Applied Economics, 1990, **22**, 1751–1759

Are monetary services indexes superior to conventional monetary aggregates as intermediate targets?

JAMES S. FACKLER[1], W. DOUGLAS MCMILLIN[2] and J. LEW SILVER[3]

Departments of Economics, [1]University of Kentucky, Lexington, KY 40506, [2]Louisiana State University, Baton Rouge, LA 70803, and *[3]Emory University, Atlanta, GA 30322, USA*

I. INTRODUCTION

The financial innovation and deregulation of recent years and the unusual behaviour of the income velocity of the conventional M1 monetary aggregate in the 1980s has heightened concern about the appropriateness of focusing on the conventional (simple-sum) monetary aggregates in implementing monetary policy. The conceptual shortcomings of the conventional monetary aggregates, which are unweighted sums of the elements of each aggregate, have been a matter of concern to monetary economists for some time. One of the primary criticisms of these aggregates is the implicit assumption that each element of a conventional aggregate is a perfect substitute for the other elements of that aggregate. Fisher (1922) regarded the simple-sum aggregates to be the worst monetary index, and Friedman and Schwartz (1970) suggested a general approach to constructing monetary aggregates in which aggregates are constructed as weighted averages of the individual elements with the value of the weights depending upon the degree of 'moneyness' of the element.

The concern about the appropriateness of the conventional aggregates has led the Federal Reserve Board to develop a set of weighted monetary aggregates referred to as the Monetary Services Indices (MSIs). These indices are based on index number theory, and are described in detail in Farr and Johnson (1985) and Lindsey and Spindt (1986). An MSI analogue to each of the conventional aggregates (M1, M2, M3, L) has been developed, and these analogues are weighted averages of the elements of the conventional aggregate. The weight given an element in an MSI is based upon the user cost of that element, and is defined as the difference between a maximum available holding period return and the own holding period return. Although the theoretial properties of the MSIs are desirable, construction of real world counterparts is not without its difficulties. Accurate measurement of the user cost is problematic, as Kennickell and DeMarco (1986) and Farr and Johnson (1985) have pointed out.

Because of the potential problems in the construction of the MSIs, their superiority to the conventional aggregates is an empirical question. Available empirical work provides mixed evidence on the desirability of the MSIs relative to the conventional aggregates. Prior work, which consists of single-equation studies, has examined the forecasting performance and coefficient stability of money demand functions for both conventional and MSI aggregates

0003-6846/90 $03.00+.12

as well as the forecasting performance and stability of St Louis-type reduced form equations (e.g. Batten and Thornton, 1985; Lindsey and Spindt, 1986; Belongia and Chalfant, 1989).

This paper compares the MSI and conventional aggregates on the basis of their ability to forecast nominal income out-of-sample. If the MSI aggregates are to be seriously considered as alternatives to the conventional aggregates as intermediate targets and/or indicators of monetary policy, there must be a close, predictable relationship between these aggregates and economic activity. It is common practice in the intermediate target/monetary indicator literature to focus on the money–nominal income relationship, and this convention is followed here. The present goal is to determine whether the MSIs can more accurately forecast nominal income than the conventional aggregates and hence whether the MSIs are good candidates to serve as intermediate targets and/or indicators of monetary policy.[1]

The ability of the conventional aggregates (M1, M2, M3 and L) and their MSI counterparts (MSI1, MSI2, MSI3, MSI4, respectively) to forecast nominal income out-of-sample is examined primarily within three variable 'Bayesian' vector autoregressive (BVAR) models.[2] These BVAR models are in the spirit of the St Louis equation in the sense that each model contains a nominal income measure (either nominal GNP or nominal personal income), a monetary aggregate, and a fiscal policy variable (cyclically adjusted federal expenditures). However, unlike the single-equation St Louis approach, the monetary aggregate and the fiscal variable are not assumed to be exogenous. The BVAR model avoids these types of potentially spurious *a priori* constraints and instead treats all variables as jointly determined. The BVAR model approach, developed by Litterman (1979, 1980), is chosen since it has been demonstrated to be competitive in terms of forecasting with large, expensive-to-build structural models (e.g. Litterman, 1986). It should be noted that although the vector autoregressive methodology has been criticized on a variety of grounds (e.g. Cooley and LeRoy, 1985), the use of vector autoregressions for forecasting is accepted by these critics. Further, the BVAR approach primarily employed here is apparently embraced by some proponents of Bayesian techniques in general (e.g. Zellner, 1985).

