Federal Reserve Bank of Minneapolis
Research Department Working Paper

A STATISTICAL APPROACH TO ECONOMIC FORECASTING

Robert B. Litterman*

Working Paper 287

September 1985

*Federal Reserve Bank of Minneapolis.

Abstract: A recently developed statistical model, called Bayesian Vector Autoregression, has proven to be a useful tool for economic forecasting. Such a model today forecasts a strong resurgence of growth in the second half of 1985 and in 1986.

In recent years a statistical time series model has been developed which for the first time appears to generate forecasts that compare favorably in terms of accuracy with those generated by the best judgment of economic forecasters. This economic outlook talk will be a little unusual in that it will focus primarily on this new statistical model rather than my view of the economic outlook. At the end I will describe the model's current, unadjusted forecast, and then I will briefly discuss the extent to which my own subjective view is different from the model forecast.

The model that produced the forecast I will talk about today is based on the statistical technique called Bayesian Vector Autoregression (BVAR). The separate concepts of Vector Autoregression (a projection of each element of a vector on its own lags and lags of each of the other elements of the vector) and of Bayesian statistics are not new. Their combination in the BVAR technique has led to models that have proven to be accurate, and yet inexpensive and relatively easy to use. Moreover, the fact that BVAR models forecast well without judgemental adjustment means that they can be used to objectively estimate the probabilities of future events, to measure uncertainty, and to answer other questions that traditional econometric models that rely on judgemental adjustment as a crucial input cannot realistically address.

Let me first briefly describe the general motivation for the BVAR technique. I will then describe the particular model that we use at the Federal Reserve Bank in Minneapolis. In the mid-1970's the economics profession, and the research department at the Minneapolis Fed in particular, became disenchanted with the large structural econometric models then commonly in use. This disenchantment was based not only on the perceived theoretical weaknesses in the standard Keynesian models then in use, but also on a number of practical problems that are generic in using large structural models. The main problems are that such models are expensive to run and that the forecasts that they generate are so often appear to be unreasonable that the standard operating procedure is to significantly adjust the entire paths of all the key variables in the model. Another problem is that the probability distributions of likely outcomes generated by such models are unreasonably narrow, and in any case cannot be taken seriously with respect to the adjusted forecast since the adjustments are not part of the statistical model.

At the Federal Reserve Bank our somewhat naive view at the time was that we could probably forecast as well, if not better, and save money, by specifying a small vector autoregression of the type that Professor Christopher Sims was working with at the University of Minnesota at that time (see Sims 1980). There are a number of seductive features about using VAR models. The VAR is a very general representation with which to approximate the stochastic process generating a multivariate time series, especially if you consider the class of VAR's with time varying coefficients. The VAR model is inexpensive and easy to specify, to estimate, and to use in generating forecasts. Unfortunately, it does have one serious problem which became obvious to us when we took a close look at the properties of our initial VAR forecasts.

The problem, which we refer to as overparameterization, is simply that when too many parameters are estimated from too few data points, one finds a very good in-sample fit, and a very bad out-of-sample forecast performance. It is very common for this to happen in the context of an unrestricted VAR where the number of parameters can easily reach or even exceed the number of observations. What happens is that parameters fit not only the systematic relationships in the data which are useful for forecasting, but also the random variation. In economic data the systematic component of the movements in a variable is often only a small part of its total variation. Many economic variables behave very much like random walks, in which case the ratio of systematic to unsystematic variation is close to zero. Thus, the basic problem of economic forecasting is to filter out the weak signal generated by the systematic correlations in the data from the din of random noise. Because of its overparameterization, the unrestricted VAR is not a particularly useful tool for accomplishing this task.

On the other hand, the Bayesian VAR which we subsequently developed at the Federal Reserve Bank to overcome this problem has proven to be a flexible tool that can, in effect, be tuned so as to become a very sensitive filter for capturing the systematic relationships in economic data. In fact, there are good theoretical reasons to believe, as well as empirical evidence to suggest, that the Bayesian VAR technique can generate forecasts that are more accurate than any of the other standard forecasting techniques.

Ultimately, there are only two sources of information useful in forecasting. One is the historical data. The other is the modeler's knowledge about the structure of the system generating the data. In the case of economic systems, there are many reasons to mistrust both sources. Reasons to mistrust the evidence in the data include measurement error, changing structures in the economy, and the fact that most of the systematic relationships are masked by random noise. On the other hand, economic theory is often unreliable as well.

