Federal Reserve Bank of Atlanta

A QUARTERLY BAYESIAN VAR MODEL
OF THE U.S. ECONOMY

by

William Roberds

Working Paper 88-2

May 1988

Working Paper Series

A QUARTERLY BAYESIAN VAR MODEL OF THE U.S. ECONOMY

by

William Roberds

Working Paper 88-2

May 1988

William Roberds is an economist in the macropolicy section of the Atlanta Fed's Research Department. He is grateful to many of his former colleagues at the Federal Reserve Bank of Minneapolis for useful discussions concerning the BVAR approach to modeling. He would especially like to thank Robert Litterman for his patient tutoring and for generously making available his forecasting programs.

FRB Atlanta Working Paper 88-2
May 1988

A QUARTERLY BAYESIAN VAR MODEL OF THE U.S. ECONOMY

Abstract

One of the more significant developments in econometric modeling over the past decade has been the invention of the forecasting technique known as Bayesian vector autoregression (BVAR). This paper provides a detailed description of the process of specifying a BVAR model of quarterly time series on the U.S. macroeconomy. The postsample forecasting performance of the model is evaluated at an informal level by comparing the model's performance to certain naive forecasting methods, and is evaluated at a formal level by means of efficiency tests. Although the null hypothesis of efficiency is rejected for the model's forecasts, the accuracy of the model exceeds that of naive forecasting methods, and seems comparable to that of commercial forecasting firms for early quarter forecasts.

Please address all correspondence to:

William Roberds
Research Department
Federal Reserve Bank of Atlanta
104 Marietta Street, NW
Atlanta, Georgia 30303-2713

A QUARTERLY BAYESIAN VAR MODEL OF THE U.S. ECONOMY

I. Introduction

The technique of Bayesian vector autoregression (BVAR) provides a useful tool for forecasters, especially forecasters with limited time and financial resources at their disposal. In particular, BVAR methodology allows forecasters to construct inexpensive and internally consistent multivariate forecasts of essentially arbitrary sets of economic time series. BVAR techniques also offer significant advantages over competing methodologies for many applications. For example, because the model produced by BVAR is multivariate, it is relatively easy to construct conditional forecasts or scenarios using such a model. This offers an advantage over univariate methods in answering the "what if" sorts of questions often faced by applied forecasters. Also, BVAR techniques require a less extensive model specification than do traditional simultaneous equation techniques. However, considerably more judgment is called for in the use of BVAR techniques than is required in univariate approaches such as Box-Jenkins.[1] In this working paper, a detailed description of the process of specifying a BVAR macro model is provided, with the goal of offering guidance to forecasters interested in using this methodology. The model, which is suitable for implementation on a microcomputer, uses nine postwar macroeconomic time series on the U.S. economy. The BVAR methodology employed is a variation on that employed by Doan, Litterman, and Sims (1984).

BVAR Defined

Let x(t) be a vector of time t observations on N series that are to be forecast. For example, in a quarterly macroeconomic model, x(t) would contain the quarter t observations on series such as GNP, the price level, the money supply, and so forth. Then a K-th order vector autoregressive model of x(t) consists of the N equations:

$$x_i(t) = \sum_{j=1}^{N} \sum_{k=1}^{K} a_{ijk}(t)\ x_j(t-k) + c_i(t) + u(t) \tag{1}$$

where $a_{ijk}(t)$ represents the time t coefficient on the kth lag of variable j in equation i, $c_i(t)$ is the time t constant term in equation i, and u(t) is a generic white noise error term. In other words, a vector autoregression or VAR consists of N equations, in which each element of the vector x(t) is regressed on K lags of itself, on K lags of all other elements of x(t), and a constant. The form of equation (1) evidently reflects a very conservative modeling strategy. In contrast to traditional simultaneous equations models, equation (1) excludes no variable in x(t) a priori from directly influencing any other variable in x(t). Nor are any exact restrictions placed on the lag structure of equation (1). In fact, the initial impetus for the formulation of BVAR models came from Sims's (1972, 1980) forceful arguments that the exact restrictions employed in traditional simultaneous equations techniques were likely to lead to serious specification error. These difficulties are avoided in equation (1) by including in each equation all possible lags of all variables of interest.

If the coefficients in equation (1) are assumed to be constant over time, one valid method of estimating these coefficients is by ordinary least squares (OLS). However, there are some serious practical difficulties with using equation (1) as a forecasting equation when its parameters are estimated by OLS.[2] In the parlance of the Box-Jenkins literature, the model represented by equation (1) is "overparametrized"--that is, because the model has so many coefficients, too much sample information is used up in estimation of so many different effects. For example, each equation of the model described in this paper incorporates six lags of nine variables plus a constant term, resulting in 55 coefficients in each equation of the VAR. Unfortunately, only about 150 quarterly observations are available on the postwar U.S. economy, leaving relatively few degrees of freedom.

In the Bayesian VAR methodology, the problem of overparametrization of VAR models is solved by incorporating additional information on the coefficients in equation (1). Essentially this information is incorporated by means of augmenting available data on economic time series with pseudo-observations that are consistent with some generally recognized characteristics of economic time series. For example, in the model described below, a minimum of 64 pseudo-observations are added to each equation of the model. The information contained in these fake data is then incorporated into the coefficient estimates using a technique that draws heavily on Theil's (1963) technique of mixed estimation. Alternatively, these fake data may be interpreted as approximate prior restrictions on the coefficients of the VAR. The word "approximate" is key here: unlike exact restrictions, the influence of the approximate prior restrictions on estimates of the model coefficients vanishes asymptotically. This is easy to see because the approximate prior restrictions can be represented as a set of fake observations of a finite, fixed size. As the size of the real data set grows, the influence of the fake data eventually diminishes to almost nothing. Strictly speaking, BVAR is not a Bayesian technique, since it does not minimize a well-specified expected loss function over the model parameters, given the posterior distribution implied by the approximate prior restrictions and the data.[3] However, BVAR does allow forecasters to make use of prior information in VAR models where the large number of parameters would make more formal Bayesian methods impractical.

There exists a fairly large literature on VAR models and on the BVAR technique, much of it due to the seminal work of Sims and Litterman. Many, although certainly not all, of the relevant articles are cited in the reference list at the end of this paper. Interested readers are referred to those articles for further background material on BVAR. Of particular interest to applied forecasters would be McNees (1986a) and Litterman (1986), which summarize the real time forecast accuracy of various BVAR models constructed by Litterman. The rest of this paper is organized as follows.

Section I briefly describes the data series used in this particular application. Section II describes some benchmark methods of forecasting that will be used to measure the marginal contribution of BVAR to postsample forecasting accuracy. Section III describes the specification of the BVAR model. Section IV compares the BVAR forecasts with those of the benchmark models and presents efficiency tests of forecasts from the final model specification. Section V concludes and offers directions for future research.

II. Data

The model presented below tracks nine quarterly data series that reflect the aggregate state of the U.S. economy. The relatively small number of series and the quarterly data frequency were chosen so as to construct a model that could be implemented using currently available microcomputer technology. All series except those associated with GNP are available at least on a monthly basis. In the model, these monthly series were converted to quarterly series by taking geometric averages of monthly average figures.

The basic purpose of the model is to provide forecasts of real GNP, unemployment rates, and inflation. The other series in the model serve the secondary role of providing information useful in predicting future movements in the three "core" series. In addition, forecasts of these series serve as a check on the internal consistency of the model forecasts (e.g., the model should not simultaneously predict a stock market boom and a sharp downturn in real activity) and provide comparability to other sets of forecasts. The following is a list of the data series, their units of measurement, and, where appropriate, a brief discussion of some difficulties likely to be encountered in forecasting the series. Sources are also listed for nonstandard series. The dataset used in the model specification search covered the period 1948:1-1987:1.

Core series:

Series 1: Real GNP, billions of 1982$, SAAR (Symbol GNP82)

A serious problem associated with real time forecasts of GNP data is the size of the revisions to preliminary data. McNees (1986b) and Mankiw and Shapiro (1986) note that such revisions have historically been quite large. For example, Mankiw and Shapiro calculate that a historical 80 percent confidence interval for the revision from the preliminary real GNP figure to the most recently available figure would be about ±2.8 percent measured in annual growth rates, but claim these revisions are unforecastable. Mork (1987), using more advanced statistical techniques, finds a significant downward bias in preliminary real GNP announcements, apparently on the order of .5 percent measured at annual growth rates. Despite the magnitude of this problem, Lupoletti and Webb (1986) conclude that revisions in current GNP figures have little informational content for forecasts of future GNP. The specification search described below ignores this issue for reasons of tractability. The "postsample" forecasts reported below are made from data that incorporate all revisions in the GNP accounts publicly available as of May 20, 1987.

