Forecasting M1 with an Autoregression Model: Some Preliminary Results

James N. Duprey

and

Robert B. Litterman

April 1979

Working Paper #: 128

Pacs File #: 2070

NOT FOR DISTRIBUTION

Forecasting M1 with a Vector Autoregressive Model: Some Preliminary Results

1. Introduction

Forecasting economic variables today usually involves a complicated, and often expensive, combination of output from large simultaneous equation econometric models, projections of univariate time-series approaches, and significant allowance for judgmental adjustment.

Sims [1977] has recently suggested as an alternative forecasting strategy the use of vector autoregressions. In these specifications, each element of a vector of economic variables of interest is projected on its own lagged values and the lagged values of every other variable in the system. The resulting models can lead to improved policy through improvements in forecasting accuracy. They are not, however, tools for finding either an optimal rule or determining the impact of a given change in a policy instrument. Policy variables such as the funds rate or M1 may appear, but not as exogenous variables that can be manipulated at the users' convenience. Rather, they appear as endogenous variables whose past history helps to determine the forecasting characteristics of the models, and whose projected time paths can be interpreted as the policymakers' most likely course of action given no change in the policy rule.

The vector autoregressive specification is inexpensive and quite general. It is capable of modeling arbitrarily well any covariance stationary stochastic process. Indeed, the main weakness of this specification, and the reason it has not often been used for forecasting, is, in a sense, its over-generality. The number of free parameters in a system increases quadratically with the number of variables, and for even moderately-sized systems the model becomes highly overparameterized. Estimation of such models leads to a very good fit of the data, but also to ex-post forecasts with large mean square errors.

The problem of overgenerality can be solved by imposing restrictions. In fact, it is possible to view several types of macroeconomic models commonly used for forecasting as special cases of the vector autoregression model derived by applying particular classes of restrictions. The reduced forms of traditional simultaneous equation econometric models are essentially very large vector autoregressions specified with huge numbers of exclusionary restrictions implied by economic theory and the categorization of variables into exogenous and endogenous. The equilibrium solutions of rational expectations models are another special case. Here the assumption of optimizing behavior of agents in the economy generally leads to a complicated set of cross-equation restrictions.

In contrast to these approaches, where the restrictions are often derived from economic theory and always applied with certainty, Litterman [1979] has developed a set of techniques, based on the use of Bayesian priors, to apply instrumental restrictions[1/] in the form of probability distributions. They are applied not for the purpose of identifying a specified economic structure, but expressly to minimize the mean square prediction error of the model. In this last sense, they are similar to the method of ridge regression.

In work reported in Litterman [1979], these techniques were employed to develop and test several quarterly, macroeconomic, vector autoregressive models. The post sample forecasting performances dominated those of the same models in unrestricted form. Further, they dominated the performance of models composed of univariate autoregressive equations.

The work reported in this paper was undertaken to see if the same kind of success could be achieved with a monthly model, in particular, with a monthly

---

[1/] By instrumental, we mean restrictions that are not derived from a particular economic model or theory.

"money market" model in which M1 and, possibly, other aggregates would appear. We were motivated by the possibility of providing an inexpensive yet good (possibly better) forecasting tool in an area of high policy significance. We were also prompted by the availability of several recent studies of the M1 forecasting record of the Board staff, studies that could be used, we believed, to judge the success of our own efforts.

In Section 2 we briefly describe our method of using priors with vector autoregressions for forecasting. Section 3 presents the particular models used to forecast M1 and their forecasting performance. Section 4 suggests several reasons for caution in comparing these results with those of real-time forecasters.

2. A Bayesian Approach to Restricting Vector Autoregressions

Autoregressive specifications often lead to multicollinearity problems and large sampling errors in estimation. This is particularly true in vector specifications which leave relatively few degrees of freedom. Several classical procedures in the form of matrix-weighted averages, such as ridge regression and Stein rule estimators, have been devised to overcome this type of problem in contexts not including lagged dependent variables. These procedures are justified on the grounds that they can generate estimators which, though biased, have smaller mean square error than OLS estimates. The same estimators have a Bayesian interpretation, which amounts to specifying the implicit prior distribution for which the particular estimator is the posterior mean.

Estimators of a similar form may be adapted to the vector autoregressive specification. While these estimators require a Bayesian justification, they are motivated, at least in part, by the above mentioned classical results which suggest that the instrumental restrictions incorporated into the priors will decrease the mean square errors of forecasts generated by these models. The experiments we report support that conclusion.

The first prior which we have used treats all equations symmetrically and specifies stochastically independent lag coefficients which, except for the first lag of the dependent variable, each have mean zero and standard errors which decrease with the length of the lag. The lag distributions are, thus, given the prior restriction that they fade away gradually. The coefficient on the first lag of the dependent variable in each equation is given a prior mean of one so that, in the limit, a tight prior, that is one in which the standard errors approach zero, corresponds to a random-walk process for the indeterministic component of each variable in the system. Each equation includes a constant, for which there is no prior.

The symmetric prior also allows for specification of less weight on other vs. own lagged variables in the sense of imposing smaller standard errors around the zero means of the coefficients on other variables.

Two parameters determine the exact form of the prior. Standard errors of lag coefficients decrease in a harmonic manner according to a given value, γ_1. The prior standard error of the k^{th} lag of each variable is $k^{-\gamma_1}$ times the prior standard error of the first lag. The second parameter, γ_2, scales the standard errors of the other variables' lags relative to those of the dependent variable.

The prior standard errors of other variables' lag coefficients are also scaled according to the relative size of their variable's innovations as measured by the sample standard error of their OLS regression. Let the standard error of the first lag of the dependent variable be λ. The prior standard error of the i^{th} coefficient, which is the j^{th} lag on the m^{th} variable, is δ_i, where

$$\delta_i = \begin{cases} j^{-\gamma_1} \text{ if the } m^{th} \text{ is the dependent variable} \\ j^{-\gamma_1}\frac{\sigma}{\sigma m}{}_2 \text{ otherwise,} \end{cases}$$

σ is the standard error of innovations in the dependent variable and σ_m the standard error of innovations in the m^{th} other variable.

Our procedure has been to act as if we were given, in 1972-1, the task of defining a mechanical procedure which would be used for the next five years to make monthly forecasts of the growth of seasonally adjusted M1 over two-month intervals at annual rates. We chose this task for this period so that we would be able to compare our results with judgemental forecasts of M1 growth made by the Federal Reserve Board as compiled by Porter, Farr, and Perea [1978]. In fact, as discussed in Section 4, we have found this a difficult comparison to make.

We have constructed a small (four variable) and a large (eight variable) vector autoregressive model, along with univariate specifications for comparative purposes. The parameters of the models, such as lag lengths, choice of variables, and, to some extent, prior parameters were picked on the basis of forecasting performance on the data up to 1972.

Given the chosen specifications, each model was used to make forecasts of the levels of M1 and the other variables in the systems each month of the projection period, 1972-1 through 1977-11. The growth rates associated with these forecasts were calculated and error statistics generated. The models were reestimated each month of the projection period so that the forecasts were based only on information available at the time the forecasts were made. One qualification needed here is that the data used were the final seasonally adjusted numbers currently available at the time this study was made (February 1979). In fact, because of the two-sided nature of the seasonal adjustment procedure, these numbers incorporate information not available to forecasters operating in real time.

