Federal Reserve Bank of Minneapolis
Research Department Staff Report 78

Revised October 1982

A USE OF INDEX MODELS IN MACROECONOMIC FORECASTING

Robert B. Litterman

Federal Reserve Bank of Minneapolis

ABSTRACT

This paper illustrates the application of observable index models to the problem of macroeconomic forecasting. In this context, a Bayesian prior is used to describe a class of models which impose the index structure with more or less weight. An out-of-sample forecasting experiment is used to measure the possible benefits of this approach. In addition, impulse response functions and the decomposition of forecast variance are analyzed to suggest a possible separation of real and nominal shocks into separate channels.

This paper was prepared for the October 1981 ASA-CENSUS-NBER conference on applied time series analysis of economic data. This draft benefited from helpful comments on an earlier version by participants in the NBER summer institute. Able computational assistance was provided by Paul O'Brien.

This paper illustrates the application of a Bayesian procedure to the estimation of a small system of macroeconomic time series subject to nonlinear restrictions on the parameters. The system is an observable index model of the type described by Sargent and Sims (1977). The ability of Bayesian estimators to incorporate nonsample information in a flexible manner is shown in this context to have a potential use in time series forecasting.

The methodology of this paper is very similar to that of Litterman 1980. That paper shows that the incorporation of prior information in the estimation of vector autoregressions can greatly improve the out-of-sample forecasting performance of such models. Here the focus is on the forecasting potential of an index model and the usefulness of prior information on the size of the effects of the index.

The idea of explaining the cross-correlations of a set of variables as the effects of one or more indexes has a long history in psychology and other social sciences. In recent years, a number of papers have applied index or factor models to economic time series.[1/] Most of these papers have used an index model to test economic theory which suggests that such a model should fit well. Unfortunately, more than one interpretation of such an index structure is often possible.

The purpose of this paper is simply to test the forecasting potential of one type of index model. The imposition of

---

[1/] Geweke (1977) describes a type of frequency domain factor analysis which has been investigated by Sargent and Sims (1977) and Singleton (1981), among others. Sims (forthcoming) also uses the observable index model.

the index structure can be viewed as a set of nonlinear cross-equation constraints on a vector autoregressive representation of the time series. If these restrictions are true, or nearly so, then their imposition during the estimation of a model may be expected to improve its out-of-sample forecasting performance. By examining the out-of-sample forecasting performance of an autoregressive model with various priors, I hope to measure the potential of this type of index model for forecasting economic time series.

Setting Up the System

The Observable Index Model

The observable index model considered here has the form

$$(1) \qquad Y(t) = \left[\sum_{i=1}^{L_d} D_i Y(t-i)\right] + \left[\sum_{j=1}^{L_a} A_j Z(t-j)\right] + C + U(t)$$

for t = 1,...,T, where Y(t) is an nx1 vector of endogenous variables, the D_i are nxn diagonal matrices of own lag coefficients, the A_j are nxq matrices of loadings on the indexes, C is an nx1 vector of constants, and the U(t) are nx1 vectors independently and identically distributed as $N(0,\Sigma)$, where Σ is a positive-definite symmetric matrix.

The qx1 index series, Z(t), is defined by

$$(2) \qquad Z(t) = \sum_{k=1}^{L_b} B_k Y(t-k+1)$$

for $t = -L_a+1,\ldots,T-1$, where the B_k are qxn matrices of coefficients.

Conditional on the initial observations, $Y(t)$ for $t < 0$, the likelihood function is given by

$$(3) \qquad L(\theta, \Sigma \mid Y) \propto |\Sigma|^{-\frac{1}{2}T} \exp\left\{-\frac{1}{2}\sum_{t=1}^{T} U(t)'\Sigma^{-1}U(t)\right\}$$

where a vector θ includes all of the free parameters in the D's, A's, and B's and in C.

Assuming little is known, a priori, about the elements of Σ, I use Jeffrey's diffuse prior for Σ

$$(4) \qquad p(\Sigma) \propto |\Sigma|^{-(n+1)/2}$$

and an independent informative prior, $p(\theta)$, for θ. The marginal posterior density for θ is then given by

$$(5) \qquad p'(\theta|Y) \propto |S|^{-\frac{1}{2}T} p(\theta)$$

where

$$S = \frac{1}{T}\sum_{t=1}^{T} \hat{U}(t)\,\hat{U}(t)'$$

and the $\hat{U}(t)$ are the residuals derived from substituting the values of θ into (1) and (2) above.

Identification

In order to identify the coefficients in the A and B matrices, some normalization is required. Since multiplication of all the elements of the A's and division of all the elements of the B's by the same constant would not affect the fit, the scale of Z is fixed by setting the elements of the first row of A_1 equal to 1. In a one-index model with finite lag lengths--that is, given L_a and L_b--the above normalization is sufficient to identify the A's and B's. However, if lag lengths increase as the number of observations increases or if the model has more than one index, then a further normalization is required. This is because in the limit the fit would not be affected by multiplication of the lag polynomials defined by the A_j, A(L), by some qxq lag polynomial, g(L), and multiplication of the lag polynomials defined by the B_k, B(L), by $g(L)^{-1}$. For a one-index model, Sargent and Sims (1977) suggest setting the elements of the first row of each of the A_j equal to zero for $j = 2,\ldots,L_a$.

With a two-index model, the natural extension of this normalization is to set some 2x2 submatrix of A_1 equal to the identity and set the same submatrix of A_j for $j = 2,\ldots,L_a$ equal to zero. Since the lag lengths estimated here are actually quite short, such restrictions may not be innocuous. For the purpose of this paper, the indexes do not have to be interpreted directly. Thus, the reliance on lag lengths to achieve identification is less of a concern than the overidentification imposed by such zero restrictions.

The basic condition required for identification of the structural coefficients, θ, is that the transformation from θ to the implied reduced-form autoregressive representation be of full rank. That will be true when the A's or B's are restricted so that there is no g(L) with

$$A^*(L)\ B^*(L) \equiv [A(L)g(L)]\ [g(L)^{-1}B(L)] = A(L)B(L)$$

such that A* and B* also satisfy the restrictions. When A(L) and B(L) are both known to be one-sided and to have known orders--that is, when A_{L_a} and B_{L_b} have full rank--then such a g(L) must have $g_j = 0$ for all $j \neq 0$. Thus, identification may be achieved by restricting A_1 so as to insure that the only such g(L) is the identity. One sufficient condition, the one I adopt here, is to set a qxq submatrix of A_1 equal to the identity. This restriction does not impose any overidentification, and any other index representation with the same orders can be transformed to its form.

A Bayesian Approach

Time series modeling of macroeconomic phenomena presents a difficult task. Economic theory suggests that many complex interactions among variables are likely to occur. Data, however, are too limited to allow estimation of more than a few relationships. Regardless of its complexity, the linear structure of a vector of covariance stationary economic time series can be approximated arbitrarily well in a mean square sense by a vector autoregressive model if the lag length is made large enough. However, the number of parameters in such a system grows with the square of the length of the vector and quickly exceeds the number for which the data are able to give accurate estimates. Moreover, the paucity of data, which prevents more accurate identification of macroeconomic relationships, is not likely to ease significantly over time. Because economic systems grow and structures change, the relevant sample size is limited. Thus, multivariate modeling of economic time series must confront this small sample problem. The essential question is how to use prior information to reduce the dimensionality of the estimation.

The type of prior information which is available is problem specific, and no one estimation strategy can be expected to dominate in every problem. The index model has been suggested by Sargent and Sims (1977) as a useful approach toward the estimation of business cycle dynamics. The fact that many aggregate macroeconomic variables seem to show high coherence at periodicities around a few years suggests that there may be one underlying factor which drives the cycle. Current business cycle theories suggest that a number of phenomena may be involved. Shocks may come from such diverse things as the financial sector, international events, relative prices, and labor markets. Independent shocks may be propagated through inventory, capital stock, and labor adjustment costs; information lags; long-term contracts; or other institutional features of the economy. Interest rates play a key role in intertemporal allocation of consumption and production, and certain sectors such as housing may be especially sensitive. Monetary and fiscal policy reactions also contribute to the dynamics of the cycle.

Given all of the possible complexity of business cycles which theory suggests, it is no small task to condense the important features into a system of variables to which time series methods can be applied. For this reason, I choose to focus here on the manufacturing sector. While this focus means that important aspects of the economy will clearly be ignored, it also means that, to the extent that the business cycle is generated primarily in the manufacturing sector, I should be able to capture the cycle's most salient features. To this end, I consider the manufacturing component of industrial production and the producer price index, the corresponding inventory-to-sales ratio, and the three-month Treasury bill rate.

Having identified a set of variables which measure important aspects of the business cycle, I face a number of alternative modeling strategies. Multivariate ARMA models have received a great deal of attention in recent years. Reviews of these methods have appeared in Granger and Newbold 1977 and Nerlove, Grether, and Carvalho 1979. Both of these texts emphasize the difficulty of specifying the form of multivariate time series models, with the latter text concluding (p. 260) that "the case for multiple time-series modeling is far from good, especially in view of the substantially greater effort and cost involved in formulating and estimating such models."

