

Improving Forecasts of the Federal Funds Rate in a Policy Model

John C. Robertson and Ellis W. Tallman

Federal Reserve Bank of Atlanta
Working Paper 99-3
March 1999

Abstract: Vector autoregression (VAR) models are widely used for policy analysis. Some authors caution, however, that the forecast errors of the federal funds rate from such a VAR are large compared to those from the federal funds futures market. From these findings, it is argued that the inaccurate federal funds rate forecasts from VARs limit their usefulness as a tool for guiding policy decisions. In this paper, we demonstrate that the poor forecast performance is largely eliminated if a Bayesian estimation technique is used instead of OLS. In particular, using two different data sets we show that the forecasts from the Bayesian VAR dominate the forecasts from OLS VAR models—even after imposing various exact exclusion restrictions on lags and levels of the data.

JEL classification: E44, C53

Key words: forecasting, Bayesian vector autoregression, federal funds rate

The author thank Jerry Dwyer and Tao Zha for helpful comments and suggestions. The views expressed here are the authors' and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the authors' responsibility.

Please address questions regarding content to John C. Robertson or Ellis W. Tallman, Research Department, Federal Reserve Bank of Atlanta, 104 Marietta Street, NW, Atlanta, GA 30303-2713, 404/521-8782 (Robertson), john.c.robertson@atl.frb.org, 404-521-8915 (Tallman), ellis.tallman@atl.frb.org.

To receive notification about new papers or to order copies of printed papers, contact the Public Affairs Department, Federal Reserve Bank of Atlanta, 104 Marietta Street, NW, Atlanta, Georgia 30303-2713, 404/521-8020. The full text of this paper may be downloaded (in PDF format) from the Atlanta Fed's World-Wide Web site at www.frbatlanta.org/publica/work_papers/.

INTRODUCTION

The application of Vector AutoRegression (VAR) models to policy analysis is extensive (see Christiano, Eichenbaum, and Evans 1997 for a recent survey). Yet, there is some concern that VAR models are of limited usefulness as policy tools because they generate inaccurate forecasts of the monetary authority's response to economic conditions. For example, Rudebusch (1998) and Evans and Kuttner (1998) report that VARs generate errors in the forecast of the federal funds rate that are much larger and are not highly correlated with errors from a forecast inferred from the federal funds futures market. Those critical of VARs point to a number of sources for the observed forecast inaccuracy, in particular: 1) the use of a time-invariant linear structure; 2) the limited scope of the information set; and 3) the tendency to specify VARs with long lag structures. We do not dispute that these factors characterize standard VAR practice. However, we suggest that the poor forecasting performance relates more to the choice of estimation technique than to the concerns listed above. Using an alternative estimation technique, we show that a VAR can, in fact, produce federal funds rate forecasts that are as accurate as those from the futures market.

The effect of over-parameterization due to the inclusion of numerous lagged observations is a legitimate concern for VAR forecasts when the coefficients are estimated by Ordinary Least Squares (OLS) (see for example, Fair and Shiller 1990 and Wallis 1989). In this paper, we examine the one and two month-ahead funds rate forecast performance of a VAR using two data sets taken from the literature. For each data set we estimate the coefficients of a VAR by OLS and confirm the poor forecast performance that has been reported previously. We then experiment with various exclusion restrictions on the VAR coefficients prior to estimation by

OLS and find that imposing exact restrictions improves the forecast performance of an unrestricted VAR somewhat, although the summary forecast error measures are significantly larger than from the futures market in all cases. Finally, we estimate the VAR coefficients using the Bayesian method presented in Kadiyala and Karlsson (1997), and Sims and Zha (1998). We use this method because it is explicitly multi-variate, and hence can be applied to identified VAR models. For both data sets the resulting Bayesian VAR (BVAR) forecasts are more accurate than the forecasts from any model estimated by OLS. For one of the data sets we find that the monthly real-time out-of-sample funds rate forecasts of the BVAR are statistically as accurate (in a mean squared error sense) as those implied by the futures market.

The remainder of this paper is organized as follows. The first section describes the data, estimation techniques, and the statistical test used in our comparisons of forecast accuracy. The subsequent section presents the results of our forecast comparisons. The last section gives some conclusions and discusses the relevance of our results.

DATA AND TECHNIQUES

2.1 Data

To examine the impact of estimation technique on VAR forecasts, we employ two sets of monthly data, both comprised of variables that would be of explicit concern to monetary policy makers. The first is a set of six variables used by Waggoner and Zha (1999) and Zha (1998). The data series are the levels of the effective federal funds rate and the unemployment rate, together with the natural logs of commodity price index, the Consumer price index, real GDP (distributed monthly), and the M2 monetary aggregate.¹ The second data set is comparable to the one used in Christiano, Eichenbaum and Evans (1997), and consists of the funds rate, and the natural

logarithms of total reserves, M1, payroll employment, the personal consumption expenditure deflator, and commodity prices. The only series common to both data sets are the federal funds rate and the commodity price index. However, many of the series in the two data sets show similar trends, so each data set does not represent completely independent information. This idea appears to be verified in the results.

Rudebusch (1998) criticized the reality of the data sets that are typically used in VAR applications for policy analysis because they use data that would not have been available at the time a forecast and policy decision was being made. We examine the quantitative significance of this claim by using real-time data for one of our data sets. To match our end-of-month futures market data we assume that the VAR forecasts are always formed at the end of any given month.² For the first data set we are able to match exactly the vintage of data that would have been available to forecasters in real-time. However, we find that using the real-time version of the data has very little impact on the resulting funds rate forecasts. For the second data set we employ only a 1998 vintage of historical data.