Series 2: Implicit GNP Deflator, 1982=100, SA (Symbol PGNP)

See the discussion under real GNP.

Series 3: Unemployment Rate, All Workers, SA (Symbol UNEMP)

Other series:

Series 4: Business Fixed Investment, billion 1982$, SAAR (Symbol BFI82)

See the discussion under real GNP. Standard theories of the business cycle assign an important role to investment as a determinant of future GNP.

Series 5: Monetary Base, billion $, SA (Symbol MBQ)

The monetary base series used in this study is the one published by the Board of Governors of the Federal Reserve System. This series is adjusted for reserve requirements in a different fashion from the more widely followed St. Louis Fed monetary base. Substituting the St. Louis base for this series made very little difference in the empirical performance of the model. Christiano (1986) presents evidence that among the various monetary aggregates, the relationship between the (St. Louis) monetary base and other macroeconomic variables has been the most stable over the postwar time period.

Series 6: Annualized Yield on 3-Month Treasury Bills, secondary market (Symbol TBILLS)

Series 7: Dollar Index, 1980 = 100 (Symbol ATL$)

Starting in 1973, this index is identical to the Atlanta Fed dollar index. For details on the construction of this index, see Rosensweig (1986). Koch, Rosensweig, and Whitt (1986) find that the Atlanta dollar index has substantial predictive power for the CPI. For this data set the Atlanta Fed index was extended backward to 1959 using the Board of Governors' dollar index. The series were spliced by multiplying the pre-1973 values of the Board index by the ratio of the average 1973 value of the Atlanta index to the average 1973 value of the Board index. Prior to 1959, the dollar series was extended backward using a backward extension of the Board index devised by Robert Litterman.

Series 8: Commodity Price Index, 1948:1 = 100, SA (Symbol PR28)

This series is a level index series constructed from BEA series 98, Percent Changes in Producer Prices for 28 Sensitive Crude and Industrial Materials. See U.S. Dept. of Commerce, Bureau of Economic Analysis, Handbook of Cyclical Indicators, 1984.

Series 9: Standard & Poor's 500 Stock Index, 1941-43 = 10 (Symbol STOCKS)

Data Transformations

Only one transformation was applied to the data series before estimation, i.e., all series except for the T-bill yield series were transformed by taking natural logarithms. The specification search described in Section IV below can therefore be interpreted as a search intended to produce approximate minimum mean square error forecasts of percentage changes in the logged model series. Jensen's inequality implies that these forecasts will be suboptimal in the mean square error sense when forecasting levels of the same series.[4]

It is perhaps instructive to note that in contrast to other modeling methodologies, BVAR does not generally require prefiltering of the data other than seasonal adjustment (which is done by the agency reporting the data in the case of the model series). For example, BVAR techniques do not require predifferencing of the data, even when nonstationarity is apparent in the data series. The major reason this prefiltering is not done is that such transformations are inimical to the "no exact restrictions" modeling philosophy described above. In the case of Box-Jenkins style predifferencing of the data, such predifferencing is mathematically equivalent to placing exact restrictions on the sums of the model coefficients. The BVAR technology, in contrast, allows the researcher to approximately impose the constraints implied by predifferencing of the data, with the closeness of the approximation determined by the postsample forecasting performance of

the resulting model. The final specification of a BVAR model in level data can very closely resemble a specification implied by predifferencing, but only to the extent that the forecasting performance of the model is improved by such a specification.

For similar reasons, the BVAR model described below does not attempt to exploit any cross-equation restrictions on the model coefficients implied by any possible co-integration of the model series.[5] Since co-integration of the model series would again imply exact restrictions on the model coefficients, direct imposition of these restrictions is avoided. Although it might be desirable to impose approximate co-integration restrictions on the model coefficients, the cross equation nature of these restrictions makes this computationally impractical. The incremental value of these restrictions for short-term predictive accuracy is in any case somewhat dubious. Published numerical exercises by Watson (1987) suggest that for the series under consideration, only a very slight reduction of coefficient uncertainty can be attained by imposing cointegration restrictions on a VAR system.

III. Forecasting Benchmarks

In the process of searching for the most accurate specification of the BVAR forecasting model, several different measures of postsample forecast accuracy were used. The first was the most commonly used measure of forecast accuracy, root mean square error (RMSE). The RMSE of a set of forecasts $\{P_t\}_{t=1}^{T}$ of a series $\{A_t\}_{t=1}^{T}$ is given by the simple formula:

$$\text{RMSE} = \left[\sum_{t=1}^{T}(P_t - A_t)^2\right]^{\frac{1}{2}}. \tag{2}$$

A drawback of RMSE is that it can be difficult to interpret because it is not unit free. Hence, a second measure of forecast accuracy that does not depend on units is also reported. This is the Theil U-statistic, defined as:

$$U = \text{RMSE}/\text{NRMSE}, \tag{3}$$

where NRMSE equals the RMSE of the "naive" forecast of no change in the series being forecast. Thus, a Theil U-statistic less than one indicates that the model forecast would have historically outperformed a random walk forecast in a mean squared error sense. Unfortunately, comparisons based on Theil U-statistics can also be misleading, depending on the nature of the series being forecast. For example, if the series has a strong upward trend, Theil U-statistics well below one will obtain very easily. Also, even if a random walk model is a good characterization of point samples of a given series, it will be relatively easy to construct a forecasting model that delivers Theil U-statistics less than one for time averages of the series. For this reason, a simple benchmark forecasting model using exponential smoothing techniques was constructed for each of the model series. In most cases, these models yielded forecasts more accurate than simple random walk forecasts.

The first sort of smoothing model considered was based on Working's (1960) characterization of time averages of random walk processes. From Working's results it can be shown that unit averages of a continuous time random walk have a discrete time integrated moving average or IMA(1,1) representation with moving average parameter α equal to $\sqrt{3} - 2$. As discussed in Harvey 1984 and Gardner 1985, such an IMA process can be forecast using the exponential smoothing algorithm:

$$S(t) = S(t-1) + \alpha e(t), \tag{4}$$

$$\hat{y}_t(m) = S(t), \tag{5}$$

where y(t) = the series being forecast, S(t) = the smoothed level of the series, $\hat{y}_t(m)$ = the m step ahead forecast of series y(t), and e(t) = the one step ahead forecast error using last period's forecast for y(t). If the series under consideration is correctly modeled as a unit time average of a continuous time random walk, then the Theil

U-statistics from the exponential smoothing algorithm given above will tend to the population values given below:

Forecast Horizon (Quarters Ahead)	Population Theil U-Statistic
1	.93301270
2	.97320508
3	.98325318
4	.98782049
5	.99043039
6	.99211914
7	.99330127
8	.99417502

Thus, even if changes in point samples of a given series are unforecastable, time averaging of the data will make possible a 7 percent reduction in one step ahead RMSE over a random walk forecast of the averaged series.[6]

A second smoothing model was considered for series with a strong secular trend. This smoothing model was derived from the simple state space model

$$S(t) = S(t-1) + T(t-1) + \varepsilon(t) + \delta(t), \tag{6}$$

state equations

$$T(t) = T(t-1) + \delta(t), \tag{7}$$

$$y(t) = S(t) + u(t), \quad \text{obs. equation} \tag{8}$$

where the ratios $(\sigma_\varepsilon/\sigma_\delta)^2 = 1$ and $(\sigma_u/\sigma_\delta)^2 = 2$ were assumed. This model allows the series y(t) to have a permanent component S(t) that follows a linear trend T(t) over time. The trend component varies as a simple random walk. The variance ratios were chosen so as to allow fairly rapid variation both in the permanent and trend components relative to the noise in the observation equation. Assuming all the shocks in this model to be uncorrelated, y can be forecast as follows using the "Holt" smoothing algorithm:[7]

$$S(t) = S(t-1) + T(t-1) + \alpha e(t) \tag{9}$$

$$T(t) = T(t-1) + \beta e(t) \tag{10}$$

$$\hat{y}_t(m) = S(t) + mT(t) \tag{11}$$

where e(t) and $\hat{y}_t(m)$ are as before, $\alpha \doteq 1.10395$, and $\beta \doteq .352582$.

