Gordon Schlegel is a research assistant at the Federal Reserve Bank of Cleveland. *The author would like to thank K.J. Kowalewski, Owen Humpage, Mark Sniderman, James Hoehn. and William Gavin for their helpful comments and suggestions.*

1. One cannot, however, perform policy simulations using vector autoregressive models. Lucas (1976) pointed out that under alternative policies, agents will have different views about the way exogenous shocks affect the system. Therefore, one cannot use the same set of parameters for all alternative policies one may wish to examine. This implies that the coefficients obtained through in-sample estimation may not accurately reflect policy changes.

Vector Autoregressive Forecasts of Recession and Recovery: Is Less More?

by Gordon Schlegel

Economic forecasts are valuable tools for decision makers in many different areas. When used with discretion, forecasts can help guide the strategic plans of businesses and corporations. A reasonably sharp picture of the future is also important in the formation of sound fiscal and monetary policy.

Forecasts are particularly important when the economy has just entered a recessionary or expansionary period. Policies that are useful in expansionary periods must often be adjusted before and during contractions, and vice versa. To get an idea of the degree to which policies must change, one needs to forecast the extent of the expansion or contraction to come.

Many economists are turning to the use of vector autoregressive (VAR) models for forecasting. A number of studies have indicated that VARs forecast as well as, if not better than, many large structural models; one such study is that of Lupoletti and Webb (1984). However, the forecast periods used in these studies are not differentiated into expansionary and recessionary periods. An economist using VARs might want to ask the question: "What VAR specification will do the best job in predicting the length and intensity of recessions and recoveries?"

This paper provides a possible answer to this important question. The first section discusses the reasons that VARs are gaining in popularity among forecasters and describes the methodology of VARs. Section II discusses the pros and cons of VARs. Section III describes the various model specifications compared in the study and the measures of forecast accuracy employed in the comparison. Section IV looks at the estimation results for the specified models, while section V considers a more recently developed VAR technique. Finally, section VI sums up the overall results of the study and mentions several cautions concerning the interpretations of the results.

I. VARs: Why and How?

In their never-ending search for the perfect crystal ball, economic forecasters try to obtain high forecast accuracy and, at the same time, use as simple a technique as possible. This is particularly true of business economists who work under significant time and resource constraints which, in turn, limit the degree of sophistication they can apply to their forecasts.

However, the forecasts must still be accurate enough to give a fairly sharp picture of the environment that firms and consumers will be facing in the immediate future. A forecast is, obviously, not useful if it does not predict with an "acceptable" degree of accuracy. However, even if the technique exists to produce a perfect forecast, the method is worthless if it is too complex for a practitioner to apply properly.

VAR techniques have been proposed as a means through which one can have the best of both worlds: simplicity and accuracy? In a VAR system with n lags, each variable being forecast is regressed against its own values in each of the n preceding periods, against the values in each of the n previous periods of all of the other variables being forecast, and against a constant term. For example, a VAR system with three variables, X, Y, and Z, and with two lags would consist of the following equations:

$$X = c_1 + a_{11}X_{-1} + a_{21}X_{-2} + b_{11}Y_{-1} + b_{21}Y_{-2} + c_{11}Z_{-1} + c_{21}Z_{-2} + e_1,$$

$$Y = c_2 + a_{12}X_{-1} + a_{22}X_{-2} + b_{12}Y_{-1} + b_{22}Y_{-2} + c_{12}Z_{-1} + c_{22}Z_{-2} + e_2,$$

$$Z = c_3 + a_{13}X_{-1} + a_{23}X_{-2} + b_{13}Y_{-1} + b_{23}Y_{-2} + c_{13}Z_{-1} + c_{23}Z_{-2} + e_3,$$

2. The in-sample fits of the various specifications are not considered. We only want to predict future values of the variables in the system, not explain their past values.

where

X_{-n} = the value of X n periods before the current period,
e_r = the error term of equation r, distributed as a normal random variable with mean 0 and constant variance, and
c_r = the constant term of equation r.

The equations are estimated individually to yield estimates for all parameters and constant terms. One can then calculate the reduced form of the system and predict the values of all variables in the current time period. These values can, in turn, be used as regressors in predicting the next period's values for the variables. The process can be continued indefinitely, enabling one to produce dynamic, out-of-sample forecasts as far into the future as desired, given the information available in the present period.

The regression equations are commonly estimated in one of two ways. With ordinary least squares, the parameters are completely unconstrained and can assume whatever values best fit the data. Bayesian techniques enable a forecaster to explicitly include, in the model, subjective judgment or other objective evidence concerning the values of the parameters, as well as the degree of confidence he has in his judgment. A general discussion of the techniques is given in Todd (1984), while Litterman (1979) approaches the topic from a more technical basis.

In this paper, we first search for the optimal ordinary least squares (OLSQ) specification, where the "optimal" specification is the one that provides the most accurate forecasts, the measures of accuracy being described below.[2] We then compare this specification to one derived through a Bayesian procedure.

II. Advantages and Disadvantages of VAR Models

VARs have a number of characteristics that make them convenient for those who make economic forecasts on a regular basis. Of these characteristics, the following five seem especially worthy of note:

1) It is relatively easy to write a computer program to perform a VAR. A programmer with a moderate amount of skill and a package of standard regression techniques should be able to implement such a program without much trouble.