2.2 The VAR and use of prior restrictions

The VAR that will be the basis for our empirical analysis can be written as

$$y_t = v + B_1 y_{t-1} + \dots + B_p y_{t-p} + u_t, \quad (1)$$

where y_t denotes an $m \times 1$ vector of observations for month t , and $u_t = A_0^{-1} \varepsilon_t$ where ε_t is assumed to be a random $m \times 1$ vector having zero mean and covariance matrix equal to an identity, with A_0 being a non-singular $m \times m$ matrix, and the covariance of u_t is $\Sigma = A_0^{-1} (A_0^{-1})'$.

¹ Robertson and Tallman (1999) describe the construction of a monthly US real GDP series.

Getting imprecise OLS estimates in a high-dimensional VAR is likely to be a common practical problem because the number of parameters is often quite large relative to the available number of observations. Various solutions to this problem have been proposed in the forecasting literature, which all amount to putting *a priori* constraints on the values of the model's coefficients, i.e., using non-data information regarding coefficient values.

One approach to reducing the coefficient uncertainty is to set some coefficients to zero or other pre-assigned values that may or may not have been determined on the basis of an initial fitting of models to the data. For example, one might pre-specify a maximum lag order p_{\max} for the VAR ($p_{\max} = 13$ in our empirical analysis), and select the $p \leq p_{\max}$ that minimizes a specific criterion. To compare with the unrestricted model performance, we investigate how alternative methods to restrict the VAR model affect the resulting forecast performance. We consider two commonly used criteria in our empirical analysis: The Akaike Information Criteria (AIC) and the Schwarz Information Criteria (SIC) (see e.g. Lutkepohl, 1991 for a discussion). Similarly, on the basis of an initial examination of the data one might choose to pre-difference series that appear to exhibit trends and/or quite persistent local levels over time, prior to fitting the VAR. This is mathematically the same as imposing restrictions on the sum of the coefficients of a VAR specified in the levels of the data.

An alternative, Bayesian approach to reducing parameter uncertainty is to use prior restrictions to influence the VAR coefficient estimates without requiring that the restrictions hold exactly (see for example, Litterman (1980, 1986), Doan, Litterman and Sims (1984), Sims (1992), Kadiyala and Karlsson (1997), and Sims and Zha (1998)). The Bayesian approach is to place a prior distribution on the coefficients and let the user adjust the “tightness” or “spread” of

² In real-time there are usually missing observations because of staggered release dates. We handle this problem using the conditional forecasting technique of Doan, Litterman and Sims 1984.

the distribution by setting the values of certain so-called “hyper-parameters”. Exact restrictions can be obtained as special cases of the priors by increasing the tightness of the restriction. The standard Bayesian VAR method is the one developed by Litterman (1980) and is based on a prior in which each series follows an independent random walk. However, the Litterman framework imposes independence between equations in the posterior as well by assuming the error covariance, Σ , is fixed and diagonal.

Doan, Litterman and Sims (1984) extended the base Litterman set-up by including a sum of coefficients prior that serves to push the estimated model toward a specification in first differences of the data as a prior weight (μ_1) is increased. However, in order to allow for the possibility of fewer trends than variables in the VAR Sims (1992) suggested incorporating a type of cointegration prior. In the limit, as a prior weight (μ_2) is increased, each series shares a single stochastic trend and the intercept terms will approach zero. Both these long-run priors are implemented by “mixing” a specific set of $m+1$ dummy observations into the data set.

In this paper, we use the sum of coefficients and co-integration priors, but replace the basic Litterman prior with a Normal-Wishart prior (see for example, Drèze and Richard, 1983, pp.539-541) of the type described in Kadiyala and Karlsson (1993,1997) and Sims and Zha (1998). This prior applies the random-walk aspect of Litterman, but does not maintain the diagonality of the error covariance as an exact restriction. Moreover, because, Σ and A_0 are related via $\Sigma^{-1} = A_0 A_0'$, this method can be adapted for Bayesian analysis of VAR models requiring identification of the elements of A_0 . In particular, one could allow priors to be formed directly on the individual elements of A_0 rather than implicitly via Σ - see Sims and Zha (1998, p.961).

The OLS estimator of the coefficients of the p^{th} -order VAR model in (1) has the form

$$\hat{B}^{OLS} = (X'X)^{-1} X'y,$$

where y is a $T \times m$ matrix with T -th row $(y_{1,T}, \dots, y_{m,T})$, and X is a $T \times (mp+1)$ matrix with T -th row $(1 \ y_{1,T-1} \dots y_{m,T-1} \ y_{1,T-2} \dots y_{m,T-2} \dots y_{1,T-p} \dots y_{m,T-p})$. The large-sample estimator of the error covariance is

$$\hat{\Sigma}^{OLS} = T^{-1}(y'y - \hat{B}'X'X\hat{B}).$$

In contrast, the coefficient estimator (the mean of the posterior distribution) based on the Normal-Wishart prior has the form

$$\hat{B}^{NW} = (\bar{H}^{-1} + X'X)^{-1}(\bar{H}^{-1}\bar{B} + X'y),$$

where \bar{B} is the prior mean of the coefficient matrix B , and \bar{H} is a diagonal matrix with elements defined as in Sims and Zha (1998, pp.954-955).³ The corresponding estimator of the error covariance is

$$\hat{\Sigma}^{NW} = T^{-1}(y'y - \hat{B}'(X'X + \bar{H}^{-1})\hat{B} + \bar{B}'\bar{H}^{-1}\bar{B} + \bar{S}),$$

where \bar{S} is the diagonal scale matrix in the prior inverse-Wishart distribution for Σ . As a

specific example, if $m=2$ and $p=2$, then we set $\bar{B} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{pmatrix}'$; \bar{S} has $(\sigma_i/\lambda_0)^2$ along the

diagonal, where σ_i is the scale term for each variable; and \bar{H} has diagonal elements $(\lambda_0\lambda_3)^2$,

$(\lambda_0\lambda_1/\sigma_1)^2$, $(\lambda_0\lambda_1/\sigma_2)^2$, $(\lambda_0\lambda_1/(2\lambda_2\sigma_1))^2$, $(\lambda_0\lambda_1/(2\lambda_2\sigma_2))^2$. Notice that λ_0 acts as an

overall tightness parameter; λ_1 controls the tightness of the random walk prior; λ_2 controls the

³ Here X and y_i contain the $m+1$ dummy observations for the sum of coefficients and co-integration priors (see Sims and Zha (1998, p.956)). GAUSS code for estimating this BVAR is available from the authors upon request.

rate of decay in the variance of the coefficients on lagged variables; and λ_3 determines the tightness of the zero prior mean on the intercepts. Smaller values for $\lambda_k > 0$ imply that the prior restrictions hold more tightly.