Jenkins and Alavi (1981) have presented a more optimistic case for multivariate ARMA modeling with a bivariate mink and muskrat example and a four-variable example with quantities and prices of butter and margarine. An important aspect of such models, however, is that much is known about their structure, a priori.

When very little is known about the structure of the interaction among variables, a difficult specification search is necessary. As Leamer (1978) emphasizes, that search is essentially a Bayesian process of combining prior information with the data. Litterman (1980) demonstrates a procedure for incorporating into the estimation of vector autoregressions the information that each of a set of economic variables (after suitable transformation) is likely to follow a random walk. That approach is potentially applicable to the set of variables studied here, but it ignores the index structure which is assumed to exist in the generation of business cycles. A combination of the above type of prior information with that incorporated here is one possible extension which has not yet been attempted.

An index structure generates an autoregressive representation of less than full rank. A number of methods may be used to estimate such a structure. Brillinger (1975) suggests a frequency domain canonical analysis, and Joreskog and Goldberger (1975) suggest a model with multiple indicators and multiple causes. I apply the observable index model described above because it allows a general form of dynamic response of the indexes to the observables and of the observables to the indexes. In addition, it allows the imposition of prior information in a convenient form.

All causal influences among variables in the index model must occur through their impact on the indexes. The prior information which is used concerns how important the impact of the indexes is thought to be. The purpose of imposing a prior density in this context is to bridge the gap between a set of univariate equations and the index model. A prior on θ allows one to specify in a probabilistic sense how important the impact of the index is thought to be.

Substituting (2) into (1) and collecting terms, the effect of the indexes on the k^{th} component of Y can be represented by

$$(6) \qquad \sum_{i=1}^{n} \sum_{j=1}^{L} \delta^{k}_{ij} \; Y_i(t-j)$$

where $L = L_a + L_b - 1$ and

$$(7) \qquad \delta^k_{ij} = \sum_{r=1}^{q} \sum_{\ell_1+\ell_2=j+1} a^r_{k\ell_1} b^r_{i\ell_2}$$

where $a^r_{k\ell_1}$ is the (k,r) element of A_{ℓ_1} and $b^r_{i\ell_2}$ is the (r,i) element of B_{ℓ_2}.

Given the difficulty of identifying the indexes individually, it is natural to formulate the prior information on the importance of the index effects directly in terms of the δ's. Thus, the information that the effects of the index are small is conveyed by a prior which puts small weight on large squared values of δ's. A prior density which incorporates this idea is

$$(8) \qquad p(\Delta) \propto \exp\left\{-\lambda^{-2} \sum_{i,j,k} \delta^k_{ij}(\theta)^2\right\}$$

where Δ is the stacked vector of δ's.

This prior is a normal density over the n^2L dimensional space spanned by the components of Δ. When the conditions for econometric identification of θ are met, this prior induces a prior on θ given by $\tilde{p}(\theta) = p[\Delta(\theta)]$, where Δ is explicitly written as a function of θ.

Equation (8) is actually a slight simplification of the prior used here, since it does not take account of the scale of the variables. If the units of a particular component of Y were to be changed by a factor of 10, for example, the prior should also be adjusted, since the δ's on that component in the equations of other components would also change by a factor of 10. In order to account for this effect, the prior includes scale factors, σ_i, associated with each component of Y. Each σ_i is the standard error of the residuals in a univariate autoregression of the i^{th} component on three lags and a constant. The adjusted prior is then given by

$$(9) \qquad p(\Delta) \propto \exp\left\{-\lambda^{-2} \sum_{i,j,k} [\delta^{k}_{ij}(\theta)(\sigma_k/\sigma_i)]^2\right\}.$$

Varying the tightness parameter, λ, will generate different priors representing, with more or less certainty, that the index effects are small. In the limit, as λ approaches zero, the index representation approaches that of a set of univariate equations. As λ gets larger, the prior becomes flat relative to the likelihood function.

The combination of index models and priors of the type described in this section does not lead to posterior densities of a computationally tractable form. For this reason, the estimation strategy described here is rather ad hoc. My intent is to describe in a meaningful sense a mapping from priors into a relevant measure of forecasting performance--in this case, a function of out-of-sample forecast errors.

In any given forecasting problem, it is desirable in a Bayesian analysis to find the predictive density of the future observation. With a squared error loss function, the posterior mean of this density is the optimal point forecast. If only one forecast were desired, it might be practical to perform the requisite numerical integration. The approach of Kloek and van Dijk (1978) might be useful in such a case. In this study, however, the desire is to look at several priors and many forecasts, so that such an approach is impractical.

Instead, I follow a procedure of maximizing the posterior density and using the estimate, θ, so derived to generate a set of forecasts. This procedure can be considered full information maximum likelihood if the information contained in the prior is viewed as a set of independent observations added to the likelihood function.

Examining the System's Performance

A Forecasting Experiment

My test of the usefulness of the observable index model is applied to the four macroeconomic variables mentioned above as those generally viewed as playing important roles in the business cycle. Again, the four variables are industrial production in manufacturing, the producer price index for manufacturing, the three-month Treasury bill rate, and the inventory-to-sales ratio for manufacturing and trade. (A fuller description of the data is given in Appendix 1.) Natural logarithms of industrial production and the producer price index are taken.

The one- and two-index models are specified with $L_d = L_a = L_b = 3$, which leads to a total of 39 free parameters in the one-index model and 60 in the two-index model. The experiment consists of choosing, for each model, a set of priors as described above with the parameter λ ranging from very loose to very tight. In each case, the model is first estimated over the monthly data from June 1948 to January 1970 (from 48:6 to 70:1). These initial estimates are used to generate out-of-sample forecasts for each month from 70:1 to 70:6. The forecasts are made for horizons ranging from 1 month to 36 months, with multistep forecasts based on the chain rule of forecasting. The model is then reestimated with data through 70:7, and another set of 6 forecasts is made. The sequence of adding observations, reestimating, and making forecasts is continued throughout the data period, which ends in 80:9. Thus, there are 128 one-step forecasts, 127 two-step forecasts, and so on. The procedure of reestimating every 6 months, rather than every month, is followed simply to reduce the computational expense.

For comparison, I also follow this forecasting procedure for a set of univariate equations (with three lags and a constant) estimated by the method of seemingly unrelated regressions, a set of univariate ARIMA models (see Appendix 2), index models with no prior, and an unrestricted vector autoregression (VAR).

Results of the Experiment

Results of this forecasting experiment are presented in two different ways in Tables 1 and 2. [Tables and figures follow page 22 of this paper.] In both tables, the results are displayed in terms of a standard forecast measure, Theil's U-statistic. This statistic is defined as the root mean square error of the forecast divided by the root mean square error of a no-change forecast.[2/] The purpose in presenting these averages of Theil statistics is simply to summarize the data. Tables 1 and 2 are

---

[2/]Under some strong distributional assumptions, one could use measures similar to these to test for significant improvements in forecasting. Rothemberg 1981 gives an example. Here, however, I am interested in the pattern of the results, not the significance of a particular statistic.

designed to make patterns in the forecast results easier to perceive than they would be in tables of simple mean square errors.

Overall, the results suggest at least a potential usefulness of index models for business cycle forecasting. Improvement in forecasting performance is not uniform, however, and the mixed nature of the results suggests that, at least for out-of-sample forecasting, it is important to damp the estimates in otherwise loosely parameterized models.

The main result of interest is that in both index specifications there is consistent improvement at the margin as the index effect is allowed to enter. As seen in Tables 1 and 2, this improvement at the margin occurs for all horizons and all variables and in both the one- and two-index models. In all cases, however, the results eventually deteriorate as the prior is loosened and the effect of the index is allowed to increase. While the extent of the improvement and the prior tightness at which the best results occur vary considerably across the different variables, the consistency of the general pattern suggests that further consideration of this type of estimator may be desirable.

Table 1 shows the average forecast performance across variables at different horizons. At the one-month forecast horizon, the results improve as the prior is loosened, up to the point where $\lambda = .5$ in the one-index model and $\lambda = .05$ in the two-index model. In both models, the performance then worsens and later improves as the prior is loosened further.

Similar patterns are evident at longer horizons, although the better λ's tend to be smaller for longer horizons. It is not obvious why priors which give more weight to the index should perform better at shorter horizons. One possibility is that the looser priors allow the models to better fit structural changes which persist only for a limited time. If this is true, then a time-varying coefficient specification in which parameters tend toward a time-invariant mean might prove beneficial. Another improvement would be to use the posterior mean for multistep forecasts rather than the chain rule, which I used to minimize computational expense.

Results for each variable averaged across horizons, shown in Table 2, follow essentially the same pattern as the averages across variables, but the point at which forecasts begin to deteriorate varies. The forecasts of the price index gain the least and quickly worsen as the index effect continues to enter. In the two-index model, industrial production also shows little improvement before worsening. Industrial production in the one-index model and the other variables in both models show substantial gains as the index effects are allowed to enter. The overall average of the Theil statistics improves from .875 in the univariate system to .818 in the one-index model with $\lambda = .2$ and to .829 in the two-index model with $\lambda = .05$. While there are no known distributions for these statistics by which to judge their significance, they do show considerably more improvement than those for the best ARIMA models.