2) The commands needed to perform an OLSQ VAR can be implemented in virtually any programming language. This would make it unnecessary to buy a specialized package to run VARs and would enable a forecaster to avoid this type of expense. The Bayesian VAR can be implemented with a little more effort, provided that matrix capabilities are available.

3) Since VARs can be programmed fairly easily, it might not be necessary to buy forecasting services from an outside data vendor. Subscriptions to the major econometric forecasting services can cost from $16,000 to $20,000 per year, no bargain if, as Lupoletti and Webb (1984) suggest, the simpler VAR models can perform as well as, or better than, the large models.

4) Because VARs only use a relatively small number of variables, it is easy to update and revise the data series as needed.

5) In their pure form, VARs require no subjective add factors. Large models contain a number of arbitrary constants that a forecaster might be unable to estimate sufficiently well for his purposes, due to a lack of necessary specialized information or expertise. The VAR gets around this problem by avoiding it.

No forecasting technique, however, is without its problems. VAR models have two major disadvantages:

1) Since most aggregate economic time series are highly correlated with their own previous values and with present and past values of other time series, multicollinearity can become a serious problem as more and more series and lagged values of series are added to the model. As the system expands, it can become very difficult to separate the effects of the explanatory variables, and the

parameter estimates can become highly sensitive to the combination of variables used in the model.

Also, a high degree of multicollinearity will make it difficult to determine which explanatory variables are significant, since the standard errors of the coefficient estimates will tend to be large. A forecaster considering

