast inefficiency are not apparent, however. A number of possible reasons may be worth considering.

First, it may be the case that the true forecasts are indeed efficient, while the published forecasts are not. This might arise because forecasting organizations are bureaucracies that need to reach consensus, and an easy way to reach consensus is to move slowly from last period's consensus to an emerging reality. Surely the high degree of forecast inefficiency of international institutions must contain some element of bureaucratically based forecast inefficiency.

Moreover, some forecasters might smooth their forecasts as a service to customers. One forecaster told me that he smoothed his forecasts because a more accurate but jumpy forecast would ''drive his customers crazy.'' President Carter indeed complained about the ''inconsistency'' of his economic advisers, stating that he was tempted to prefer the fortune teller at the Georgia State fair. Another reader commented that too-quick forecast revisions would entail reversing decisions about investment plans too often.

Each of these rationalizations for smoothing appears unjustified in a completely rational world (the refutations are left to the reader to devise). Whether they are reasonable in an imperfect world where decisionmaking is costly, information is scarce, and misperceptions are rampant is an open question. In any case, at least some of the revealed tendency to smooth forecast revisions probably arises from customer-oriented forecasts.

A final possibility is that the tendency to smooth one's forecasts is rooted in the way people think about the future. There is a suggestion from this study as well as from other evidence that people tend to smooth their forecasts too much. That is, we break the good or bad news to ourselves slowly, taking too long to allow surprises to be incorporated into our forecasts. Is this a pervasive phenomenon — akin to the well-known finding that people overestimate the precision of their knowledge (see Tversky and Kahnman [1981] and Arrow [1982])? Is there perhaps a conservative tilt inside the human brain, a tilt toward wishing to hold on to old notions that are relatively familiar rather than to have to adjust to surprises? Did such a conservative tilt serve good purposes for more stable primitive societies, selecting out people who become fearful about every poorly understood thunderstorm or eclipse? Does such a tilt serve poorly those who must forecast in a rapidly changing world? If in fact there is an internal bias toward excessive smoothing of surprises, perhaps we will have to train ourselves to adapt more rapidly to new circumstances. A motto to remember might be, ''Get all the bad news out, and get it out fast.''

REFERENCES

Arrow, K. J. [1982]. ''Risk Perceptions in Psychology and Economics,'' Economic Inquiry, Vol. 20, pp. 1-9.

Central Intelligence Agency [1980]. OECD Countries: Prospects for Nuclear Power in the 1980's, NFAC/OER/M/IE, Washington, D.C., July 1980.

DRI Review. Data Resources, Inc. Review of the U.S. Economy, Data Resources, Inc., various years.

Fama, Eugene F. [1976]. Foundations of Finance, Basic Books, New York.

Haerter, Manfred [1985]. Guete- und Erfolgsbeurteilung Zukunftsbezogener Aussager, Doctoral Dissertation, University of Stuttgart, 1985.

Nordhaus, William and Steven Durlauf [1984]. ''Empirical Tests of the Rationality of Economic Forecasters: A Fixed-Horizons Approach,'' Cowles Foundation Discussion Paper No. 717-R, November 1984.

Shiller, Robert [1984]. ''Stock Prices and Social Dynamics,'' Brookings Papers on Economic Activity, 1984:2.

Tversky, A. and D. Kahnman [1981]. ''The Framing of Decisions and the Psychology of Choice,'' Science, 211, pp. 453-458.