Still, there is an obvious limitation to the practical application of this type of procedure. In formulating a prior, there is no way to know ahead of time what the appropriate tightness should be. If the forecasts were not too sensitive to this choice, it would not matter much how the tightness was chosen. In fact, however, the results demonstrate an unfortunate degree of sensitivity to this choice. For example, in the two-index model, the overall average of the Theil statistics jumps from .829 to .896 as λ moves from .05 to .1. In the one-index model, the three-year horizon average jumps from 1.020 to 1.735 to 1.035 as λ moves from 1 to 2 to 5. Thus, this class of models has the potential for generating bad, as well as good, forecasts.

One possible explanation for the mixed nature of these results is that it is an example of the well-known phenomenon that with limited observations highly parameterized models, even when correctly specified, generate out-of-sample estimates with large mean square errors. This interpretation suggests that the two-index restrictions on the vector autoregressive representation are not strong enough for forecasting purposes. If this is true, it is interesting to note that with less than 16 parameters per equation and the number of observations ranging from 265 to 391, the two-index model is not outside the normal range in terms of parameter-to-observation ratios for econometric models. This result would seem to suggest that there may be many contexts similar to this one in which Bayesian priors have a potential value for damping estimates in order to improve forecast performance.

Examining the Influences Among Variables

Testing the Index Fit

The observable index specification imposes cross-equation restrictions on the vector autoregressive representation which can be tested using the classical likelihood ratio test. The restrictions which the one-index structure imposes on the vector autoregressive model are rejected. The two-index model's restrictions on the vector autoregression, however, cannot be rejected at conventional significance levels.[3/] Detailed results of these likelihood ratio tests are presented in Table 3.

Decomposition of Variance

A useful way to summarize the channels of influence among variables in an econometric model is through a decomposition of forecast variance. This technique is described in Appendix 3, and the results of applying it here are shown in Table 4. The decomposition highlights two aspects of the index models. First, it reveals that there is a natural separation of effects into one real index and one nominal index. Second, it demonstrates the extent to which the two indexes are able to capture the full extent of dynamic interaction which is obtained in the unrestricted vector autoregression.

Given the ordering chosen here, in the univariate system the decomposition of industrial production must attribute 100

[3/] A direct test of the one-index model against the two-index model does not satisfy the conditions for the likelihood ratio test. Under the null hypothesis that the one-index model is correct, the parameters of the two-index model are not identified.

percent of the forecast variance at all horizons to own innovations. This result does not follow for the one-index model, wherein cross-effects are to be expected. However, in my one-index model, industrial production does in fact turn out to be almost completely exogenous. (Actually, at the four-year horizon, other innovations account for .18 of 1 percent of production's forecast variance.)

As seen in Table 5 (which shows the improvements in fit of individual equations), the inventory-sales ratio is the variable for which the fit is most improved by the addition of an index. The index allows the impact of industrial production on the forecast variance of this ratio to increase from slightly less than 30 percent, the amount due to contemporaneous correlation, to 62 percent at a one-year horizon (see Table 4d). The index also picks up substantial effects of both of these real variables on the nominal variables, prices and interest rates. Notice, however, that the index does not pick up any appreciable influence of the nominal variables on either of the real variables, nor does it capture any interaction between prices and interest rates.

The one-index model is a representation in which real variables are block exogenous and alone determine a real index. Innovations in industrial production and, to a lesser extent, in the inventory-sales ratio drive the index, and it in turn helps to predict the inventory-sales ratio, producer prices, and nominal interest rates. Again, industrial production itself is completely exogenous.

We have seen from Table 3, however, that the one-index model is strongly rejected by the data. This lack of fit is reflected in the decomposition of variance by the absence of channels of influence in the one-index model which do appear in the two-index model and the vector autoregression. The data do not reject the two-index model, and its decomposition of variance is very similar to that of the vector autoregression.

Figures 1 and 2 illustrate the main differences between the one- and two-index models. Arrows show the significant causal channels, and the width of the arrows and the numbers indicate the impacts measured by the percentage of forecast variance at a four-year horizon attributable to innovations in each of the variables.

Changes resulting from the addition of the second index are consistent with the interpretation that one index represents real effects and the other nominal effects. The addition of the second index primarily improves the fit of the nominal variables, but it also seems to open a channel of influence from nominal variables to real variables. In the two-index model, instead of being completely exogenous, industrial production obtains 28.56 percent of its variance at the four-year horizon from nominal interest rates. Prices and the inventory-sales ratio, however, still have very little influence on the variance of production. With the addition of the second index, nominal variables explain a substantial portion of the inventory-sales variance. At the four-year horizon, prices explain 7 percent and nominal interest rates 21 percent.

The addition of the second index causes a minor increase in the interest rate effect in the decomposition of prices, and it provides a channel for some additional impact of prices and industrial production on interest rates. Aside from this last impact of industrial production, which occurs only at horizons greater than one year, all effects of innovations in real variables which appear in the two-index model also appear in the one-index model. Yet most of the influences of nominal variable innovations show up only in the two-index model. Thus, although two indexes with real and nominal components are necessary to fit the data, when only one index is allowed, the fit is maximized by defining an index of the real components alone.

Impulse Response Functions

Some additional insight into the interpretation of these index effects can be obtained from the moving average representations of the index models. These response functions are plotted in Figures 3 through 6 along with the response function of an unrestricted vector autoregressive representation.

As is to be expected, the sizes of the responses, scaled in terms of own innovation standard deviations, correspond roughly to the relative importance in the decomposition of variance. In addition, the response functions reproduce the exogeneity of industrial production in the one-index model, the introduction of nominal cross-equation effects in the two-index model, and the close correspondence between the two-index model and the vector autoregression.

The main new information contained in these response functions concerns the direction and timing of different cross-equation effects. A positive innovation in industrial production generates the same basic pattern of responses in each model. This set of responses can be interpreted as an anticipated increase in demand, in the standard sense that price and quantity move up together. The industrial production innovation is met immediately by a persistent increase in inflation. Nominal interest rates build up over a period of four months, and the inventory-sales ratio falls immediately, then is slowly built back up to normal.

The response to the orthogonalized innovation in the inventory-sales ratio--that is, that part of the innovation uncorrelated with innovations to the other variables--also leads to the same responses in all three models. This set of responses can be interpreted as an unanticipated decrease in demand. The increase in the inventory-sales ratio is gradually reduced as industrial production dips and then returns to normal. There is an immediate deflationary impact which persists, and interest rates drop and remain lower than otherwise for approximately two years. Between the first and second year, industrial production increases to a point above where it would have been without the shock.

Price level innovations are persistent, but have small cross-equation effects in all the models. In each case, industrial production rises for about six months and then falls below normal, inventory-sales dip and then rise above normal, and interest rates rise and then slowly return to normal. The responses are basically the same in all three models, but are somewhat larger in the two-index than in the one-index model and are larger still (but still small) in the vector autoregression.

The only noticeable difference in the response patterns across the three models appears in the set of responses to an interest rate innovation, between the one-index and the two-index model. The one-index model does not appear to provide a channel for this innovation to have an effect, whereas the two-index model does. There the interest rate innovation causes an immediate rise in the inventory-sales ratio, which is followed shortly by a prolonged decline in industrial production. Prices respond with a small decline within a few months, followed by a sustained drop after about a year. The fact that the inventory-sales ratio responds positively and somewhat before industrial production while prices respond negatively suggests that these reactions reflect a reduction in demand in response to tighter credit rather than a correct anticipation by borrowers and lenders of future negative supply effects. One possible interpretation is that these innovations largely represent actions of the monetary authority.

Taken together, the decompositions of variance and the impulse response functions suggest an interesting separation of the causes of manufacturing cycles into real shocks, which are primarily shifts in demand, and financial shocks, which may be policy induced.

Conclusion

The use of Bayesian procedures allows the estimation of models with varying degrees of influence from observable indexes. Relative to a system of univariate equations, the introduction of the index effect at the margin improves out-of-sample forecasting performance for a set of economic time series related to the business cycle. There appears to be a general pattern of improvement and then deterioration in forecasts as the index effect is increasingly allowed to enter. Inspection of the unrestricted one- and two-index models and their associated decompositions of forecast variance suggests that there may be a separation of real and nominal influences into different indexes. Such a separation might provide a useful distinction between real shocks to tastes and technology and shocks generated in financial markets or by monetary policy. Such an identification of the observable indexes might be achieved by testing exclusionary restrictions in the context of a larger set of variables. That will be the subject of future research.

