Federal Reserve Bank of Minneapolis
Research Department Staff Report 92

March 1984

SPECIFYING VECTOR AUTOREGRESSIONS FOR MACROECONOMIC FORECASTING

Robert B. Litterman

Federal Reserve Bank of Minneapolis

ABSTRACT

This paper describes a Bayesian specification procedure used to generate a vector autoregressive model for forecasting macroeconomic variables. The specification search is over parameters of a prior. This quasi-Bayesian approach is viewed as a flexible tool for constructing a filter which optimally extracts information about the future from a set of macroeconomic data. The procedure is applied to a set of data and a consistent improvement in forecasting performance is documented.

This paper was prepared for the volume "Bayesian Inference and Decision Techniques with Applications: Essays in Honor of Bruno de Finetti."

1. Introduction

This paper describes a Bayesian solution to the problem of specifying a vector autoregression for forecasting macroeconomic variables. Although the specification process described here is unorthodox, it does have antecedents in two areas: the work of Hoerl and Kennard [4], Stein [14], Shiller [10] and Leamer [5, 6] on shrinkage estimation and its Bayesian interpretation; and the work of Sims [11-13] and Litterman [7, 8] on specifying loosely parameterized vector autoregressions. The motivation for this approach is described first, followed by an example that demonstrates how it can be used to improve macroeconomic forecasting.

2. The Specification Problem

Although it does not usually get much attention, the specification process is a critical stage in empirical analysis. In specifying a model, a researcher constructs the instrument that will be used to filter information from the data. Bayesian statistics focuses on the process of combining information in a prior with information from a set of data, and thus a Bayesian approach to the specification process naturally gives an important role to the construction of an informative prior. A Bayesian specification search is described here in which the construction of a prior is viewed as a flexible means of generating a filter for the optimal extraction of information from a set of data.

The particular specification described here generates a set of loosely parameterized equations designed for forecasting macroeconomic variables. It differs significantly from the usual approaches to the specification of macroeconomic models, which adopt a much narrower focus. A specific behavior, often the solution to a particular optimization problem, is the guide to finding the right specification for an equation in most macroeconometric applications. This is a reasonable approach to explain many individual economic phenomena; but it seems likely to fail to explain aggregate economic behavior, which reflects the solutions of many different agents to a multitude of different problems. With respect to the cause and propagation of business cycles, for example, numerous explanations have been suggested, most probably containing a degree of truth and many probably having had empirical relevance at one time or another. It is doubtful, however, that any one paradigm can explain more than a very small part of the behavior of any macroeconomic time series.

For forecasting purposes, it seems quite reasonable to suppose that small bits of useful information concerning the macroeconomy are scattered throughout the data, and a narrowly focused approach is thus unlikely to find much useful information. The problem of econometric models is to filter out as much of the information as possible from the data and to give each little bit an appropriate weight. The proposed solution for the problem of filtering information useful for forecasting is motivated by two assertions, which I take to be self-evident, concerning the location and amount of information available for macroeconometric forecasting. The first assertion is that there is a very low signal-to-noise ratio in macroeconomic data, by which I mean that the predictable movement in such variables based on their own past is only a small fraction of their total variation. The second assertion is that the current state of macroeconomic theory leaves a great deal of uncertainty concerning which economic structures are useful for macro-forecasting, particularly at the level of detail that is necessary for short-run prediction.

From a Bayesian perspective, the latter assertion is that economic theory gives small positive probability to a large number of economic structures, each of which can be represented as an equation with a flat prior distribution over a wide range of parameter values. If that is true, then the true nature of prior information may be badly misrepresented by the usual approach of focusing on a particular economic theory and (unless it is rejected by the data) imposing the restrictions it suggests. Nonetheless, such restrictions are used because they are considered necessary to reduce the number of estimated parameters.

Unfortunately, testing economic hypotheses is especially difficult given the sample sizes of macroeconomic time series and the low signal-to-noise ratio in such data. When this is the case, the data alone cannot reject many restrictions, nor can they provide useful estimates, for forecasting purposes, of coefficients in an unrestricted specification. Faced with these problems, the specification process boils down to a set of decisions concerning what functional form to adopt and which explanatory variables to include. The choices that are made in specifying an equation determine, in effect, a filter that is used to extract information from the data. When an important variable is left out of the equation, a potential channel for extracting information in the data is closed. When too many variables are included in the equation, the filter is too wide and noise in the data will obscure the relatively weak signal. Thus, the choice of which variables to include involves a crucial tradeoff between oversimplification and overparameterization.

Neither of the two strategies commonly used in specifying macroeconomic forecasting equations addresses this tradeoff adequately. The structural approach relies on the restrictions suggested by one or more economic theories. The specification in this case includes a few of the many possible explanatory variables suggested by different economic theories. The alternative approach is to use a time series representation, which usually means inclusion of a few autoregressive and moving-average parameters. When such specifications are univariate, as is almost always the case, they cannot capture the interactions among variables. In this case, the focus on generating a parsimonious representation will lead to oversimplification of multivariate processes. In the standard approach to either structural or time series modeling, the tradeoff is crudely struck, essentially by restricting oneself to only a few parameters.