2.3 A test of the difference in forecast accuracy

Diebold and Mariano (1995) present several statistics that could be used for testing relative multi-step forecast accuracy. In our empirical work we employ their general asymptotic test statistic, using the small sample correction suggested by Harvey et al. (1997). The test statistic is based on assessing the size of the deviation in some function of two sets of forecast errors, such as the average difference in the squared errors. Specifically, for comparing the h step-ahead forecasts of two methods i and j , the modified Diebold and Mariano statistic has the form:

$$DM = \left(\frac{k + 1 - 2h + n^{-1}h(h-1)}{k} \right) \frac{\bar{d}}{\sqrt{\hat{f}_h/k}}$$

where $\bar{d} = k^{-1} \sum_{t=1}^k [g(e_{it}) - g(e_{jt})]$ is the average difference of some pre-specified function of the k forecast errors from each method, and \hat{f}_h is a consistent estimate of the asymptotic variance of \bar{d} (see Diebold and Mariano (1995) for more details). The first term in parentheses in the DM statistic is an adjustment suggested by Harvey et al. (1997) that tries to account for the fact that the raw DM test statistic can have an inflated size in small samples when $h > 1$. We make this correction although for the number of forecast observations in our application ($k = 96$), the impact is negligible.

2.4 Forecasts – Accuracy and Comparison

We report results from comparing the implicit futures market forecasts of the federal funds rate against forecasts from the following seven alternative schemes:

1. The *no change*, or *random walk* forecast.
2. An unrestricted VAR specification in the levels of the data, estimated by OLS, and with $p = 13$ imposed. This is denoted as the *VAR* forecast.
3. A VAR specification in the levels of the data and where the lag length is chosen on the basis of the Schwarz Information Criterion (SIC), with p_{\max} set at 13. This is denoted as the *VAR-SIC* forecast.
4. A VAR specification with the sum of coefficients restriction imposed exactly (the data are first differenced prior to estimation), and with the lag length chosen on the basis of the SIC, with p_{\max} set at 13. This is denoted as the *DVAR-SIC* forecast.⁴
5. A univariate AR specification for the funds rate denoted as the *DAR* forecast.
6. A univariate AR specification for the funds rate with the sum of coefficients restriction imposed exactly, and where the lag length is chosen on the basis of the SIC with p_{\max} set at 13. This is denoted as the *DAR-SIC* forecast.
7. The BVAR specification described in the previous section. Following Waggoner and Zha (1998) the parameters controlling the priors are set at $\lambda_0 = 0.6$, $\lambda_1 = 0.1$, $\lambda_2 = 1$, $\lambda_3 = 0.1$, $\mu_1 = 5$ and $\mu_2 = 5$. The scale parameters, σ_i , are obtained from the residuals of univariate AR regressions. The forecast is denoted as the *BVAR* forecast.

⁴ Results are not reported for the VAR forecasts based on the AIC since although they are better than an unrestricted VAR they are strictly dominated by those based on the SIC. The fact that specifications in which the lag length is selected based on a penalty function outperform those that do not suggest that down-weighting distant lags is advantageous to forecast performance.

It should be stressed that our empirical exercise asks whether there is empirical evidence that a VAR model can generate accurate forecasts of the federal funds rate. We are not investigating whether a VAR model actually did produce accurate forecasts. There are two reasons for this. First, our BVAR uses an estimation technology that was not available until quite recently and hence unavailable to forecasters in 1988 when the futures market first began trading. Second, the values for the prior weights (the hyper-parameters) are those chosen by Zha (1998) in order to maximize a multivariate, one year ahead, forecast accuracy criterion over the 1979-1996 period. While we are not looking at annual forecasts, the results may give a misleading guide as to how well a forecaster using the BVAR would have done if they had also been faced with a task of selecting hyper-parameter values in real-time. We do not undertake a real-time parameter search here, but we can get a feel for the sensitivity of our empirical results by varying the base hyper-parameter values. We find that for a reasonably large portion of the parameter space the empirical results are robust –see the Appendix for the details.

For each VAR specification a maximum of $p = 13$ lagged observations of the six variables are included in each equation. The VAR is first fit to data for the period February 1960 to December 1989, with the 13 pre-sample values being those for January 1959 to January 1960. The VAR is re-estimated (and the lag length is reselected in the models applying either the AIC or the SIC) each month through to September of 1997. In this model re-estimation process, the coefficient estimates can vary in response to new data and revisions to existing data. Each time the model is re-estimated, we generate funds rate forecasts for the next two months. Thus, for example, with the first data set at the end of March of 1997 we have available data on the funds rate and commodity prices for March; CPI, M2 and unemployment data for February, and GDP data for December of 1996. In this example the coefficients are estimated using the data through December 1996, and the missing observations on CPI, M2 and unemployment for March,

together with the January to March values of GDP are estimated using the conditional forecasting technique. Point forecasts of the funds rate are then constructed for April and May in the standard way. Applying this process starting in December 1989 and ending in November 1997 generates a sample of 96 one-month-ahead forecasts, and 95 two-month-ahead forecasts to analyze.