The estimators in a given equation may be written as

$$b = (X'X+k(R'R))^{-1}(X'Y+kR'r),$$

where the prior is given by $R\beta = r+v$ with $v \sim N(0,\lambda^2 I)$ where $k = \hat{\sigma}^2/\lambda^2$, $\hat{\sigma}^2$ being the estimated variance of the regression residuals in the OLS regression without a prior. To put the prior in this form, we set R = diag[$1/\delta_i$] and r = [0...010..0] where the one corresponds to the first lag of the dependent variable. Here Y represents the vector of observations on the dependent variable and X the matrix of observations on all lags of all variables in the system.[2/]

Given the coefficient estimates of the autoregressive representation, projections are made according to Wold's "chain rule of forecasting." In January 1972, for example, the current and lagged values of the variables in the system are used to forecast levels for February. These values are then used to forecast March, and so on.

The one-step forecast of the two-month growth rate of M1 then refers to the quantity $[\hat{M}(t+1)-M(t-1)]600/M(t-1)$ where M(t-1) is the value of M1 for December 1971 and $\hat{M}(t+1)$ is the forecast for February 1972. Similarly, the two step forecast is $[\hat{M}(t+2)-M(t)]600/M(t)$.

As the size of the vector autoregressive system increases, it becomes increasingly implausible that the procedure of treating all variables in the prior symmetrically is optimal. In an N variable system, one must, at least implicitly, specify N-1 parameters in each equation which determine relative weights of other vs. own variables. Treatment of each parameter separately seems arbitrary, while the symmetric approach which reduces the specification to a single parameter seems overly restrictive.

We have used the symmetric approach in the four variable system, but in the eight variable system have used a circle-star type prior, described in Litterman [1979]. Variables in a central star are those which are assumed to have a strong, but equal impact on all other variables in the system. The other

---

[2/] Our estimation technique of treating each equation separately is not fully efficient. In unrestricted vector autoregressions, separate estimation is justified because the right-hand side variables are the same in all equations. This result is no longer true when prior restrictions are added; however, we believe the loss in efficiency in our estimation is small because the covariance matrixes of the residuals for these models are nearly diagonal.

variables are arranged in a circular ordering with related variables placed close together in the ordering. Relative tightness of the prior on coefficients of variables in an equation is then made a function of the relative positions of the variables. Details of the priors for the particular models are given in the following section.

3. Particular Models and Their Performance

The techniques just described were employed to develop two monthly money market models--one a small (and inexpensive) four variable model, and the second an expanded version with eight variables--that could serve as forecasting tools for M1 and other aggregates. In both models, estimation was done with and without the application of priors in order to demonstrate the value of prior restrictions. Univariate autoregressive models were estimated and employed to generate forecasts that could serve as a basis of comparison. The data begin in 1953-1.

The four variables of the small model are M1, personal income, the rate on 4- to-6-month commercial paper, and the consumer price index. Income and the commercial paper rate were chosen because of the frequency with which they serve as key arguments in studies of the demand for nominal money balances. The consumer index appears in level form; at a later stage we plan to substitute the rate of change of the index for its level as a measure of the commodity opportunity cost of holding money.

The parameters of the prior in the four variable system were $\gamma_1=.5$, $\gamma_2=.5$, and $\lambda=.5$. Thus, the coefficient on the first lag of the dependent variable in each equation has a mean of 1 and a standrd error of $\lambda=.5$. The coefficient on the second lag has a mean of zero and a standard error of $\lambda k^{-\gamma_1} = .5(2)^{-.5} = .3536$. The coefficient on the first lag of the m^{th} other variable is $\lambda\gamma_2(\frac{\sigma}{\sigma_m}) = .25(\frac{\sigma}{\sigma_m})$. Table 1 shows the prior on the M1 equation in the four

Table 1
M1 Equation in the 4 Variable Model

Coefficient	(Lag)	Prior		Estimation	
		Mean	Standard Error	OLS	Restricted
Constant		--	--	4.84423	5.08283
M1	(1)	1.	.5000	1.05925	1.03916
M1	(2)	0.	.3536	-.14065	-.11526
M1	(3)	0.	.2887	.31659	.27803
M1	(4)	0.	.2500	-.45494	-.37066
M1	(5)	0.	.2236	.30059	.22980
M1	(6)	0.	.2041	-.07045	-.05063
Income	(1)	0.	.0474	.02911	.02426
Income	(2)	0.	.0336	-.02870	-.02094
Income	(3)	0.	.0274	-.00319	-.00674
Income	(4)	0.	.0237	-.01413	-.00720
Income	(5)	0.	.0212	.04215	-.02703
Income	(6)	0.	.0194	-.01722	-.00801
Paper Rate	(1)	0.	.6229	-.48636	-.49440
Paper Rate	(2)	0.	.4405	.03236	.09655
Paper Rate	(3)	0.	.3597	.44900	.24374
Paper Rate	(4)	0.	.3114	-.42979	-.19723
Paper Rate	(5)	0.	.2785	.07504	-.01970
Paper Rate	(6)	0.	.2542	.13497	.13338
Prices	(1)	0.	.6434	-.23289	-.15940
Prices	(2)	0.	.4550	.46715	.32752
Prices	(3)	0.	.3715	-.10268	-.04816
Prices	(4)	0.	.3217	-.07283	-.10614
Prices	(5)	0.	.2877	-.34708	-.20324
Prices	(6)	0.	.2626	.19081	.08791

variable system along with the unrestricted (OLS) and the restricted (posterior mean) estimates of the coefficients.

The eight variable model includes the four listed above and adds to them other time and savings deposits (which can be used with M1 to find M2), the currency component of M1, nonborrowed reserves, and the ratio of the commercial paper rate to the Q ceiling rate on savings deposits at commercial banks. Nonborrowed reserves was selected to represent supply-side reserve factors, and the ratio of market to ceiling rates included to capture the effects of disintermediation.

The prior on the eight variable system differs in two important ways from that on the smaller system. First, the prior is considerably tighter, that is, has smaller standard errors. Also, it includes structure in the sense that standard errors are adjusted on other variables' coefficients not only according to the size of their innovations, but also with respect to their position in the circle-star prior.

The parameters of the prior are $\gamma_1=.5$, $\gamma_2=.05$, and $\lambda=.2$. The central variables are the four variables in the smaller system. Table 2 shows the prior, unrestricted and restricted estimates of the coefficients in the M1 equation in this system. The tighter prior in this specification is required by the larger number of parameters in each equation relative to the number of degrees of freedom. The value of .05 for γ_2 was chosen on the basis of a comparison of forecasting performance of several models with different values of γ_2 on the period 1968-71.

