

NBER WORKING PAPER SERIES

THE INFORMATIONAL CONTENT
OF EX ANTE FORECASTS

Ray C. Fair

Robert J. Shiller

Working Paper No. 2503

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 1988

We are indebted to Stephen K. McNees for providing us with the data on the ex ante forecasts used in this paper. This research was supported by the National Science Foundation. The research reported here is part of the NBER's research program in Economic Fluctuations. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

The Informational Content
of Ex Ante Forecasts

ABSTRACT

The informational content of different forecasts can be compared by regressing the actual change in a variable to be forecasted on forecasts of the change. We use the procedure in Fair and Shiller (1987) to examine the informational content of three sets of ex ante forecasts: the American Statistical Association and National Bureau of Economic Research Survey (ASA), Data Resources Incorporated (DRI), and Wharton Economic Forecasting Associates (WEFA). We compare these forecasts to each other and to "quasi ex ante" forecasts generated from a vector autoregressive model, an autoregressive components model, and a large-scale structural model (the Fair model).

Ray C. Fair
Cowles Foundation
Box 2125, Yale Station
New Haven, CT 06520

Robert J. Shiller
Cowles Foundation
Box 2125, Yale Station
New Haven, CT 06520

The Informational Content of Ex Ante Forecasts

by

Ray C. Fair and Robert J. Shiller

I. Introduction

In a previous paper, Fair and Shiller (1987), we proposed a procedure for examining the informational content of forecasts. The procedure involves running regressions of the actual change in the variable forecasted on forecasts of the change, along the lines described in the literature on encompassing tests of non-nested hypotheses. We used this procedure to examine forecasts for the period 1976 III - 1986 II from the Fair (1976) model, two autoregressive (AR) models, six vector autoregressive (VAR) models, and eight "autoregressive components" (AC) models. The procedure requires that forecasted changes for a period be based only on information available in the period prior to the first period of the forecast, and we were careful to impose this requirement. All models were estimated only through period $t-s$ for a forecast of the change between periods $t-s$ and t . Also, autoregressive equations for all the exogenous variables in the Fair model were added to the model, and these equations were used to predict the exogenous-variable values. (The other models examined contain no exogenous variables.) Finally, a version of the Fair model was used that existed in 1976 II, which insures that no information after this date was used for the specification.

Although all the forecasted changes between period $t-s$ and t were based only on information through period $t-s$, the forecasts were not ex ante forecasts in the sense of having been forecasts that were actually made at the time. In this paper we use our procedure to examine three sets of ex ante forecasts: the American Statistical Association and National Bureau of

Economic Research Survey (ASA), Data Resources Incorporated (DRI), and Wharton Economic Forecasting Associates (WEFA). The data on the forecasts were provided us by Stephen K. McNees, who has been collecting forecasts on a consistent basis from the forecasters as the forecasts were made. He is thus able to verify the exact date when the forecast became available. McNees has done a number of studies comparing the accuracy of the different forecasts -- see, for example, McNees (1981), (1985), (1986).

It is well known that forecasts from models like DRI and WEFA are subjectively adjusted. One interpretation of the adjustment procedure is that the model builders use all the information available to them at the time of the forecast, much of it outside the model, in deciding how to adjust the model. In other words, the forecasts are an aggregation of a considerable amount of information as sifted through the minds of the model builders.

We are interested in two sets of questions. The first is whether, say, the DRI forecasts contain information not in the WEFA forecasts and vice versa. The second is whether the forecasts generated in our previous paper (based only on information through the period prior to the first period of the forecast) contain information not in the ex ante forecasts and vice versa. We will call the forecasts generated in our previous paper "quasi ex ante" forecasts to distinguish them from the true ex ante forecasts of ASA, DRI, and WEFA.