Fig. 1 Dynamic Out-of-Sample Root Mean Squared Error

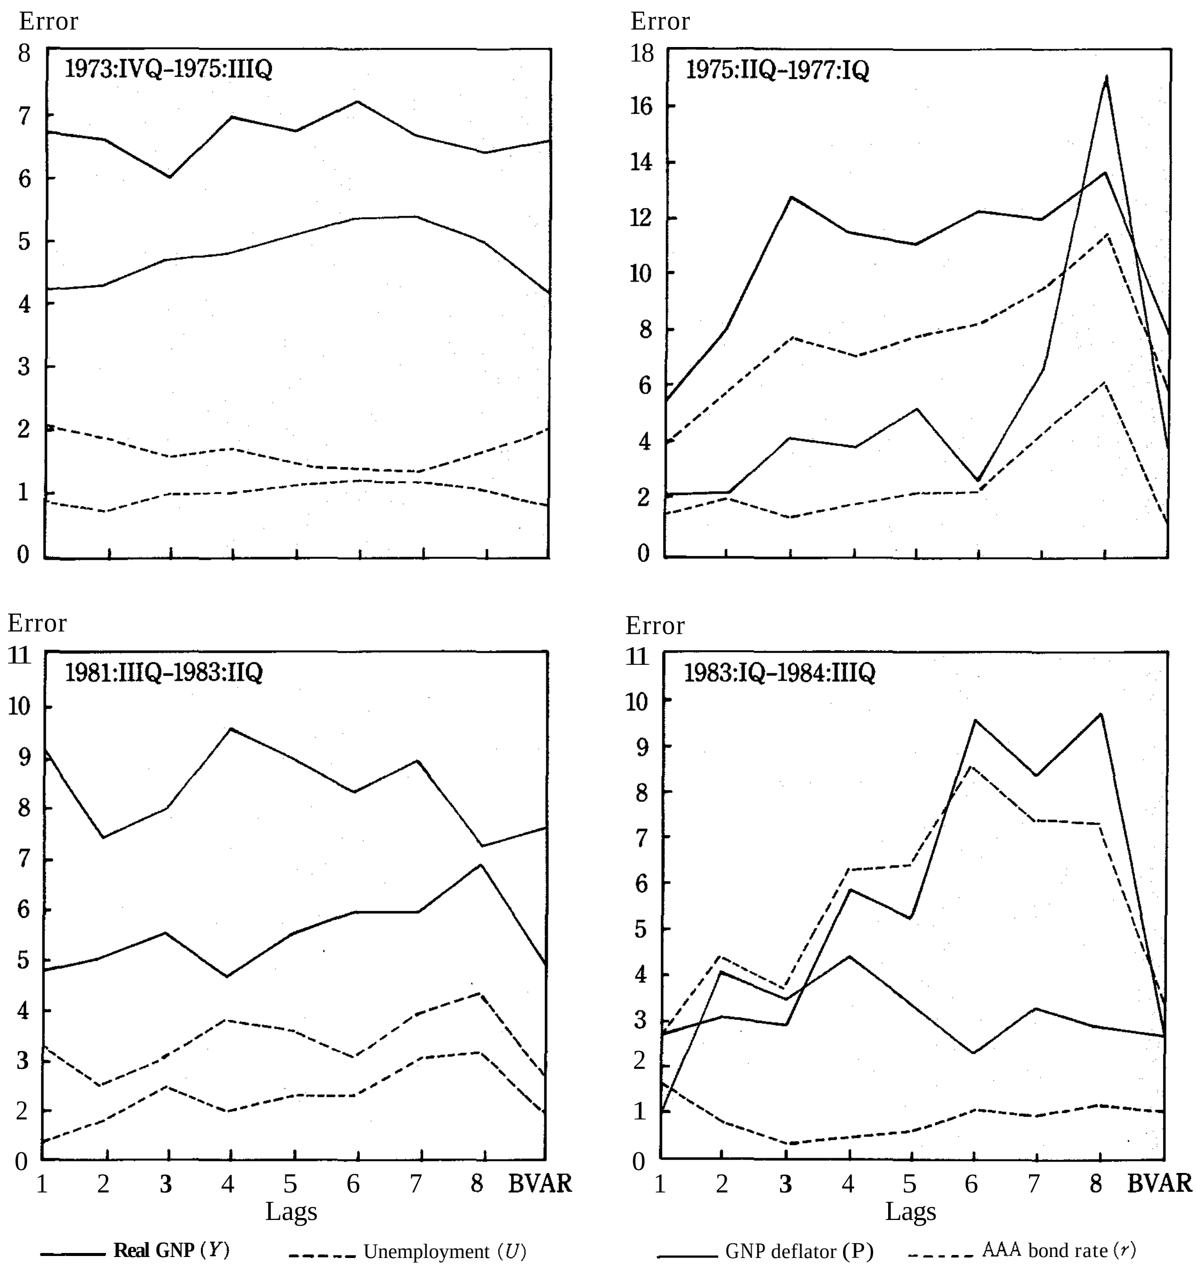

3. We choose the growth rate of real GNP instead of a measure of the level of this variable. This implies that we are interested in the pattern of GNP growth over our forecast horizon, not just the proportion by which output will have grown seven or eight quarters hence.

4. Implicit in this methodology is the assumption that turning points are recognized when they occur. In practice, there may be a time lag ofseveral months between the occurrence of a turning point and its recognition by forecasters.

a certain lag structure might want to ask if certain lagged variables can be dropped from the system without sacrificing forecast accuracy. A detailed discussion is found in Intriligator (1978), among others.

As far as the forecasting aspects of multicollinearity are concerned, Christ (1966) points out that if the joint distribution of the regressors changes during a forecasting period, multicollinearity between regressors will affect the accuracy of the forecasts. Given the increasing volatility of aggregate measures of economic activity over the past 10 years, particularly interest rates, it would appear that such changes have taken place. Multicollinearity, therefore, seems to present a problem for VAR forecasting.

2) As the number of variables of a VAR model increases, the number of parameters to be estimated goes up rapidly. If a variable is added to the model, each equation has n more coefficients to be estimated, where n is the number of lags for each variable.

If a lag period is added, each equation has *r* more parameters, where *r* is the number of variables in the system. As the number of coefficients increases relative to the amount of available data, random events of the past, as well as systematic relationships, are increasingly reflected in the coefficients. If these coefficients are used in out-of-sample prediction, a set of future random events that differs from the shocks of the past would be expected to result in less accurate forecasts. This problem is discussed in Todd (1984).

Table 1 Rankings of Root Mean Squared Errors of Dynamic Out-of-Sample Forecasts

Forecast period	Variables	Number of lags 1	2	3	4	5	6	7	8
1973:IVQ–1975:IIIQ	P	1	2	3	4	6	7	8	5
	Y	6	4	1	7	5	8	3	2
	r	2	1	3	4	6	8	7	5
	U	8	7	4	6	3	2	1	5
1975:IIQ–1977:IQ	P	1	2	5	4	6	3	7	8
	Y	1	2	7	4	3	6	5	8
	r	2	4	1	3	5	6	7	8
	U	1	2	4	3	5	6	7	8
1981:IIIQ–1983:IIQ	P	7	2	3	8	5	4	6	1
	Y	2	3	5	1	4	6	7	8
	r	4	1	2	6	5	3	7	8
	U	1	2	6	3	5	4	7	8
1983:IQ–1984:IIIQ	P	1	7	6	8	5	2	4	3
	Y	1	3	2	5	4	7	6	8
	r	1	3	2	4	5	8	7	6
	U	8	4	1	2	3	6	5	7
Total ranking		47	49	55	72	75	86	94	98

III. Model Specification

The model contains four variables: the growth rate of the GNP deflator (P), the growth rate of real GNP (Y), Moody's AAA corporate bond rate (r), and the civilian unemployment rate (*U*).[3] All variables are expressed as percentages—the growth rates being annualized. We wanted to examine how well the various model specifications estimate the scope of the expansion or recession to come because, as mentioned before, once an expansion or contraction begins, an economist needs an idea of how long the new phase of the business cycle will last.[4]

One-quarter- and eight-quarter-ahead, *ex post*, dynamic, out-of-sample forecasts were produced from two cyclical peaks: the fourth quarter of 1973 and the third quarter of 1981, and from one cyclical trough: the second quarter of 1975. For the period beginning in the first quarter of 1983, a cyclical trough, a seven-quarter-ahead forecast was made rather than one for eight quarters ahead, since revised data for the fourth quarter of 1984 were not available at the time this paper was written.

The first step in our estimation process was to perform a multivariate time series analysis on the four variables for each in-sample period. Using the techniques described in Box and Jenkins (1976) and Tiao and Box (1981), it was found that, in each period, an AR(1) or, at most, an AR(2) specification provided an adequate in-sample fit.[5] Since these models contain no moving average or lagged error terms, they closely approximate a standard VAR with one or two lags of each explanatory variable.[6] This makes our use of VAR techniques to solve the model under consideration

Fig. 2 One-Step-Ahead Errors (Absolute value)