A Bayesian approach allows one to generate a class of estimators that highlight the tradeoffs between oversimplification and overparameterization. Moreover, for forecasting purposes, the use of out-of-sample prediction errors provides a natural measure for use in picking a specification that balances these tradeoffs.

Consider the following type of question, which typically arises in the context of a specification search: Should past values of variable x be included in the equation for another variable, y? Such a question is asked when one or more theories suggest that variations in x should contain information about future movements in y. There are probably many such x variables, however, only a few of which can reasonably be included in the equation for y. Economic theory does not provide enough information to make the decision. The usual approach is to choose the variables to include on the basis of statistics generated from classical tests of the hypotheses that the coefficients on variables in question are zero.

A Bayesian interpretation of the above procedure suggests that it is an inadequate way to combine prior information with evidence in the data. The procedure leads to a specification that can be viewed as reflecting very strong priors that coefficients are zero on all the excluded variables and as reflecting flat priors for the coefficients on the variables that are included in the equation. This result does not reflect the original symmetrical nature of the prior information.

As an alternative, I suggest a procedure in which the number of included explanatory variables is large, determined by computational expense considerations rather than the limited information content of the data; and in which the relatively symmetrical prior information on all of the included variables is used to balance the tradeoff between oversimplification and overparameterization. The basic idea is to specify a relatively unrestricted vector autoregression and a prior that can be varied along one or more dimensions affecting this tradeoff. Consideration of out-of-sample forecast errors as a function of movement in these dimensions is then used to find the optimal balance for forecasting purposes. In effect, these parameters are used to fine-tune the prior, which then acts as a filter to extract as much information from the data as possible. Implementation of this specification process is described in the context of the following example.

3. The Model

The specification search described here is designed to generate a model for short-run forecasting of monthly values of a set of macroeconomic variables. Included are four variables of primary interest: measures of output, prices, interest rates, and money; along with three informational variables: a stock price index, the flow of total nonfinancial debt, and the change in business inventories. The data are described fully in the Appendix. Observations begin in 1948:1 and end as of 1981:12. All variables, except changes in business inventories and the interest rate, are logged.

As a first step in the specification process, a set of benchmark univariate autoregressive forecasting equations is estimated by least squares. This benchmark, and the measures of fit that are used throughout, are based on out-of-sample forecast errors. Because I search over many forecasting models and choose among them based on this measure, it is true that the ultimate specification is, in a sense, fit "in-sample." It is important to recognize, however, that what is being fit is a balance between oversimplification and overparameterization based on the out-of-sample performance of estimators derived from a particular prior.

For each of the model specifications described in this paper, the following procedure is used to generate the out-of-sample forecast errors. Estimation is carried out one observation at a time using a Kalman filter algorithm. The posterior distribution from a given set of observations is used as the prior distribution to be combined with the next observation. An approximation to the posterior mean is then used to generate a set of forecasts for 1 through 12 steps ahead. In making forecasts at each point in time, for computational reasons I followed the usual procedure of treating the coefficient estimates as fixed, even though that procedure does not generate the Bayesian posterior mean predictions that minimize mean square error. The updating and forecasting procedure is continued for each observation in the sample. The computations were carried out at the University of Minnesota using the Regression Analysis of Time Series package, RATS, which is available from the author. Each cycle through the estimation and forecasting procedure required 20 seconds of computer time on a Cray 1 and cost about $10.

The out-of-sample forecast errors form the raw data for the measures of fit. The overall measure of fit is the log of the determinant of the sample covariance matrix of the one-step-ahead forecast errors. This criterion is suggested by Sims [13] when the desire is to study a procedure's predictive accuracy and one does not "believe literally in the fine detail of a model's nominal probability structure," as is the case here. In addition to the log-det measure, I also calculate the standard errors of the forecasts for each variable. Both the logdet and the standard errors are calculated for forecast errors 1, 3, 6, and 12 steps ahead and for each of three subperiods as well as the full forecast sample. The purpose in considering these different measures is to be able to observe any systematic changes in the predictive accuracy of different models over time or with respect to different variables. The forecasts are begun in 1951:1, with the subperiods defined as follows:

1. 1952:1 through 1961:12
2. 1962:1 through 1971:12
3. 1972:1 through 1981:12

The full sample consists of the sum of the three subperiods and thus contains 360 forecast errors for each variable and for each horizon, a total of more than 10,000 forecast errors. The overall pattern of results was quite stable across subperiods, variables, and forecast horizons. To save space, few of these results are reported in this paper; the complete set of results is available from the author on request.

The main purpose in estimating a set of univariate autoregressive models is to provide a benchmark against which to compare other specifications. Another natural choice would be to use ARIMA models as a benchmark. I rejected this approach primarily because of the computational expense it would have involved. Because they are nonlinear, even parsimoniously parameterized ARIMA models are expensive relative to the models estimated here. Moreover, the judgmental process involved in identifying ARIMA models makes them less suitable as benchmarks and raises difficult issues during implementation. For example, does one use the same specification of (p,d,q) throughout, or how often does one identify the equation? Finally, as shown in Litterman [8], low-order autoregressive models often perform better out of sample without the addition of moving average terms.