If the series under consideration is correctly modeled by the state space model given above, then the population Theil U-statistics for levels forecasts from the Holt exponential smoothing algorithm will all be zero. This is because the state space model above implies that first differences of the series are themselves only difference stationary. However, the exponential smoothing forecasts of first differences (i.e., of growth rates for lagged series) of such a series are easily shown to have the following population Theil U-Statistics:

Forecast Horizon (Quarters Ahead)	Population Theil U-Statistic
1	.68297511
2	.80498798
3	.83577935
4	.85817308
5	.87519231
6	.88856456
7	.89934864
8	.90822964

Again, when the series is correctly modeled by the state space model given above, even a simple smoothing algorithm can result in a large reduction in RMSE over a random walk forecast of growth rates for that particular series.

Each of the above measures of forecast accuracy is univariate in character, since each of these measures only considers forecasts of one series at a time. As a system measure of forecast accuracy of the BVAR model against the benchmark system of smoothing models, the unit free multivariate measures suggested by Doan, Litterman, and Sims (1984) were used:

$$J_K = \log \{\det [\frac{1}{(T-K)} \sum_t \nu_K(t) \nu_K(t)']\} \tag{12}$$

where $\nu_K(t)$ = the vector of k step ahead forecast errors from forecasts made at time t. Comparing J_K from the BVAR model to the analogous quantity for the system of smoothing models gives a sense of the approximate percentage improvement in RMSE at forecast horizon K averaged over the nine model series.[8]

IV. Specification Search[9]

The specification search described in this section makes extensive use of a mathematical algorithm known as the Kalman filter. For readers unfamiliar with this algorithm, a brief nontechnical summary follows.[10] The Kalman filter is an algorithm for estimating the value of an unobservable state vector, i.e., in this case the model coefficients, given observations dated time 1,2, ..., t on a vector of related variables, i.e., the model variables x(t). A convenient feature of the Kalman filter is its recursive character. That is, time t estimates of the coefficients can readily be calculated from time t-1 estimates without having to replicate all the time t-1 calculations. Hence, this feature is especially useful for performing searches over out of sample forecasting accuracy.

To describe the current application of the Kalman filter more formally, let the vector b(t) represent the actual time t values of the coefficients in a typical model equation (1).[11] Following Doan, Litterman, and Sims (1984), assume that the Kalman state vector b(t) evolves over time as a random walk, i.e.,

$$b(t) = b(t-1) + e(t), \tag{13}$$

where e(t) is a white noise error term assumed to be independent of the error term u(t) of equation (1). The usual assumption that b(t) is constant over time corresponds to the special case of (13) where e(t) has zero variance. The unobservable state vector b(t) is

assumed to be related to the vector of observations on the model series x(t) through the typical model equation (1). Then for a given initial mean value of b(t), b(-1); a known initial covariance matrix of b(t), S(-1); a known covariance matrix of the coefficient shocks e(t), W; and a known value of the variance of u(t), the Kalman filter allows for recursive calculation of the linear projection of the coefficient vector b(t) on information available at time t. Since estimation of the second moment matrices and the initial mean values needed for Kalman filtering is not computationally practical in most BVAR applications, an essential part of the BVAR forecasting methodology involves use of heuristic techniques that yield some informed guesses about the value of these moments.

Initial Specification of the Model

In the initial specification for this model, the standard "Minnesota" prior mean b(-1) for b(t) was chosen by setting the mean of the coefficient on the first own lag equal to one, and all other coefficients equal to zero.[12] Lag length for all variables was truncated at six quarters. A fixed coefficient version of the model was then estimated using the data from 1948:1 to 1966:4. Some adjustments were then made to this model based on the model's postsample forecast accuracy over the period 1967:1-1987:1. For the coefficients of the kth lag of variable j in equation i, a prior standard error of the form was assumed:

$$S(i,j,k) = g\ f(i,j)\ s(j)/[s(i)\ k], \tag{14}$$

where g is a parameter reflecting the overall tightness of the prior covariance matrix; f(i,j) is a weight chosen to reflect the importance of variable j in explaining variable i; and s(i) is the standard error of a six lag univariate autoregression of variable i. The value of the parameter g was taken as .2 and the matrix F = [f(i,j)] was taken as equal to:

In Equation For	Weight of								
	GNP82	BFI82	UNEMP	MBQ	PGNP	PR28	TBILLS	STOCKS	ATL$
GNP82	1.0	0.3	0.3	0.1	0.1	0.1	0.8	0.1	0.1
BFI82	0.3	1.0	0.3	0.3	0.3	0.3	0.5	0.3	0.3
UNEMP	0.2	0.2	1.0	0.2	0.2	0.2	0.2	0.2	0.2
MBQ	0.3	0.3	0.3	1.0	0.3	0.3	0.3	0.3	0.3
PGNP	0.3	0.3	0.3	0.3	1.0	0.3	0.3	0.3	0.3
PR28	0.3	0.3	0.3	0.3	0.3	1.0	0.3	0.3	0.3
TBILLS	0.3	0.3	0.3	0.3	0.3	0.3	1.0	0.3	0.3
STOCKS	0.3	0.3	0.3	0.3	0.3	0.3	0.3	1.0	0.3
ATL$	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	1.0

The initial covariance matrices S(-1) were then assumed to be diagonal matrices with the squared values of S(i,j,k) along their main diagonal (except for the last element along the main diagonal, which was left free, corresponding to a diffuse prior). Alternatively, the prior distributions given above could be characterized as pseudo-observations on equation (1). For example, to record the information that the coefficient on the first lag of real GNP in the equation for real GNP has prior mean 1 and prior standard deviation S(1,1,1), one could construct a pseudo-observation in which the observation on current GNP82 is set equal to 1/S(1,1,1), the observation on GNP82 last period to 1/S(1,1,1), and observations on all other explanatory variables to zero (assuming the variance of u(t) is normalized to unity).

The values for g and F given above were arrived at in the following manner. The value g=.2 was found to work fairly well for U.S. postwar macro time series in a previous study (Miller and Roberds 1987). The values for f(i,j) were initially taken as f(i,j)=1 for i=j and f(i,j)=.3 for $i \neq j$. The off diagonal weights for the GNP and unemployment equations were then reduced to .1 and .2, after such reductions were found to increase

postsample forecast accuracy for these variables. Some further experimentation with the GNP equation resulted in loosening the priors for investment and unemployment back to the original specification of .3, and a substantial loosening of the prior on T-bill yields to 0.8. Loosening of the prior on T-bill yields in the investment equation also yielded a small increase in GNP forecast accuracy.

Modifications for the Random Coefficient Model

The next step in the specification search was to consider a random coefficient version of the model, i.e., a version of the model in which e(t) in equation (13) has nonzero variance. For this version of the model, the initial means of the coefficients in equations (1) were set to zero, except for coefficients on own lags. For own lags, the coefficient on the kth own lag was assumed to have initial mean $(1-\alpha)\alpha^k$, where again $\alpha = \sqrt{3} - 2$, corresponding to the first K lags of an autoregressive representation for discrete unit time averages of a continuous time random walk (see Section III). To specify the initial parameter covariance matrices S(-1), I began by taking these to be equal to the estimated covariance matrices for the fixed coefficient model as of 1966:4.[13] Four adjustments were then made to these matrices.

The first adjustment applied was to multiply all of the S(-1) matrices by a constant scale factor $\pi_1 > 1$.[14] This upward scaling of the initial covariance matrices seemed necessary to avoid "double counting" the data dated 1948:1-1966:4. The value of π_1 was chosen by experimenting with a few values of this scale factor and observing the impact of these values on the model's forecasting performance over the interval 1967:1-1987:1. The second adjustment consisted of specifying initial first and second moments for the sums of coefficients on lags of the same variable. Following the strategy of Doan, Litterman, and Sims (1984), the initial mean for the sum of coefficients on own lags was set equal to one, and on lags of other variables to zero. The net effect was to shrink the matrices S(-1) towards zero by forcing the model toward a first differences

specification. The initial second moments for the sums were set by scaling the initial variances of the summed coefficients. Following the procedure in Miller and Roberds (1987), two scaling factors were used. The factor π_2 was applied to all sums of coefficients on real GNP, investment, unemployment, and the monetary base. The factor π_3 was used to weight sums of coefficients on the deflator, commodity prices, T-bill rates, stock prices, and the dollar. In subsequent searches over the values of π_2 and π_3, large values of π_2, but relatively small values of π_3, were found to be useful in improving the model's out of sample forecasting properties. Except for the inclusion of the monetary base in the first group of variables, this finding roughly corresponds to the notion of long-run neutrality of nominal effects on real variables. In the case of the monetary base, putting a stronger weight on the sum of coefficient priors resulted in somewhat better forecasting properties in the long run, but somewhat poorer forecasts at shorter horizons.