TABLES AND FIGURES

Table 1

Average Forecasting Performance
of All Variables Over Each Horizon

(Averages of Theil U-Statistics)

Model	Forecast Horizon in Months						
	1	2	3	6	12	24	36
Univariate							
ARIMA*	.851	.843	.857	.880	.887	.891	.870
SUR**	.869	.869	.882	.903	.886	.840	.847
One-Index							
λ = .001	.869	.868	.881	.901	.883	.837	.843
.010	.867	.866	.878	.897	.880	.832	.839
.100	.855	.843	.849	.850	.816	.736	.698
.200	.847	.832	.834	.834	.806	.755	.752
.500	.844	.829	.836	.848	.819	.775	.779
1.000	.850	.843	.858	.893	.938	.952	1.020
2.000	.862	.863	.892	.971	1.097	1.302	1.735
5.000	.855	.847	.864	.901	.944	.949	1.035
No Prior	.847	.833	.848	.886	.922	.938	1.039
Two-Index							
λ = .001	.869	.867	.880	.900	.882	.835	.841
.010	.859	.851	.858	.863	.837	.774	.740
.050	.855	.843	.847	.847	.821	.770	.744
.100	.864	.863	.877	.907	.928	.918	.978
.200	.847	.840	.851	.878	.914	.964	1.130
.500	.835	.828	.837	.856	.859	.882	1.032
1.000	.830	.825	.834	.844	.821	.825	.957
2.000	.830	.825	.833	.845	.823	.834	.965
No Prior	.854	.861	.879	.927	.978	.993	1.120
VAR	.849	.843	.856	.880	.888	.914	1.098

*The selection of ARIMA models is described in Appendix 2.

**The SUR line reports results for a set of univariate autoregressions with three lags and a constant, estimated by the seemingly unrelated regression method. This model is the limiting case for either index model as λ approaches zero.

Table 2

Average Forecasting Performance
of Each Variable Over All Horizons

(Averages* of Theil U-Statistics)

Model	Industrial Production	Producer Price Index	Treasury Bill Rate	Inventory-Sales Ratio	All Four Variables
Univariate					
ARIMA	.860	.604	.997	1.014	.869
SUR	.900	.602	1.020	.975	.875
One-Index					
λ = .001	.900	.601	1.018	.972	.873
.010	.898	.610	1.005	.965	.870
.100	.899	.642	.893	.857	.823
.200	.898	.641	.888	.845	.818
.500	.842	.664	.934	.862	.826
1.000	.928	.681	1.079	.885	.893
2.000	1.350	.669	1.085	.962	1.017
5.000	1.029	.675	1.004	.885	.898
No Prior	1.009	.636	1.007	.886	.885
Two-Index					
λ = .001	.900	.602	1.018	.969	.872
.010	.896	.610	.964	.885	.839
.050	.897	.607	.949	.865	.829
.100	1.033	.655	.950	.945	.896
.200	.982	.684	.938	.955	.890
.500	.888	.702	.935	.896	.855
1.000	.840	.691	.933	.874	.835
2.000	.848	.677	.930	.893	.837
No Prior	1.075	.667	.926	1.017	.921
VAR	.993	.705	.928	.888	.878

*Each average is weighted across different horizons with weight proportional to k^{-1} at the k-step horizon for $k = 1,\ldots,36$.

Table 3

Results of the Likelihood Ratio Tests*

Hypothesis	Log Determinants	χ^2	Marginal Significance Level
I: One-Index vs. VAR	-29.9171 -30.1910	$\chi^2(45) = 100.54$	$< 10^{-5}$
II: Two-Index vs. VAR	-30.1280 -30.1910	$\chi^2(24) = 23.14$	0.51

*The likelihood ratio tests follow Sims' (1980) suggestion of correcting for degrees of freedom. The tests are performed on the entire sample, 388 observations. The degrees of freedom correction is 21 for both hypotheses. Thus, the statistic reported is $(T-k)(\log|\Sigma^u| - \log|\Sigma^c|)$, where $T-k = 367$ and Σ^u and Σ^c are the covariance matrices of the unrestricted (VAR) and constrained (index) models, respectively. The marginal significance level is $\Pr[\chi^2(\cdot) > c]$, where c is the value of the test statistic.

Table 4

The Decomposition of Forecast Variance

Table 4a Industrial Production

MODEL/Horizon	Standard Error	% of Variance in Production Due to Innovations in Industrial Production	Producer Price Index	Treasury Bill Rate	Inventory-Sales Ratio
UNIVARIATE					
1	.0122	100.00	.00	.00	.00
2	.0200	100.00	.00	.00	.00
3	.0267	100.00	.00	.00	.00
6	.0420	100.00	.00	.00	.00
12	.0629	100.00	.00	.00	.00
24	.0914	100.00	.00	.00	.00
36	.1132	100.00	.00	.00	.00
48	.1315	100.00	.00	.00	.00
ONE-INDEX					
1	.0120	100.00	.00	.00	.00
2	.0205	99.89	.00	.00	.10
3	.0286	99.70	.01	.00	.29
6	.0484	99.61	.00	.00	.38
12	.0748	99.65	.03	.01	.31
24	.1075	99.73	.06	.01	.20
36	.1303	99.79	.06	.00	.15
48	.1485	99.82	.06	.00	.12
TWO-INDEX					
1	.0117	100.00	.00	.00	.00
2	.0196	99.84	.09	.00	.06
3	.0267	99.52	.17	.08	.24
6	.0443	97.92	.14	1.28	.67
12	.0673	92.01	.34	7.27	.38
24	.0961	78.22	.89	20.10	.79
36	.1144	70.95	.87	26.29	1.89
48	.1276	67.89	.80	28.56	2.74
VAR					
1	.0116	100.00	.00	.00	.00
2	.0194	99.73	.23	.00	.04
3	.0263	99.29	.49	.08	.13
6	.0431	97.12	.35	1.34	1.18
12	.0641	90.44	2.24	6.23	1.09
24	.0895	74.62	5.98	18.47	.93
36	.1053	66.35	6.08	25.50	2.07
48	.1168	63.16	5.64	28.15	3.05

Table 4b The Producer Price Index

MODEL/Horizon	Standard Error	% of Variance in Price Index Due to Innovations in: Industrial Production	Producer Price Index	Treasury Bill Rate	Inventory-Sales Ratio
UNIVARIATE					
1	.00434	.40	99.60	.00	.00
2	.00708	.40	99.60	.00	.00
3	.01042	.40	99.60	.00	.00
6	.01991	.40	99.60	.00	.00
12	.03641	.40	99.60	.00	.00
24	.03641	.40	99.60	.00	.00
36	.08304	.40	99.60	.00	.00
48	.10299	.40	99.60	.00	.00
ONE-INDEX					
1	.00422	.18	99.82	.00	.00
2	.00688	1.05	98.10	.00	.84
3	.00999	1.47	97.65	.00	.88
6	.01861	3.87	94.36	.01	1.75
12	.03330	10.26	85.88	.06	3.79
24	.05737	22.46	70.19	.16	7.18
36	.07880	30.01	60.73	.25	9.02
48	.09919	34.45	55.27	.31	9.96
TWO-INDEX					
1	.00407	.80	99.20	.00	.00
2	.00644	2.54	95.53	.17	1.76
3	.00911	4.05	93.47	.35	2.13
6	.01769	6.92	89.79	.50	2.78
12	.03460	12.36	82.37	.39	4.88
24	.06341	21.47	69.73	.73	8.07
36	.08665	26.96	61.70	2.06	9.32
48	.10673	29.91	56.94	3.71	9.44
VAR					
1	.00402	.91	99.09	.00	.00
2	.00642	3.44	94.90	.56	1.10
3	.00912	5.50	92.55	.79	1.17
6	.01768	9.68	87.70	1.29	1.33
12	.03409	17.14	78.67	.73	3.45
24	.06013	27.81	63.77	.78	7.63
36	.08020	33.27	55.03	1.96	9.74
48	.09781	35.66	50.56	3.48	10.29

Table 4c The Three-Month Treasury Bill Rate

MODEL/Horizon	Standard Error	% of Variance in Bill Rate Due to Innovations in: Industrial Production	Producer Price Index	Treasury Bill Rate	Inventory-Sales Ratio
UNIVARIATE					
1	.382	3.37	2.98	93.65	.00
2	.669	3.37	2.98	93.65	.00
3	.837	3.37	2.98	93.65	.00
6	1.107	3.37	2.98	93.65	.00
12	1.509	3.37	2.98	93.65	.00
24	2.033	3.37	2.98	93.65	.00
36	2.391	3.37	2.98	93.65	.00
48	2.658	3.37	2.98	93.65	.00
ONE-INDEX					
1	.370	2.82	2.05	95.13	.00
2	.643	4.09	2.23	93.34	.34
3	.811	6.77	2.42	88.84	1.97
6	1.114	13.76	2.10	80.77	3.37
12	1.590	18.96	1.19	76.21	3.64
24	2.307	22.49	.58	73.07	3.86
36	2.893	23.62	.38	72.16	3.84
48	3.416	23.94	.31	72.01	3.74
TWO-INDEX					
1	.354	1.99	2.22	95.79	.00
2	.604	2.90	3.83	93.02	.25
3	.742	4.86	5.05	88.41	1.69
6	.938	9.38	7.30	80.74	2.59
12	1.105	16.70	6.24	73.03	4.03
24	1.218	27.92	5.19	61.61	5.28
36	1.281	31.44	4.78	58.85	4.93
48	1.327	32.91	4.87	57.62	4.60
VAR					
1	.352	2.07	2.09	95.84	.00
2	.600	3.74	4.16	91.94	.15
3	.738	5.52	5.87	86.98	1.63
6	.926	10.26	8.43	78.76	2.54
12	1.113	17.36	6.50	71.38	4.76
24	1.235	26.92	5.70	60.15	7.22
36	2.285	29.57	5.49	57.91	7.03
48	1.320	30.66	5.34	57.33	6.67