These univariate autoregressive models also show how important the tradeoff is between oversimplification and over-parameterization. For these models the tradeoff arises in the context of choosing a lag length, and the best forecasting performances were generated with a surprisingly small number of lags. There is a large literature addressing the problem of choosing lag length by associating penalties with longer lags and balancing these against the improvements in fit obtained as lags are increased. (See, for example, Geweke and Meese [3].) The best one-step-ahead forecasts, using the overall log-det measure, were generated with only four lags. These results are also not sensitive to the subperiod or forecast horizon. One might suppose, for example, when starting with more than 250 monthly observations and wanting to forecast a year ahead, that it would help to include at least a year's worth of lags; however, in general that is not the case. For all 12 of the combinations of subperiods and horizons, the four-lag specification by this measure performs better than 12 lags. In fact, for 10 of these 12 measures, even the one-lag specification performs better than 12 lags.

TABLE 1
Univariate autoregression forecast performance
as measured by log determinants of forecast errors

Number of lags in model	Overall	Subperiods 1	2	3
	1-Step horizon results			
1	-42.312	-44.804	-45.936	-41.481
2	-42.643	-45.077	-46.489	-41.701
3	-42.615	-44.966	-46.543	-41.718
4	-42.702	-45.095	-46.746	-41.783
5	-42.676	-45.043	-46.693	-41.783
6	-42.626	-44.940	-46.673	-41.772
9	-42.481	-44.405	-46.584	-41.738
12	-42.190	-44.017	-46.307	-41.458
	12-Step horizon results			
1	-23.395	-23.781	-28.593	-25.519
2	-24.084	-24.878	-28.506	-25.171
3	-24.295	-24.696	-28.433	-25.459
4	-24.564	-24.730	-28.631	-25.967
5	-24.078	-24.016	-28.620	-25.863
6	-24.001	-23.637	-28.668	-25.756
9	-23.994	-22.939	-28.462	-25.876
12	-22.960	-21.168	-28.042	-25.809

1-Step standard error results

(All units are percentages except change in business inventories (Inven), which is expressed in billions of '72 dollars.)

Number of lags in model	GNP	M1	Stocks	Tbills	Debt	Prices	Inven
1	.8673	.3817	3.330	.5367	14.46	.2073	5.582
2	.8550	.3808	3.275	.5237	14.39	.1862	5.493
3	.8552	.3823	3.290	.5239	14.37	.1856	5.488
4	.8450	.3818	3.291	.5274	13.67	.1855	5.499
5	.8435	.3800	3.295	.5313	13.69	.1862	5.516
6	.8457	.3812	3.306	.5381	13.73	.1866	5.506
9	.8549	.3847	3.300	.5261	13.99	.1893	5.617
12	.8687	.3891	3.322	.5432	14.07	.1940	5.715

Given a set of univariate benchmark models, it is natural to seek improvement in forecast performance by allowing multivariate interaction. This is the motivation for the second step in the specification process, which is to define a class of estimators that are derived as approximations to the posterior means associated with prior distributions. The priors are allowed to vary along a dimension that determines how likely coefficients on lags of other variables are to deviate from a prior mean of zero. Movement along this dimension includes the univariate specification at one end and an unrestricted vector autoregression at the other.

The representation to which this prior applies is a sixth-order autoregression of the vector of current observations, X_t, given by

$$X_t = \sum_{s=1}^{6} A_s X_{t-s} + C + \varepsilon_t, \qquad (1)$$

where the A_s are 7x7 coefficient matrices and C is a vector of constant terms. The choice of six lags is motivated by computational considerations, as is the limitation of the estimation procedures to single-equation techniques. It would be prohibitively expensive to form the more accurate approximation of the posterior mean generated by stacking the coefficient vectors of the multiequation system. A viable alternative suggested by Sims [13] would be to form a recursive system for which single-equation methods are appropriate. Although we would expect some improvement when applying such a system to the exercises considered here, the results would probably not be very different because the forecast errors are not based on contemporaneous observations. In a practical forecasting problem, however, some variables are observed before others, and when that is the case this alternative approach may have significant benefits. On the other hand, such an approach is not simply a different estimation technique; it also forces one to make substantive decisions about an ordering of the variables.

At this stage in the specification process, the prior distribution is flat with respect to own-lag coefficients and has a mean of zero on lags of other variables in each equation. In the prior, the distributions for each coefficient are treated as independent and normal. The variances of the distributions for coefficients on lags vary as a function of the lag number, being tighter around lags further back in the distribution. The variance of the prior distribution for the coefficient on lag ℓ of variable j in equation i is given by the formula

$$\frac{\pi_1 \sigma_i^2}{\ell \sigma_j^2}. \tag{2}$$

For this prior, and all others described below, the constant term in each equation is given a prior with a mean of zero and a variance 10^6 times π_1 to represent lack of knowledge about the means of the variables. The forecast results were not significantly affected by order-of-magnitude changes in this variance.

The prior distribution is parameterized by π_1, which determines the tightness of the prior around zero for each of the coefficients on variables other than own lags in each equation. The σ's are the standard errors of errors in six-lag univariate regressions over the full sample; they are included to scale the prior to make it independent of the units of the variables. The computation of the posterior mean is an application of Theil's [15] mixed estimation technique as described in Litterman [7]. The basic idea is to formulate a prior as a set of dummy observations that is added to the data. Ordinary least squares is then applied, producing an approximation of the posterior mean. The updating of estimates is accomplished using the Kalman filter algorithm, a description of which can be found in Bertsekas [1].