Table 1 displays the results from the real-time forecast application for one and two-month-ahead forecasts.⁵ The root mean squared errors (RMSE) and mean absolute errors (MAE) of the various VAR specification forecasts are listed together with the corresponding summary statistic for the federal funds futures forecast errors. Also, the table reports the DM statistic for testing the equality of forecast accuracy based on squared error and absolute error criteria. For one-month-ahead forecasts, the table also presents the correlation between the futures market forecast errors and the errors from each alternative scheme. The results in Table 1 suggest several general inferences regarding the forecast accuracy of the various specifications of the VAR models estimated by OLS. First, the least accurate forecasts of the federal funds rate are produced by the unrestricted VAR in levels estimated by OLS. The RMSE is almost three times higher than for the futures market at the one-month horizon, and about two times higher than the corresponding RMSE of the no-change forecast. Also, the forecast accuracy of the VAR generally improves as more restrictions are used. In fact, the simple no change forecast appears to set a rather high hurdle for the models reflecting the fact that over the forecast sample period the federal funds rate did not vary much. The best performing OLS-based VAR has the lag length chosen via the SIC (usually, $p=2$ is selected) and uses first differences of the data. However, the RMSE of this specification is still significantly larger than that of the futures

⁵ As in the previous research we also make the somewhat ad-hoc bias correction to the futures market forecasts by subtracting off the sample average bias. For out-of-sample analysis the average bias estimate is updated each time a forecast is formed.

market forecasts at both forecast horizons.⁶ Specifically, the null hypothesis of equal MSE and MAE between the VAR and the futures market forecasts is rejected in each case at the 1 percent significance level. In contrast, the estimated BVAR specification generates forecast errors that are not significantly larger (in a mean square error sense) than those implicit in the federal funds futures market at either the one or two-month horizons.⁷

Figure 1 presents the one-month-ahead forecast errors of four of the models plotted against the errors from the federal funds futures market. The first panel plots the forecast from the unrestricted VAR in levels. The second panel gives the no-change forecast errors. The third panel presents the forecast errors from the DVAR-SIC model, while the BVAR forecast errors are plotted in the fourth panel. Figure 2 presents the corresponding errors for the two-step-ahead forecasts. Clearly, the errors from the BVAR are less variable than from the other VAR-based procedures. Notice however, that at the one-month horizon the BVAR generates more moderately sized errors than does the futures market. This feature underlies the result that the forecast accuracy of the BVAR and futures market is less similar at the one-month horizon when based on the MAE rather than the MSE that tends to penalize large errors much more than moderate sized errors.⁸

As Figures 1 and 2 make clear, despite the relatively similar average out-of-sample forecasting accuracy, the futures market and the BVAR still generate somewhat different forecasts. Thus, to the extent that the futures rate can be taken to represent the market's

⁶ For one-month-ahead forecasts there are additional accuracy gains if one employs a univariate $AR(p)$ in first differences of the funds rate and with p selected by the SIC (usually $p = 3$ is selected). However, this gain is less obvious at a two-month-horizon.

⁷ In the case of mean absolute error the conclusion is not as strong at the one-month horizon, but at the two month horizon there is again no statistically discernable differences in performance between the futures market and the BVAR.

⁸ We also generated forecast errors from the same selection of models in Table 1 using a data set comprised of the latest revisions of the data series. The inferences are essentially unaffected by the use of latest available data.

expectation of the short-term path for the federal funds rate the BVAR is not the model used by the market. However, the correlation between the one-month-ahead forecast errors reported in Table 1 suggests that the BVAR is a much closer approximation than are unrestricted VAR specifications. In particular, the correlation between the errors from the futures market and the VAR forecasts is only 0.21, while the correlation between the futures market errors and those from the BVAR is 0.63.

The forecast error results listed in Table 1 are conditional on the data set used together with the settings for the prior hyper-parameters. To get some idea of the importance of these choices we examine the forecast performance of a VAR for a different data set, but keeping the hyper-parameters fixed at the values used previously. Specifically, we use the seven-variable data set comparable to that from Christiano, Evans, and Eichenbaum (1997), and used in Rudebusch (1998) and Evans and Kuttner (1998). We then undertake the same experiment that we performed with the six-variable data set, except that we use only the latest data vintage throughout. The results are summarized in Table 2. Again, the BVAR has smaller summary error measures than the other VAR specifications, and generates forecast errors that are more highly correlated with those from the futures market. However, the statistical tests of forecast error differences are less compelling. For both the one- and two-step forecast horizon, the null hypothesis that the variability of the federal funds futures and the BVAR forecast errors are the same can be rejected at the 5 percent level. Hence, it appears that the relative performance of a VAR is not independent of the variables included in the model, and/or the selection of the prior hyper-parameters. The seven-variable VAR generates better forecasts than the six-variable VAR estimated by OLS, but the seven-variable BVAR generates worse funds rate forecasts than the six-variable BVAR for the same hyper-parameter settings.

CONCLUSION

A requirement for a model to be useful in policy making is that its forecasts should provide a reliable baseline. Recent criticisms of VAR models call into question the accuracy of VAR forecasts, and hence, raise concern about the viability of using these models in the policy arena. The main concerns stem from three potential sources of limitations of VARs 1) linear time-invariant structure, 2) the limited scope of the information set, and 3) long lag structures. In this paper we demonstrate that even in the presence of these three well-known limiting characteristics, a specific multi-variate Bayesian VAR estimation technique can improve substantially the accuracy of federal funds rate forecasts. These Bayesian forecasts are more accurate than forecasts from both unrestricted and restricted VARs estimated by OLS. Hence, choice of estimation technique appears to be an important consideration when implementing VAR models designed for policy analysis and forecasting. Our results appear to refute the argument that VAR models forecast poorly, and thus, are not good tools for policy analysis.