Table 4d The Inventory-Sales Ratio

MODEL/Horizon	Standard Error	% of Variance in I-S Ratio Due to Innovations in Industrial Production	Producer Price Index	Treasury Bill Rate	Inventory-Sales Ratio
UNIVARIATE					
1	.0225	28.76	1.61	.00	69.63
2	.0303	28.76	1.61	.00	69.63
3	.0362	28.76	1.61	.00	69.63
6	.0470	28.76	1.61	.00	69.63
12	.0565	28.76	1.61	.00	69.63
24	.0617	28.76	1.61	.00	69.63
36	.0627	28.76	1.61	.00	69.63
48	.0629	28.76	1.61	.00	69.63
ONE-INDEX					
1	.0210	29.18	1.68	.01	69.13
2	.0288	38.47	1.38	.12	60.02
3	.0357	48.65	1.00	.32	50.04
6	.0476	58.63	1.39	.43	39.54
12	.0583	61.98	3.34	.46	34.22
24	.0623	61.65	4.91	.47	32.97
36	.0626	61.31	5.27	.47	32.96
48	.0627	61.35	5.36	.47	32.83
TWO-INDEX					
1	.0207	28.19	1.55	.00	70.26
2	.0282	35.93	2.09	.42	61.57
3	.0344	44.22	1.66	.87	53.25
6	.0443	49.90	1.26	2.51	46.33
12	.0529	46.89	3.40	9.84	39.86
24	.0587	39.30	6.75	20.80	33.14
36	.0602	39.35	7.10	21.62	31.93
48	.0609	40.16	7.11	21.30	31.43
VAR					
1	.0207	27.76	1.58	.00	70.67
2	.0280	35.43	2.28	.42	61.87
3	.0340	44.01	1.92	.98	53.09
6	.0440	48.36	1.88	2.68	47.08
12	.0536	42.34	8.16	8.25	41.26
24	.0605	34.26	14.58	17.80	33.36
36	.0623	34.81	14.15	18.80	32.25
48	.0629	35.48	13.91	18.61	32.00

Table 5

How Much Adding an Index
Improves the Fit of Individual Equations

(% Reductions in Standard Errors)

	Industrial Production	Producer Price Index	Treasury Bill Rate	Inventory-Sales Ratio
From Univariate to One-Index	1.35	2.87	3.10	6.48
From One-Index to Two-Index	2.26	3.52	4.35	1.07
From Two-Index to VAR	1.12	1.29	0.46	0.52

Figure 1

Economic Interpretation of the One-Index Model: Which Variables Determine Which

(Numbers refer to variance decomposition at a four-year horizon.)

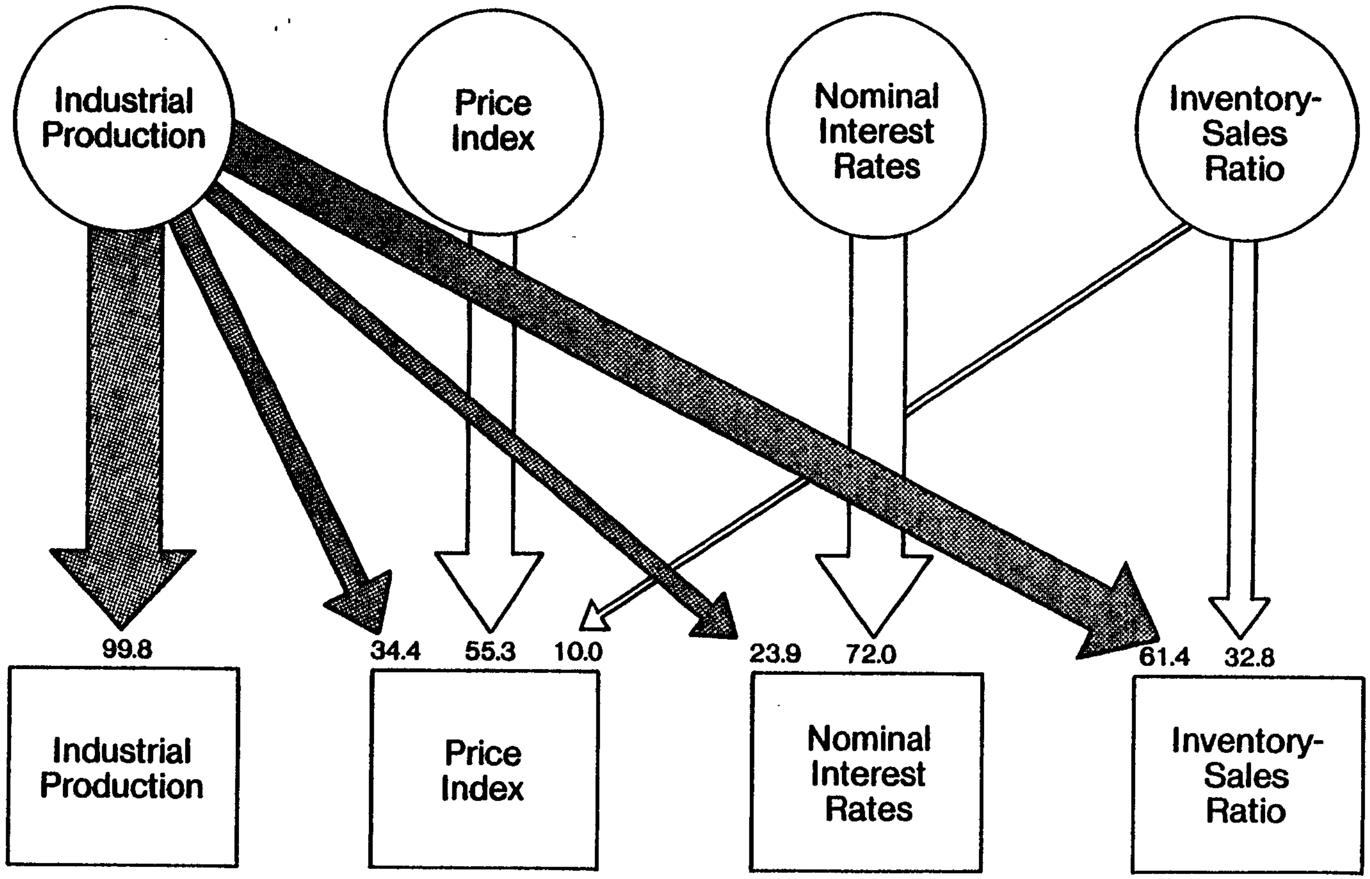

Figure 2

Economic Interpretation of the Two-Index Model: Which Variables Determine Which

(Numbers refer to variance decompostion at a four-year horizon.)