By varying π_1 and following the out-of-sample forecasting procedure described above, it is possible to map different prior distributions, which vary according to how much multivariate interaction is allowed, into a measure of predictive accuracy. The results of such an experiment are shown in table 2. This mapping is the first example illustrating the assertion made above that a Bayesian procedure could generate a class of models in which a balance could be struck between oversimplification and overparameterization.

TABLE 2
Forecast results from the search along a dimension of the prior allowing varying amounts of multivariate interaction

1-Step log-det results

Values of $_1$	Overall	Sub-periods 1	2	3
Univariate	-42.626	-44.940	-46.673	-41.772
.00005	-42.692	-45.091	-46.809	-41.849
.0001	-42.696	-45.079	-46.787	-41.904
.0005	-42.689	-45.005	-46.702	-42.019
.001	-42.678	-44.982	-46.646	-42.022
.005	-42.615	-44.891	-46.417	-41.927
.01	-42.535	-44.736	-46.269	-41.832
.05	-42.140	-43.989	-45.834	-41.455
.1	-41.885	-43.524	-45.632	-41.250
1.0	-40.881	-41.612	-45.103	-40.712

1-Step standard error results

Values of $_1$	GNP	M1	Stocks	Tbills	Debt	Prices	Inven
Univariate	.8457	.3812	3.305	.5380	13.72	.1866	5.506
.00005	.8394	.3790	3.328	.5344	13.45	.1837	5.513
.0001	.8357	.3780	3.329	.5339	13.40	.1843	5.518
.0005	.8272	.3740	3.336	.5298	13.35	.1885	5.536
.001	.8257	.3713	3.343	.5264	13.36	.1911	5.545
.005	.8325	.3644	3.378	.5148	13.55	.1986	5.590
.01	.8423	.3626	3.402	.5092	13.71	.2029	5.628
.05	.8844	.3640	3.494	.4995	14.32	.2176	5.790
.1	.9111	.3671	3.552	.4974	14.71	.2258	5.894
1.0	1.0361	.3831	3.779	.4966	16.22	.2589	6.376

For this mapping, the results show that such a balance requires tight priors around zero on coefficients of other variables in each equation. In the model that forecasts best, overall, the variance around the first lag (aside from a scale factor) is only .0001. Although the improvement is not large in absolute magnitude, it is quite consistent across variables, subperiods, and forecast horizons. Notice also that the largest improvement is apparent in subperiod 1, followed by subperiod 2 and then subperiod 3. This result is consistent with the expectation that prior information is most helpful when there is the least amount of data.

To get a rough guide to the interpretation of the magnitude of changes in the log-det statistics, notice that (aside from the covariance terms) changes in the log-det approximate a sum of the percent changes in the variance of forecast errors from each equation. Thus, to get a rough estimate of the average percentage of reduction in forecast standard errors in each equation, multiply the change in the log-det by 100/(7x2), approximately 7. By this guide the reduction in the one-step log-det, from -42.626 in the univariate case to -42.696 for the best value of π_1, represents an average reduction of about .5 percent. The actual average reduction in standard errors in this case was .78 percent.

Starting from the π_1 = .001 model of the previous exercise, it is natural to ask what possible improvement could be obtained by restricting the own-lag coefficients in some manner. This value of π_1 was chosen despite the fact that the overall log-det was minimized with π_1 = .0001. There are several reasons for this choice. First, notice that this value of π_1 is best in the most recent subperiod. Also, the differences in log-dets between the values of π_1 in this range of values are not large. Finally, these results were obtained with a flat prior on own-lag coefficients and no differentiation among coefficients of variables, except own versus other. When, in later steps in the specification process, such additional prior information is added, one might expect more room for other variables to enter in each equation. In effect, this search over π_1 is serving several roles, one of which is to describe the tradeoff along this dimension and another to serve as a rough guide to the location in a larger dimensional space where the best forecasting performances can be found. If the computational expense were not excessive, one would prefer to search directly over the larger space. The prior mean for the own lags is taken to be a random walk specification; that is, all coefficients are given a prior mean of zero except the own first lag, which is given a prior mean of one. Another parameter, π_2, determines a variance of the prior distributions around this mean. The variance of own lag ℓ is given by the formula

$$\frac{\pi_2}{\ell}. \qquad (3)$$

Again, by varying π_2 a class of estimators is generated that balances tradeoff between oversimplification and overparameterization, and again there is an interior minimum with respect to movements along this dimension. The results of this search are given in table 3.