The improved forecast accuracy for BVARs is robust for VARs estimated using two different data sets. For one data set, the BVAR generates forecasts whose accuracy is largely indistinguishable from the accuracy of the forecasts implicit in federal funds futures market data. For another data set the BVAR forecasts are more accurate than those of alternative specifications, but are still statistically inferior to the forecasts implied by futures market data. That is, the BVAR's good forecast performance is not completely independent of the data set used in the model. The empirical findings in this paper suggest that using probabilistic prior information can provide an effective way of introducing a degree of parsimony into a high-dimensional VAR specification without at the same time imposing too many strong, and possibly false, exact *a priori* restrictions on the model.

APPENDIX 1

The BVAR forecast results in the paper employ a specific set of hyper-parameter values chosen by Zha (1998) to minimize a forecast error criterion. In this Appendix, we present forecast results from the BVAR using a variety of hyper-parameter settings. For instance, the overall tightness parameter λ_0 is set at 0.6 in baseline BVAR. We examine forecasts when the value of λ_0 is tighter at 0.2 and looser at 1.0. In the base specification, the parameter controlling the tightness of the random walk prior is set to 0.1; we investigate how the forecasts respond to increasing this value to 0.3. Finally, the sum of coefficients and the co-integration priors are set to 5 in the base specification. We compare the forecast accuracy when using the base values to forecasts generated when we increase the weights 10 and decrease them to 1.0. There is a clear interaction between the overall tightness and the random walk parameter as seen in Tables A1 and A2.

The BVAR formulation generates improvement in forecasting accuracy over the unrestricted VAR in all cases. Across the two data sets, certain parameter settings generate more noticeable improvement than others. For example, for the Waggoner-Zha data set, the most accurate forecasts come from a VAR with parameter settings of (0.6,0.1, 1,0.1, 5, 5) and (0.2,0.3,1, 0.1,5,5). In contrast, for the Christiano, Eichenbaum, and Evans data set, the most accurate forecasts come from a parameter setting of (1,0.1,1,0.1,10,10). Hence, the choice of the “best” set of hyper-parameters depends on the given data set. Also, tightening the priors in some cases worsens the forecast performance. Overall, the results in this appendix suggest that while the Bayesian VAR improves forecast accuracy across a range of parameter settings, the degree of forecast accuracy improvement for a given data set is dependent on the choice of hyper-parameter values.

TABLE A1: Real-time Out-of-Sample Forecast Accuracy Statistics (Six-Variable VAR)

A Sensitivity Analysis

	One-month-ahead forecasts					Two-month-ahead forecasts			
	RMSE	DM-test	MAE	DM-test	Correlation	RMSE	DM-test	MAE	DM-test
Futures	13.69		9.55			20.99		16.00	
No Change	19.28	-2.65	13.51	-2.74	0.58	33.11	-2.77	22.42	-2.42
VAR	41.58	-6.68	32.62	-8.64	0.21	77.72	-5.73	64.74	-7.26
VAR-SIC	27.57	-6.49	22.94	-10.18	0.48	55.32	-5.81	47.23	-9.18
DVAR-SIC	19.05	-2.86	14.79	-3.71	0.54	31.52	-2.64	24.97	-3.45
AR	19.26	-3.52	14.92	-5.17	0.58	35.34	-3.75	26.11	-4.54
DAR-SIC	18.55	-3.01	14.31	-3.99	0.59	32.40	-3.33	22.96	-3.33
BVAR (.2,.1,1,.1,1,1)	17.54	-2.53	12.82	-2.78	0.59	30.06	-2.62	22.39	-3.00
BVAR (.2,.1,1,.1,5,5)	16.28	-1.79	11.59	-1.81	0.63	26.65	-2.11	18.43	-1.29
BVAR (.2,.1,1,.1,10,10)	16.28	-1.76	11.50	-1.73	0.64	26.55	-2.11	18.29	-1.28
BVAR (.2,.3,1,.1,1,1)	16.05	-1.71	12.57	-2.69	0.56	25.85	-2.09	19.66	-1.67
BVAR (.2,.3,1,.1,5,5)	15.04	-1.00	11.44	-1.75	0.64	22.92	-0.91	16.04	-0.02
BVAR (.2,.3,1,.1,10,10)	15.29	-1.15	11.29	-1.66	0.64	23.54	-1.13	16.27	-0.15
BVAR (.6,.1,1,.1,1,1)	16.05	-1.71	12.58	-2.69	0.56	25.85	-2.09	19.66	-1.67
BVAR (.6,.1,1,.1,5,5)	15.04	-1.00	11.44	-1.75	0.64	22.92	-0.91	16.04	-0.02
BVAR (.6,.1,1,.1,10,10)	15.28	-1.15	11.29	-1.66	0.64	23.53	-1.10	16.26	-0.15
BVAR (.6,.3,1,.1,1,1)	18.00	-3.10	14.18	-3.74	0.51	30.40	-3.05	23.96	-3.14
BVAR (.6,.3,1,.1,5,5)	17.23	-2.02	12.75	-2.73	0.59	28.09	-2.07	20.86	-2.51
BVAR (.6,.3,1,.1,10,10)	17.95	-2.29	13.21	-3.21	0.58	30.01	-2.24	22.54	-3.18
BVAR (1,.1,1,.1,1,1)	16.53	-1.91	13.21	-3.09	0.54	26.79	-2.23	20.67	-2.04
BVAR (1,.1,1,.1,5,5)	15.72	-1.38	11.93	-2.18	0.62	24.34	-1.30	17.47	-0.81
BVAR (1,.1,1,.1,10,10)	16.21	-1.63	12.03	-2.32	0.61	25.63	-1.57	18.59	-1.43
BVAR (1,.3,1,.1,1,1)	19.98	-3.36	15.64	-4.71	0.49	35.14	-3.75	28.06	-4.24
BVAR (1,.3,1,.1,5,5)	19.22	-2.66	14.19	-3.65	0.56	32.85	-2.59	24.81	-3.74
BVAR (1,.3,1,.1,10,10)	19.94	-2.84	14.79	-4.11	0.56	34.71	-2.61	26.36	-4.13