II. The Procedure

Let \hat{Y}_{1t} denote a forecast of Y_t (in our application, log real gross national product at time t) made by forecaster 1 (or model 1 with its associated estimation procedure and forecasting method) at time $t-s$, $s > 0$. The foundation of the empirical work that follows (as in Fair and Shiller (1987)) is the regression equation:

$$(1) \quad Y_t - Y_{t-s} = \alpha + \beta(\hat{Y}_{1t} - Y_{t-s}) + \gamma(\hat{Y}_{2t} - Y_{t-s}) + u_t.$$

If neither forecast 1 nor forecast 2 contains any useful information for s -period-ahead forecasting of Y_t , then the estimates of β and γ should both be zero. In this case the estimate of the constant term α would be the average s -period-change in Y . If both forecasts contain independent information for s -period-ahead forecasting, then β and γ should both be nonzero. If both forecasts contain information, but the information in, say, forecast 2 is completely contained in forecast 1 and forecast 1 contains further relevant information as well, then β but not γ should be nonzero. (If both forecasts contain the same information, then they are perfectly correlated, and β and γ are not separately identified.)

The procedure we have proposed is to estimate equation (1) for different forecasts and test the hypothesis H_1 that $\beta = 0$ and the hypothesis H_2 that $\gamma = 0$. H_1 is the hypothesis that forecast 1 contains no information relevant to forecasting s periods ahead not in the constant term and in forecast 2, and H_2 is the hypothesis that forecast 2 contains no information not in the constant term and in forecast 1.

Our testing procedure is similar to the C-test of Davidson and MacKinnon (1981) -- which is a special case of the "Wald encompassing test"

of Mizon and Richard (1986)¹ -- but it differs from this procedure in a number of important ways.

First, in our procedure the tests will be done for s equal to four as well as one. Davidson and MacKinnon, along with many others, have focussed attention exclusively on one-period-ahead forecasts.² The information content of forecasts may differ depending on forecast horizon, as we will see below. Second, the C-test restricts β and γ to sum to one.³ In our application this restriction does not seem sensible. As noted above, if both models' forecasts are just noise, the estimates of both β and γ should be zero. Third, the C-test restricts the constant term α to be zero.⁴ Again, in our application this restriction does not seem sensible. If, for example, both forecasts were noise and we estimated equation (1) without a constant term, then the estimates of β and γ would not generally be zero when the mean of the dependent variable is nonzero.

Fourth, we require that forecasts beginning in period t contain only information through period $t-1$. Davidson and MacKinnon do not require this. The *ex ante* forecasts obviously satisfy this requirement, and we have made

¹See also Hendry and Richard (1982) and Chong and Hendry (1986). Nelson (1972) and Cooper and Nelson (1975) are early examples of the use of encompassing-like tests.

²Their doing so was dictated by their setup of the model, wherein multi-period forecasts are not defined.

³Granger and Newbold (1986) in their discussion of combining forecasts also speak of constraining the coefficients to sum to one, without presenting an argument why one should do so. In their work, constraining the coefficients to sum to one and setting the constant term to zero makes possible some simple theorems that offer interpretations of the single parameter estimated in their regression.

⁴Chong and Hendry's (1986) formulation of (1) also does not contain a constant term, although they do not constrain β and γ to sum to one.

sure that the quasi ex ante forecasts also satisfy it. Forecasts that are based on rolling estimation of a model may have different properties from those made with a model estimated with future data. If the model is misspecified (e.g., parameters change through time), then the rolling estimation forecasts (where estimated parameters vary through time) may carry rather different information from forecasts estimated over the entire sample. Also, some models may use up more degrees of freedom in estimation than others, and with varied estimation procedures it is often very difficult to take formal account of the number of degrees of freedom used up. In the extreme case where there were so many parameters in model 1 that the degrees of freedom were completely used up when it was estimated, it would be the case that $Y_t = \hat{Y}_t$ and there would be a spurious perfect correspondence between the variable forecasted and the forecast. This would cause $\beta = 1$ in (1) whether or not model 1 were a good model. One can guard against this degree of freedom problem by requiring that no forecasts be within-sample forecasts.⁵