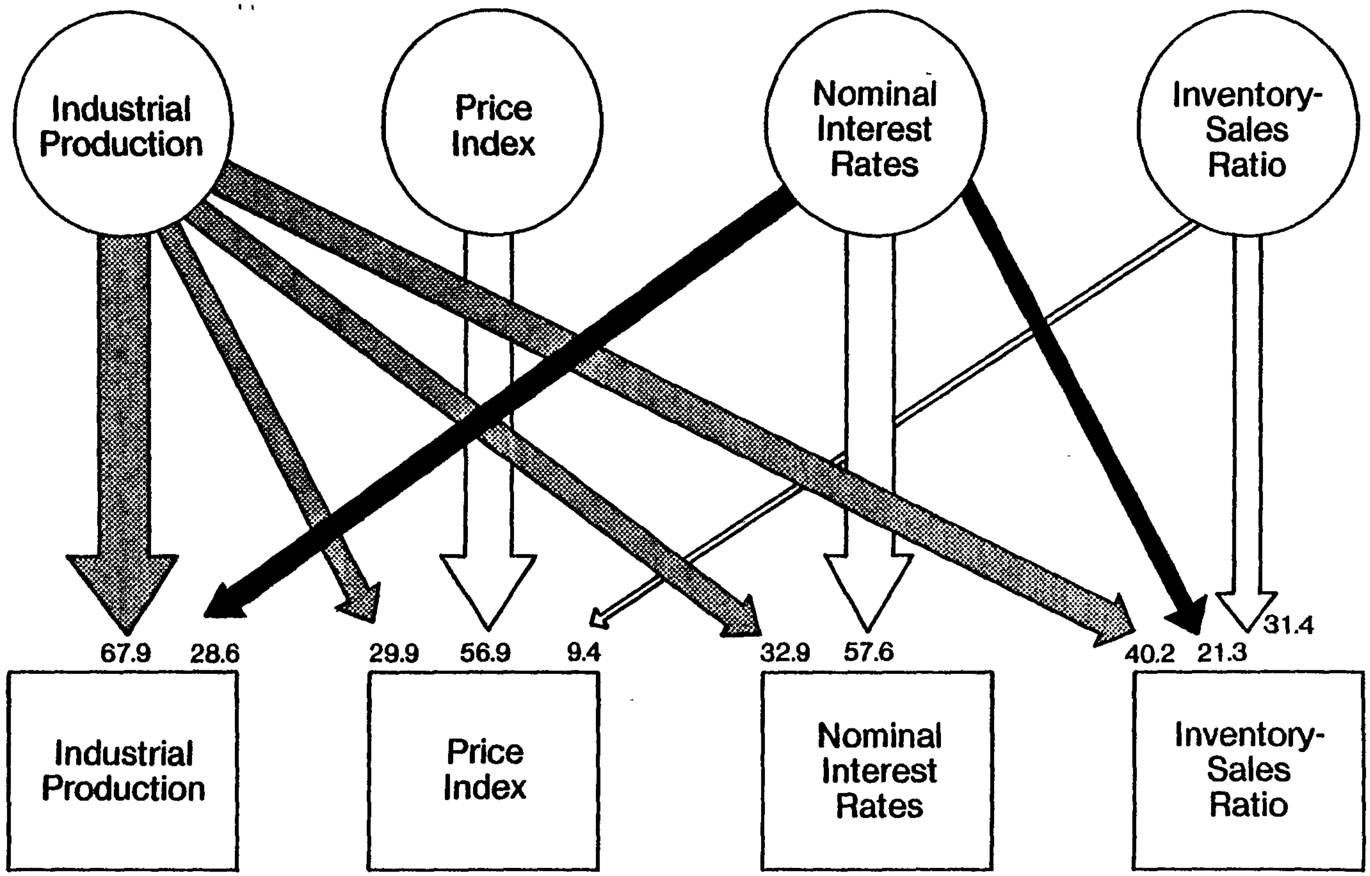

Figure 3

Response to an Innovation in Industrial Production

Key: Response of ——— Industrial Production
- - - - - - Price Index
— - — - — Nominal Interest Rates
— — — — Inventory-Sales Ratio

Vertical scale: Standard Deviations
Horizontal scale: Months

One-Index Model

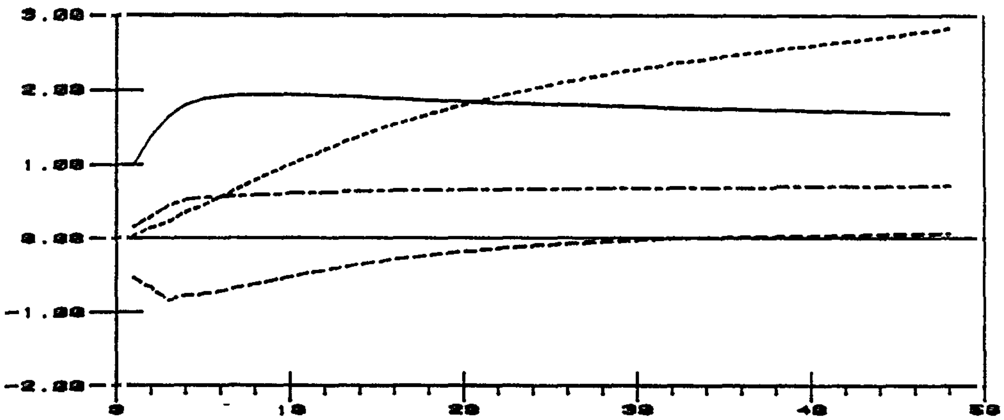

Two-Index Model

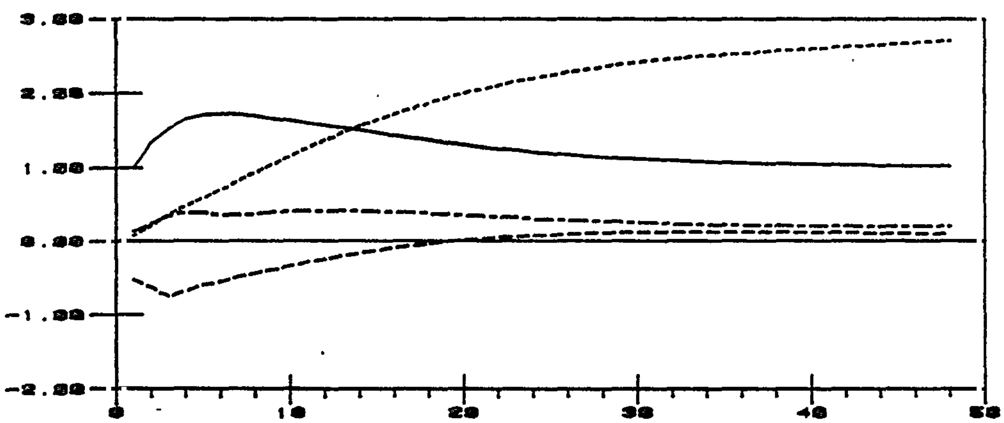

Vector Autoregression

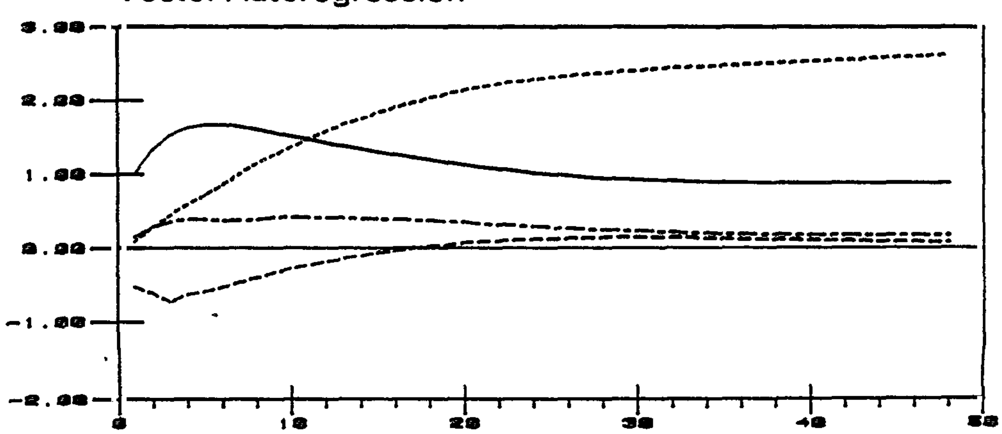

Figure 4

Response to an Innovation in the Price Index

Key: Response of
- ———— Industrial Production
- - - - - - Price Index
- — — — Nominal Interest Rates
- — - — Inventory-Sales Ratio