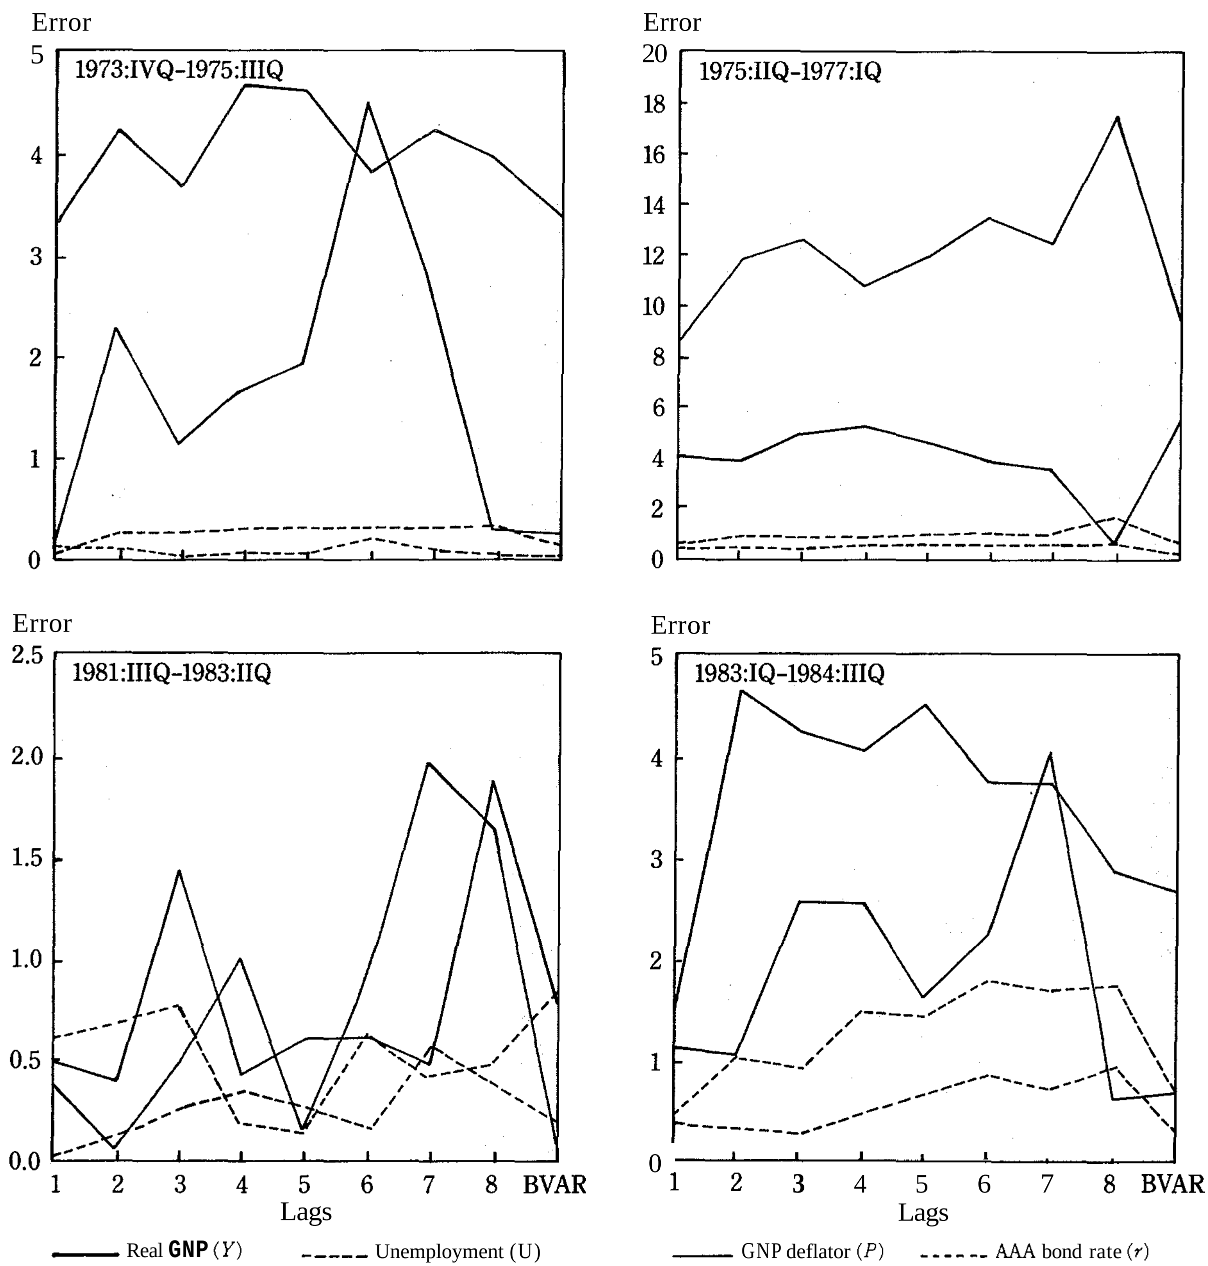

justified by these more general time series analysis procedures.

Each specification of the model consists of four OLSQ regressions. In the equations, each variable at period t is regressed against the values of all four variables at times t–1 through t–n, as well as a constant. For this paper, the lag length n ranged from one to eight. Despite the multicollinearity problems and estimation difficulties mentioned above, OLSQ estimation has been used in such seminal VAR models as that of Sims (1980). Our goal is to compare the different lag specifications to see which size of OLSQ VAR model provides the best out-of-sample forecasts of recession and recovery.