The empirical results presented below provide no evidence supporting the superiority of the MSIs in terms of their abilities to foreshadow changes in economic activity. Rather, it is found that in terms of forecasting nominal income, the conventional aggregates outperform the MSIs. This conclusion is based on a set of models which, approximately, simulate real-time forecasting exercises.[3] Specifically, for each money measure and each income variable investigated, an optimal forecasting model is selected and then a sequence of out-of-sample forecasts of nominal income is made. Quarterly data for the period 1971:1–1984:3 are employed. Since the base period estimation is over 1971:1–1979:3 and since the data end in 1984:3, we have 31 four-quarter-ahead forecasts. The forecasting results based on our

[1]It is well known that a close link between a monetary aggregate and economic activity is a necessary, but not sufficient, condition for the aggregate to be a good intermediate target/indicator candidate. The variable must also be controllable by the central bank and information about the aggregate must be available on a timely basis. However, the only concern here is with the links between the aggregate and nominal income.

[2]It might also be interesting to include in such a comparison Spindt's (1985) MQ measure of money. It is not examined here, since the narrow conventional aggregate (M1) and the narrow Monetary Services Index (MSI1) perform relatively poorly and since MQ excludes non-medium-of-exchange assets such as those in the broader conventional and MSI aggregates.

[3]The term 'approximately' is used because, for example, use is made of recent revisions to the data which would not have been available to the real-time forecaster. Data used in this project are available on request.

experiments appear unique since other quarterly forecasting exercises which have been performed using data of this type have estimated a base period regression and then simulated several quarters into the future, so that the forecast comparisons are based only upon one realization for each aggregate.

The organization of the paper is as follows. Section II presents a brief discussion of the models. Section III presents the empirical results and Section IV contains concluding comments.

II. ESTIMATED MODELS: AN INTRODUCTION

Methodology

This section introduces the case for comparing the forecasting abilities of four conventional and four MSI measures in terms of a 'real-time' forecasting experiment. Because a real-time forecasting exercise is simulated the results presented below should be substantially more reliable than previous comparisons of forecasts using data sets similar to those employed here. Specifically, what has typically been done before is to estimate a set of models, one corresponding to each monetary measure, over a base period. Then, each model is simulated, using the parameters estimated over the base period, over some particular forecast period. For example, Lindsey and Spindt (1986) report that St Louis-type equations were estimated from 1971:1–1979:4 and simulated thereafter. While results from such a procedure may be suggestive, they are suspect for a number of reasons.

First, this procedure suggests that there is no apparent 'learning' about the model over time. Even in the absence of a time-varying parameter estimator, updating parameter estimates as time passes is essentially improving the efficiency with which the presumably fixed parameter is estimated. It seems quite doubtful whether, for policy purposes, a model estimated through the end of 1979 would still be used to make forecasts nearly a decade later. Second, comparison of the forecasting abilities of the various monetary measures based on one realization of forecasts is suboptimal. For instance, suppose that the policymaker is interested in using one-year-ahead forecasts in setting policy. In the procedure using parameters estimated in one time period and then forecasting over several years into the future, there is only one realization of a one-year-ahead forecast error. Third, there is no guarantee that the comparisons across monetary aggregates are based upon the 'best' forecasting equation for each.

In the forecasting exercises reported below, an attempt is made to overcome each of these problems. First, our procedure is to estimate the forecasting equations over a base period (1971:1–1979:3) and make a one-year-ahead forecast. Then, the data set is updated and the equations re-estimated before another forecast is made, and so on until the end of the data set is reached. Thus, we simulate a real-time forecaster in that all available information at the time a forecast is made is utilized.[4] Note that the process of updating and re-estimating forecasting models is a standard procedure for both reduced-form and structural models. This procedure implies, of course, that parameter estimates may change over time, so that

[4]Strictly speaking, this is only approximately true. As stated in footnote 3, a recent update of each series is used. An implication of this is that, with seasonally adjusted data, the data employed here have more information contained in them than would actually be available to the forecaster since the seasonal factors are typically two-sided moving averages around the date in question.

model forecast errors through time are, in part, the result of estimated parameters changing through time. In fact, it is possible that a model using one monetary aggregate with rapidly changing parameters through time could produce forecast error statistics for income which are marginally lower than those associated with a model using another monetary aggregate with essentially constant parameters over time. If the former model is statistically unstable, then its use for policy evaluation is suspect; its superior forecast performance could be due solely to chance. However, for the models presented below, splitting the sample period approximately in half (i.e. splitting in mid-1977) and performing Chow tests suggests that instability of the income equations is generally not a problem. Second, by updating our parameter estimates period by period, a series of one-year-ahead forecasts is obtained. As noted earlier since the base period estimation is over 1971:1–1984:3 and since the data end in 1984:3, there are 31 such forecasts. Third, as discussed briefly in the next section, a concerted effort has been made to obtain the 'best' forecasting results for each monetary measure.