Fifth, we do not assume that u_t is identically distributed, as do Davidson and MacKinnon. It seems quite likely that u_t is heteroskedastic. If, for example, $\alpha = 0$, $\beta = 1$, and $\gamma = 0$, then u_t is simply the forecast error from model 1, and in general forecast errors are heteroskedastic. Also, we will be considering four-period-ahead forecasts in addition to one-period-ahead forecasts, and this introduces a third-order moving-average

⁵Nelson (1972) and Cooper and Nelson (1975) do not require the forecasts to be based only on information through the previous period. Chong and Hendry (1986) do, however, require this. In their procedure the models that give rise to the forecasts are estimated using sample period 1 through T and their regression analogous to equation (1) is run using sample period beginning in T+1.

process to the error term in equation (1).⁶ We correct for both heteroskedasticity and the moving average process in the estimation of the standard errors of the coefficient estimates. For the one-quarter-ahead forecasts we use the method of White (1980), and for the four-quarter-ahead forecasts we use the method suggested by Hansen (1982), Cumby, Huizinga, and Obstfeld (1983), and White and Domowitz (1984). The exact formula that we used for the covariance matrix of the coefficient estimates is presented in Fair and Shiller (1987).

III. The Forecasts and Models

Any comparison of ex ante forecasts must confront the problem of data revisions. The data for GNP are revised back three years every year, and from time to time the data are revised back to the very beginning of the sample. Let $Y_{t,T}$ represent the value of time t log real GNP that is the latest available from the U.S. Commerce Department at time T , $T \geq t$. (It is understood that when the second subscript T is omitted, we mean T - end of the full sample available now.) Let ${}_{t-s}Y'_t$ be the ex ante forecast of log GNP for quarter t that existed at time $t-s$ (the ' replacing the ^). The problem is how to compare ${}_{t-s}Y'_t$ and Y_t , given that $Y_{t-s,t-s}$ and Y_{t-s} may be quite different because of data revisions? There is obviously no right answer to this problem. What we have done is to adjust ${}_{t-s}Y'_t$ to make the forecasted change (from Y_{t-s}) be the same as the ex ante forecasted change.

⁶The error term in equation (1) could, of course, be serially correlated even for the one-period-ahead forecasts. Such serial correlation does not appear to be a problem with any of the models we study here, however, and we have assumed it to be zero. One should not, of course, uncritically apply procedures such as ours to all models, as Granger and Newbold (1986) have warned.

In other words, we have taken the new value of the forecasted level of log real GNP for quarter t , ${}_{t-s}\hat{Y}_t$, to be:

$$(2) \quad {}_{t-s}\hat{Y}_t = {}_{t-s}Y'_t + Y_{t-s} - Y_{t-s,t-s}$$

Adjustments of this type are fairly common when dealing with ex ante forecasts -- see, for example, McNees (1981).

We will now briefly discuss the three models whose quasi ex ante forecasts we are comparing to the actual ex ante forecasts.

The Fair Model (FAIR)

The first version of the Fair model was presented in Fair (1976) along with the estimation method and method of forecasting with the model. This version was based on data through 1975 I. One important addition that was made to the model from this version was the inclusion of an interest rate reaction function in the model. This work is described in Fair (1978), which is based on data through 1976 II. The version of the model in Fair (1976) consists of 26 structural stochastic equations, and with the addition of the interest rate reaction function, there are 27 stochastic equations. There are 106 exogenous variables, and for each of these variables an eighth order autoregressive equation with a constant and time trend was added to the model. This gave a model of 133 equations, and this is the version that was used.

For the rolling estimations, the first estimation period ended in 1976 II, which is the first quarter in which the model could definitely be said to exist. This allowed the model to be estimated 40 times (through 1986 I).