TABLE 3
Forecast results from the search along a dimension of the prior that imposes a random walk specification for own lags

1-Step log-det results

Values of π_2	Overall	Subperiods 1	2	3
Infinite	-42.678	-44.982	-46.646	-42.022
1.0	-42.717	-45.150	-46.668	-42.034
.1	-42.790	-45.311	-46.694	-42.094
.05	-42.808	-45.344	-46.683	-42.127
.01	-42.772	-45.261	-46.547	-42.182

1-Step standard error results

Values of π_2	GNP	M1	Stocks	Tbills	Debt	Prices	Inven
Infinite	.8257	.3713	3.343	.5264	13.37	.1912	5.546
1.0	.8251	.3686	3.340	.5245	13.34	.1909	5.539
.1	.8225	.3672	3.328	.5201	13.30	.1901	5.519
.05	.8221	.3672	3.324	.5185	13.32	.1902	5.519
.01	.8270	.3671	3.325	.5166	13.50	.1932	5.588

The specification of π_1 and π_2 has generated a prior that is represented in the accompanying schematic. Although this prior contains no economic theory and treats each equation identically, it can lead to improvements in forecasting. The symmetry of the prior is limiting and unnecessary, however, and a reasonable next step is to ask whether further improvements can be obtained by specifying prior information that is more specific to each equation. One might want to include in a prior, for example, one's knowledge that stock prices are more likely to respond like a random walk than are changes in business inventories. A way to include this information is to vary the variances of distributions