Comparing Forecast Accuracy

There are many measures of forecasting accuracy that one may use to compare different models that propose to explain the same phenomena. For this study, the following techniques were chosen:

1) To compare the one-step-ahead forecasts for each lag specification, we simply compare the absolute values of the one-step-ahead forecast errors. Here, we assume that it is just as undesirable to overestimate the actual values of the variables being forecast as to underestimate them, since either type of error can cause problems. We only want to know the degree to which the forecasts miss the mark.

2) For the seven- or eight-quarter-ahead forecasts, we look at the root mean squared errors of the forecasts for each variable. This seems to be an appropriate procedure, since we are not directly comparing forecasts of different variables.

Also, in business, as forecasts become more inaccurate, the fallout from decisions based on these forecasts increases even faster than the inaccuracy of the forecasts. The more inaccurate a forecast, the more sectors of a business' operation are affected by decisions made on the basis of the incorrect prediction. Thus, we seem justified in using a squared error measure, as opposed to a measure based on the simple difference between the actual and predicted values. Again, this implies that it is equally important to avoid overprediction and underprediction.

3) It would also seem useful to know if the longer-term forecasts consistently overestimate or underestimate the actual values of the variables we are interested in. If forecasts constantly miss the mark in the same direction, the problems caused by the decisions based on the forecasts will be compounded over time, rather than being compensated for by mistakes in the other direction. The measure used here is the bias component of the Theil U decomposition described in Theil (1961). This bias component is calculated as:

$$Bias = (Y - \overline{Y})^2/MSE,$$

5 *There might, however, be significant moving average terms in the ARIMA specification which provides the best out-of-sample fit.*

6. *Box and Jenkins (1976) show that, for moderate or large samples, the ordinary least squares estimates of the parameters of a VAR equation differ only slightly from those obtained through the Yule Walker equations used in ARIMA type analyses.*

Table 2 Rankings of Absolute Values of One-Step-Ahead Forecast Errors

Forecast period	Variables	Number of lags 1	2	3	4	5	6	7	8
1973:IVQ–1975:IIIQ	P	1	5	2	8	7	3	6	4
	Y	1	6	3	4	5	8	7	2
	r	1	2	3	5	7	6	4	8
	U	6	7	1	4	3	8	5	2
1975:IIQ–1977:IQ	P	5	3	7	8	6	4	2	1
	Y	1	3	6	2	4	7	5	8
	r	1	3	2	4	6	5	7	8
	U	1	4	3	2	5	7	6	8
1981:IIIQ–1983:IIQ	P	3	1	4	6	2	5	8	7
	Y	4	1	7	2	5	6	3	8
	r	5	7	8	2	5	6	3	8
	U	1	2	5	6	4	3	8	7
1983:IQ–1984:IIIQ	P	1	8	6	5	7	4	3	2
	Y	3	2	7	6	4	5	8	1
	r	1	3	2	5	4	8	6	7
	U	3	2	1	4	5	7	6	8
Total ranking		38	59	65	73	75	94	87	85