The VAR Model

We considered six VAR models in Fair and Shiller (1987), but here we consider only the VAR model that gave the best results. This model is the same as the model used in Sims (1980) except that we have added the three-month Treasury bill rate to the model. There are seven variables in the model: real GNP, the GNP deflator, the unemployment rate, the nominal wage rate, the price of imports, the money supply, and the bill rate. All but the unemployment rate and the bill rate are in logs. Each equation consists of each variable lagged one through four times, a constant, and a time trend, for a total of 30 coefficients per equation. We have imposed Bayesian priors on the coefficients of the model. We imposed the Litterman prior that the variables follow univariate random walks. The standard deviations of the prior take the form

$$(3) \quad S(i,j,k) = \gamma g(k) f(i,j) (s_j / s_i),$$

where i indexes the left-hand-side variable, j indexes the right-hand-side variables, and k indexes the lag. s_i is the standard error of the unrestricted equation for variable i . The parameter values chosen imply fairly tight priors: 1) $f(i,j) = 1$ for $i = j$, $f(i,j) = .5$ for $i \neq j$, 2) $g(k) = k^{-1}$, and 3) $\gamma = .1$. These are the values used by Litterman (1979, p. 49).

The VAR model was estimated 40 times using the same sample periods as were used for the Fair model. The model was then used to make 40 forecasts of real GNP.

The AC Model

Eight AC models were considered in our earlier paper, but again we consider only the one that gave the best results. An AC model is one in which each of the components of real GNP is determined by a simple autoregressive equation and GNP is determined as the sum of the components (i.e., by the GNP identity). The version we use here has 17 components. Each equation for a component contains the first eight lagged values of the component, a constant, a time trend, and the first four lagged values of real GNP itself. The equations are not in log form. The same sample periods and procedures were used for the AC model as were used for the Fair and VAR models.

IV. The Results

The results of estimating equation (1) are presented in Table 1. The sample period used for the one-quarter-ahead results is 1976 III - 1986 II, for a total of 40 observations. The sample period for the four-quarter-ahead results is 1977 II - 1986 II, for a total of 37 observations. As mentioned above, for the quasi ex ante forecasts each forecast observation is based on a different set of coefficient estimates of the model -- rolling estimation is used. Also, for the Fair model all exogenous variable values are generated from the autoregressive equations; no actual values are used. Finally, the one-quarter-ahead regressions White's (1980) correction for heteroskedasticity has been used and for the four-quarter-ahead regressions the method of Hansen (1982), Cumby, Huizinga, and Obstfeld (1983), and White

TABLE 1

Comparison of the Forecasts: Estimation of Equation (1)

One-Quarter-Ahead Forecasts
Dependent variable is $Y_t - Y_{t-1}$
Sample period - 1978III-1986II

Other Model	ASA			DRI			WEFA		
	Const	ASA	Other	Const	DRI	Other	Const	WEFA	Other
DRI	.0000 (0.04) [.00673]	2.36 (4.88) (2.47)	-1.23 (2.47)						
WEFA	.0018 (1.28) [.00841]	2.49 (8.24) (3.29)	-1.36 (3.29)	.0014 (0.70) [.00820]	.92 (1.48)	.01 (0.01)			
VAR	-.0008 (0.50) [.00700]	1.30 (8.01) (1.58)	-.33 (1.58)	-.0016 (0.77) [.00814]	.98 (4.09) (0.37)	-.09 (0.37)	-.0007 (0.30) [.00851]	.89 (3.54) (0.30)	-.08 (0.30)
AC	-.0000 (0.62) [.00720]	1.07 (4.39) (0.27)	.06 (0.27)	-.0003 (0.15) [.00806]	.81 (2.88) (1.08)	.29 (1.08)	-.0010 (0.40) [.00839]	.73 (2.45) (1.03)	.31 (1.03)
FAIR	-.0021 (0.95) [.00703]	.95 (4.08) (1.40)	.32 (1.40)	-.0025 (0.98) [.00778]	.73 (3.18) (2.38)	.50 (2.38)	-.0034 (1.21) [.00802]	.84 (2.90) (2.57)	.55 (2.57)