A Schematic Representation of the Prior

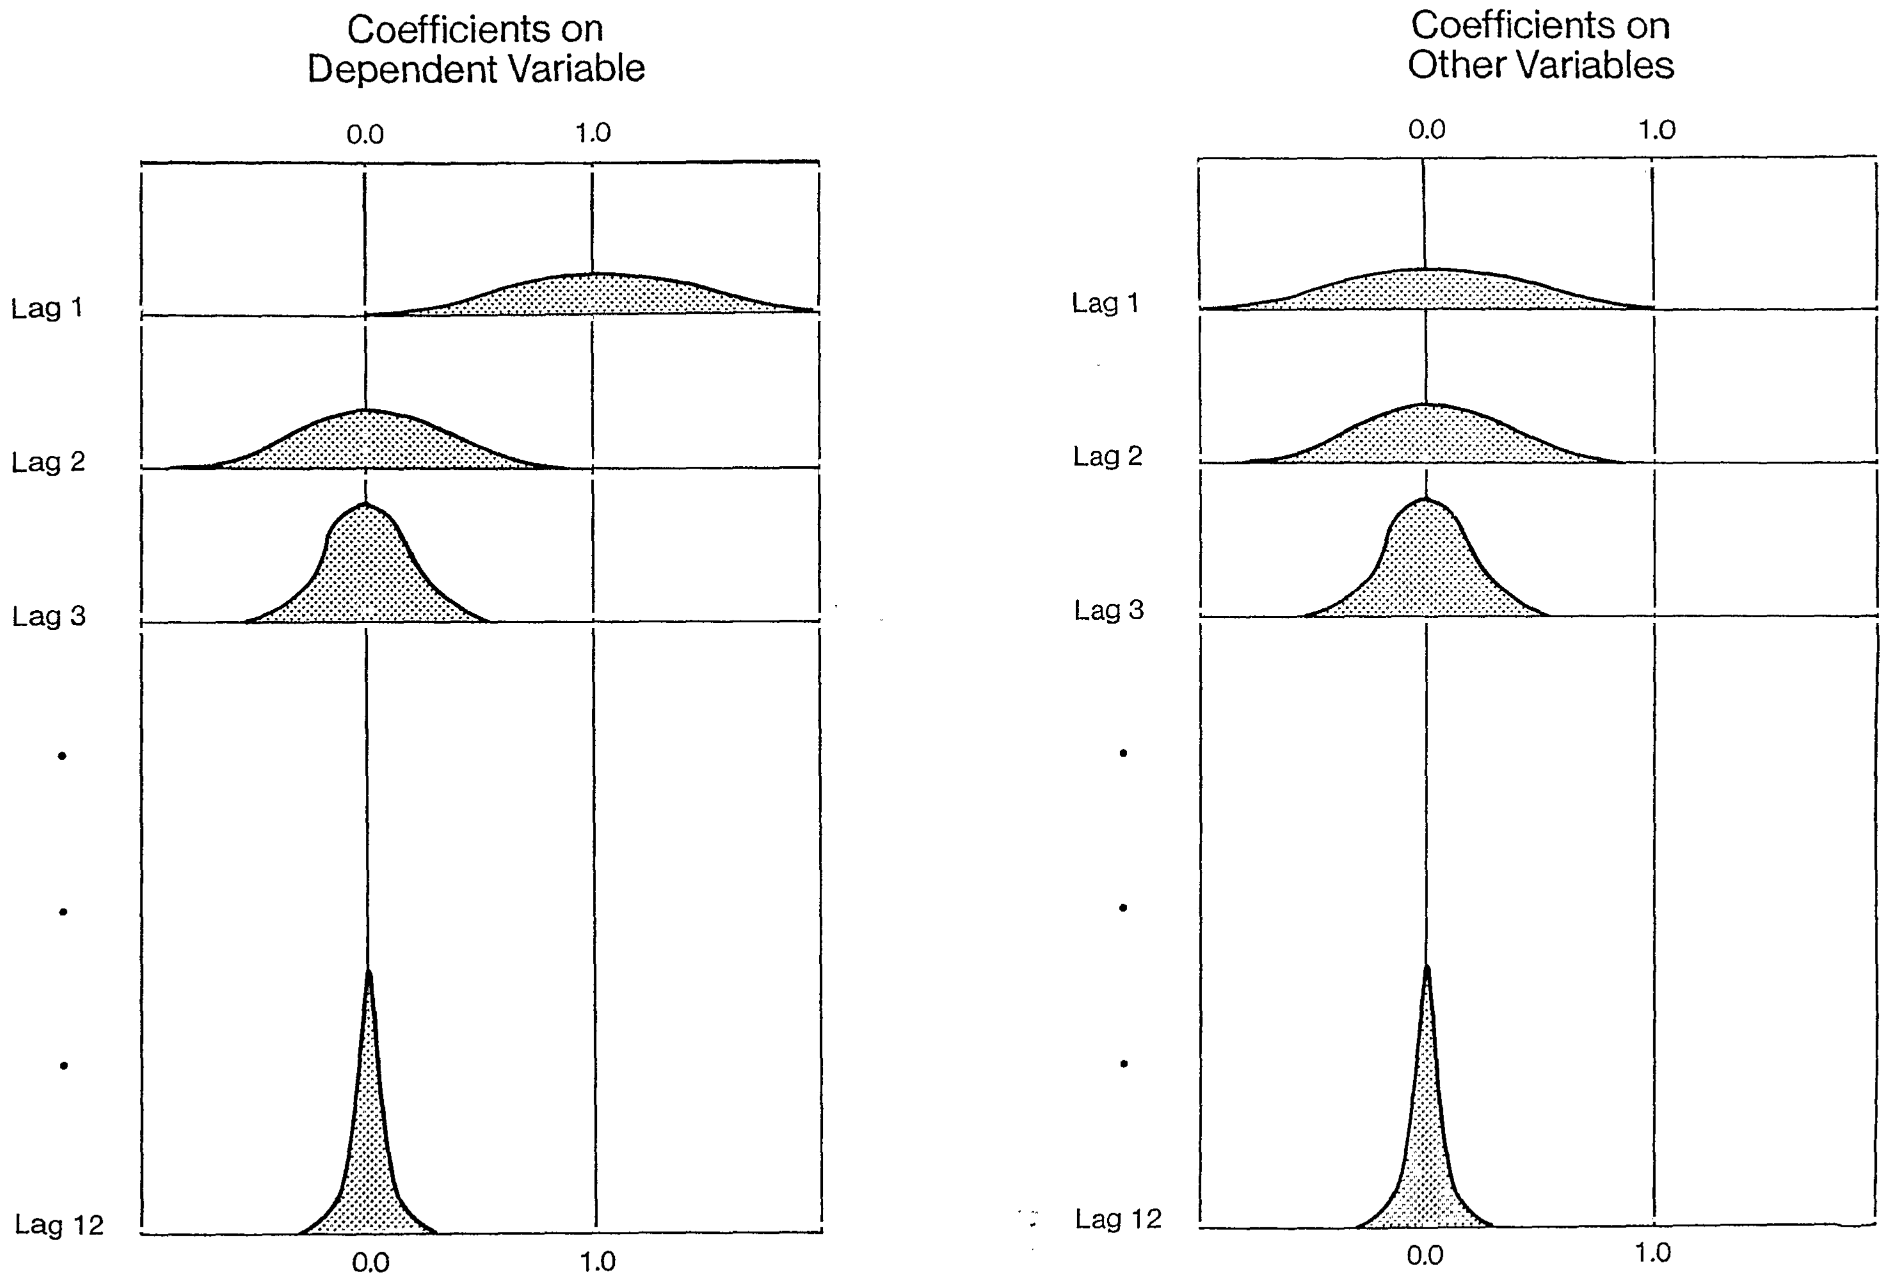

for coefficients on different variables in each equation. For an n variable vector autoregression, there are n decisions to be made for each equation. Another way to differentiate different equations is to change the prior mean; in particular, for variables that are clearly stationary, it may be appropriate to use a prior mean on the first own-lag coefficient that is somewhat less than 1.

Although we could use intuition guided by economic theory to suggest more or less weight for each variable in a particular equation, it is much more challenging to translate those qualitative feelings into a quantitative prior. One way to proceed is to specify a set of relative weights for the purpose of scaling the variances for each variable in each equation, and then to search over a dimension that determines how much effect these weights have. Such a search is demonstrated here by specifying a set of weights, w_{ij}, and defining the variance of the distribution for lag ℓ of variable i in equation j to be

$$\frac{\pi_1 \ \sigma_i^2 \ \exp(-\pi_3) \ w_{ij}}{\ell \ \sigma_j^2}. \tag{4}$$

The search over π_3 defines a dimension of more or less differentiation among variables. When π_3 is zero, all variables are treated symmetrically; as π_3 increases, the limiting specification has zero restrictions on all coefficients for which w_{ij} is positive. The weights used in this search are given in table 4. Note that the larger the weight, the faster the decrease of the variance on that coefficient as π_3 increases. The positive diagonal elements for stock prices and interest rates, for example, reflect a prior that those variables are more likely to follow a random walk process.