The third modification was applied to the coefficients on lags of the deflator and the monetary base, as well as to coefficients of all variables in the equations for those two variables. In this modification, priors were specified for the sums of the accumulated coefficients of the indicated variables. Sums of accumulated coefficients were assumed to have initial mean equal to one for own lags, and zero for other lags. Initial second moments again were determined by the weighting factor π_6. The net effect again is to shrink the matrices S(-1) toward zero by forcing the model toward a second differences specification in the money and price variables.

The fourth and final modification was to inflate the prior variance of the constant term in each equation by a factor π_5, allowing for nonzero drift terms in the model equations. The covariance matrices W of the coefficient shocks were then set by scaling the modified S(-1) matrices by the factor π_4.

Final Specification

To calibrate the scaling factors π_1 through π_6, an informal search was conducted over various values of these factors, using the measures of postsample forecast accuracy discussed in Section III. After some experimentation, a final specification was selected and represented by the values:

$$\pi_1 = 4.0$$
$$\pi_2 = 500.$$
$$\pi_3 = 1.0$$
$$\pi_4 = .00001$$
$$\pi_5 = 10$$
$$\pi_6 = 1.$$

These values can be given the following interpretations. (Criteria used in arriving at this specification are described in the next section.) The value of π_1 is relatively large and reflects a fairly loose overall prior. The large value of π_2, however, indicates that this prior can be tightened quite a bit in the direction of a first differences specification for real GNP, investment, unemployment, and the monetary base. The relatively low value of π_3 indicates that the opposite is true for the other variables of the model. The value of π_5 is also relatively small, and indicates that tightening the prior towards a second differences specification for the deflator and the monetary base yields only a marginal increase in forecast accuracy. The minuscule value of the time variation parameter π_6 indicates that the final model is extremely close to a fixed coefficient specification. In summary, the final specification implied by the scale factors $\pi_1 - \pi_6$ is very close to a fixed coefficient VAR model in which real GNP, investment, unemployment, and the monetary base all enter the model in first difference form.

Variance Decomposition

A standard method of describing the dynamics of VAR models is to report a decomposition of forecast error variance, as described in Sims (1980). Accordingly, one such decomposition is reported below, based on the end-of-sample model coefficient estimates and the variance-covariance matrix of postsample one step ahead forecast errors for the last ten years of the sample. The standard method of decomposition assumes that the final coefficient estimates are known and nonstochastic, while we have assumed neither. Since the amount of time variation in the parameters is quite small, randomness in the model parameters is of little concern in interpreting the variance decomposition. However, calculations by Runkle (1987) suggest that the sampling error of the coefficient estimates may be quite large even after imposition of the model priors. In interpreting the figures given below, following a suggestion of Runkle's, a cutoff point of 10 percent is adopted for "significant" explanation of forecast error variance.

Calculation of the variance decomposition requires that the one-step ahead forecast errors be orthogonalized in some fashion. In the table below this is done by constructing an "ordering" of the model forecast errors, as described in Sims 1980. The ordering assumed was: the dollar, S&P 500, the T-bill rate, commodity prices, the deflator, the monetary base, unemployment, investment, and output (exactly the reverse of the list given in the table). This ordering is arbitrary. Since the primary purpose of the model is to provide short-run predictions, the decomposition is reported at a horizon of eight quarters instead of the standard sixteen quarters.

Decomposition of Variance, Eight Quarters Ahead

Explaining	Innovations in								
	GNP82	BFI82	UNEMP	MBQ	PGNP	PR28	TBILLS	STOCKS	ATL$
GNP82	69	3	18	2	1	4	2	2	0
BFI82	1	61	19	1	1	7	7	2	1
UNEMP	1	0	81	1	0	10	5	2	0
MBQ	5	0	3	79	1	2	2	5	2
PGNP	0	4	7	0	58	24	4	0	1
PR28	0	1	0	0	0	79	5	3	11
TBILLS	0	5	3	0	0	8	70	10	4
STOCKS	0	2	6	0	0	8	6	74	3
ATL$	0	0	0	0	0	0	16	1	83

Given the general difficulty of placing an economic interpretation on these decompositions, any such interpretation of the table above is left to the reader. However, some of the interesting statistical features of the table may be summarized as follows. Taking 10 percent as the threshold of significant feedback, most variables in the model seemed to be explained largely by their own innovation and to a much lesser extent by innovations in one other variable. This underscores the generally weak nature of relationships between economic time series at this level of aggregation. Some exceptions to this general pattern are the monetary base, stock prices, and the commodity price index. Forecast errors in the monetary base are apparently explained only by the innovations in that series, which in turn seem to contribute very little to forecast errors of other series. Stock price forecast errors also appear to be explained significantly only by innovations in stock prices, which also account for a significant proportion of forecast errors in interest rates. The most surprising aspect of the decomposition reported above is the role of commodity prices. Apparently forecast

errors in the commodity price index are significantly accounted for only by its own innovations and those in the dollar. However, innovations in commodity prices significantly account for forecast errors in the deflator and in the unemployment rate. If the threshold of significant feedback is lowered to 5 percent, then innovations in the commodity price index significantly impact all of the other model variables except for GNP and the monetary base. The relative importance of this series accords with the causality results reported in Hamilton (1983).

Compatibility Tests

Following Theil (1963), one can examine the compatibility of the model priors with the sample information by constructing for each equation Wald tests of the restrictions implied by the means of the priors. When this was first done for the priors described above, the prior restrictions were rejected for all equations. However, the value of the compatibility statistics fell drastically when the restrictions on the constant terms were dropped from the tests and when more recent information was used in constructing the model "priors." Making use of all available data and dropping the constant term restrictions, the compatibility statistics were as follows:

Variable	Compatibility of Prior Means [χ^2 (54)]	Significance
GNP82	31.3	99%
BFI82	44.9	80%
UNEMP	63.5	17%
MBQ	233.	0%
PGNP	176.	0%
PR28	92.8	.08%
TBILLS	56.4	38%
STOCKS	44.4	81%
ATL$	42.5	86%

These statistics indicate that the "continuous time random walk" priors described above are largely compatible with all the series in the model, with the exception of nominal series that are measured in levels.

V. Model Performance

In searching over the values of the model scaling factors $\pi_1 - \pi_6$, a number of somewhat inexact criteria were adopted for the model's postsample forecast accuracy over the model evaluation period 1967:1-1987:4. These are, roughly in order of importance:

(1) The model should generate forecasts of real GNP growth that improve on both the random walk and smoothing model forecasts.

(2) The model should generate forecasts of inflation and unemployment that are at least as good as the random walk and smoothing model forecasts.

(3) The forecasts of the other variables of the model should be at least as good as the random walk and smoothing model forecasts of those variables.

(4) The multivariate measures of forecast error J_k should be minimized, and should not exceed the corresponding values for a system of smoothing models.

The model's out of sample forecasting performance is summarized on a variable by variable basis in Tables 1 and 2. Table 1 presents the forecasting performance of the model over the past twenty years, i.e., the period used in setting the scale factors $\pi_1 - \pi_6$. For purposes of comparison, Table 2 presents the forecasting performance of the model over the most recent ten-year period.[15] In these tables, the RMSEs of the final BVAR specification are compared with the RMSEs of random walk forecasts and of the smoothing models described in Section III. For all variables except the T-bill interest rate and the unemployment rate, statistics are reported for both the levels and annualized rates of change. Only levels are reported for the latter two series.