Forecast errors for all models were computed by taking the one- and two-step forecasts of two-month growth rates and subtracting the actual growth rates as determined from the data file. These errors were summarized in terms of root mean square errors (RMSE) and Theil U statistics (TU). It should be noted that the TU's were computed according to the formula

Table 2
M1 Equation in the 8 Variable Model

Coefficient	(Lag)	Prior: Mean	Prior: Standard Error	Estimates: OLS	Estimates: Restricted
Constant	--	--	--	4.18932	1.54905
M1	(1)	1.	.2000	1.11031	1.11072
M1	(2)	0.	.1414	-.17583	-.14242
M1	(3)	0.	.1155	.35059	.16628
M1	(4)	0.	.1000	-.39750	-.21223
M1	(5)	0.	.0894	.21068	.11437
M1	(6)	0.	.0816	-.07115	-.02125
Income	(1)	0.	.0018	.02718	.00042
Income	(2)	0.	.0013	-.02245	-.00010
Income	(3)	0.	.0010	-.00521	.00002
Income	(4)	0.	.0009	-.01252	.00016
Income	(5)	0.	.0008	.03626	.00035
Income	(6)	0.	.0007	-.02539	.00017
Paper Rate	(1)	0.	.0248	-.75560	-.06726
Paper Rate	(2)	0.	.0175	.41837	-.02117
Paper Rate	(3)	0.	.0143	.40391	-.00755
Paper Rate	(4)	0.	.0124	-.22772	-.00402
Paper Rate	(5)	0.	.0111	-.51101	-.00009
Paper Rate	(6)	0.	.0101	.69717	.00291
Prices	(1)	0.	.0259	-.36161	-.02551
Prices	(2)	0.	.0183	.64408	-.00695
Prices	(3)	0.	.0149	-.18161	-.00515
Prices	(4)	0.	.0129	-.03977	-.00397
Prices	(5)	0.	.0116	-43853	-.00235
Prices	(6)	0.	.0106	.32377	.00002
Reserves	(1)	0.	.0055	-.48387	.00083
Reserves	(2)	0.	.0039	.36840	.00069
Reserves	(3)	0.	.0032	.08682	.00057
Reserves	(4)	0.	.0028	.31065	.00043
Reserves	(5)	0.	.0025	-.38671	.00027
Reserves	(6)	0.	.0023	.02339	.00018
Currency	(1)	0.	.0116	-.96217	-.00360
Currency	(2)	0.	.0082	-.04182	-.00130
Currency	(3)	0.	.0067	.78621	-.00040
Currency	(4)	0.	.0058	-.68306	-.00053
Currency	(5)	0.	.0052	1.25497	-.00004
Currency	(6)	0.	.0047	-.43878	-.00005
Other Time	(1)	0.	.0021	.16788	.00253
Other Time	(2)	0.	.0015	-.23366	.00096
Other Time	(3)	0.	.0012	.23715	.00061
Other Time	(4)	0.	.0011	-.32410	.00044
Other Time	(5)	0.	.0010	.06423	.00036
Other Time	(6)	0.	.0009	.11812	.00033
r/Q	(1)	0.	.0139	1.8360	-.00956
r/Q	(2)	0.	.0098	-2.0798	-.00356
r/Q	(3)	0.	.0080	.5267	-.00151
r/Q	(4)	0.	.0070	-1.0630	-.00088
r/Q	(5)	0.	.0063	2.0622	-.00024
r/Q	(6)	0.	.0057	-1.7302	.00014

$$TU = \left[\frac{\sum_{1}^{m} (P_t - A_t)^2}{\sum_{1}^{m} (P_t^{NC} - A_t)^2} \right]^{1/2}$$

where

P = predicted two-month growth rate;

P^{NC} = no change prediction (the two-month growth rate obtained by extrapolating the actual growth rate of the preceding two months);

A = actual two-month growth rate; and

m = number of forecasts.

A TU statistic of zero indicates perfect forecasts; a value of one indicates that the forecaster has done as well as a simple extrapolation procedure; and a value greater than one shows that a simple extrapolation procedure would have done better. For the forecast period 1972-77, both RMSE's and TU's were compiled.

On the basis of the RMSE measure of forecast accuracy the eight variable system with prior restrictions performed best overall. Table 3 shows the one- and two-step RMSE's and TU statistics for each variable in each of the models in which it appears.[3/] Relative to the univariate autoregressions, the eight variable system without a prior generated one-step forecast error RMSE's, which were in all cases worse, and on average were 15.1 percent worse. The addition of the prior, however, reduced the errors to the extent that for seven of the eight variables, the RMSE's were smaller than the univariate specification. On average, there was a reduction of 1.2 percent in the RMSE's.

In general, the four variable systems did not forecast quite as well as the univariate systems. The average increase of RMSE's in the unrestricted system was 4.7 percent, while the average increase with prior restrictions was

---

[3/] The RMSEs and TUs for the commercial paper rate and the ratio of the paper rate to the Q ceiling on passbook savings are based on levels rather than growth rates.

Table 3
Error Statistics for the Period 1972.1 - 1977.11

Root Mean Square Errors

Model	M1	Income	Paper Rate	Prices	Reserves	Currency	Other Time	r/Q
Univariate								
one-step	2.284	3.074	.512	1.749	7.978	1.747	1.521	.099
two-step	3.474	4.199	.928	2.709	12.525	2.318	2.857	.169
4 variable - no prior								
one-step	2.276	3.226	.526	1.944	--	--	--	--
two-step	3.469	4.312	.936	2.946	--	--	--	--
4 variable - prior								
one-step	2.244	3.126	.512	1.882	--	--	--	--
two-step	3.404	4.224	.923	2.890	--	--	--	--
8 variable - no prior								
one-step	2.480	3.508	.527	1.951	9.753	2.104	1.843	.118
two-step	3.960	4.560	.960	2.819	15.436	2.792	3.638	.209
8 variable - prior								
one-step	2.245	3.002	.497	1.737	7.909	1.731	1.531	.097
two-step	3.328	4.074	.886	2.713	12.512	2.283	2.905	.164

Theil U Statistics

Model	M1	Income	Paper Rate	Prices	Reserves	Currency	Other Time	r/Q
Univariate								
one-step	.694	.745	.929	.856	.853	.779	.716	.917
two-step	.753	.761	.996	.945	.905	.794	.865	.945
4 Variable - no prior								
one-step	.692	.782	.956	.951	--	--	--	--
two-step	.752	.781	1.004	1.028	--	--	--	--
4 Variable - prior								
one-step	.682	.757	.929	.921	--	--	--	--
two-step	.738	.765	.990	1.009	--	--	--	--
8 Variable - no prior								
one-step	.754	.850	.957	.954	1.043	.939	.868	1.100
two-step	.858	.826	1.030	.984	1.116	.956	1.109	1.165
8 Variable - prior								
one-step	.683	.727	.903	.850	.846	.772	.721	.906
two-step	.722	.738	.950	.947	.904	.782	.885	.913

1.9 percent. It should be noted, however, that the value of γ_2 in the prior specification was chosen specifically with M1 in mind, and for M1 alone the forecast errors of the four variable systems do show improvement over the univariate projections.

The results we have reported are, of course, conditioned on the particular prior we used. There are undoubtedly other sets of prior restrictions which would generate smaller forecast errors on our projection period than the prior we have chosen. We did not search for these prior restrictions. Such a task might be useful for other purposes, but as a basis of comparison of forecasting performance, it amounts to data mining. What we have attempted to do is to choose our prior without using the data from the projection period and then test its performance on that period. Despite the particular selection of variables, projection period and prior restrictions we have used, we suggest that the results of that test support the following conclusions:

A. There exists a class of relatively inexpensive estimators, generated as posterior means of priors incorporating instrumental types of restrictions, which greatly improve the forecasting performance of vector autoregressive models relative to unrestricted OLS estimators.