Vertical scale: Standard Deviations
Horizontal scale: Months

One-Index Model

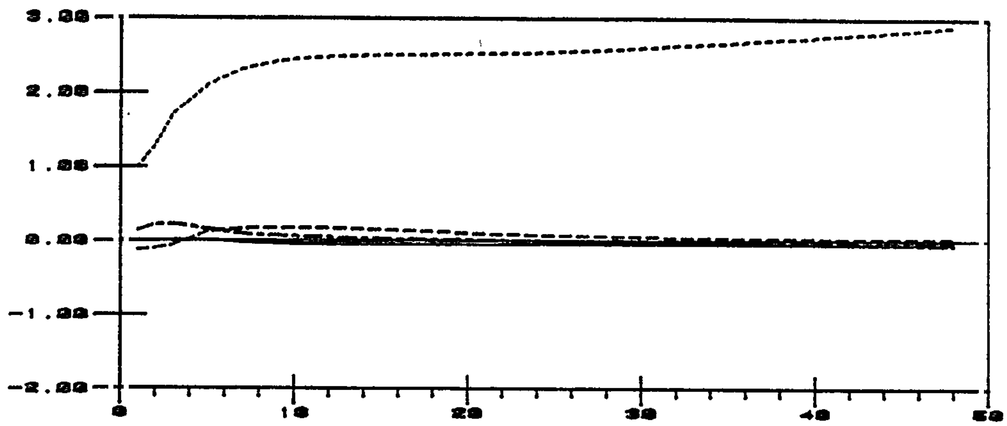

Two-Index Model

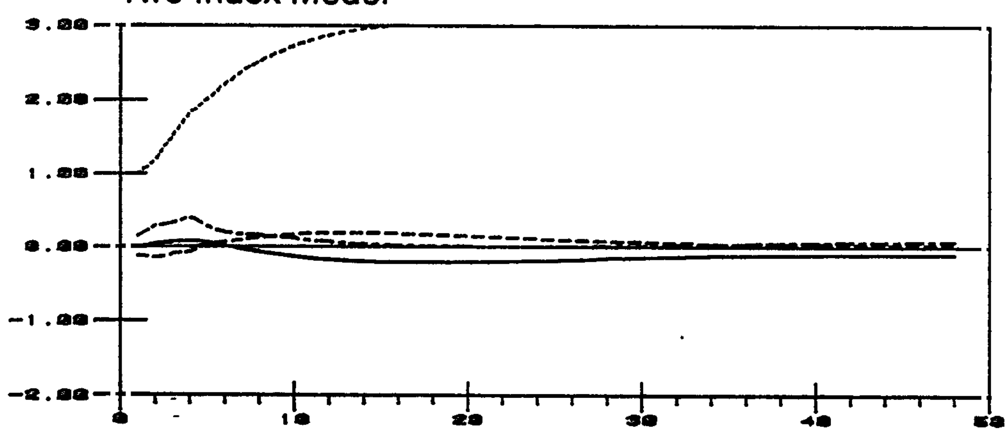

Vector Autoregression

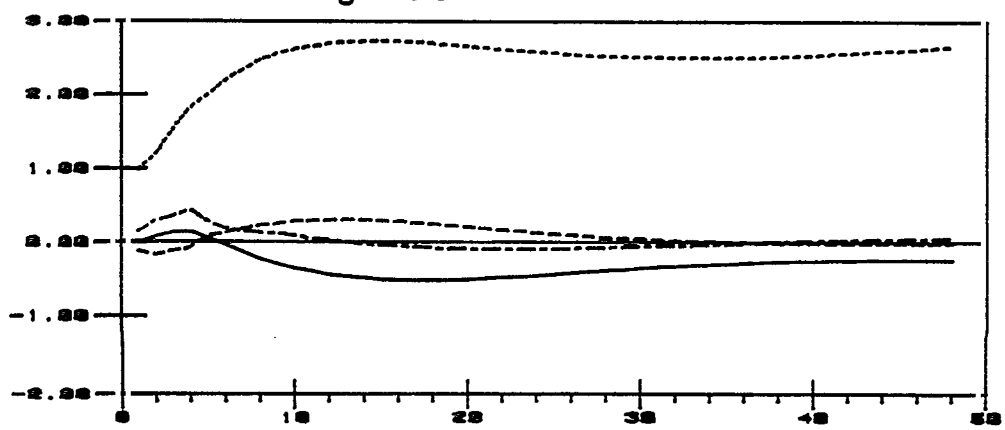

Figure 5
Response to an Innovation in Nominal Interest Rates

Key: Response of ——— Industrial Production
- - - - - - Price Index
— - — - Nominal Interest Rates
— — — — Inventory-Sales Ratio

Vertical scale: Standard Deviations
Horizontal scale: Months

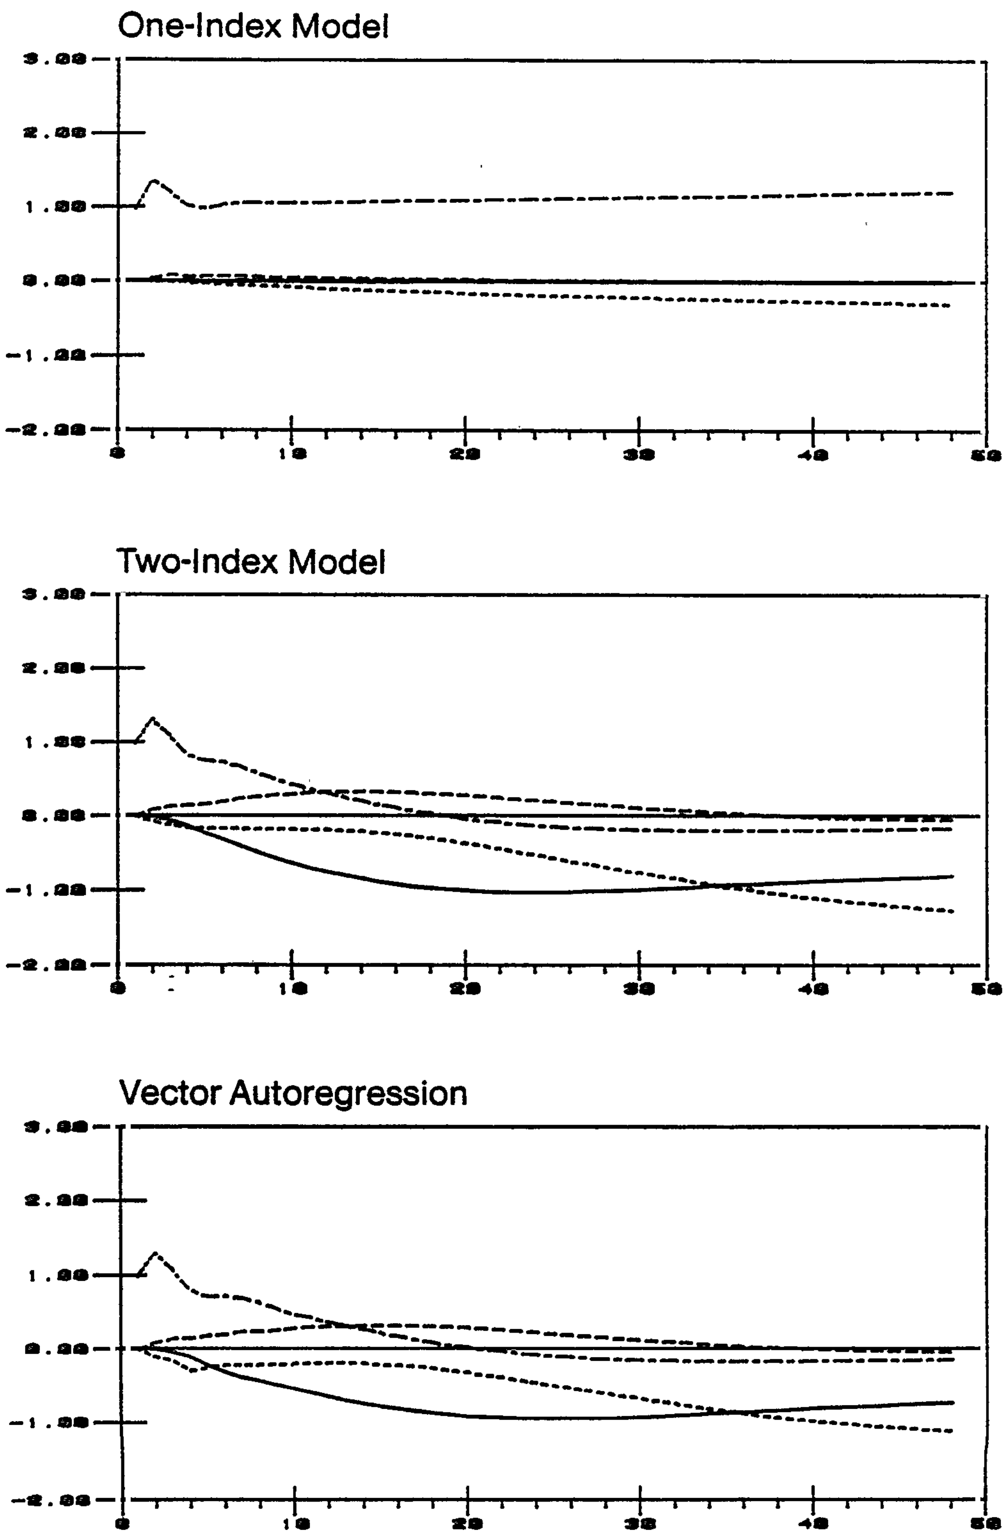

Figure 6

Response to an Innovation in the Inventory-Sales Ratio

Key: Response of
——— Industrial Production
------ Price Index
—-—- Nominal Interest Rates
— — — Inventory-Sales Ratio