Tables 1 and 2 indicate that the model is generally successful in delivering improved forecasts of real GNP growth and unemployment. The RMSEs for these variables are well below one and also below those for the smoothing models. The model is somewhat less successful in forecasting inflation. Table 1 shows that over the past 20 years, the model would have generated inflation forecasts with Theil U-statistics below one, but that these forecasts would have been outperformed slightly by a simple smoothing model. However, Table 2 shows that the model's inflation forecasting performance over the most recent ten-year period is substantially better than the smoothing model's performance, especially at longer horizons. This suggests that the model specification is flexible enough to recover from an episode as disruptive as the 1973-74 oil shock and still deliver useful forecasts.

Turning to some of the model's secondary variables, the model's performance on business fixed investment is somewhat mixed. In terms of annual growth rates, the model seems to slightly dominate the simple smoothing model in RMSE terms. On the other hand, the smoothing model seems to dominate the BVAR at short horizons in levels. However, the extremely volatile nature of this GNP component suggests that level forecasts for investment should not be taken too seriously without first applying log normal corrections to these forecasts. The model's record on forecasting the monetary base closely mirrors its record with inflation. One gets a slight overall increase or decrease in forecast accuracy over naive methods, depending on whether the first oil shock is included in the forecast evaluation period. In contrast, the model's forecasting performance for the commodity price series seems consistent over the past 20 years. That is, the model delivers improved forecasts of commodity price movements in the short run, but erratic forecasts in the long run.

Of the three financial market series included in the model, the model seems to have the most success tracking the (time averaged) T-bill rate. In both Table 1 and Table 2, the model's predictive record for T-bill rates exceeds both the theoretical and empirical

performance of the time averaged random walk model. For the S&P 500 series, the model delivers somewhat poorer forecasts than the smoothing model over the past twenty years, but somewhat better forecasts over the past ten years. Finally, the model's forecasts for the value of the dollar are largely dominated by those of the smoothing model. The values of the Theil U-statistics obtained for the last two series using the smoothing model suggest that these series are well modeled as time averages of continuous time martingales (compare Section III). For the dollar series, this result is in agreement with Meese and Rogoff's (1983) finding that point samples of exchange rates are better forecast as random walks than by conventional structural models.

Tables 3 and 4 present some multivariate summaries of the model's forecast accuracy over the periods 1967:1-1987:1 and 1977:1-1987:1, respectively. Table 3 shows that the over the past 20 years, the average gain in forecast accuracy of the BVAR methodology over simple smoothing is relatively modest, with an average reduction in RMSE on the order of 1.5 to 2.5 percent at horizons of four quarters or less. This suggests that much of the forecasting power of the BVAR model comes from its ability to capture relatively simple patterns in the data such as trends and time averaging. Table 4 shows that over the past ten years, the average gain in forecast accuracy using the BVAR methodology is somewhat more substantial, with an average reduction in RMSE on the order of 5-12 percent at horizons of four quarters or less.

Finally, Table 5 offers some comparisons of the forecasting accuracy of the BVAR model with the real time forecasting records of forecasters using traditional forecasting techniques. This is done by taking the median RMSEs of the "early quarter" forecasts reported in McNees (1986a) for the period 1980:2-1985:1, and comparing these to the "out of sample" RMSEs of the BVAR model over this period.[16] This comparison is not a strictly valid one due to the short time period covered and problems with data revisions, but does provide a rough idea of the relative accuracy of the BVAR forecasts. Table 5 suggests that the BVAR model presented above, on the whole, is only slightly less

accurate than the median commercial forecaster. A notable weak point of the BVAR forecasts over this time period is the relatively poor performance of the model at predicting real GNP at horizons of more than two quarters.

Efficiency Tests

In addition to comparisons against other forecasters, it is of interest to compare the accuracy of the BVAR forecasts against a theoretical limit implied by the mathematics of linear projections.[17] Since the theoretical lower bound on the RMSE of forecasts of the model variables is unknown, such tests must be conducted in an indirect fashion. Let x(t) again represent the 9×1 vector of actual data on the model variables, and let p(t) be the corresponding vector of time t-1 forecasts of x(t). If p(t) is the predictor that minimizes the sum of the MSEs of the model variables, the mathematics of linear projections implies that in the multivariate regression

$$x(t) = p(t)\beta + \alpha + \varepsilon(t) \tag{15}$$

the population value of the matrix β is the identity matrix, and the population value of the intercept vector α is zero.[18]

One potential difficulty in testing the hypothesis $\beta = I$, $\alpha = 0$ in equation (15) occurs because we cannot rule out conditional heteroskedasticity of the vector of error terms $\varepsilon(t)$. Hence the usual OLS inference procedures cannot be applied to equation (15), even though (15) may be estimated by OLS. However, using results from Hansen 1982, one can construct asymptotic standard errors for OLS estimates of α and β as follows.[19] Augmenting p(t) by a constant term and augmenting β with α, rewrite equation (15) as

$$x(t) = p(t)^{*}\beta^{*} + \varepsilon(t). \tag{16}$$

Then the OLS estimate of β^* has asymptotic variance-covariance matrix given by

$$\text{var}\ (\text{vec}\ (\beta^*)) \approx \hat{\Omega} = \tag{17}$$

$$[I \otimes (\sum_t p(t)^* p(t)^{*\prime})^{-1}][\sum_t \text{vec}(p(t)^* \hat{\varepsilon}(t)') \ \text{vec}(p(t)^* \hat{\varepsilon}(t)')']$$

$$\times\ [I \otimes (\sum_t p(t)^* p(t)^{*\prime})^{-1}]$$

where I has dimension equal to nine (the number of model variables), and $\hat{\varepsilon}(t)$ is the time t vector of OLS residuals. The hypothesis $\beta=I$, $\alpha=0$, may then be tested by means of Wald tests utilizing $\hat{\Omega}$.

Let W be the Wald test statistic associated with the test of the hypothesis H_0: $\beta=I$, $\alpha=0$. It is easy to show that W can be represented as the sum of two Wald statistics W_1 and W_2, where W_1 is the Wald statistic associated with the test of the null hypothesis H_A: β is diagonal. If R_1 is the selector matrix that picks out off-diagonal elements of β and elements of α from vec β^*, and R_2 selects diagonal elements of vec β^*, then the asymptotic variance-covariance matrix of diag $\beta = R_2\text{vec}\ (\beta^*)$, conditioned on H_A is

$$\hat{\Omega}_C = R_2[\hat{\Omega} - \hat{\Omega}\ R_1'[R_1\hat{\Omega}\ R_1']^{-1}R_1\hat{\Omega}\ ]R_2'. \tag{18}$$

The statistic W_2 is then the Wald statistic associated with the test of H_B: diag $\beta=\iota$, and $\alpha=0$, conditioned on H_A · W_2 is constructed in the usual fashion using $\hat{\Omega}_C$.

The statistics W_1 and W_2 have a natural interpretation if one considers the system of equations

$$x_i(t) = \gamma_i + \delta_i p_i(t) + \eta(t). \tag{19}$$

If equation (19) is estimated by instrumental variables, taking p(t) and a constant as instruments, then W_1 is the test statistic associated with a test of the overidentifying orthogonality conditions imposed by efficiency on equation (19). W_2 is then the test statistic associated with a test of unbiasedness of the predictions p(t), conditional on the validity of the overidentifying conditions.

To implement the tests described above, equation (19) was estimated over the 30-year period 1957:1-1987:1. The values of W_1, W_2, and W were computed as

$$W_1 = 842.1$$

$$W_2 = 1366.4$$

$$W = 2208.5.$$

W_1, W_2, and W should be distributed as χ^2 random variables with degrees of freedom equal to 72, 18, and 90, respectively. The null hypotheses H_0, H_A, and H_B can therefore be rejected at even minuscule significance levels. The rejection of H_B calls into question even the unbiasedness of the model forecasts.

The tests described above are powerful in the sense that they are joint tests of the numerous restrictions imposed by the hypothesis of efficiency of the model forecasts. But they are also relatively uninformative in the sense that they offer little insight into possible sources of inefficiency in the model forecasts. However, some insight into this problem is available if one estimates the equations

$$x_i(t) = \gamma_i + \delta_i p_i(t) + \emptyset_i(x_i(t-1) - p_i(t-1)) + \eta_i(t). \qquad (20)$$

If the forecasting model's predictions $p_i(t)$ are minimum MSE forecasts of $x_i(t)$, then the coefficient vector $(\gamma_i, \delta_i, \emptyset_i)$ should equal (0, 1, 0). To test these hypotheses, equation (20) was estimated by OLS. Wald tests of the hypotheses were then conducted using White's (1980) heteroskedasticity-consistent estimator of the variance-covariance matrix of $(\hat{\gamma}_i, \hat{\delta}_i, \hat{\emptyset}_i)$. For the period 1957:1-1987:1, equation-by-equation

efficiency was rejected for the GNP, unemployment, and monetary base forecasts at the 5 percent level. In no case was efficiency rejected at the 1 percent level. For the period 1967:1-1987:1, equation-by-equation efficiency is rejected only at the 5 percent level for the dollar forecasts only. For the period 1977:1-1987:1, efficiency is not rejected for any equation. These results are consistent with the general tendency of the model's forecasts to improve over time.