B. The forecasting errors of unrestricted vector autoregressions tend to increase as the size of the system grows relative to the number of degrees of freedom. This phenomenon is not true of the restricted models we have used.

Unfortunately, our results do not seem to shed much light on the question of whether the restricted vector autoregressions represent a possible improvement of forecasting over univariate autoregressions or judgmental forecasts. The latter question will be considered in Section 4. As to the former, our four variable system tends to favor the univariate, while the eight variable system favors the vector specification.

We emphasize that these results are strongly conditioned on our particular priors. In retrospect, we suspect that a tighter prior on the smaller system may have led to smaller forecast errors. Our experiments with different priors on the 1968-1971 projection period showed that in moving from the univariate to a vector specification, that is, letting γ_2 increase from 0 toward 1, there was a region of improvement followed by worsening. This phenomenon is also found in a different system described in Litterman [1979]. In the four variable system the improvement region for the earlier projection period seemed quite broad for the measure which we looked at, namely the RMSE on projections of M1. These results led us to the loose prior represented by the choice of γ_2=.5 and λ=.5. The eight variable system, on the other hand, showed a much smaller region of improvement, leading to our choice of γ_2=.05, and λ=.2.

The question of whether there are restricted vector specifications which dominate univariate autoregressions in terms of forecasting performance, seems to us to rest on whether one can generate restrictions which can be expected to consistently lead to a region of improvement. The results of this test suggest to us that such regions of improvement probably do exist. On the other hand, in this system they appear not to have been as large, as consistent, or as easily exploitable as we had hoped.

In addition to our projections on the 1972-1977 period, we have generated forecasts for the months from then to the present. In Figure 1 we present the graphs of the errors generated by our restricted four variable models along with the forecast errors made by the Federal Reserve Board at the time that most closely corresponds to our forecasts, that is during the week for which estimates of the level of M1 for the previous month first became available. The board forecast errors and those of the vector autoregression are very similar through most of 1978. More recently both do very poorly, with the autoregression making especially large errors.

Figure 1

Shown here are the recent forecast errors of the Board and the restricted four-variable vector autoregression (VAR).

Errors in Forecasts of Two-Month Growth Rates of M_1

(Growth from months indicated)