where

Y = mean of the forecast values of Y,

$\overline{Y}$ = mean of the actual values of Y, during the forecast period, and

Fig. 3 Theil U Biases

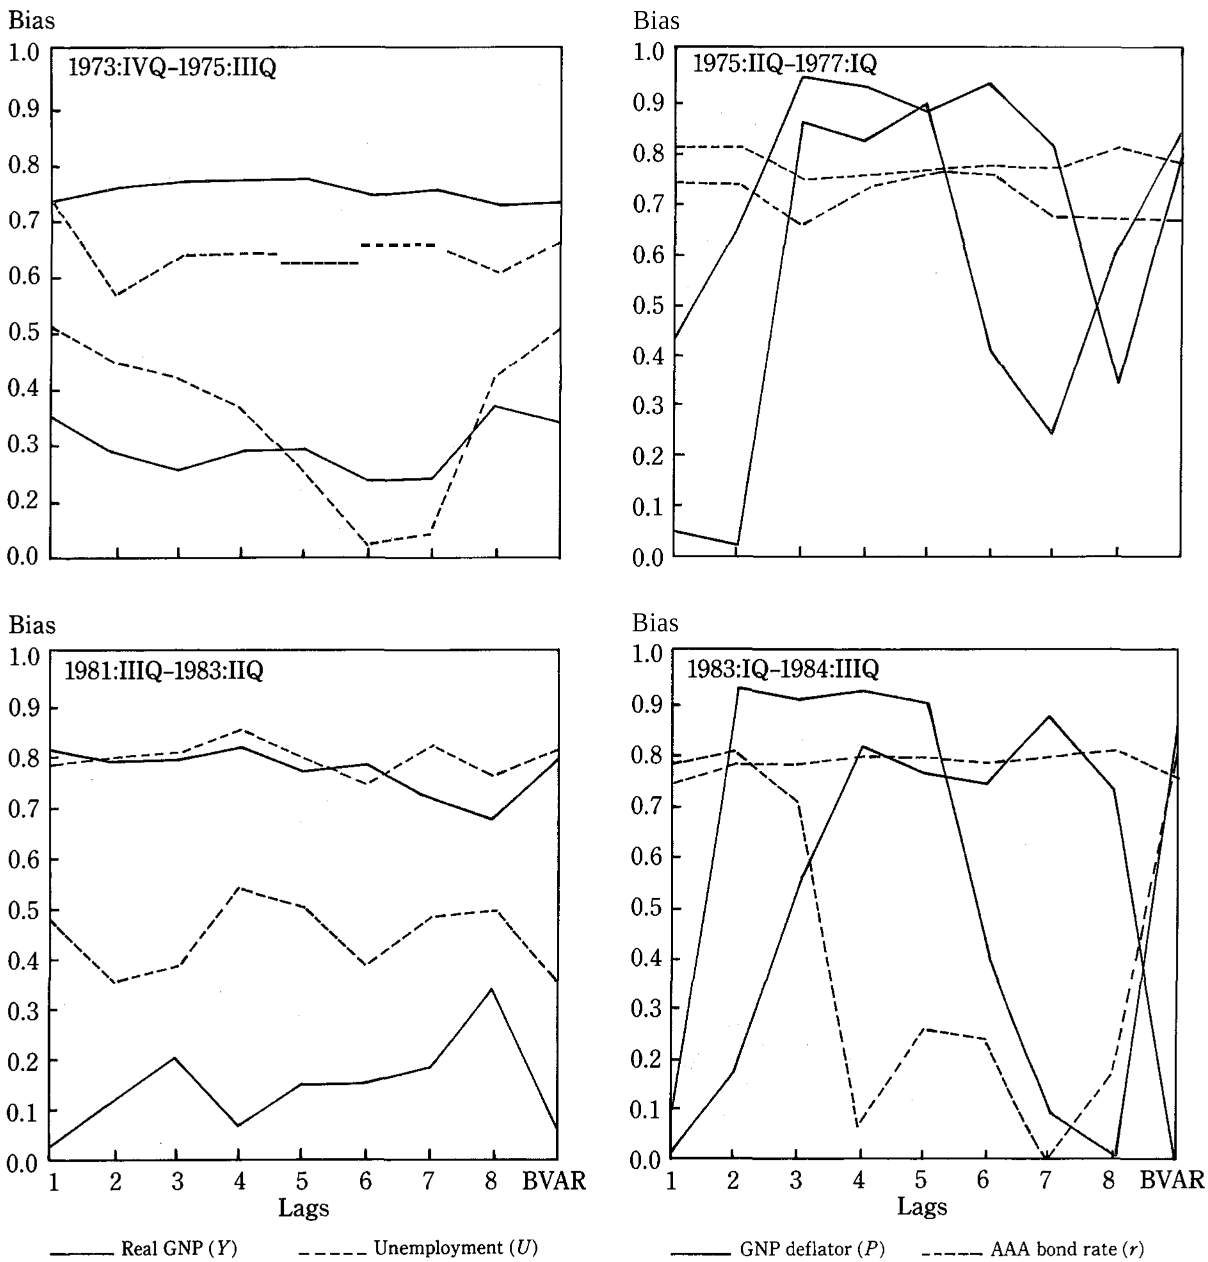

MSE = mean squared error of the forecast.

It must be noted that all of these measures of accuracy are subject to McNees's (1975) comments concerning the use of *ex post* forecasts to compare the predictive power of different models. However, McNees's critique does not apply to the VAR models examined here as much as it does to the large models he studies. With VARs, we have no exogenous variables and no subjective adjustments—two factors that McNees feels present a strong case for the use of *ex ante* forecasts when judging the comparative performance of econometric models. For our purposes, the *ex post* forecasts would seem to be appropriate.

To evaluate the rankings of the forecasts, we used the following techniques:

1) For each variable in each forecast period, the smallest error or bias is given a rank of one. The next smallest is given a rank of two and so on, the largest error or bias being assigned a rank of eight. If there is a tie, say, for the third smallest error, the tied errors are each given a rank of three, while the next largest error gets a five ranking. Since there are four forecast periods and four variables involved, we have 16 sets of rankings for each of the three accuracy measures.

2) The 16 sets of rankings for each measure are then added for each of the eight lag lengths. We thus obtain the totals of all the ranks for each lag length, one through eight. The lag length with the smallest total ranking is considered the one that forecasts the best, the length with the second smallest total ranking is considered the one that forecasts second best, and so on.

Several assumptions are implicit in this type of ranking scheme. We assume that all variables and all time periods are equally important. We also assume that the quantitative differences in error measures between forecasts are not important; we only want to know which forecast does better. It must be noted that even if two forecasts have different quantitative error measures, the difference between the measures may not be statistically significant. Ashley, Granger, and Schmalensee (1980) suggest a technique through which one can test the squared errors of forecasts from various models for such significance. However, our methodology generates only four forecasts of a given number of steps ahead for each variable in each model specification. Therefore, we do not have enough forecasts to utilize their method for comparing prediction errors. No test is currently available to examine the Theil U biases of different models for statistical significance.

Table 3 Rankings of Theil U Bias Statistics

Forecast period	Variables	Number of lags 1	2	3	4	5	6	7	8
1973:IVQ–1975:IIIQ	P	2	5	6	7	8	3	4	1
	Y	7	4	3	5	6	1	2	8
	r	8	1	3	6	5	4	7	2
	U	8	7	5	4	3	1	2	6
1975:IIQ–1977:IQ	P	2	1	7	6	8	4	3	5
	Y	2	3	8	6	5	7	4	1
	r	5	6	1	4	7	8	3	2
	U	6	8	1	2	3	5	4	7
1981:IIIQ–1983:IIQ	P	7	5	6	8	3	4	2	1
	Y	1	3	7	2	4	5	6	8
	r	4	1	2	8	7	3	5	6
	U	3	5	6	8	4	1	7	2
1983:IQ–1984:IIIQ	P	2	8	6	7	5	4	3	1
	Y	1	2	3	7	6	5	8	4
	r	1	3	2	6	5	4	7	8
	U	7	8	6	2	5	4	1	3
Total ranking		66	70	72	88	84	*63*	68	*65*