Vertical scale: Standard Deviations
Horizontal scale: Months

One-Index Model

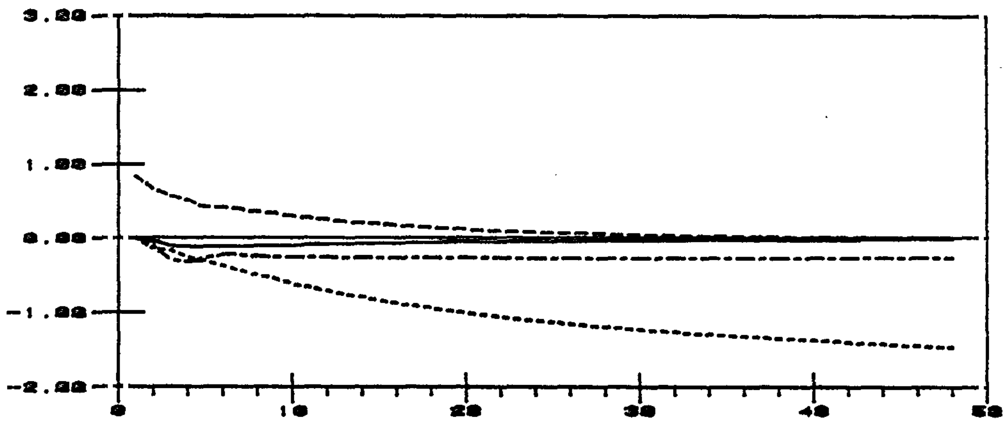

Two-Index Model

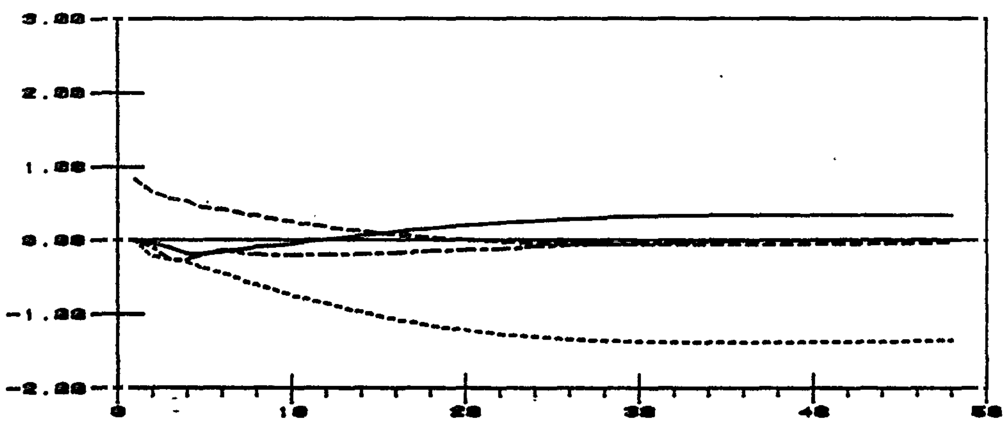

Vector Autoregression

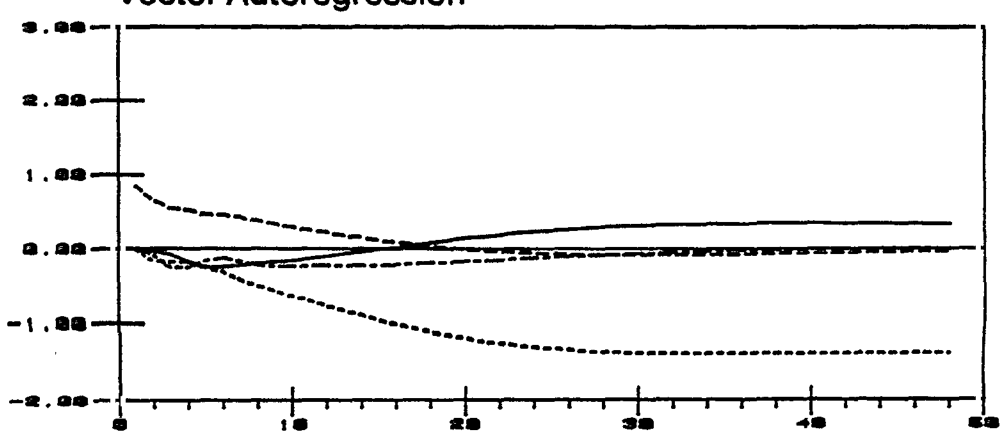

APPENDIXES

1	The Data
2	The ARIMA Models
3	The Decomposition of Variance

Appendix 1

The Data

The series used in this study are taken from the Citibank data base available from the TROLL computer system at MIT. The data are monthly from January 1948 through September 1980. The series are as follows:

IPMFG: Industrial Production: Manufacturing (1967 = 100) seasonally adjusted

PWMSA: Producer Price Index: Manufactured Goods (1967 = 100) seasonally adjusted

FYGM3: U.S. Government Security Yield Three-Month Bill, Market Yield (percent per annum)

IVT72: Manufacturing and Trade, constant '72 dollar inventory to sales ratio seasonally adjusted.

Appendix 2

The ARIMA Models

The benchmark ARIMA models are estimated by standard techniques.* I choose a number of alternative specifications on the basis of the autocorrelation and partial autocorrelation functions. For each series, the choice is the specification which produces the best out-of-sample forecasts. The resulting equations are as follows:

Industrial Production	(2, 1, 0)
Producer Price Index	(3, 1, 0)
Treasury Bill Rate	(2, 1, 0)
Inventory-Sales Ratio	(0, 1, 0)

where (i, j, k) gives the order of the autoregression, differencing, and moving average, respectively. Notice that this procedure essentially provides a lower bound on the size of errors for ARIMA models. It may give a distorted picture of the performance one would get if models were chosen by conventional methods.

---

*Alternative specifications are selected essentially as suggested in Chapters 7 and 10 of Nerlove, Grether, and Carvalho 1979.

Appendix 3

The Decomposition of Variance

The decomposition of variance is derived from the moving average representation associated with a time series model. The moving average representation gives the response functions of each variable to innovations in each of the endogenous variables in the system. An innovation is defined as the difference between the projection of a variable on information available one period earlier and the variable's actual value. (Sargent 1978, 1979 give a more detailed discussion.) The moving average representation for an index model can be easily generated by forming the autoregressive representation associated with an index model and inverting it.

Here, let the moving average representation be given by

$$Y_t = \sum_{j=0}^{\infty} M_j \, \varepsilon_{t-j}$$

where M_0 is normalized to be the identity matrix and ε_t is the vector of innovations at time t. The variance matrix of a k-step-ahead prediction is then given by

$$\Omega_k = \sum_{j=0}^{k-1} M_j \, \Sigma \, M_j'$$

where Σ is the variance matrix of the innovation vector. If the components of the innovation vector were orthogonal (that is, if Σ were diagonal), then the forecast variance could be decomposed as

$$\Omega_k = \sum_{i=1}^{n} \sum_{j=0}^{k-1} M_j \Sigma_i M_j'$$

where Σ_i is the matrix with the i^{th} diagonal element equal to the variance of the innovation in the i^{th} variable and with zeros elsewhere.

Since the innovations are not, in general, orthogonal, a decomposition of forecast variance relies on some orthogonalization procedure. One way to accomplish this orthogonalization is to choose an ordering and define the i^{th} orthogonalized innovation to be that part of the i^{th} innovation which is orthogonal to all previous innovations.

The ordering I use is industrial production, producer price index, Treasury bill rate, inventory-sales ratio. In any of the models I test, the only correlation between innovations which is greater than .20 is that between industrial production and the inventory-sales ratio. Thus, the major impact of this ordering is that the forecast variance of the inventory-sales ratio which is correlated with industrial production is attributed to the latter in the decomposition. This attribution is apparent in the decomposition at a one-step horizon in which all cross-variable effects are due to contemporaneous correlations of innovations.

References

Brillinger, David R. 1975. Time series: data analysis and theory. New York: Holt, Rinehart and Winston.

Geweke, John F. 1977. The dynamic factor analysis of economic time series. In Latent variables in socio-economic models, ed. D. J. Aigner and A. S. Goldberger, pp. 365-83. Contributions to Economic Analysis, vol. 103. Amsterdam: North-Holland.

Granger, C. W. J., and Newbold, Paul. 1977. Forecasting economic time series. New York: Academic Press.

Jenkins, Gwilym M., and Alavi, Athar S. 1981. Some aspects of modelling and forecasting multivariate time series. Journal of Time Series Analysis 2 (No. 1): 1-47.

Joreskog, Karl G., and Goldberger, Arthur S. 1975. Estimation of a model with multiple indicators and multiple causes of a single latent variable. Journal of the American Statistical Association 70 (September): 631-39.

Kloek, T., and van Dijk, H. K. 1978. Bayesian estimates of equation system parameters: an application of integration by Monte Carlo. Econometrica 46 (January): 1-19.

Leamer, Edward E. 1978. Specification searches: ad hoc inference with nonexperimental data. New York: Wiley.

Litterman, Robert B. 1980. Techniques of forecasting using vector autoregressions. Ph.D. diss., University of Minnesota. Also 1979. Research Department Working Paper 115. Federal Reserve Bank of Minneapolis.

Nerlove, Marc; Grether, David M.; and Carvalho, Jose L. 1979. Analysis of economic time series: a synthesis. New York: Academic Press.

Rothemberg, Julio J. 1981. Selection of econometric models with out-of-sample data. Massachusetts Institute of Technology. Manuscript.

Sargent, Thomas J. 1978. Estimation of dynamic labor demand schedules under rational expectations. Journal of Political Economy 86 (December): 1009-44.

________________. 1979. Macroeconomic theory. New York: Academic Press.

Sargent, Thomas J., and Sims, Christopher A. 1977. Business cycle modeling without pretending to have too much a priori economic theory. In New methods in business cycle research: proceedings from a conference, ed. C. A. Sims, pp. 45-109. Minneapolis: Federal Reserve Bank of Minneapolis.

Sims, Christopher A. 1980. Macroeconomics and reality. Econometrica 48 (January): 1-48.

________________. Forthcoming. An autoregressive index model for the U.S., 1948-1975. In Large-scale macro-econometric models: theory and practice, ed. J. Kmenta and J. B. Ramsey. Contributions to Economic Analysis, vol. 141. Amsterdam: North-Holland.

Singleton, Kenneth J. 1981. Real and nominal factors in the cyclical behavior of interest rates, output, and money. Conference Paper 123. National Bureau of Economic Research.