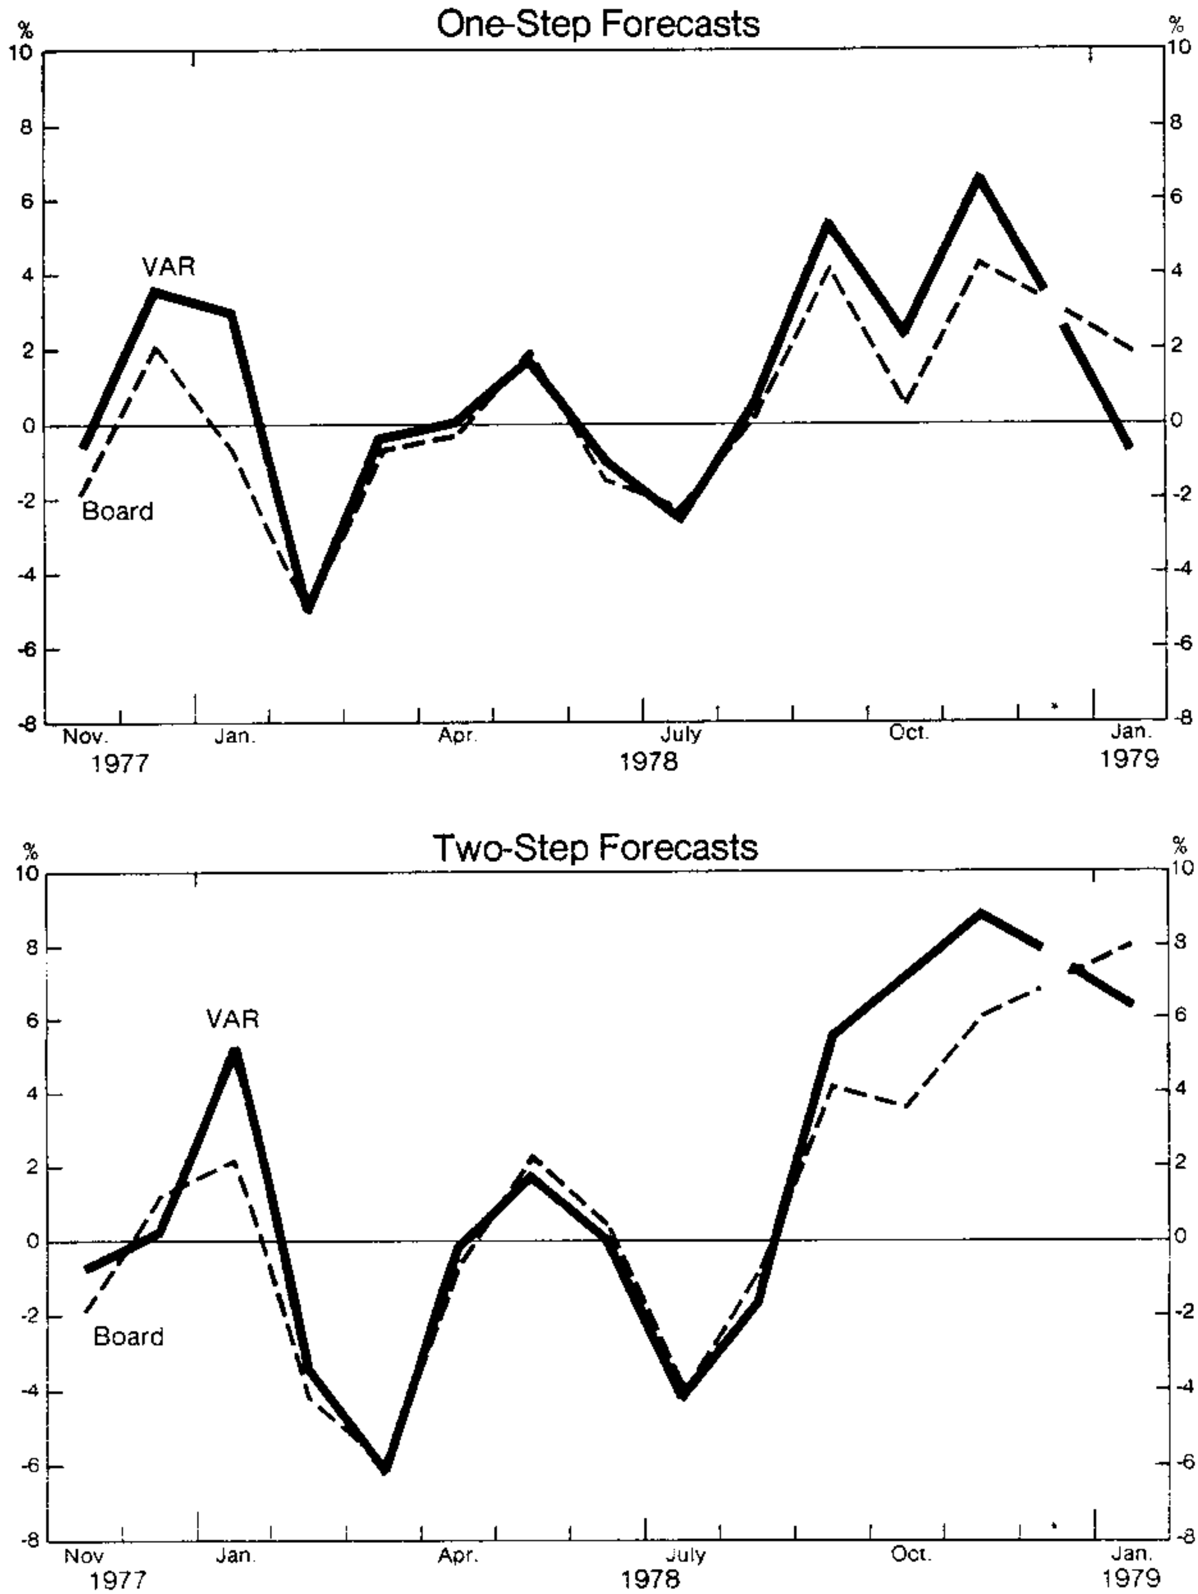

*No forecast

4. Real-Time Forecasting and Seasonal Adjustment[4/]

We had hoped in this paper to present a comparison of the forecasts of growth rates of M1 generated by the mechanical vector autoregressive models with the judgmental forecasts produced by the Board and reported by Porter, Farr, and Perea. There are several problems with such a comparison. First, forecasts made in real time, necessarily involve only past data and therefore, as we discuss below, require the use of a one-sided seasonal adjustment procedure. This is not true of the seasonally adjusted data series we have used. In this section we estimate this difference and suggest a correction to adjust for it. The Board, however, has not computed its forecast errors from the final two-sided seasonally adjusted numbers, but rather from the one-sided numbers as they became available. We have tried to recreate these numbers with the use of the one-sided seasonal adjustment procedure defined below, but we have not been able to regenerate the growth rate series which the Board attempted to forecast. The possible differences which remain include benchmark revisions of the data based on nonmember bank reports, other definitional changes in the M1 series, and differences in the seasonal adjustment procedures. As we show here, the series forecast by the Board shows more variation than the series we have forecast. Thus, the smaller forecast errors which we have generated do not necessarily represent an improvement over those of the Board.

Consider now the problem of seasonal adjustment. The standard procedure for obtaining a "final" seasonally adjusted M1 series requires the application of a two-sided filter to the raw data. Letting the two-sided seasonally adjusted data be $M1_s$, the unadjusted data be $M1_N$, then a two-sided filter is represented by

$$M1_s(t) = F[M1_N(t-k), M1_N(t-k+1), \ldots, M1_N(t), \ldots, M1_N(t+k)].$$

[4/] The results in this section were generated using the Regression Analysis of Time Series, RATS, computer program written by Thomas A. Doan at the Federal Reserve Bank of Minneapolis.

Thus, future, as well as past, values of the unadjusted $M1_N$ series are required for some length, k, depending on the filter.

The results presented in Section 3 are based on data current as of late February 1979; for purposes of this discussion they can be considered final. A forecaster operating in real time, however, faces a more difficult problem than we have posed for ourselves. At the time of the forecast, t, the forecaster does not know the final seasonally adjusted numbers for the last k periods, and he must first estimate $M1_s(t-s)$, s=0, 1, ..., k, based on $M1_N(t-s)$, s=0, 1, ..., and then project future values of $M1_s$. We refer to the estimation of $M1_s(t-s)$, s=0, 1, ..., k, on the basis of $M1_N(t-s)$, s=0, 1, ..., as one-sided seasonal adjustment, and let $M1_s^t(t-s)$, s=0, 1, ..., be the estimates. Notice that for s > k, $M1_s^t(t-s) = M1_s(t-s)$. Several methods have been suggested to accomplish this one-sided seasonal adjustment. We use the procedure suggested by Geweke [1978] because it has the property of minimizing expected subsequent revision in the seasonal factors.

In theory we could face the problem of real-time forecasting by starting with unadjusted data and incorporating a one-sided seasonal adjustment procedure explicitly into the estimation process at each point in time during the projection period. We have not done this because of the large computing expense which would be involved. Not only would there be the cost of seasonal adjustment each period, but more importantly, because the data series themselves change, use of the Kalman filter updating algorithm would no longer be possible and a set of matrix inversions would, thus, have to be performed each period to form the desired projections.[5/]

Instead of using this costly procedure throughout, we have followed it once in a univariate system as an experiment in order to estimate the magnitude

---

[5/] The Kalman filter is a recursive algorithm which, given current coefficient estimates and a set of additional observations, generates the new coefficients for the enlarged data set.

of the difference between using two-sided and one-sided seasonally adjusted data.

The two-sided seasonal adjustment procedure we used consisted of the following. Given the entire $M1_N$ series, the seasonally adjusted series, $M1_S$, is formed using a simplified version of the multiplicative, ratio-to-moving average method used in the Census X-11 program as described in Shiskin, Young, and Musgrave [1967]. The first step is calculation of a centered 23-term moving average of $M1_N$. S-I (seasonal-irregular) ratios are formed by dividing $M1_N$ by the moving average. Seasonal factors are generated as a (3x5) moving average of the S-I ratios individually for each month. The seasonally adjusted series $M1_S$ is $M1_N$ divided by the seasonal factors. This procedure requires values of the unadjusted series 47 steps ahead and previous to each period in order to calculate the seasonally adjusted value.

This mechanical procedure generates a series which closely approximates the published seasonally adjusted M1 series, M1-SA. A comparison of the two during our projection period, along with the unadjusted data, is given in Figure 2. The series labeled "Filtered M1-NSA" is generated using the above method and is based on unadjusted data available through 1978-11 and projections of $M1_N$ beyond that date. Shown are deviations from constant and trend of the logarithms of each of the three series.

The following is our one-sided seasonal adjustment procedure. Given $M1_N(t-s)$, s=0, 1, ..., a forecast of $M1_N(t+s)$, s=1, 2, ..., k is generated. The forecast is made in the manner suggested by Geweke, that is, using data up to time t the regression of $M1_N(t)$ on $M1_N(t-1)$, $M1_N(t-12)$, and $M1_N(t-13)$ is computed. The residuals are then regressed on themselves in a sixth order autoregression. They are projected k steps ahead using the chain rule. These forecasts of residuals can then be plugged into the first regression allowing

Figure 2

This graph compares the seasonally adjusted M_1 series generated by the two-sided seasonal filter defined in the text (Filtered M1-NSA) with the published seasonally adjusted series (M1-SA) and not seasonally adjusted series (M1-NSA). All three are shown as deviations from constant and trend of logarithms of the data.