Four-Quarter-Ahead Forecasts
Dependent variable is $Y_t - Y_{t-4}$
Sample period - 1977II-1986II

Other Model	ASA			DRI			WEFA		
	Const	ASA	Other	Const	DRI	Other	Const	WEFA	Other
DRI	.0001 (0.01) [.0206]	.34 (1.06) (2.12)	.86 (2.12)						
WEFA	-.0088 (0.80) [.0201]	.12 (0.31) (1.37)	.96 (1.37)	-.0078 (0.47) [.0200]	.17 (0.32)	.89 (1.14)			
VAR	.0017 (0.18) [.0200]	.70 (2.33) (1.19)	.38 (1.19)	.0024 (0.23) [.0194]	.89 (2.79) (1.36)	.38 (1.36)	-.0054 (0.44) [.0191]	.79 (2.91) (1.75)	.35 (1.75)
AC	-.0081 (0.70) [.0200]	.79 (5.21) (1.86)	.44 (1.86)	-.0082 (0.53) [.0199]	.77 (5.32) (1.47)	.42 (1.47)	-.0135 (1.14) [.0195]	.88 (4.84) (1.59)	.35 (1.59)
FAIR	-.0182 (2.92) [.0134]	.38 (2.23) (4.94)	1.00 (4.94)	-.0171 (2.62) [.0136]	.33 (1.84) (4.33)	1.00 (4.33)	-.0204 (3.64) [.0134]	.38 (1.75) (4.50)	.97 (4.50)

Notes: Y = log of real GNP.
t-statistics in absolute value are in parentheses.
Estimated standard errors of the regressions are in brackets.
See text for discussion of the estimation methods.
For the one-quarter-ahead results, estimates are of the coefficients of $t-1$ Y_{t-1}
For the four-quarter-ahead results, estimates are of the coefficients of $t-4$ Y_{t-4}

and Domowitz (1984) has been used (with a moving average of order 3).⁷

Consider now the results in Table 1. For the one-quarter-ahead forecasts, the ASA forecast dominates in the sense that it has the largest coefficients. It is the case, however, that the ASA forecasts are made later in the period than the others, which gives them a considerable advantage for the one-quarter-ahead results. (McNees (1985) classifies the ASA forecasts as "mid quarter," whereas the DRI and WEFA forecasts are classified as "early quarter.") What the present results show is that by the time the ASA forecasts are made, they contain substantial information not in the other forecasts. For the four-quarter-ahead results, both the DRI and FAIR forecasts appear to contain information not in the ASA forecasts -- the t-statistic for the DRI forecast is 2.12 and the t-statistic for the FAIR forecast is 4.94.

It is interesting to note that when the DRI and WEFA one-quarter-ahead forecasts are compared to the ASA one-quarter-ahead forecast, the DRI and WEFA forecasts are significant at the 5 percent level, but with negative weights. The negative coefficient estimates do not mean, however, that the DRI and WEFA forecasts are not necessarily optimal forecasts given their (early quarter) information set. Consider the following example. Let X_1 be the optimal forecast given the early quarter information set, and let $X_1 + X_2$ be the optimal forecast given the mid quarter information set, where X_1 and X_2 are uncorrelated. Assume that only a third of the ASA respondents

⁷In one case for the four-quarter-ahead results -- ASA versus DRI -- the estimated covariance matrix of the coefficients estimates was nearly singular and the results were not sensible. In this case we assumed a second order MA process for the error term instead of a third order, which solved the problem. Had this been a more wide spread problem, we would have used one of the estimators in Andrews (1987), but this seemed unnecessary given only one failure.

use the new information available after the date of the early quarter forecast. If the DRI forecast is X_1 and the ASA forecast is $X_1 + (1/3)X_2$, a regression of the actual value on the two forecasts will give a coefficient of 3 for ASA and -2 for DRI, thus achieving the optimal forecast $X_1 + X_2$. The DRI forecast is in effect "correcting" the ASA forecast for using only a 1/3 weight on X_2 .