VI. Conclusion

The analysis presented in Section V suggests that the following conclusions may be drawn for the BVAR model presented in Section IV:

(1) The model generally produces better forecasts than naive forecasting methods for all model variables with the exception of the dollar and possibly stock prices. This is true even when the benchmark naive methods exploit such features of the data series as time aggregation and trends. Moreover, the relative forecasting accuracy of the model has increased in recent years.

(2) The forecasting accuracy of the model seems comparable to that of conventional econometric forecasts, at least for early quarter forecasts.

(3) Over the past 20-year period, it is difficult to reject efficiency of the model forecasts on a series-by-series basis. However, system efficiency is strongly rejected for the past 30 years, indicating a large potential for improvement in the model's forecasting performance.

Given that the current model is very close to a time-invariant parameter specification, one possibly fruitful avenue for future research would be to pursue methods by which more time variation could be introduced in the model parameters. In the engineering literature, a number of filtering algorithms have been proposed that allow for more systematic feedback from the realizations of observable model variables to the model state vector, i.e., the parameter vector b(t) in the current application.[20]

Unfortunately most of these "adaptive" algorithms appear not to be useful in a macroeconomic context due to dimensionality/degrees of freedom considerations. Some modification of these techniques, however, may yield some noticeable improvement in the model's forecast accuracy.

A more conventional path for future research would be to explore ways in which additional information could be used to increase the forecasting accuracy of the BVAR model. In particular, the analysis above has not investigated the forecasting performance of the model conditional on any information that might be available between quarterly GNP announcements. Also, the potential for combination of forecasts from the BVAR model and other forecasts has not been explored.[21] These and other modifications of the BVAR methodology present a rich agenda for future research.

Footnotes

1. See Hoehn, Fomby, and Gruben (1984).

2. On this issue see Fair (1979) and Lupoletti and Webb (1986).

3. See Doan, Litterman, and Sims (1984) and Litterman (1979). Litterman shows that (non-time-varying) BVAR corresponds to Theil's mixed estimation procedure. Thus, BVAR on an equation-by-equation basis can be given the usual Bayesian rationale for mixed estimation, i.e., that it maximizes a normal approximation to the posterior likelihood of the coefficients aijk, assuming a diffuse prior on the variance of the equation error term.

4. See for example, DeGroot (1970).

5. Two first difference stationary processes are said to be co-integrated of order (1,1) if some nontrivial linear combination of the two is a stationary process. For example, if two processes x and y have a bivariate AR(1) representation in which every coefficient on lagged x and y are equal to .5, then x and y are co-integrated with co-integrating vector (1,-1), since x-y is white noise. See Engle and Granger (1987) on co-integration and its implications for VAR models.

6. That is, a 7 percent reduction in RMSE is possible using only the past history of the time averaged values of the series. Further reductions will in general be possible if past observations on other series are available.

7. Again see Gardner (1985) and Harvey (1984).

8. This approximation is exact only if the forecast errors of the model series are mutually uncorrelated.

9. The model specification described in this section is derived from the specification in Miller and Roberds (1987), which is a slight modification of that described in Doan, Litterman, and Sims (1984). All computations were performed using the RATS (Regression Analysis of Time Series) program, written by Thomas Doan and Robert Litterman. A diskette containing RATS input and data files for the model is available on request from the author.

10. More rigorous descriptions of the Kalman filter can be found in numerous engineering texts, for example, Goodwin and Sin (1984). The updating formulas used in BVAR applications of the Kalman filter (see Doan and Litterman (1986) differ slightly from the formulas given in most engineering presentations of the Kalman filter. The difference results from the fact that in BVAR, the Kalman filter is used to estimate the current value of the Kalman state vector (i.e., the equation coefficients) rather than to estimate next period's state vector.

11. In BVAR applications, the Kalman filter is applied one equation at a time. System Kalman filtering seems impractical at the current time due to dimensionality considerations.

12. See Todd (1984) or Litterman (1986).

13. Real time applications of the model would use all available information in the calibration of the S(-1) matrices. In the specification search for the model, this was not done, so as to avoid overstating the model's postsample forecast accuracy.

14. Such scaling factors are also known as "hyperparameters"; see Doan, Litterman, and Sims (1984).

15. For the simulation results reported in Table 2, the model priors were recalibrated using data through 1976:4.

16. In calculating the median RMSEs for Table 5, the forecasting record attributed to Robert Litterman's BVAR model was omitted. The interested reader is referred to McNees (1986a) for an assessment of the forecast accuracy of Litterman's BVAR models. The model described above generally seems to be more accurate than Litterman's models for nominal series, but less accurate than Litterman's model for real series. The same conclusion holds for comparisons of the present model with Lupoletti and Webb's (1986) non-Bayesian VAR model. For the simulation results reported in Table 5, the model priors were recalibrated using data through 1980:1.

17. See Luenberger (1969, chapters 3 and 4) for a summary of linear projections theory. As explained at the beginning of Section V, the model is not intended specifically to minimize the sum of MSEs of the model series, but rather to meet a number of inexact criteria. It is nonetheless informative to run efficiency tests on the model's forecasts to give some idea of how far the forecasts are from this absolute criterion.

18. If the RHS of equation (15) contained terms involving lagged prediction errors, e.g., x(t-1) - p(t-1), then efficiency would also imply that the coefficients on such variables must equal zero. Unfortunately, there are insufficient degrees of freedom in the model data set to allow testing of these additional restrictions.

19. Hansen's results require stationarity of x(t) and p(t). Accordingly, the series for GNP, investment, the deflator, the monetary base, and commodity prices were first differenced before estimation of equation (15).

20. See Goodwin and Sin (1984, chapter 9) for a survey of these algorithms.

21. Lupoletti and Webb (1986) present evidence suggesting that affine combinations of forecasts from a (non-Bayesian) VAR and those of commercial forecasting firms may outperform any single set of forecasts.

References

Christiano, L.J. 1986. "Money and the U.S. Economy in the 1980's: A Break From the Past?" Federal Reserve Bank of Minneapolis Quarterly Review 10 (Summer): 2-13.

DeGroot, M.H. 1970. Optimal Statistical Decisions. New York: McGraw-Hill.

Doan, T.A., and R.B. Litterman. 1986. RATS Version 2.00 User's Manual. Evanston, Ill.: VAR Econometrics.

Doan, T.A., R.B. Litterman, and C. Sims. 1984. "Forecasting and Conditional Projection Using Realistic Prior Distributions." Econometric Reviews 3: 1-100.

Engle, R.F., and C.W.J. Granger. 1987. "Co-integration and Error Correction: Representation, Estimation, and Testing." Econometrica 55 (March): 251-76.

Fair, R.C. 1979. "An Analysis of the Accuracy of Four Macroeconomic Models." Journal of Political Economy 87: 701-18.

Gardner, E.S. 1985. "Exponential Smoothing: The State of the Art." Journal of Forecasting 4: 1-28.

Goodwin, G.C., and K.S. Sin. 1984. Adaptive Filtering Prediction and Control. Englewood Cliffs, N. J.: Prentice-Hall.

Harvey, A.C. 1984. "A Unified View of Statistical Forecasting Procedures" Journal of Forecasting 3: 245-75.

Hansen, L.P. 1982. "Large Sample Properties of Generalized Method of Moment Estimators." Econometrica 50 (July): 1029-54.

Hoehn, J.G., W.C. Gruben, and T.B. Fomby. 1984. "Some Time Series Methods of Forecasting the Texas Economy." Federal Reserve Bank of Dallas Research Paper No. 8402, April.

Koch, P.D., J.A. Rosensweig, and J.A. Whitt. 1986. "The Dynamic Relationship between the Dollar and U.S. Prices: An Intensive Empirical Investigation." Federal Reserve Bank of Atlanta Working Paper 86-5, June.