TABLE A2: Out-of-Sample Forecast Accuracy Statistics (Seven-Variable VAR)

A Sensitivity Analysis

	One-month-ahead forecasts					Two-month-ahead forecasts			
	RMSE	DM-test	MAE	DM-test	Correlation	RMSE	DM-test	MAE	DM-test
VAR	27.79	-5.70	23.14	-7.18	0.38	44.52	-4.03	36.04	-5.25
VAR-SIC	18.34	-2.79	14.57	-4.13	0.53	33.51	-4.30	27.87	-6.10
DVAR-SIC	19.23	-3.15	14.89	-3.73	0.53	30.67	-3.03	23.93	-3.16
BVAR (.2,.1,1,.1,1,1)	19.19	-3.11	14.01	-3.66	0.59	33.40	-3.12	24.98	-3.82
BVAR (.2,.1,1,.1,5,5)	18.76	-2.96	13.80	-3.64	0.60	32.19	-3.02	24.14	-3.56
BVAR (.2,.1,1,.1,10,10)	18.33	-2.74	13.31	-3.24	0.61	31.20	-2.82	22.79	-3.02
BVAR (.2,.3,1,.1,1,1)	19.31	-2.86	14.46	-3.77	0.55	34.39	-2.86	26.57	-3.64
BVAR (.2,.3,1,.1,5,5)	18.22	-2.42	13.61	-3.41	0.57	31.55	-2.37	23.83	-2.93
BVAR (.2,.3,1,.1,10,10)	17.33	-2.04	12.93	-3.08	0.60	29.25	-2.00	20.58	-1.81
BVAR (.6,.1,1,.1,1,1)	19.31	-2.86	14.46	-3.77	0.56	34.39	-2.86	26.56	-3.64
BVAR (.6,.1,1,.1,5,5)	18.22	-2.42	13.61	-3.40	0.57	31.55	-2.37	23.83	-2.93
BVAR (.6,.1,1,.1,10,10)	17.32	-2.04	12.93	-3.07	0.59	29.24	-1.99	20.57	-1.81
BVAR (.6,.3,1,.1,1,1)	17.94	-2.37	13.63	-3.04	0.54	29.99	-2.21	22.46	-2.37
BVAR (.6,.3,1,.1,5,5)	17.26	-2.12	13.32	-3.06	0.55	28.00	-1.86	20.61	-1.92
BVAR (.6,.3,1,.1,10,10)	17.29	-2.26	13.55	-3.40	0.55	28.09	-2.12	21.49	-2.47
BVAR (1,.1,1,.1,1,1)	18.32	-2.45	13.49	-3.00	0.55	31.75	-2.44	23.57	-2.62
BVAR (1,.1,1,.1,5,5)	17.40	-2.07	12.90	-2.81	0.57	29.15	-1.97	20.77	-1.83
BVAR (1,.1,1,.1,10,10)	16.96	-1.95	12.97	-3.05	0.58	27.86	-1.85	19.93	-1.68
BVAR (1,.3,1,.1,1,1)	18.35	-2.68	16.00	-3.25	0.52	30.29	-2.40	23.60	-2.81
BVAR (1,.3,1,.1,5,5)	17.79	-2.54	13.95	-3.32	0.53	28.84	-2.22	22.28	-2.64
BVAR (1,.3,1,.1,10,10)	18.05	-2.84	14.39	-3.69	0.52	29.78	-2.74	23.84	-3.36

REFERENCES

- Christiano, Lawrence, J., Martin Eichenbaum, Charles L. Evans 1997 “Monetary Shocks: What Have We Learned, and to What End?” Forthcoming in the *Handbook of Macroeconomics*.
- Diebold, Francis, X., and Roberto S. Mariano. 1995 “Comparing Predictive Accuracy.” *Journal of Business and Economic Statistics*, 13: 253-263.
- Doan, Thomas , Robert B. Litterman and Christopher A. Sims. 1984. “Forecasting and Conditional Projection Using Realistic Prior Distributions.” *Econometric Reviews*. 3: 1–100.
- Drèze, J.H., and Jean-F. Richard. 1983 “Bayesian Analysis of Simultaneous Equation Systems.” In Z. Griliches and M.D. Intriligator (eds), *Handbook of Econometrics*, Vol I, Amsterdam.
- Evans, Charles L. and Kenneth N. Kuttner. 1998 “Can VARs Describe Monetary Policy?” Unpublished Manuscript, April.
- Fair, Ray C. and Robert J. Shiller. 1990. “Comparing Information in Forecasts from Econometric Models.” *American Economic Review* 80, 3: 375–389.
- Harvey, David, Stephen Leybourne, and Paul Newbold. 1997 “Testing the equality of prediction mean squared errors.” *International Journal of Forecasting*, 13: 281-291
- Kadiyala, K. Rao and Sune Karlsson. 1993 “Forecasting with Generalized Bayesian Vector Autoregressions.” *Journal of Forecasting*. 12: 365-378.
- Kadiyala, K. Rao and Sune Karlsson. 1997 “Numerical Methods for Estimation and Inference in Bayesian VAR-Models.” *Journal of Applied Econometrics*, 12, 99-132..
- Litterman, Robert B. 1980. “A Bayesian Procedure for Forecasting with Vector Autoregression.” Working Paper, Massachusetts Institute of Technology, Department of Economics.
- Litterman, Robert B. 1986 “Forecasting with Bayesian Vector Autoregressions – Five Years of Experience.” *Journal of Business and Economic Statistics*. (January)4,1:25–38.
- Lutkepohl, Helmut. 1991. *Introduction to Multiple Time Series Analysis*. Springer Verlag: New York.
- Robertson, John C. and Ellis W. Tallman. 1998. “Data Vintages and Measuring Forecast Model Performance.” Federal Reserve Bank of Atlanta *Economic Review*, Fourth Quarter: 4–20.
- Robertson, John C. and Ellis W. Tallman. 1999. “Vector Autoregressions: Forecasting and Reality.” Federal Reserve Bank of Atlanta *Economic Review* (forthcoming).
- Rudebusch, Glenn D. 1998. “Do Measures of Monetary Policy in a VAR Make Sense?” *International Economic Review*. (November) 39, 4: 907-31.
- Sims, Christopher A. 1992. “A Nine-Variable Probabilistic Macroeconomic Forecasting Model.” In *Business Cycles, Indicators, and Forecasting*, Stock, James H. and Mark W. Watson, eds. University of Chicago Press: Chicago. 179–204.
- Sims, Christopher A. and Tao A. Zha. 1998. “Bayesian Methods for Dynamic Multivariate Models.” *International Economic Review*. ” 39, 4: 949–968.
- Waggoner, Daniel and Tao A. Zha. 1998. “Conditional Forecasts in Dynamic Multivariate Models.” Manuscript, September.