All of the standard techniques, from simple smoothing algorithms to complex structural models, can be viewed as attempts to solve the filtering problem by using prior beliefs about the structure of the data generating process to restrict the parameterization of the general VAR representation. Structural models, for example, set almost all of the coefficients in the VAR representation to zero, relying on economic theory to suggest a few coefficients in each equation with which to fit the movements in the data. Time series techniques generally exclude the possibility of fitting cross variable interactions and instead rely on fitting the autocorrelations in the data via a small number of parameters. Though not usually viewed from a Bayesian perspective, these methods can be, and when viewed that way they are seen to be too rigid to allow the the forecaster to express his true prior beliefs.

The basic problem with any of the standard approaches from a Bayesian perspective is that in specifying his prior beliefs about coefficients the forecaster has essentially only two choices, to exclude a variable, which is to specify that a coefficient is exactly zero or to include the variable, which is to say that he knows nothing about its likely value. In order to avoid overfitting the forecaster is forced to rely heavily on exclusion restrictions even though that represents an unreasonably strong prior, one that will never be altered by evidence in the data. On the other hand, the only alternative to exclusion in standard approaches in inclusion of a coefficient without a prior probability distribution. Such a specification is the expression of an unreasonably vague prior, and if too many variables are included the approach rapidly leads to overfitting of the data. Notice that no matter how much specification testing is undertaken, the basic problem of not being able to specify realistic prior information cannot be avoided.

The approach taken by the Bayesian VAR technique is to solve the overparameterization problem by specifying in a Bayesian framework the likely values for all of the coefficients. Instead of setting lots of coefficients to zero, the prior that we use specifies that most coefficients are likely to be close to zero. To see the advantage of this approach, consider the problem of choosing a lag length for one variable. The standard approach relies on one of several possible ad hoc rules for choosing a value k, the number of lags to include. The motivation for making such a choice is the knowledge that more recent values of a variable are more likely to contain useful information about its future movements than older values. Yet the prior implicit in a choice k is too rigid to incorporate that information. It specifies that we know nothing about the coefficients on the first k lags, and that we know that all other lags have coefficients that are exactly zero. A Bayesian approach can directly incorporate the original information by including a long set of lags and specifying that the larger the lag, the more likely that the coefficient is to be close to zero. An example of our approach would be to specify that the jth lag has an independent normal distribution with a mean of zero and a standard deviation inversely proportional to j. The proportionality constant is a parameter we would refer to as the "overall tightness" since it specifies how close all of the coefficients are to their prior mean. We usually call such a parameter of the prior a "hyperparameter" in order to distinguish it from the estimated coefficients, the parameters of the model itself.

Specifying the value of such a hyperparameter is a difficult issue, similar in some ways to the problem of choosing a lag length k, but the important point is that the model generated by choosing any reasonable value for the hyperparameter will reflect the prior information available to the modeler much more accurately than the prior implicit in any choice of k, and therefore such a prior is likely to perform better as a filter for capturing the useful information in the data. The kind of flexibility illustrated in this simple example of choosing a lag specification is what gives the Bayesian approach its power to dominate other approaches in terms of forecast accuracy. In fact, most other approaches can be viewed as special cases of the Bayesian prior, but just as in the case of choosing the prior for a lag specification, one would rarely be led to the rigid prior implicit in the standard approach. A more complete description of the BVAR approach can be found in Litterman (1980), Litterman (1984a), Doan Litterman and Sims (1984) and Todd (1984).

The BVAR model that we currently use at the Federal Reserve Bank in Minneapolis, which is described in Litterman (1984b), has 47 variables arranged in 8 sectors. The typical equation has 15 lags on each of ten explanatory variables, a total of some 8,000 coefficients to estimate. We make the problem of specifying priors for all those coefficients manageable by specifying a general functional relationship between nine hyperparameters and the priors for all the coefficients in the VAR representation. Typically most of the coefficients are given prior means of zero, except for the first own lag. That coefficient is generally given a prior mean of one so that we often refer to a "random-walk" prior. The tightness around the means of the coefficients reflect considerations that range from the position in the lag distribution, as described above, to the importance that theory suggests one variable should have with respect to another. We also allow the parameters to vary over time, and we use a hyperparameter to control the amount of such variation.

In experimenting with different hyperparameter settings, we have paid particular attention to the out-of-sample forecasting properties of the models implied by those choices. Those experiments, which show that a wide range of hyperparameter choices lead to consistent improvements in forecast accuracy relative to other types of models, provide one form of evidence that the Bayesian VAR technique works. Another, perhaps more convincing, bit of evidence comes from the actual performance of BVAR models that have been specified and used over the past five years. One small BVAR that I have been using to forecast with mechanically each month for the past five years has been shown by Stephen McNees (1985) and myself, Litterman (1985) to have generated forecasts that for most variables and horizons have been superior to the majority of commercial judgemental forecasts. This simple model takes only ten minutes to estimate on a microcomputer.