Litterman, R.B. 1979. "Techniques of Forecasting Using Vector Autoregressions." Federal Reserve Bank of Minneapolis Working Paper No. 115, November.

__________. 1986. "Forecasting with Bayesian Vector Autoregressions: Five Years of Experience." Journal of Business and Economic Statistics 4 (January): 25-38.

Luenberger, D.G. 1969. Optimization by Vector Space Methods. New York: John Wiley.

Lupoletti, W.M., and R.H. Webb. 1986. "Defining and Improving the Accuracy of Macroeconomic Forecasts: Contributions From a VAR Model." Journal of Business 59 (April): 263-85.

Mankiw, N.G., and M.D. Shapiro. 1986. "News or Noise? An Analysis of GNP Revisions." National Bureau of Economic Research Working Paper No. 1939, June.

McNees, S.K. 1986. "Forecasting Accuracy of Alternative Techniques: A Comparison of U.S. Macroeconomic Forecasts." Journal of Business and Economic Statistics 4 (January): 5-15.

______. 1986. "Estimating GNP: The Trade-off between Timeliness and Accuracy." New England Economic Review (January/February): 4-10.

Meese, R., and K. Rogoff. 1983. "Empirical Exchange Rate Models of the Seventies: Are Any Fit to Survive?" Journal of International Economics 14 (February): 3-24.

Miller, P.J. and W. Roberds. 1987. "The Quantitative Significance of the Lucas Critique." Federal Reserve Bank of Minneapolis Staff Report 109, July.

Rosensweig, J.A. 1986. "The Atlanta Fed Dollar Index and Its Component Sub-indexes." Federal Reserve Bank of Atlanta Working Paper 86-7, August.

Runkle, D.E. 1987. "Vector Autoregressions and Reality." Journal of Business and Economic Statistics 5 (October): 437-42.

Sims, C.A. 1972. "The Role of Approximate Prior Restrictions in Distributed Lag Estimation." Journal of the American Statistical Association 67 (March): 169-75.

______. 1980. "Macroeconomics and Reality." Econometrica 48 (January): 1-48.

Theil, H. 1963. "On the Use of Incomplete Prior Information in Regression Analysis." Journal of the American Statistical Association 58: 401-14.

Todd, R.M. 1984. "Improving Economic Forecasting with Bayesian Vector Autoregression." Federal Reserve Bank of Minneapolis Quarterly Review 8 (Fall): 18-29.

Watson, M.W. 1987. "Comments on 'Vector Autoregressions and Reality,' by D.E. Runkle." Journal of Business and Economic Statistics 5 (October): 451-53.

White, H. 1980. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and Direct Test for Heteroskedasticity." Econometrica 48 (May): 817-38.

Working, H. 1960. "Note on the Correlation of First Differences of Averages in a Random Chain." Econometrica 28 (October): 916-18.

Table 1

Forecast Statistics for a BVAR Model over 1967:1-1987:1*

Series	BVAR Statistics RMSE	BVAR Statistics Theil U	Smoothing Model Statistics RMSE	Smoothing Model Statistics Theil U		Approximate Improvement in RMSE** (percent)
GNP82						
Horizon						
1 qtr.	29.2	.81	31.4	.88	{Level}	7.3
	4.17	.80	4.42	.84	{Growth}	5.8
2 qtrs.	45.2	.74	53.2	.87		16.3
	4.14	.78	4.64	.87		11.4
3 qtrs.	61.8	.72	78.1	.91		23.4
	4.47	.74	5.02	.84		11.6
4 qtrs.	74.1	.68	102.	.94		32.0
	4.34	.73	5.12	.86		16.5
8 qtrs.	106.	.57	205.	1.09		66.0
	4.52	.65	5.57	.80		20.9
BFI82						
Horizon						
1 qtr.	9.66	.97	9.62	.97	{Level}	-0.4
	10.4	.85	11.0	.90	{Growth}	5.6
2 qtrs.	15.5	.94	15.5	.94		0.0
	10.8	.90	11.7	.98		8.0
3 qtrs.	22.9	.98	22.6	.97		-1.3
	11.4	.75	11.8	.77		3.4
4 qtrs.	27.4	.94	28.1	.97		2.5
	11.2	.75	11.9	.80		6.1
8 qtrs.	35.5	.78	45.5	.99		24.8
	11.4	.61	12.0	.65		5.1

UNEMP

Horizon						
1 qtr.	.303	.75	.349	.86	{Level}	14.1
2 qtrs.	.593	.81	.683	.94		14.1
3 qtrs.	.852	.84	.971	.96		13.0
4 qtrs.	1.12	.89	1.22	.98		8.6
8 qtrs.	1.72	.92	1.85	1.00		7.3

MB

Horizon						
1 qtr.	.678	.26	.681	.26	{Level}	0.4
	2.11	.94	2.11	.94	{Growth}	0.0
2 qtrs.	1.25	.24	1.28	.25		2.4
	2.39	1.00	2.43	1.02		1.7
3 qtrs.	1.94	.25	2.03	.27		4.5
	2.76	.93	2.17	.93		-24.1
4 qtrs.	2.71	.27	2.85	.28		5.0
	2.96	.87	2.98	.88		0.7
8 qtrs.	5.94	.31	6.47	.34		8.5
	3.08	.80	3.25	.85		5.4

PGNP

Horizon						
1 qtr.	.292	.26	.298	.27	{Level}	2.0
	1.89	.93	1.87	.92	{Growth}	-1.1
2 qtrs.	.533	.24	.560	.25		4.9
	2.18	.92	2.14	.91		-1.9
3 qtrs.	.745	.25	.870	.26		15.5
	2.26	.91	2.29	.92		1.3
4 qtrs.	.971	.22	1.23	.28		23.6
	2.43	.94	2.45	.95		0.8
8 qtrs.	2.46	.27	3.15	.35		24.7
	3.25	1.00	3.06	.94		-6.0

PR28

Horizon						
1 qtr.	6.08	.74	7.08	.86	{Level}	15.2
	11.7	.95	13.5	1.09	{Growth}	14.3
2 qtrs.	12.4	.83	14.0	.93		12.1
	14.7	.91	15.6	.97		5.9
3 qtrs.	19.7	.94	20.0	.96		8.9
	17.3	.91	15.7	.83		-9.7
4 qtrs.	27.8	1.06	25.5	.97		-8.6
	18.2	.83	15.7	.72		-14.8
8 qtrs.	52.0	1.29	39.9	.99		-26.9
	16.4	.75	16.0	.74		-2.5

TBILLS

Horizon						
1 qtr.	1.03	.93	1.07	.98	{Level}	3.8
2 qtrs.	1.63	.97	1.71	1.03		4.8
3 qtrs.	1.80	.95	1.93	1.01		7.0
4 qtrs.	2.05	.94	2.20	1.01		7.1
8 qtrs.	3.06	.96	3.23	1.01		5.4

STOCKS

Horizon						
1 qtr.	7.52	.91	7.67	.93	{Level}	2.0
	27.8	.85	27.9	.85	{Growth}	0.4
2 qtrs.	12.8	.96	12.7	.96		-0.8
	31.1	.79	29.4	.75		-5.6
3 qtrs.	17.3	.99	16.9	.97		-2.3
	31.6	.77	29.5	.72		-6.9
4 qtrs.	20.9	.98	20.8	.97		-0.5
	29.2	.75	29.7	.76		1.7
8 qtrs.	33.0	1.09	30.1	1.00		-9.2
	32.0	.77	30.2	.73		-5.8

ATL$

Horizon						
1 qtr.	2.61	.98	2.48	.93	{Level}	5.1
	9.05	.92	8.51	.86	{Growth}	-6.2
2 qtrs.	4.66	1.03	4.30	.96		-8.0
	9.88	.85	9.19	.75		-7.2
3 qtrs.	6.43	1.06	5.90	.97		-8.6
	9.98	.81	9.25	.75		-7.6
4 qtrs.	7.88	1.07	7.21	.98		-8.9
	9.57	.81	9.31	.78		-2.8
8 qtrs.	13.1	1.17	11.1	1.00		-16.6
	10.8	.81	9.53	.72		-12.5

* Except for the unemployment and T-bill yield series, two sets of forecast statistics are presented for each variable at each horizon. These correspond to the models forecasts errors in (1) levels and (2) percent change at an annual rate.