IV. Estimation Results

As the lag length increased, the in-sample fits improved. This follows directly from the theory of least squares regression, which states that as more explanatory variables are added to a model, the in-sample fit should improve or stay the same. However, the graphs and the tables of rankings show that, by the methodology described above, the out-of-sample forecasts worsened as the lag lengths increased. In the case of the seven- or eight-quarter-ahead forecasts, forecast accuracy decreased over the entire range of lag lengths, with one lag giving the best forecasts and eight lags the worst. These results are shown in table 1 and figure 1. In table 2 and figure 2 we see that, in the case of the one-step-ahead forecast errors, the one-period lag gave, by far, the most accurate predictions. The forecasts got uniformly worse, as longer lags were used, until the seven-period lags, when there was a slight improvement. For the Theil U biases, shown in table 3 and figure 3, the rankings deteriorated uniformly from one lag period to four, improved slightly with five-period lags, then returned to a level very close to that of the one-period lag for lag lengths six through eight.

In sum, these results seem to indicate that, in a vector autoregressive system estimated with OLSQ, the best forecasts of recessions and recoveries are obtained by assuming that the value of each variable depends only on the values, in the immediately preceding period, of itself and all other variables in the model. A one-lag model, in essence, restricts the coefficients for all longer lags to zero.

It is possible, however, that a forecaster may have prior information—information not reflected in the data—which indicates that some of the coefficients for variables lagged two or more periods can be nonzero. To explicitly accommodate these "priors" in a statistical model in the hope of obtaining better forecasts, we can use Bayesian vector autoregression.

V. The Bayesian VAR Method

By using the Bayesian vector autoregression (BVAR) techinque, one can include, in the model, subjective estimates of the model's parameters and measures of the forecaster's confidence in his estimates?

Very briefly, the BVAR technique involves the following steps:

1) Choose the lag structure and variables of the model. Here, we use the same variables as before (P, Y, r, and U) and regress each on the past three values of all four variables and a constant term.

2) Make an estimate of the coefficient values and your confidence in the estimates. Here, we have applied what Todd (1984) calls the Minnesota prior. The Minnesota prior assumes that all variables in each equation of the model behave according to a random walk; that is, all coefficients are zero except for the coefficient of the most recent value of the endogenous variable, which is one.

In other words, it is expected that the value of a variable at any given time equals the value of that variable in the preceding period. The Minnesota prior also assumes that one has more confidence in his estimates of the coefficients as the lag lengths get longer; the longer the lag, the more certain the forecaster is that a lagged variable has no effect on the system.[8]

3) Divide the variables of each equation into own and cross variables, where the endogenous variable of any given equation is the own variable for that equation, and all other variables in the equation are cross variables. Once this is done, scale the prior variances of the cross variables to units equivalent to those of the own variable?

4) Multiply all own and cross-variances by hyperparameters H_o and H_c, respectively, to convert the weights determined in steps two and three to estimates of the absolute prior variances. For this estimation, we set H_o at 0.1 and H_c at 0.05 for all cross variables.

5) Perform a mixed estimation simulation using the method described, for example, in Theil (1970). A further discussion of points two, three, and four may be found in Todd (1984).

When we compare the results from the Bayesian VAR with those of the OLSQ estimations, we find that the BVAR performs at a level comparable to that of the non-Bayesian VAR with one lag. The ordinals of the root mean squared errors for the longer term fore-

7. This technique is being used by the Federal Reserve Bank of Minneapolis to model and forecast economic conditions in the Ninth Federal Reserve District. The forecasts are presented in District Economic Conditions, *available free of charge from the Research Department of the Federal Reserve Bank of Minneapolis, Minneapolis, MN 55480.*

8. The Minnesota prior constrains the variance of the coefficient of any n-period lagged variable to be 1/n times the variance of the coefficient of that variable when lagged once.

9. This is done by multiplying each relative prior variance of a cross variable by s_o/s_c, where s_o is the standard error of the regression in which the own variable is the endogenous variable, and s_c is the standard error of the equation in which the cross variable is the endogenous variable.

10. The ordinal scores in table 5 for the OLSQ VARs are not strictly comparable to those presented in tables 1 to 3. In table 5, we are comparing nine specifications: eight OLSQ and one Bayesian. The Bayesian model is not ranked in tables 1 to 3.

casts show that the BVAR performs slightly better that the one period VAR estimated with OLSQ. For the one-step-ahead forecast errors, the BVAR performs better than all other specifications except for the one-period non-Bayesian VAR, which does a shade better. Finally, the Theil U bias statistics show that the BVAR forecast consistently over- or under-estimates the realized values by about the same degree as the one-, six-, seven-, or eight-period, lagged non-Bayesian VAR. The breakdown of the rankings for the BVAR is shown in table 4, while table 5 compares the BVAR performance to that of the OLSQ autoregressions.[10] Figures 1 through 3 chart the BVAR performance against that of OLSQ.