The comparisons of the DRI and WEFA forecasts in Table 1 show that the two forecasts are too collinear for any strong conclusions to be drawn. None of the forecasts individually is significant. The VAR forecasts appear to contain no information not in the DRI and WEFA one-quarter-ahead forecasts (the VAR coefficient estimates are highly insignificant), but they do carry a weight of about a third for the four-quarter-ahead forecasts. The AC forecasts get a weight of about a third when compared with either the DRI or WEFA forecasts for both the one-quarter-ahead and four-quarter-ahead results. The DRI and WEFA forecasts are significant at the 5 percent level when compared with the VAR and AC forecasts, and so they appear to contain information not in the VAR and AC forecasts.

For the one-quarter-ahead results the FAIR forecasts contain information not in the DRI and WEFA forecasts and the DRI and WEFA forecasts contain information not in the FAIR forecasts. For the four-quarter-ahead results it is still true that the FAIR forecasts contain information not in the DRI and WEFA forecasts, but it is now no longer the case that the DRI and WEFA forecasts are statistically significant when compared with the FAIR forecasts. They get weights of about a third, with t-statistics for DRI and WEFA of 1.64 and 1.75, respectively.

The following points thus emerge from the results:

1. The procedure cannot discriminate well between DRI and WEFA. Both sets of model builders seem to use very similar information sets, and the two forecasts do not contain much independent information.
2. No one-quarter-ahead forecast carries as much information as in the ASA forecast, the ASA forecast being made later than the others.
3. The VAR and AC quasi ex ante forecasts appear to contain only a modest amount of information not in the ASA, DRI, and WEFA forecasts. Another way of looking at this is that the ASA forecasters and the DRI and WEFA model builders have not overlooked a lot of useful forecasting information in the variables in the VAR and AC models.⁸
4. The FAIR model quasi ex ante forecasts, on the other hand, do contain a substantial amount of information not in the ASA, DRI, and WEFA forecasts (except for the one-quarter-ahead ASA forecast). In other words, the ASA forecasters and the DRI and WEFA model builders have overlooked useful forecasting information in the FAIR model forecasts.
5. For the one-quarter-ahead results the ASA, DRI, and WEFA forecasts contain useful forecasting information not in the FAIR forecasts. The large amount of information sifted through the minds of the model builders when they make a forecast does appear to contain some useful information for forecasting one quarter ahead that is not in the FAIR quasi ex ante forecasts. On the other hand, this is much less the case for the four-

⁸McNees (1986) found that the Litterman Bayesian VAR forecasts did better than any of the other forecasts studied for the four-quarter-ahead forecast of real GNP for the sample period 1980-II to 1985-I. On the other hand, the Bayesian VAR forecasts were not relatively good at forecasting one-quarter-ahead real GNP. This sample is only a third as long as ours, and so it is of questionable relevance to our results.

quarter-ahead results except for the ASA forecast, which has a head start. In this sense the quasi ex ante FAIR forecasts look quite good.

IV. Conclusion

The procedure that we have proposed for examining the informational content of forecasts appears to be a useful alternative to the standard procedure of comparing forecasts by the size of their root mean squared errors (RMSEs) or mean absolute errors (MAEs). In many cases our procedure may be able to discriminate better. It is often the case, for example, that the RMSEs and MAEs for two forecasts are so close that one is not sure if the differences are economically meaningful. In at least some of these cases our procedure may be more informative. For example, the dominance of the one-quarter-ahead ASA forecasts in this paper is not something that is obvious from simply looking at the RMSEs and MAEs. The same is also true of the dominance of the ASA, DRI, and WEFA forecasts over the VAR and AC forecasts. There are also, of course, cases where our procedure does not discriminate either, such as the DRI versus WEFA comparison, but there appear to be fewer of these cases for our procedure than for the procedure of comparing RSMEs and MAEs.