Estimation

As noted earlier, the relationships among the variables under consideration are investigated in a reduced-form framework. There now exists a variety of ways to proceed with such a specification. In what follows, we experiment with two variants of the BVAR technique[5] and, for purposes of comparison, with an unconstrained vector autoregression (UVAR). Specifically, in the context of the BVARs, various types of 'prior' information are tried and a search is made for the prior which, when combined with the data, provides the 'best' forecasts. For the UVAR, an estimate is merely made of a set of equations without any constraints on the shape of the lag distribution. For the models with priors, a four-quarter lag was used.[6] For the UVAR models, experimentation with up to eight lags was undertaken. However, the four-lag results for the UVAR outperformed lag lengths of six or eight, and are not reported here but are available on request.

As with any modelling technique, it is assumed that the macro-economy may reasonably be described as being stable over the period of investigation. It is further presumed that policymakers have adhered to a consistent policy rule over the entire time period.[7] Further, it should be noted that all the out-of-sample forecasts are unconditional, so that no policy interventions have been imposed on the analysis.

The two variants of the Bayesian estimator augment the information from the data about the parameters of an equation with a set of prior information. Assume prior information of the form $r = R\beta + v$, where each element of r is the best guess about the mean of the corresponding element of $R\beta$, where R is a known matrix that relates the elements of β to r,

[5]The BVAR model can be viewed as an application of the Theil and Goldberger (1961) 'mixed estimator'.

[6]No experimentation with other lag lengths was done due to the expense of specifying the BVAR models. It is unlikely that adding additional lags would improve the forecasting performance, however, since the priors used here heavily discount information from earlier periods.

[7]Thus, while the Lucas critique is potentially applicable, we appeal to Sims' (1982) intuition and evidence about the behaviour of the economy: 'The US postwar data contain enough information to give a useful characterization of the conditional distribution of the future of major macroeconomic aggregates given the past. Although there is evidence that this structure changes over time, there is also evidence that it does not change suddenly, so that a model fit to the whole postwar period as if its parameters were fixed over that whole period is not badly biased because of parameter changes'.

and where the vector $\mathbf{v}$ represents uncertainty about this guess. It is assumed that there are no priors on cross-equation covariances, so that $E(vv') = V$.

Theil (1971) showed that this prior information can be combined with the data in the following way:

$$\begin{bmatrix} y \\ r \end{bmatrix} = \begin{bmatrix} X \\ R \end{bmatrix} \beta + \begin{bmatrix} \varepsilon \\ v \end{bmatrix}, E\begin{bmatrix} \varepsilon \\ v \end{bmatrix} = 0;\ VCV\begin{bmatrix} \varepsilon \\ v \end{bmatrix} = \begin{bmatrix} \sigma^2 \mathrm{I} & 0 \\ 0 & V \end{bmatrix} \quad (1)$$

The resulting estimator of β is $\beta = (X'X + \sigma^2 R'V^{-1}R)^{-1}(X'y + \sigma^2 R'V^{-1}r)$.

Note the following intuition about the estimator. First, in the absence of any prior information (so that $r=0$ and $R=0$), the estimator is just the ordinary least squares estimator. Second, as certainty about the prior restrictions is approached, the diagonal elements of V approach zero. Then the diagonal elements of V^{-1} are 'large', so the prior restrictions dominate the data in the computation of the parameter vector. If neither of these extremes holds, then the estimated parameter vector is a mixture of the data and the prior.