TABLE 4
Relative weights of variables in different equations

Equation	GNP	M1	Stocks	Tbills	Debt	Prices	Inven
GNP	0.0	1.0	1.0	0.0	1.0	1.0	1.0
M1	1.0	0.0	1.0	0.0	1.0	1.0	1.0
Stocks	5.0	5.0	1.0	3.0	5.0	5.0	5.0
Tbills	2.0	1.0	2.0	1.0	2.0	2.0	2.0
Debt	1.0	1.0	1.0	0.0	0.0	1.0	1.0
Prices	1.5	1.0	1.5	1.0	1.5	0.0	1.5
Inven	1.5	3.0	3.0	1.5	4.0	3.0	0.0

At the same time that these asymmetries in the prior variances are introduced, the mean values for the first lag of two variables--changes in business inventories and GNP--are changed in their own equation from 1 to .8 and .95, respectively. The prior means of 1 seemed reasonable for the other variables. The value of π_2 was set at .05 for this search, and π_1 remained at .001. The results of the search are shown in table 5. The effects of the changes in the prior means alone can be seen by comparing the line $\pi_3 = 0.0$, in table 5, with the line $\pi_2 = .05$, in table 3.

TABLE 5
Forecast results from the search along a dimension of the prior that allows asymmetric treatment of variables in different equations

1-Step log-det results

Values of π_3	Overall	Subperiods 1	2	3
0.0	-42.806	-45.346	-46.685	-42.122
1.0	-42.854	-45.319	-46.755	-42.169
1.5	-42.868	-45.326	-46.793	-42.164
2.0	-42.875	-45.338	-46.810	-42.159
3.0	-42.871	-45.334	-46.819	-42.122

1-Step standard error results

Values of π_3	GNP	M1	Stocks	Tbills	Debt	Prices	Inven
0.0	.8220	.3672	3.324	.5185	13.32	.1902	5.510
1.0	.8248	.3686	3.313	.5227	13.29	.1847	5.501
1.5	.8261	.3689	3.305	.5246	13.28	.1829	5.499
2.0	.8273	.3690	3.299	.5258	13.27	.1818	5.492
3.0	.8292	.3687	3.305	.5259	13.26	.1813	5.488

The final step in this specification search is to estimate a time-varying parameter representation. Up to this point, the underlying model has been one in which it is assumed that the true coefficients are constant through time. At best, this is thought to be a reasonably good approximation, which is necessary given the limitations in the information content of the data. We could, however, proceed in a manner similar to the three searches described above and generate a class of estimators that include constant-coefficient specifications as one special case and that allow time-varying parameters with various amounts of freedom. Again, by varying a parameter, π_4, a balance can be obtained between the oversimplification of a constant-coefficient specification and the overparameterization of a time-varying coefficients model. Sims [13] reports an informal search of exactly this type. Estimation of a time-varying parameters representation requires a specification of the nature of this variation and of the variances of the disturbances that cause the parameters to change each period. The process considered here follows Sims's suggestion that the parameters follow a random walk and that the variances of the disturbances for each coefficient be proportional to their variances in the prior distribution. The factor of proportionality is π_4; thus, when $\pi_4 = 0.0$ a constant-coefficients result is obtained. As π_4 increases, more parameter variation is allowed. The search takes values of π_1 and π_2 as before, and π_3 is set at 2.0. The results of this search, displayed in table 6, again indicate that an improvement in forecast performance is possible. A comparison of the forecast results between the best univariate benchmark model and those generated with the best value of π_4 shows an average reduction of 1.74 percent in the standard errors.

TABLE 6
Forecast results from the search along a dimension of the prior that allows time-varying coefficients

1-Step log-det results

Values of π_4	Overall	Sub-periods 1	2	3
0.0	-42.876	-45.338	-46.810	-42.159
$.1\times10^{-8}$	-42.884	-45.338	-46.825	-42.170
$.5\times10^{-8}$	-42.900	-45.337	-46.854	-42.192
$.1\times10^{-7}$	-42.910	-45.336	-46.870	-42.205
$.5\times10^{-7}$	-42.920	-45.324	-46.887	-42.220
$.1\times10^{-6}$	-42.913	-45.308	-46.881	-42.207
$.1\times10^{-5}$	-42.785	-45.118	-46.766	-42.031

1-Step Standard Error Results

Values of π_4	GNP	M1	Stocks	Tbills	Debt	Prices	Inven.
0.0	.8273	.3690	3.299	.5258	13.27	.1818	5.492
$.1\times10^{-8}$	.8264	.3680	3.299	.5258	13.27	.1817	5.492
$.5\times10^{-8}$	.8239	.3661	3.300	.5257	13.27	.1819	5.493
$.1\times10^{-7}$	.8222	.3651	3.300	.5257	13.27	.1821	5.493
$.5\times10^{-7}$	.8202	.3644	3.301	.5254	13.27	.1824	5.493
$.1\times10^{-6}$	.8216	.3654	3.302	.5250	13.27	.1825	5.494
$.1\times10^{-5}$	.8385	.3756	3.307	.5219	13.32	.1842	5.509

4. Conclusion

This paper has attempted to motivate and illustrate a specification search for a vector autoregressive representation to be used in short-run macroeconomic forecasting. In designing a model for use in forecasting, we encounter a particularly crucial tradeoff between oversimplification and overparameterization. A Bayesian approach to the specification process suggests several ways to balance this tradeoff. An example demonstrates that such a procedure can lead to consistent improvements in out-of-sample forecast performance.