** Approximate percentage improvement in RMSE of the BVAR over the naive model, calculated by taking differences of the logarithms of the RMSEs.

Table 2

Forecast Statistics for a BVAR Model over 1977:1-1987:1*

Series	BVAR Statistics RMSE	Theil U	Smoothing Model Statistics RMSE	Theil U		Approximate Improvement in RMSE** (percent)
GNP82						
Horizon						
1 qtr.	30.3	.76	34.1	.86	{Level}	11.8
	3.87	.77	4.32	.86	{Growth}	11.0
2 qtrs.	46.9	.68	58.9	.86		22.8
	3.85	.70	4.67	.86		19.3
3 qtrs.	62.1	.65	85.8	.90		32.3
	4.15	.71	4.95	.83		17.6
4 qtrs.	74.3	.62	102.	.94		31.7
	4.34	.73	5.04	.83		15.0
8 qtrs.	111.	.55	226.	1.11		71.0
	3.89	.55	5.27	.75		30.4
BFI82						
Horizon						
1 qtr.	12.5	1.01	12.2	.99	{Level}	-2.4
	12.6	.84	12.9	.87	{Growth}	2.4
2 qtrs.	19.2	.97	18.5	.94		-3.7
	12.6	.98	13.2	1.02		4.7
3 qtrs.	28.4	1.01	27.3	.97		-4.0
	13.0	.72	13.2	.73		1.5
4 qtrs.	33.3	.96	33.5	.97		0.6
	12.6	.76	13.2	.81		4.7
8 qtrs.	37.7	.71	52.7	.99		33.5
	12.4	.66	12.7	.68		2.4

UNEMP

Horizon						
1 qtr.	.287	.73	.345	.88	{Level}	18.4
2 qtrs.	.521	.72	.671	.93		25.3
3 qtrs.	.734	.73	.971	.96		28.0
4 qtrs.	.967	.80	1.23	.97		24.1
8 qtrs.	1.67	.87	1.90	.99		12.9

MB

Horizon						
1 qtr.	.843	.25	.853	.25	{Level}	1.2
	2.14	.96	2.16	.92	{Growth}	0.9
2 qtrs.	1.52	.23	1.59	.24		5.1
	2.38	.96	2.47	1.00		3.7
3 qtrs.	2.34	.24	2.50	.25		6.6
	2.57	.92	2.77	.99		7.5
4 qtrs.	3.28	.25	3.56	.27		8.2
	2.79	.87	3.04	.94		8.6
8 qtrs.	7.21	.29	8.64	.35		18.1
	3.13	.73	3.89	.90		21.7

PGNP

Horizon						
1 qtr.	.310	.23	.340	.25	{Level}	9.2
	1.47	.86	1.55	.91	{Growth}	5.3
2 qtrs.	.527	.20	.630	.23		17.9
	1.63	.98	1.74	1.05		6.5
3 qtrs.	.735	.18	1.01	.25		31.8
	1.53	.79	2.00	1.04		26.8
4 qtrs.	.975	.18	1.47	.27		41.1
	1.58	.76	2.22	1.06		34.0
8 qtrs.	2.49	.23	4.05	.36		48.6
	2.26	.75	3.10	1.03		31.6

PR28

Horizon						
1 qtr.	6.28	.75	7.35	.88	{Level}	15.7
	8.17	.83	10.1	1.01	{Growth}	21.2
2 qtrs.	12.1	.80	14.2	.94		16.0
	10.1	.79	11.6	.90		13.8
3 qtrs.	17.4	.83	20.1	.96		14.4
	10.8	.79	11.6	.85		7.1
4 qtrs.	23.9	.90	25.7	.97		7.3
	12.1	.75	11.8	.73		-2.5
8 qtrs.	49.2	1.23	39.6	.99		-21.7
	13.0	.79	10.6	.64		-20.4

TBILLS

Horizon						
1 qtr.	1.29	.95	1.35	.98	{Level}	4.5
2 qtrs.	2.04	1.00	2.14	1.04		4.8
3 qtrs.	2.20	.96	2.36	1.02		7.0
4 qtrs.	2.49	.94	2.70	1.01		8.1
8 qtrs.	3.71	.94	4.02	1.02		8.0

STOCKS

Horizon						
1 qtr.	8.87	.88	9.30	.93	{Level}	4.7
	28.7	.86	29.4	.87	{Growth}	2.4
2 qtrs.	14.3	.90	15.2	.95		6.1
	30.7	.83	31.2	.84		1.6
3 qtrs.	19.3	.90	20.6	.96		6.5
	32.4	.79	31.5	.76		-2.8
4 qtrs.	24.0	.89	26.1	.97		8.4
	31.4	.77	32.0	.79		1.9
8 qtrs.	33.5	.84	39.4	.99		16.2
	31.5	.84	32.8	.87		4.0

ATL$

Horizon						
1 qtr.	2.78	.90	2.80	.91	{Level}	0.7
	9.49	.90	9.66	.92	{Growth}	1.8
2 qtrs.	5.11	.96	5.03	.94		-1.6
	10.6	.85	10.7	.85		0.9
3 qtrs.	7.36	1.00	7.12	.97		-3.3
	10.9	.79	10.8	.78		-0.9
4 qtrs.	9.40	1.03	8.90	.98		-5.5
	11.1	.81	11.0	.80		-0.9
8 qtrs.	15.4	1.16	13.3	1.00		-14.7
	12.0	.78	11.1	.72		-9.6

* Except for the unemployment and T-bill yield series, two sets of forecast statistics are presented for each variable at each horizon. These correspond to the models' forecasts errors in (1) levels and (2) percent change at an annual rate.

** Approximate percentage improvement in RMSE of the BVAR over the naive model, calculated by taking differences of logarithms of the RMSEs.

Table 5

Root Mean Squared Forecast Errors 1980:2-1985:1
RMSE of BVAR Model/ Median RMSE of "Early Quarter" Forecasts
Listed in McNees (1986a)*

Variable	Horizon in Quarters			
	1	2	3	4
GNP Deflator	1.55/ 1.60	1.97/ 1.35	1.76/ 1.50	1.72/ 1.65
Real Business Fixed Investment	12.08/ 11.45	9.90/ 10.75	10.76/ 10.50	11.10/ 10.25
Real GNP	4.09/ 3.85	3.43/ 3.10	4.04/ 3.00	4.49/ 2.75
90-Day Treasury Bill Rates	1.76/ 1.40	2.52/ 2.50	2.51/ 2.80	3.03/ 3.10
Unemployment Rate**	.35/ .30	.55/ .60	.69/ .95	.95/ 1.25

* Errors are in annualized percentage growth rates except unemployment and T-bill rates in percent.

** Not strictly comparable because the model unemployment series is for all workers while the other forecasts are for the civilian unemployment rate.

Table 3

Log determinants of k-step ahead forecast error covariance matrices over 1967:1-1987:1. Forecast errors are of log levels except for the T-bill rate, which is in levels.

Forecast Horizon k	BVAR Log Determinant	Smoothing Model Log Determinant	Approximate Improvement in RMSE*
1	-65.591	-65.314	1.54
2	-56.418	-56.035	2.13
3	-50.948	-50.732	1.20
4	-47.537	-47.140	2.21
8	-39.564	-38.385	6.55

* Approximate average percentage reduction in RMSE of the BVAR forecasts over the naive model, obtained by taking the difference in log determinants and multiplying by 5.56 (divide by 18 to get standard errors for 9 variables and multiply by 100 to get percent).

Table 4

Log determinants of k-step ahead forecast error covariance matrices over 1977:1-1987:1. Forecast errors are of log levels except for the T-bill rate, which is in levels.

Forecast Horizon k	BVAR Log Determinant	Smoothing Model Log Determinant	Approximate Improvement in RMSE*
1	-67.334	-66.308	5.70
2	-59.208	-57.580	9.04
3	-54.183	-52.191	11.07
4	-50.865	-48.764	11.67
8	-43.424	-41.292	11.84

* Approximate average percentage reduction in standard error of the forecast obtained by taking the difference in log determinants and multiplying by 5.56 (divide by 18 to get standard errors for 9 variables and multiply by 100 to get percent).

FEDERAL RESERVE BANK OF ATLANTA
INCORPORATED MAY 18, 1914
6-F