Before presenting the forecasts of the model itself, I should warn you that in many cases the model's forecasts lie well outside the range of other forecasters. Because I know that the BVAR model is basically as likely to be as accurate as any of the other forecasts, I consider it a strength that it often provides information distinct from the consensus view. Nonetheless, I feel it necessary to provide this warning because to the uninitiated these forecasts often appear quite unreasonable. Luckily, from the point of view of making this presentation more provacative and interesting, this particular period happens to be one where the model is well outside the range of the other forecasters.

Briefly, the BVAR is currently extremely bullish on the economy. I will begin with the short run prospects. Based on data through July 28, the model projects real GNP growth of 6 percent at annual rates for the third quarter. That rate is, I believe, well above the current consensus view. The latest Blue Chip consensus released July 10 was for 3.9 percent growth in the third quarter. I must caution though, that when I say the model projects 6 percent, that is really a shorthand way of saying that the model projection of real GNP is a distribution with mean of 6 percent. Since the model also projects a standard deviation of 3.3 percent for that distribution, a better description would be to say that model projects that real GNP growth has about a two-thirds probability of falling between 2.7 percent and 9.3 percent growth.

That distribution for GNP is rather wide, and some might ask what good is such an uncertain forecast. My answer, of course, is that there is a huge amount of uncertainty about economic forecasts, even the best model cannot eliminate that, and it is just as important to have a realistic measure of that uncertainty as it is to have a point forecast.

The model's distribution for GNP rises and widens a little bit in the fourth quarter and peaks in the first quarter of '86 when the mean reaches 7.6 percent with a standard deviation of 3.6 percent. For the rest of '86, the mean of the forecast distribution declines rather sharply to a value of 3.6 percent for the fourth quarter.

This quarterly path leads to year-over-year mean growths of 3.0 percent in 1985 and 5.9 percent in 1986.

With repect to inflation, the model is much closer to the consensus. Its projection of the distribution for the third quarter of this year of growth in the GNP deflator has a mean of 2.7 percent with a standard deviation of 1.4 percent. Notice that there is considerably less uncertainty about inflation rates than real growth rates. The mean inflation forecast rises to 4.1 percent in the fourth quarter and continues to rise to just above 5 percent by the fourth quarter of 1986. These rates lead to year over year inflation of 3.7 percent in 1985 and 4.1 percent in 1986.

In line with the optimistic real growth forecast, the model also projects unemployment to decline significantly over the coming two years. Here again the model differs sharply with the consensus view that has the unemployment rate holding essentially flat over that horizon. In constrast, the BVAR model projects the mean of the distribution for the unemployment rate to drop at a rate of one-tenth of one percent each month for the next year, bringing it to a level of around six percent by late summer of 1986, at which point the mean remains steady for the rest of the year. That mean is about one standard deviation, that is about 1.3 percent, below the consensus view of 7.3 percent for the end of 1986.

When the model projects growth that is so much higher than other forecasters, a natural question is where is that growth going to come from, that is in what components of GNP. Let me answer that by comparing the projected growth rates for the third quarter with those measured in the second quarter. By far the largest component of GNP is personal consumption expenditures. These account for about two-thirds of GNP. Those consumption expenditures were actually quite strong in the second quarter. They grew at a 5.2 percent rate. Despite recent increases in consumer confidence, the model expects consumption to moderate a bit to a rate of 4.1 percent. Business fixed investment was also strong in the second quarter, with growth at 13.6 percent. Here again the model expects a moderation, down to a growth of 7.1 percent. Government expenditures grew moderately in the second quarter, at a rate of 3.8 percent. Here again the model expects less in the third quarter; it projects a growth of only 1.8 percent.

With growth in all the major components of GNP slowing, how can the model be projecting growth to rise from 1.7 percent to 6.0 percent? The answer is in the remaining small, but volitile components of inventory investment, net exports and housing. In total these components sum to just over 2 percent of GNP, but because they jump around, they can make a tremendous difference in terms of quarterly growth rates. First, notice what these components did to the second quarter growth rate. Inventory investment and net exports alone dropped 19 billion (1972 dollars) off of the level of real GNP. In other words, had they stayed at their first quarter levels GNP would have measured 6.3 percent growth in the second quarter. The model projects that inventory investment will stay essentially unchanged at the low second quarter level, and that net exports will become slightly less of a drag on GNP than it was at the record second quarter level. In particular, the model projects net exports to be minus 30 billion rather than minus 34 billion in 1972 dollars.