For a given r and R, the best forecasting model is found by iterating on how tightly the prior is imposed. Specifically, it is assumed that the elements of the vector $\mathbf{v}$ are set in the following way. Let the prior standard deviation on coefficient j in equation i at lag k be

$$\mathbf{v}_{i,j,k} = \text{gamma} \cdot f_{ij} \cdot (k^{-d}) \cdot (s_i/s_j)$$

gamma is the overall tightness parameter and represents the prior standard deviation on the own first lag; f_{ij} indicates the importance placed on variable j in equation i; if it is believed that variable j is relatively unimportant, then setting this parameter close to zero implies that this belief is relatively firmly held.[8] The decay parameter, d, allows us to restrict how lagged values of the independent variables enter the equation. The larger the setting for d, the tighter the variance of that particular lag coefficient around the postulated mean. The (s_i/s_j) factor scales for differences in the units of the variables.[9]

The mixed estimator used here works in the usual way. First, a base period regression was estimated, assuming a particular set of values for r, R, gamma, f_{ij}, and d. Using these parameters, a four-quarter-ahead forecast for income was made. After saving the forecast error, the sample period was updated with one data point, the equation was re-estimated, and another forecast was made. By continuing to add data one point at a time, a time series of forecast errors is generated that allows for computation of forecast error statistics like the root-mean-squared error (RMSE).

For a given set of priors about the means of the coefficients in the forecasting equations, i.e. for given r and R, the forecasting equations for each monetary aggregate were found by searching over a grid of values for gamma, f_{ij}, and d.[10] In the present analysis, the last period to be forecast is 1984:3, so that there are 31 four-quarter-ahead forecasts for each point in the grid.

[8]Note that $f_{ii} = 1$.

[9]The s_i are the estimated standard errors of univariate autoregressions.

[10]Typically a grid was searched for $0.001 < \text{gamma} < 0.20$, with f_{ij} and d at the relatively loose values of 0.50. Then, given the RMSE-minimizing value for gamma, we searched over $0.05 < f_{ij} < 0.60$. Finally, given gamma and f_{ij}, we searched over $0.0 < d < 6.0$ to find the decay parameter. Wider grid searches were undertaken occasionally when the RMSE-minimizing parameter value was at the boundary of the above regions. There is no necessary reason why the initial value selected for gamma will continue to be optimal after the selection of f_{ij} and d. However, a limited amount of experimentation did not reveal sensitivity of the first round parameter specifications to subsequent grid search.

The first BVAR model is based on the random walk prior, in which the coefficient on the own first lag in an autoregression is unity and all other coefficients are zero, as proposed by Litterman (1979, 1980).[11] As noted earlier, evidence is presented by Litterman (1986) that the random walk prior is competitive with alternative techniques in producing low-RMSE forecasts. However, the random walk prior is a relatively restrictive lag specification; it implies a discontinuous set of lag coefficients. Thus, more flexibility than is inherent in the random walk prior may capture more accurately the dynamics inherent in the data. To achieve this flexibility, we estimate a second BVAR model in which the priors are based on the Pascal lag distribution.[12] We note that the Pascal distribution is sufficiently flexible that it could be used to approximate the lag distribution associated with the polynomial distributed lag specification of the St Louis equation.

III. EMPIRICAL RESULTS[13]

Reported in this section are the results of our forecasting exericses with the various money measures, with the alternative income measures, and with alternative functional forms. Models were estimated in both log-levels and differences of log-levels. The forecasts of log-levels were then transformed into levels in order to produce the error statistics reported below;[14] error statistics associated with the growth rate formulation are in terms of actual growth rates relative to predicted growth rates. Forecasting ability is investigated in terms of both levels and growth rates in light of recent interest in nominal GNP targeting.[15] A standard explanation of the usefulness of nominal GNP targets is that GNP can be viewed either as the product of money and velocity or as the product of price and output. For example, a simple policy rule is that a change in velocity is offset by a change in the money

[11]Thus, an estimation using the mixed estimator, the vector **r** has all zero elements except for an element of unity corresponding to the coefficient for the own first lag.

[12]For the simplest case of a bivariate model, the Pascal distributed lag model may be written as

$$y_t = b \sum_{i=0}^{\infty} w_i X_{t-1-i} + u_t$$

where $w_i = (1-\delta)^n \binom{n+i-1}{i} \delta^i$, $0<\delta<1$, and n is an integer. Choice of alternative values of n and δ allows for a wide variety of lag distributions. In fact, for $n=1$ and δ arbitrarily close to zero, the Pascal distribution approximates the random walk. More generally, it allows for smoother lag coefficients. For the results reported here, the values $n=2$ and $\delta=0.1$ were used; the results were not sensitive to alternative values for these parameters for $n=(2, 3)$ and $\delta=(0.1, 0.2, 0.3)$.

[13]Data used in the results reported in the text were current when this project was initiated in autumn 1986. In May 1987, all data were updated to take account of data revisions and each model was re-estimated using the optimal priors found with the earlier data set. The results were virtually unchanged.