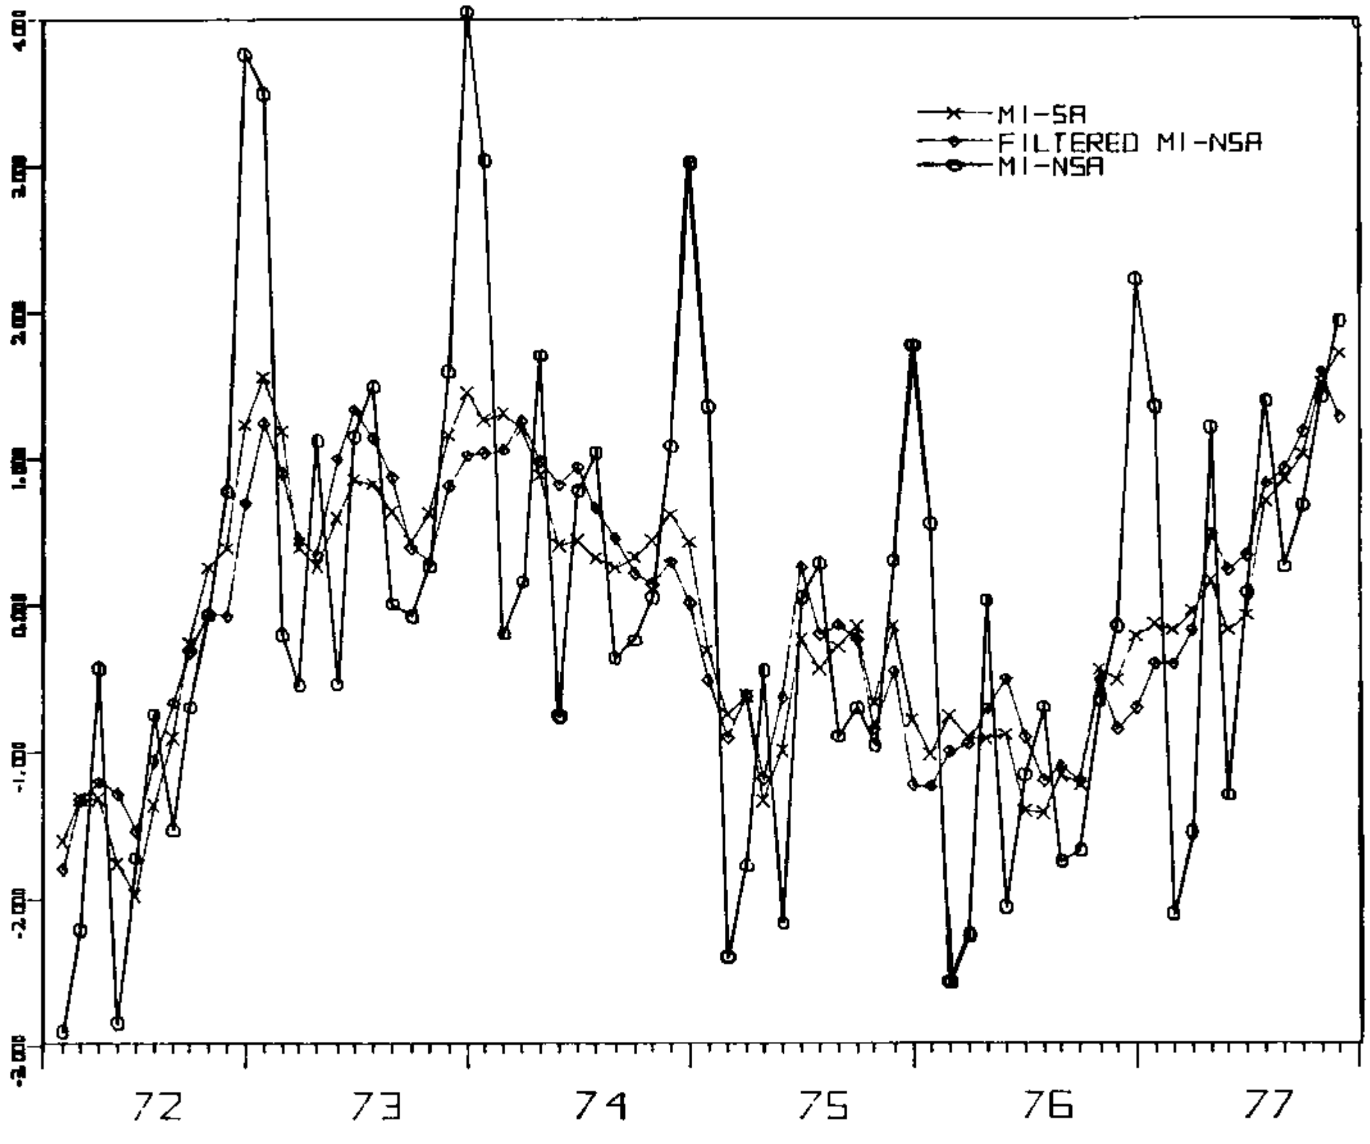

calculation of $M1_N(t+s)$, s=1, 2, ..., k. Using these values for $M1_N$ the above two-sided adjustment procedure is applied to generate $M1^t_s(t-s)$, s=0, 1, ...

The above techniques allow us to generate two series of growth rates of seasonally adjusted M1 as follows:

1-S NSA(t) = $600[M1^t_s(t)-M1^t_s(t-2)]/M1^t_s(t-2)$ that is, the two-month growth rates of seasonally adjusted M1 which become available month by month through the use of the one-sided seasonal adjustment procedure; and

2-S NSA(t) = $600[M1_s(t)-M1_s(t-2)]/M1_s(t-2)$, the two-month growth rates of $M1_s$, the series generated by the two-sided seasonal adjustment procedure applied to the entire not seasonally adjusted M1 series.

In Figure 3 we have plotted these two series along with "SA," the growth rates implied by the final published seasonally adjusted M1 data which we used in Section 3. The series "1-S NSA" closely approximates both "2-S NSA" and "SA."

We now define "real-time forecasting" as the technique of forecasting future values of $M1_s$ based on $M1^t_s$, that is, of reestimating. $M1^t_s(t-s)$, s=0, 1, ..., on the basis of $M1_N(t-s)$, s=0, 1, ..., each period and using those estimates to project $M1_s(t+1)$, $M1_s(t+2)$, ... We have calculated error statistics comparing the forecasts of growth rates of $M1_s$ generated by this method with those generated by projecting $M1_s$ on itself.

Our experiments show that there is only a rather small advantage gained by using the final seasonally adjusted numbers throughout rather than the real-time forecasting method. Recall from Section 3 that the one-step projections by a univariate sixth order autoregression of two-month growth rates of "SA" generate a root mean square error of 2.28 on our projection period. The corresponding error statistic using $M1_s$ is 2.34. The root mean square error of forecasts using the real-time forecasting method is 2.43. Thus, we have demonstrated a procedure of real-time forecasting which, in the univariate case,

Figure 3

A comparison of the growth rates of three seasonally adjusted M_1 series shows that applying the one-sided and two-sided adjustment procedures defined in the text to not seasonally adjusted data generates series with growth rates very close to those of the published seasonally adjusted data.