Table 4 Rankings of Bayesian VAR Model by Variable and Forecast Period

Forecast period	Variables	7-, 8-quarter-ahead RMSE	1-step-ahead forecast error	Theil U bias
1973:IVQ–1975:IIIQ	P	1	2	2
	Y	3	2	7
	r	2	2	8
	U	8	2	8
1975:IIQ–1977:IQ	P	5	9	7
	Y	2	2	4
	r	1	1	2
	U	3	2	6
1981:IIIQ–1983:IIQ	P	3	2	5
	Y	3	1	2
	r	2	9	2
	U	3	4	7
1983:IQ–1984:IIIQ	P	3	2	5
	Y	2	2	1
	r	2	2	2
	U	6	2	8
Total ranking		49	46	76

Table 5 Total Rankings of Bayesian and Ordinary Least Squares Models

Error measure	OLSQ (number of lags) 1	2	3	4	5	6	7	8	BVAR
1-step-ahead	44	71	79	87	89	95	101	98	46
RMSE	55	57	68	84	89	99	108	111	49
Theil U bias	75	79	81	98	93	70	76	72	76

VI. Conclusions and Caveats

The results indicate that, at least when the economy moves from an expansionary period to one of contraction, or vice versa, the forecasting ability of a VAR system deteriorates as longer lags are incorporated into the model. It also seems that a Bayesian estimation procedure does not produce forecasts that are substantially better than those of the non-Bayesian VAR with one lag per variable. Since the Bayesian method is more difficult to implement than the standard OLSQ technique, a forecaster using VAR techniques under these circumstances would probably want to stick with OLSQ.

Three important considerations must be noted, however, concerning these results. First, it may be that the comparative forecasting abilities of VARs with different lag specifications would change if the forecasts were made at points other than those considered here. For example, the one-lag model might not be superior to the others if the forecasts were being made in the middle of a cyclical expansion. Such an investigation might prove to be a useful topic for future work. If the one-lag specification is still the best method at any point of the business cycle, there is no need to use longer lags at any time. If this is not so, then we need a measure of when to change between different VAR specifications in forecasting.

The second issue is that a forecaster usually doesn't know when a recession or recovery has begun until several periods after the fact. Would the one-lag method still be best if applied when a forecaster became aware that the economy had taken a turn, rather than at the turn itself?

Finally, there is no guarantee that the Minnesota prior provides the best Bayesian VAR forecasts at the times we consider. The data may, in fact, be strongly rejecting the imposition of a random walk, producing biased coefficient estimates. A different set of estimates for the values of the coefficients and variances might yield even better predictions.

It must also be noted that, for many economists, it is more important to predict *when* the economy will turn than to forecast the *magnitude* of the turn. How well can VARs forecast the timing of the beginning and end of a recession compared to other small models and large econometric systems? Also, what VAR lag specification calls the timing of the turns most accurately? These questions must be addressed to better evaluate the usefulness of VAR forecasting methods.

As we have seen, VARs, while freeing one from the assumptions underlying a structural economic model, present problems of their own. However, since even the prototype BVARs, for instance, outperformed many commercial forecasters (see, for instance, Doan, Litterman, and Sims [1984]), further research on the models should prove very fruitful in clearing up our crystal balls.

References

Ashley, R., C.W.J. Granger, and R. Schmalensee. "Advertising and Aggregate Consumption: An Analysis of Causality," *Econometrica*, vol. 48, no. 5, pp. 1149–67.

Box, George E.P., and Gwilym M. Jenkins. ***Time Series Analysis: Forecasting and Control.*** Revised Edition, San Francisco, CA: Holden-Day Inc., 1976.

Christ, Carl F. ***Econometric Models and Methods.*** New York, NY: John Wiley and Sons, Inc., 1966.

Doan, Thomas, Robert Litterman, and Christopher Sims. "Forecasting and Conditional Projection Using Realistic Prior Distributions," ***Econometric Reviews,*** vol. 3, no. 1, pp. 1-100.

Intriligator, Michael D. ***Econometric Models, Techniques, and Applications.*** Englewood Cliffs, NJ: Prentice-Hall Inc., 1978.

Litterman, Robert B. "Techniques of Forecasting Using Vector Autoregressions," ***Working Paper*** 115, Federal Reserve Bank of Minneapolis, Revised November 1979.

Lucas, Robert E. "Econometric Policy Evaluation: A Critique," in K. Brunner and A. Meltzer, eds., ***The Phillips Curve and Labor Markets, Vol. 1,*** Carnegie-Rochester Conference Series on Public Policy. Amsterdam: North-Holland, 1976, pp. 19–46.

Lupoletti, William M., and Roy H. Webb. "Defining and Improving the Accuracy of Economic Forecasts: Contributions From a VAR Model," ***Working Paper*** 84-6, Federal Reserve Bank of Richmond, October 1984.

McNees, Stephen K. "An Evaluation of Economic Forecasts," *New England Economic* ***Review,*** Federal Reserve Bank of Boston, November/December 1975, pp. 3–39.

Sims, Christopher A. "Macroeconomics and Reality," *Econometrica*, vol. 48, no. 1 (January 1980), pp. 1–48.

Theil, Henri. ***Economic Forecasts and Policy,*** 2nd Revised Edition, Amsterdam: New-Holland Publishing Co., 1961.

. ***Principles of Econometrics,*** New York: John Wiley and Sons, Inc. 1970.

Tiao, G.C., and G.E.P. Box. "Modeling Multiple Time Series with Applications," ***Journal of the American Statistical Association,*** vol. 76, no. 376 (December 1981), pp. 802–16.

Todd, Richard M. "Improving Economic Forecasting With Bayesian Vector Autoregressions," ***Quarterly Review,*** Federal Reserve Bank of Minneapolis, Fall 1984, pp. 18–29.