[14]The k-step-ahead levels forecast, X, of logs of the data, denoted by x, is computed as:

$$X(k) = \exp[x(k) + 0.5 \times \sigma^2(k)]$$

where $x(k)$ is the k-step-ahead forecast of the log of the data and $\sigma^2(k)$ the square of the standard error of the k-step-ahead forecast.

[15]It is noted that it is unclear whether 'GNP targeting' in the literature refers to levels or growth rates. Tobin (1983) interprets such proposals as intending growth rates while Hall (1983) interprets them as intending levels.

stock so as to achieve the GNP target. Alternatively, nominal GNP targets presumably help avoid sustained inflation; a rise in prices, given nominal GNP, leads to a decline in real output which, in turn, eases prices pressures.

The results of the forecasting exercises are summarized in Table 1, which shows the one-year-ahead root-mean-squared forecast errors (RMSEs) for the log-levels and growth rate models. Precise definition of the RMSEs for the log-level and growth rate income measures are given in Table 1. Average percentage forecasting errors (which are not reported in order to conserve space) gave no indication of bias in the forecasts. Almost invariably, it was observed that the forecasts of income are best with the imposition of the random walk prior. That is, the RMSEs are typically lower for the random walk prior BVAR than for the Pascal prior BVARs and the UVARs. Income forecasts with the Pascal prior are, in general, though not always, better than those produced with UVARs. Although these results are not reported here in order to conserve space, it was noted that in terms of the ability to forecast the overall system, the Pascal prior, in some instances, produced at least as good forecasts as the random walk prior.[16]

Table 1. *Root-mean-squared errors, four-quarter ahead forecasts*[a]

	Log-level			Growth rate		
Variable	Random walk prior	Pascal prior	UVAR	Random walk prior	Pascal prior	UVAR
A. *Nominal GNP models*						
M1	3.18	4.16	5.25	4.93	5.65	6.13
MSI1	3.18	4.36	4.96	4.88	4.89	5.45
M2	3.06	3.51	3.43	4.96	5.58	5.59
MSI2	3.18	3.57	4.09	4.82	4.96	5.63
M3	3.18	3.52	3.49	4.77	4.85	4.77
MSI3	3.18	3.41	4.06	4.95	4.93	5.76
L	3.18	3.87	3.87	4.98	5.33	6.01
MSI4	3.18	3.58	4.58	4.95	5.17	5.68
B. *Personal Income Models*						
M1	2.59	3.71	3.58	3.86	4.16	5.21
MSI1	2.59	3.40	3.38	3.86	3.95	4.31
M2	2.49	2.58	2.61	3.86	3.79	4.40
MSI2	2.25	2.80	3.35	3.86	3.85	5.13
M3	2.59	2.74	2.69	3.86	3.81	3.81
MSI3	2.28	2.93	3.41	3.86	3.81	5.18
L	2.60	3.04	3.11	3.86	3.81	3.99
MSI4	2.28	3.25	3.92	3.86	4.24	5.47

[a]The root-mean-squared error for the log level models is defined as:

$$[\Sigma((F_t - A_t)/A_t)^2)/T]^{1/2}$$

where $F(A)$ is the forecast (actual value) of the level of nominal income.
The root-mean-squared error for the growth rate models is defined as:

$$[(\Sigma(F_t - A_t)^2)/T]^{1/2}$$

where $F(A)$ is the forecast (actual) of the growth rate of nominal income.

[16]The ability to forecast the overall system was measured, following Litterman (1979, 1980), by the log of the determinant of the variance–covariance matrix of the forecast errors.

For each prior and for each of the four income measures, the monetary aggregate that produces the smallest RMSEs for nominal income can be determined. There is a total of twelve columns in Table 1. In one case (the random walk prior for the growth rate of personal income models), all aggregates generate the same RMSE. This occurs because in this case the initial prior of a coefficient of one on the own first lag and coefficients of zero on the other variables in the income equation was not modified by the data. In nine of the eleven remaining cases, a conventional aggregate produces the lowest RMSE. In only two cases does an MSI generate the lowest RMSE. It was further noted that the narrow aggregates (M1 and MSI1) never generate the lowest RMSE and the same is true for the broadest aggregates (L and MSI4). Thus, in all nine cases won by the conventional aggregates, either M2 or M3 generates the lowest RMSE, and M2 generates the lowest RMSE in five of these nine cases. Only for the growth rate of GNP model does the same aggregate (M3) produce the lowest RMSEs for both BVARs and the UVAR. Similar results are found if the quality of the forecasts is measured by the average absolute percentage forecasting error rather than the RMSE.