The final component of GNP, housing construction, is projected to boom in the third quarter with growth of 30 percent at annual rates. This is up from 19 percent growth in the second quarter. A housing boom is a predictable event after the significant declines in interest rates that we have seen recently. Some, however, currently argue that special factors could cause the housing sector to not respond as it usually does. I don't put much weight on such stories, but even if the housing boom doesn't materialize, and suppose, for the sake of argument that housing grows only at the second quarter rate. That would reduce the growth rate of real GNP by only approximately one-half of one percent. Thus, when I look carefully at the components of growth forecast by the model it gives me some comfort because I don't see any areas which look particularly unreasonable.

What is this model picking up that other forecasters are ignoring? It is hard to answer such a question precisely, but this model finds that information useful for forecasting is often concentrated more in the financial indicator variables such as interest rates, stock prices, the value of the dollar and money growth; than it is in the values of current indicators of real activity. And while you cannot find evidence of strength in the real indicators, there have been significant movements downward in interest rates and upward in stock prices and money growth in recent months. Such movements are historically strongly associated with a lagged increase in real activity. The level of uncertainty is such that the response may not come as quickly as the model suggests, but it would be very unusual, given this combination of factors, not to see a strong response sometime in the near future.

People sometimes feel uncomfortable with a forecast from a statistical model--what is your gut feeling they ask? Or to put it somewhat more scientifically, if you were forced to bet on a forecast, what numbers would you want to bet on. My response is that if I had to bet on a forecast, and if the payoff were proportional to the distance between the forecast and the actual value (because if the contest is a game between my forecast and some other forecast I would behave very differently), I don't think I would adjust this forecast much. I suppose I would adjust the housing component down a few billion because the uncertainty about the value of real estate investments over the coming years may alter their historical response to interest rate movements, and perhaps I would want to lower the model's forecast of investment in business structures by a few billion because it does not know about unprecedented vacancy rates. However, on the other side, I would adjust the government spending component up several billion because defense expenditures have been running well below the pace necessary to reach this years' appropriation.

The bottom line, then, is that I agree with the model that we should be very optimistic about the economy in the short run. That optimism should not be misinterpreted to mean that I would rule out the possibility of a bad outcome. After all, the model projects that there is about a forty percent probability that we will see at least one quarter of negative real growth over this six quarter horizon. On the other hand, one quarter of negative growth is not too bad, and the model suggests that there is less than a one-in-ten chance of having two negative quarters in a row, which is a standard rule of thumb for recessions. Such a probability is unusually low, especially given that the current recovery is already of about the postwar average length.

Based on the model projection, my main concerns about the future are more long term. The model projects that we are more likely to start getting into trouble toward the end of next year. At that point, it projects declining real growth, an increasing probability of recession, continued large budget deficits, and accelerating inflation.

References

Doan, Thomas; Robert Litterman; and Christopher Sims. 1984. "Forecasting and Conditional Projection Using Realistic Prior Distributions." Econometric Reviews 3 (1), pp. 1-144.

Litterman, Robert B. 1980. "Techniques for Forecasting Using Vector Autoregressions." Ph.D. diss., University of Minnesota.

_______________. 1984a. "Specifying Vector Autoregressions for Macroeconomic Forecasting." Research Department Staff Report 92, Federal Reserve Bank of Minneapolis. Also forthcoming in Bayesian Inference and Decision Techniques With Applications: Essays in Honor of Bruno de Finetti, ed. P. Goel. Amsterdam: North-Holland.

_______________. 1984b. "Above-Average National Growth in 1985 and 1986." Federal Reserve Bank of Minneapolis Quarterly Review 4 (Fall): 3-7.

_______________. 1985. "Forecasting with Bayesian Vector Autoregressions--Five Years of Experience." Federal Reserve Bank of Minneapolis, Working Paper 274. Also forthcoming in the Journal of Business and Economic Statistics.

McNees, Stephen K., 1985. "The Forecasting Accuracy of VARs and CMs: The Evidence and Its Interpretation." Federal Reserve Bank of Boston.

Sims, Christopher A. 1980. "Macroeconomics and Reality," Econometrica 48, pp. 1-48.

Todd, Richard M. 1984. "Improving Economic Forecasting With Bayesian Vector Autoregression." Federal Reserve Bank of Minneapolis Quarter Review 4 (Fall): 18-29.