Appendix

The seven variables used in this study are monthly measures of real GNP, the GNP deflator, three-month Treasury bill yields, the M1 measure of the money supply, the flow of total nonfinancial debt, the Standard and Poor's stock price index, and change in business inventories. Only three of the series are published on a monthly basis; the other four are constructed by interpolating the quarterly series using related monthly series. All series are seasonally adjusted except the interest rate and stock price index. The money series is extended back from 1948:1 through 1958:12 using the old M1 series scaled by a constant factor.

The interpolation is accomplished using the Chow-Lin [2] first-order Markov (M) procedure, which in a few cases is modified to be a random walk Markov (RWM) procedure as suggested in Litterman [9]. Monthly real GNP is constructed as the sum of eight real components, each of which is either available monthly or is interpolated. For each component and for the other interpolated series, the following list shows which procedure--M or RWM--was used and which monthly series were used as explanatory variables. In several cases, only nominal monthly series were available for interpolating real quarterly data. In these cases, flagged by the letters DFL, the monthly interpolated GNP deflator was used to deflate the monthly series. In one case, flagged SA, the monthly series required seasonal adjustment, which was accomplished using a regression of the logged data on monthly seasonal dummies. All equations included a constant and trend. In some cases where monthly series were not available for the entire sample, additional constant and trend dummies were added for those periods.

Real Consumption (available on a monthly basis)

Real Residential Fixed Investment (RWM)
Personal income DFL
New private construction DFL
Expenditures on private construction DFL
Total private construction put in place SA

Real Nonresidential Fixed Investment (M)
Contracts and orders for plant and equipment
Capital investment commitments 1967=100
New orders, nondefense capital goods
Three-month Treasury bill yields
Industrial production index
Commercial and industrial loans DFL

Real Change in Business Inventories, Durable Goods (M)
Net change in inventories on hand and on order
Change in wholesale inventories, durable goods DFL
Change in inventories, durable goods DFL
Change in finished goods inventories, durable goods DFL

Real Change in Inventories, Nondurable Goods (M)
Net change in inventories on hand and on order
Change in wholesale inventories, nondurable goods DFL
Change in inventories, nondurable goods DFL
Change in finished goods inventories, nondurable goods DFL

Real Government Purchases of Goods and Services (RWM)

Real Exports
Exports, excluding military aid shipments DFL

Real Imports
General imports, F.A.S. basis DFL

GNP Price Deflator (RWM)
Consumer price index
Producer price index

Flow of Total Nonfinancial Debt (M)

Change in the Level of Consumer Installment Debt
- Standard and Poor's stock price index
- Three-month Treasury bill yield
- Consumer price index
- Commercial and industrial loans

References

[1] Dimitri P. Bertsekas, Dynamic programming and stochastic control. New York, Academic Press, (1976).

[2] Gregory Chow and An-loh Lin, Best linear unbiased interpolation, distribution and extrapolation of time series by related series. Review of Economics and Statistics 53 (1971), 372-375.

[3] John Geweke and Richard Meese, Estimating regression models of finite but unknown order. International Economic Review 22, no. 1 (1981), 50-70.

[4] Arthur E. Hoerl and Robert W. Kennard, Ridge regression: Biased estimation for nonorthogonal problems. Technometrics 12, no. 1 (1970), 55-67.

[5] Edward E. Leamer, A class of informative priors and distributed lag analysis. Econometrica 40 (1972), 1059-1081.

[6] Edward E. Leamer, Specification searches, ad hoc inference with nonexperimental data. New York, Wiley, (1978).

[7] Robert B. Litterman, Techniques for forecasting with vector autoregressions. Doctoral thesis, University of Minnesota, (1980).

[8] Robert B. Litterman, A use of index models in macroeconomic forecasting. In: Applied Time Series Analysis of Economic Data. Arnold Zellner, ed. Proceedings of the ASA-Census-NBER Conference on Applied Time Series Analysis of Economic Data, October 13-15, 1981, Arlington, Virginia, U.S. Bureau of the Census (1983).

[9] Robert B. Litterman, A random walk, Markov model for the interpolation of time series. Journal of Business and Economic Statistics 1 (1982), 169-173.

[10] Robert J. Shiller, A distributed lag estimator derived from smoothness priors. Econometrica 41 (1973), 775-788.

[11] Christopher A. Sims, Macroeconomics and reality. Econometrica 48 (1980), 1-48.

[12] Christopher A. Sims, Scientific standards in econometric modeling. University of Minnesota Working Paper, (1981) (forthcoming in the proceedings of the January 1982 conference in honor of the Rotterdam Econometrics Institute).

[13] Christopher A. Sims, Policy analysis with econometric models. Brookings Papers on Economic Activity, 1 (1982), 107-154.

[14] Charles M. Stein, Multiple regressions. Chap. 37 in Contributions to probability and statistics: Essays in honor of Harold Hotelling edited by I. Olkin. Stanford University Press, (1974).

[15] H. Theil, Principles of econometrics. New York, Wiley, (1971).