Our results also suggest that combining forecasts may be useful. Although there is not much point in combining the DRI and WEFA forecasts, since they are so similar, some gain may be achieved by combining either of them with the FAIR model forecast for one-quarter-ahead forecasting. There is, of course, no assurance that such combined forecasts will work well. The forecasts that go into a regression may have changing stochastic properties through time. For example, as time progresses and a model is

reestimated, the forecast from a model is based on more and more data. Thus, a model estimated using rolling estimation methods may forecast much better now, at the end of the sample, than it did on average over the entire sample. The ex ante forecasts are also updated using new data, and the model builders who put their judgment into the forecasts are themselves learning from past errors, just as we are with our regression analysis. They may have already in effect combined the forecasts. One must thus be cautious in combining forecasts from regressions like those in Table 1.

References

- Andrews, Donald W. K., "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," October 1987, mimeo.
- Chong, Yock Y., and David F. Hendry, "Econometric Evaluation of Linear Macro-Econometric Models," Review of Economic Studies, 53, 671-90, August, 1986.
- Cooper, J. Phillip and Charles R. Nelson, "The Ex-Ante Prediction Performance of the St. Louis and FRB-MIT-Penn Econometric Models and Some Results on Composite Predictions," Journal of Money, Credit and Banking, 7:1-32, February, 1975.
- Cumby, R. E., J. Huizinga, and M. Obstfeld, "Two-Step Two Stage Least Squares in Models with Rational Expectations," Journal of Econometrics, 21, 333-355, 1983.
- Davidson, R., and J. G. MacKinnon, "Several Tests of Model Specification in the Presence of Alternative Hypotheses," Econometrica 40, 781-93, 1981.
- Fair, Ray C., A Model of Macroeconomic Activity, Vol. 2. The Empirical Model, Ballinger Publishing Co., 1976.
- Fair, Ray C., "The Sensitivity of Fiscal Policy Effects to Assumptions About the Behavior of the Federal Reserve," Econometrica, 46:1165-79, 1978.
- Fair, Ray C., and Robert J. Shiller, "Econometric Modeling as Information Aggregation," December 1987, mimeo.
- Granger, C. W. J., and Paul Newbold, Forecasting Economic Time Series 2nd. Ed., Academic Press, 1986.
- Hansen, Lars Peter, "Large Sample Properties of Generalized Method of Moments Estimators," Econometrica, 50, 1029-1054, 1982.
- Hendry, David F., and Jean-Francois Richard, "On the Formulation of Empirical Models in Dynamic Economics," Journal of Econometrics, 20, 3-33, October, 1982.
- Litterman, Robert B., "Techniques of Forecasting Using Vector Autoregression," Federal Reserve Bank of Minneapolis Working Paper No. 115, November 1979.
- McNees, Stephen K., "The Recent Record of Thirteen Forecasters," New England Economic Review, pp. 5-21, September/October 1981.
- McNees, Stephen K., "Which Forecast Should You Use?" New England Economic Review, pp. 36-42, July/August 1985.

- McNees, Stephen K., "The Accuracy of Two Forecasting Techniques: Some Evidence and Interpretation," New England Economic Review, pp. 20-31, March/April 1986.
- Nelson, Charles R., "The Prediction Performance of the FRB-MIT-Penn Model of the U. S. Economy," American Economic Review, 62:902-17, December 1972.
- Sims, Christopher A., "Macroeconomics and Reality," Econometrica, Vol. 48, No. 1, pp. 1-48, January 1980.
- White, Halbert, "A Heteroskedasticity Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," Econometrica Vol. 48, No. 4, pp. 817-38, May 1980.
- White, Halbert, and Ian Domowitz, "Nonlinear Regression with Dependent Observations," Econometrica, 52, 143-61, 1984.