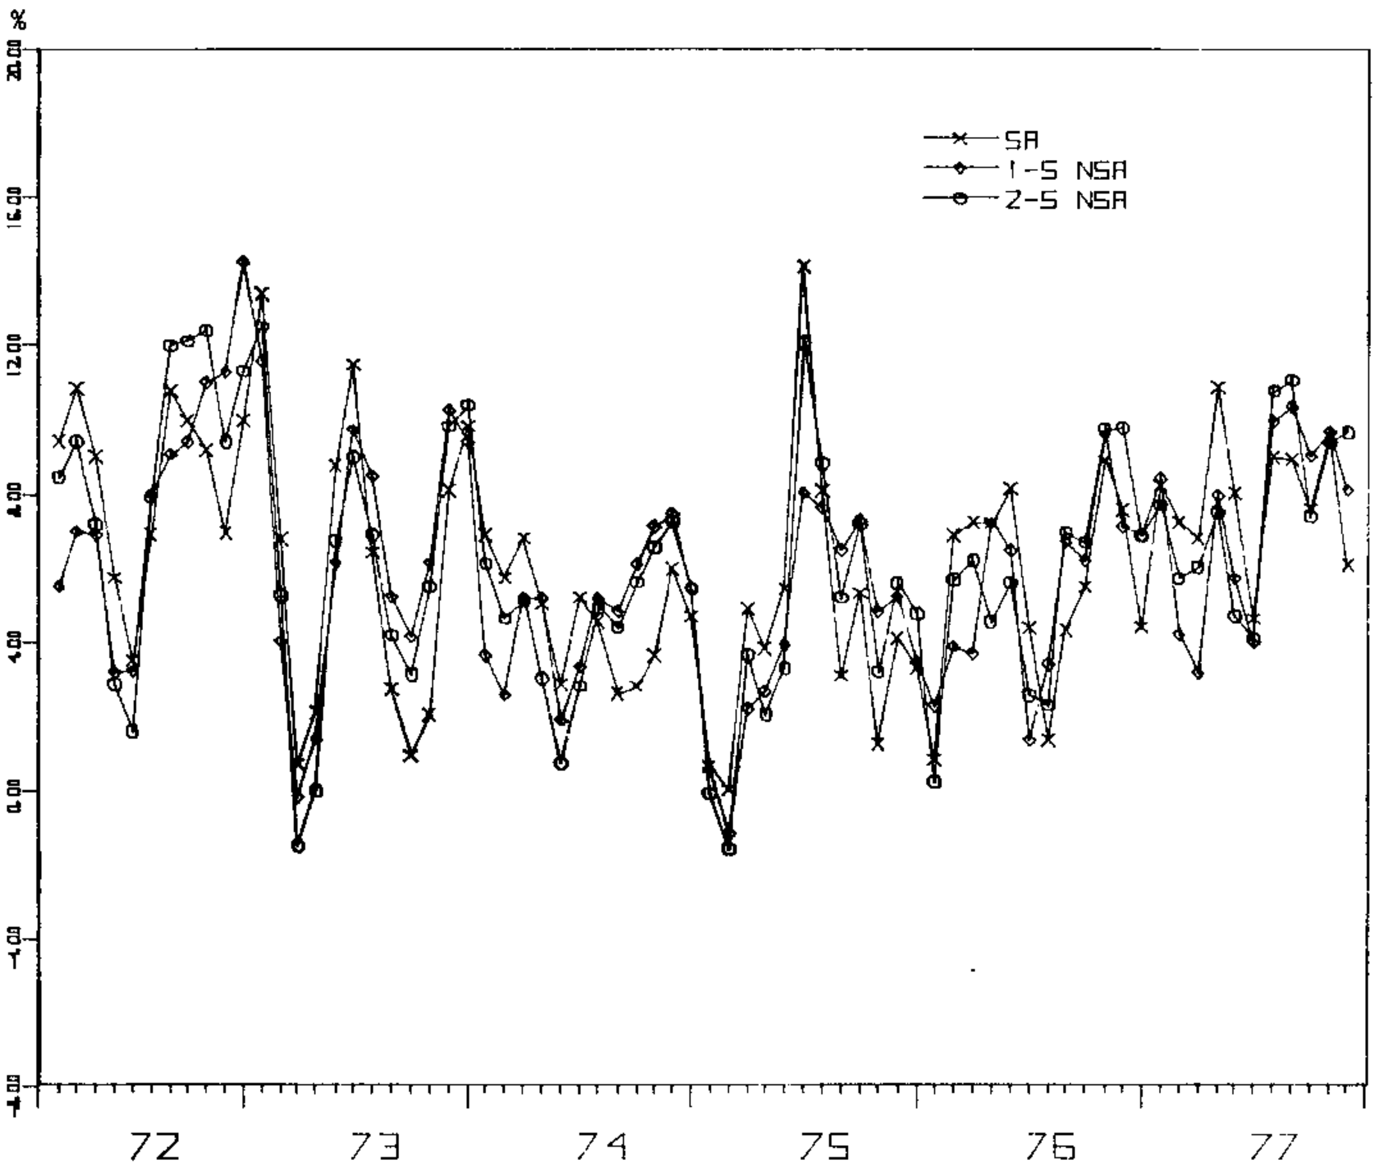

generates errors only slightly larger than our original method of projecting final seasonally adjusted data. These results show that our root mean square error statistics on final seasonally adjusted series should be adjusted upward by approximately .09 (2.43-2.34) in order to account for errors which would have been made in real-time forecasting over this period using our one-sided seasonal adjustment procedure.

On the other hand, the Board does not attempt to forecast growth rates which correspond to our final seasonally adjusted data. Their procedure corresponds to using current one-sided seasonal numbers to project the next period's one-sided seasonal growth rate. Following Porter, Farr, and Perea [1978], we call the realized one-sided seasonally adjusted growth rates, "actuals." We could not directly attempt to forecast this series because we do not have its values prior to 1972-2. Instead, we performed the corresponding experiment, that is, forecast future values of "1-S NSA" based on current $M1_s^t$ with a univariate system, and obtained one-step forecast errors with a RMSE of 2.10. This would indicate a substantial improvement over the Board results, which are summarized below.

However, as shown in Figure 4, our "1-S NSA" is not a particularly close approximation to the "actual" series used by the Board. One difference between the two is attributable to the revisions in the data after the "actual" number is generated. To this extent the "actual" series represents a noisy measurement of the "1-S NSA" series. Some of the difference is also caused by the Board's use of a one-sided seasonal adjustment procedure which does not minimize subsequent revision. In any case, it is clear that the "actual" series has more variation than the series we have used. We have quantified the increased variation by computing the standard errors of the different series about their means, and by computing the standard errors of the residuals in

Figure 4

The series "actual," which forms the basis of the real-time forecast errors of the Board, shows more variation than the growth rates of 1-S NSA, the series being forecast in this study which most closely approximates it.