- Wallis, Kenneth. 1989. "Macroeconomic Forecasting: A Survey," *Economic Journal*. 99, 394: 28–61.
- Zha, Tao A. 1998 "A Dynamic Multivariate Model for Use in Formulating Policy." Federal Reserve Bank of Atlanta *Economic Review*. 83, 1: 16–29.

TABLE 1: Real-time Out-of-Sample Forecast Accuracy Statistics (Six-Variable VAR)									
	One-month-ahead forecasts					Two-month-ahead forecasts			
	RMSE	DM-test	MAE	DM-test	Correlation	RMSE	DM-test	MAE	DM-test
Futures	13.69		9.55			20.99		16.00	
No Change	19.28	-2.65	13.51	-2.74	0.58	33.11	-2.77	22.42	-2.42
VAR	41.58	-6.68	32.62	-8.64	0.21	77.72	-5.73	64.74	-7.26
VAR-SIC	27.57	-6.49	22.94	-10.18	0.48	55.32	-5.81	47.23	-9.18
DVAR-SIC	19.05	-2.86	14.79	-3.71	0.54	31.52	-2.64	24.97	-3.45
AR	19.26	-3.52	14.92	-5.17	0.58	35.34	-3.75	26.11	-4.54
DAR-SIC	18.55	-3.01	14.31	-3.99	0.59	32.40	-3.33	22.96	-3.33
BVAR	15.24	-1.12	11.63	-1.93	0.63	23.73	-1.00	16.37	-0.21

TABLE 2: Out-of-Sample Forecast Accuracy Statistics (Seven-Variable VAR)									
	One-month-ahead forecasts					Two-month-ahead forecasts			
	RMSE	DM-test	MAE	DM-test	Correlation	RMSE	DM-test	MAE	DM-test
VAR	27.79	-5.70	23.14	-7.18	0.38	44.52	-4.03	36.04	-5.25
VAR-SIC	18.34	-2.79	14.57	-4.13	0.53	33.51	-4.30	27.87	-6.10
DVAR-SIC	19.23	-3.15	14.89	-3.73	0.53	30.67	-3.03	23.93	-3.16
BVAR	17.83	-2.24	13.25	-3.14	0.57	30.51	-2.24	22.45	-2.42