As a check on the robustness of these results, log-level monthly models were also estimated for the two income measures and the monetary aggregates.[17] Monthly growth rate models were not estimated due to the time and expense involved so there are six cases for the monthly data. The results (not reported here) are consistent with the quarterly results. In five of six cases, the conventional aggregates generated the lowest RMSEs. Again the narrowest aggregates (M1 and MSI1) and the broadest aggregates (L and MSI4) did not generate any of the lowest RMSEs.

IV. CONCLUSIONS

Taken together, the results presented in this paper suggest that the conventional aggregates dominate the monetary services indices in terms of out-of-sample forecasting of nominal income. These conclusions appear to hold regardless of the nominal income variable employed, regardless of data frequency, and regardless of functional form used in the estimation. Thus, the monetary services indices do not appear, on these grounds, to be better candidates than the conventional aggregates for intermediate targets of monetary policy or as indicators of monetary policy.

These results are, however, subject to a variety of caveats, most of which relate to the heavy computational burden of the exercises undertaken. First, no attempt has been made to evaluate the usefulness of either deterministic or stochastic trends in the equations. Second, experimentation with only a limited array of Pascal priors has been undertaken. Third, the forecasting results have only been examined using the most recently revised data, not the data which would have been available to a 'real-time' forecaster. Nonetheless, given these constraints, care has been taken to specify the 'optimal' models for each monetary aggregate, and the results suggest a preference for the conventional over the monetary services indices.

[17]Since nominal GNP and cyclically adjusted federal expenditures are not available monthly, these series have had to be interpolated. A more detailed description of the interpolation procedure is available on request.

REFERENCES

Batten, D. S. and Thornton, D. L. (1985) Weighted monetary aggregates as intermediate targets, Working Paper 85-010, Federal Reserve Bank of St Louis.

Belongia, M. T. and Chalfant, J. A. (1989) The changing empirical definition of money: some estimates from a model of the demand for money substitutes, *Journal of Political Economy*, **97**, 387–97.

Cooley, T. F. and LeRoy, S. F. (1985) A theoretical macroeconomics: a critique, *Journal of Monetary Economics*, **16**, 283–308.

Farr, H. T. and Johnson, D. (1985) Revisions in the monetary services (Divisia) indexes of monetary aggregates, Special Studies Paper 189, Board of Governors of the Federal Reserve System.

Fisher, I. (1982) *The Making of Index Numbers: A Study of Their Varieties, Tests, and Reliability*, Houghton Mifflin, Boston.

Friedman, M. and Schwartz, A. J. (1970) *Monetary Statistics of the United States: Estimates, Sources and Methods*. Columbia University Press, New York.

Hall, R. E. (1983) Macroeconomic policy under structural change, in *Industrial Change and Public Policy*, Federal Reserve Bank of Kansas City.

Kennickell, A. B. and DeMarco, G. B. (1986) Availability of data, sensitivity of calculation and possible improvements in data collection for the MSI, MT and MQ indexes, Special Studies Paper 197, Board of Governors of the Federal Reserve System.

Lindsey, D. E. and Spindt, P. (1986) An evaluation of monetary indexes, Special Studies Paper 195, Board of Governors of the Federal Reserve System.

Litterman, R. B. (1979) Techniques of forecasting using vector autoregressions, Working Paper 115, Federal Reserve Bank of Minneapolis.

Litterman, R. B. (1980) Specifying vector autoregressions for macroeconomic forecasting, Staff Report 92, Federal Reserve Bank of Minneapolis.

Litterman, R. B. (1986) Forecasting with Bayesian vector autoregressions – Five Years of Experience, *Journal of Business and Economic Statistics*, **4**, 25–38.

Sims, C. A. (1982) Policy analysis with econometric models, *Brookings Papers on Economic Activity*, **1**, 107–52.

Spindt, P. A. (1985) Money is what money does: monetary aggregation and the equation of exchange, *Journal of Political Economy*, **93**, 175–204.

Theil, H. (1971) *Principles of Econometrics* John Wiley, New York.

Theil, H. and Goldberger, A. S. (1961) On pure and mixed statistical estimation in economics, *International Economic Review*, **2**, 65–78.

Tobin, J. (1983) Monetary policy: rules, targets, and shocks, *Journal of Money, Credit and Banking*, **15**, 506–18.

Zellner, A. (1985) Bayesian econometrics, *Econometrica*, **53**, 253–69.