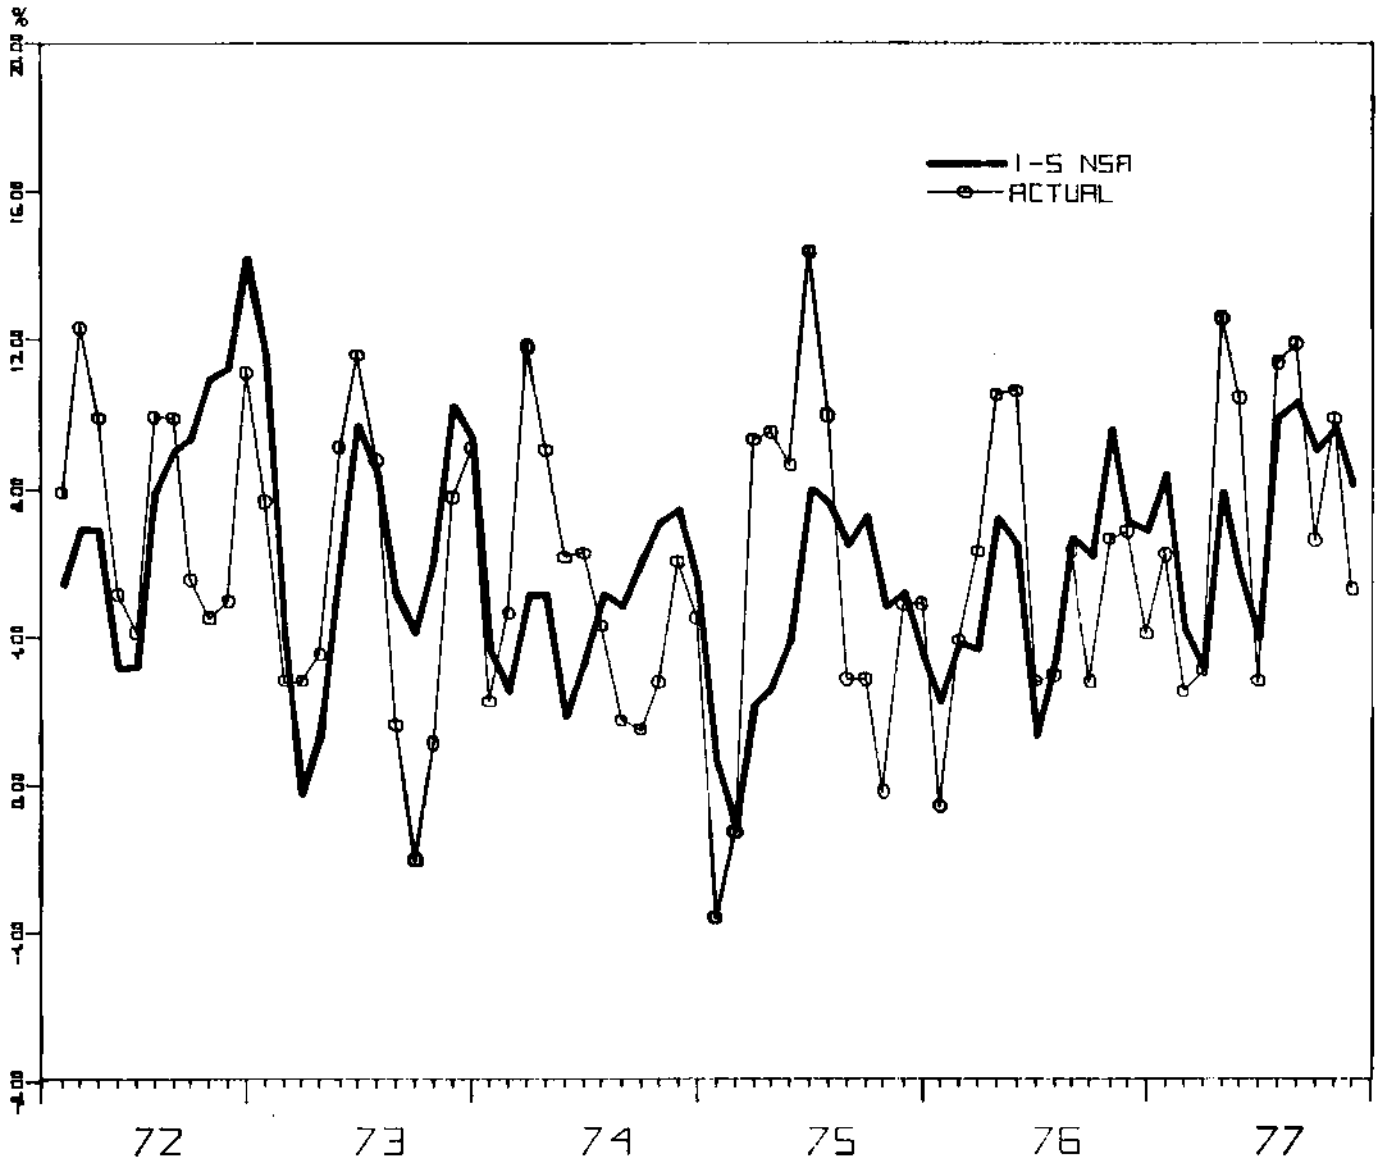

third-order autoregressions with constants for the 72-5 to 77-11 period. The results are shown in Table 4.

Table 4

Variation of Four M1 Growth Series

	Actuals	SA	1-S NSA	2-S NSA
Standard Deviation of Series	3.915	3.182	3.062	3.431
Standard Deviation of Residuals of 3rd Order Auto-regression	3.369	2.668	2.299	2.739

In the study by Porter, Farr, and Perea, Board forecast errors are reported as a function of the number of weeks after the FOMC meeting the forecasts were made. Since the timing of the FOMC meeting varies to some extent from month to month, none of their error statistics can exactly represent forecasts made with information sets which match those implicit in our procedure.

We suggest that the information available two weeks after FOMC is the best approximation to the information used in our one-step forecasts. The RMSE of those forecasts is 2.62. The other error RMSE's, ranging from one week prior to FOMC to three weeks after, were 3.49, 3.18, 2.92, 2.62, and 2.18 respectively.

Relative to these numbers, the RMSE's generated by both the univariate and vector specifications we have tested appear rather small. Unfortunately, the considerations which we raise above force us to conclude that such a comparison must be viewed with extreme caution.

5. Conclusions

This paper presents the results of an experiment in which we have tried to test the forecasting performances of mechanical procedures and compare them with the compiled errors of the judgmental forecasts of the Board. In general, our forecasting results by the mechanical procedures have appeared encouraging relative to judgmental results. However, we have found many difficulties in making such a comparison.

We feel more useful comparisons can be made amoung different mechanical forecasting procedures using the techniques we have described. Our tests on univariate, small and larger vector autoregressive specifications suggest that instrumental restrictions in the form of Bayesian priors can generate significant improvements in forecast performance over unrestricted models.

Bibliography

Geweke, John. (1978) "An Efficient Method For Revising Seasonally Adjusted Time Series," Discussion Paper No. 78-22, Social Science Institute, University of Wisconsin.

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

Porter, Richard D., Helen T. Farr, and Juan Perea. (1978) "An Evaluation of Board Staff Forecasts of M1," Federal Reserve Board.

Shiskin, Julius, Allan H. Young, and John C. Musgrave. (1967) "The X-11 Variant of the Census Method II Seasonal Adjustment Program," U.S. Department of Commerce, Bureau of the Census, Technical Paper No. 15.

Sims, Christopher A. (1977) "Macro-economies and Reality," Discussion Paper No. 77-91, Center for Economic Research, University of Minnesota.