FIGURE 1

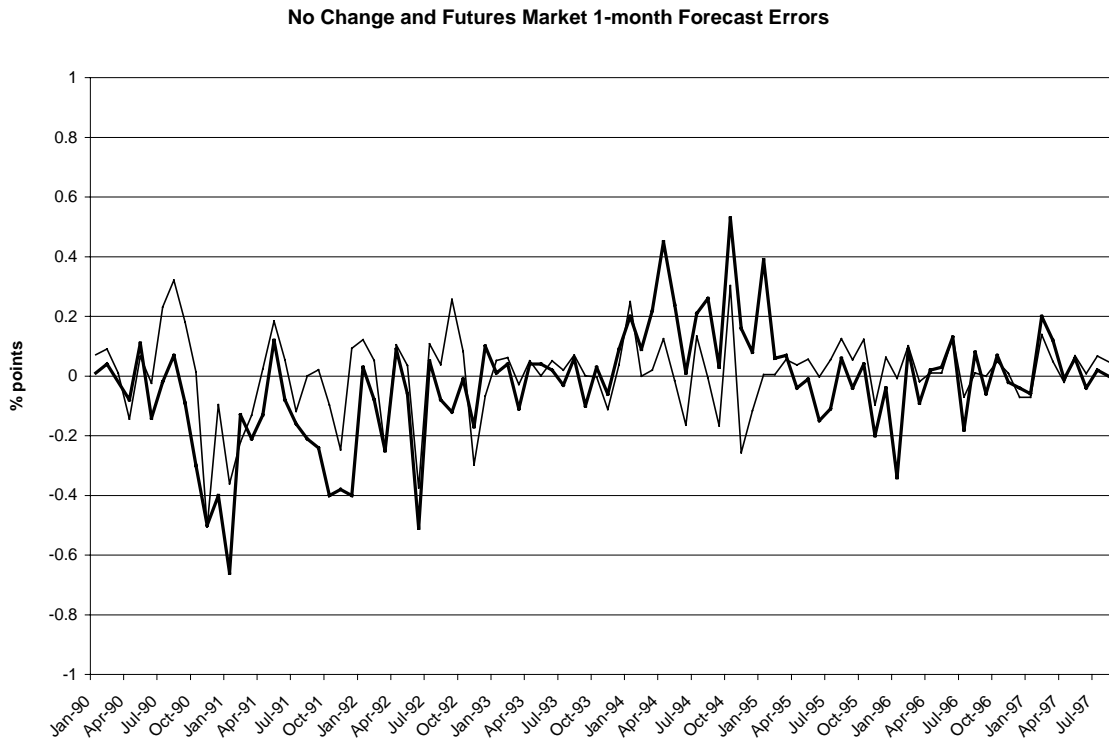
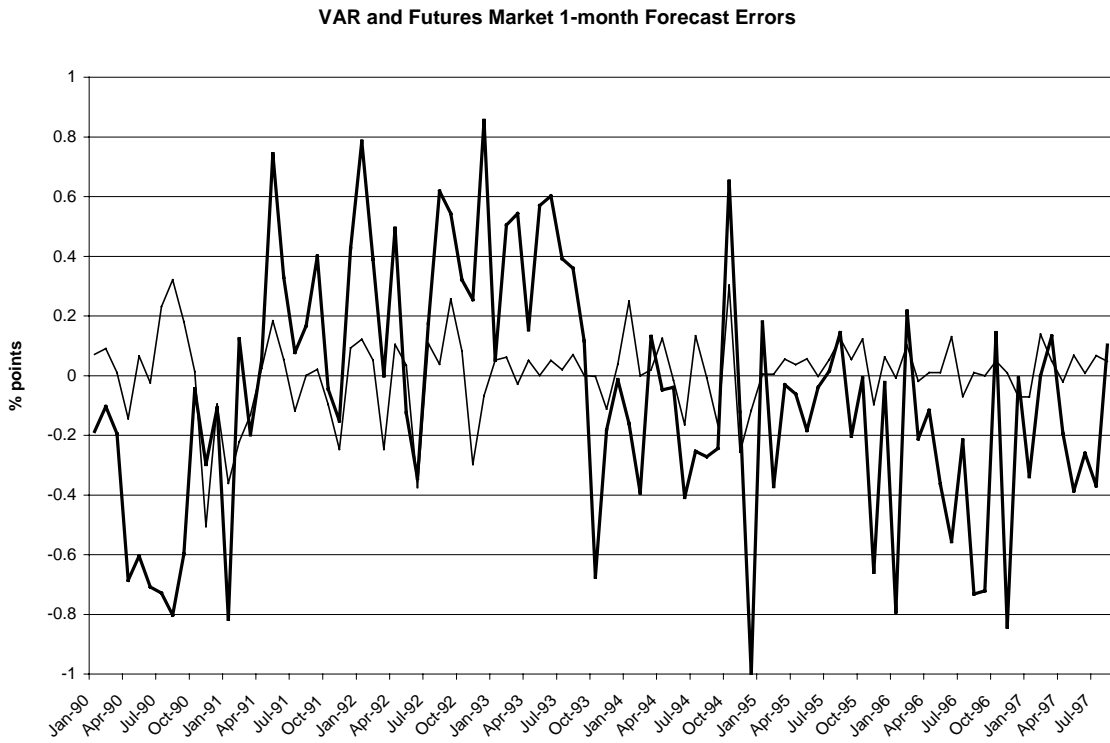
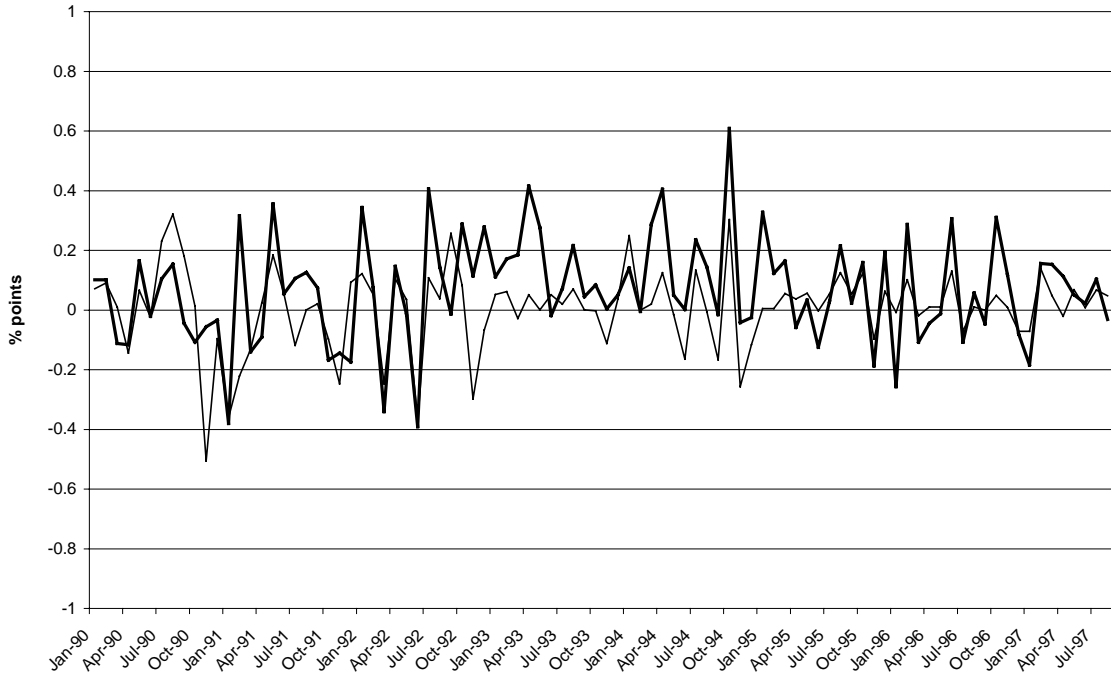


FIGURE 1 CONTINUED

DVAR-SIC and Futures Market 1-month Forecast Errors



BVAR and Futures Market 1-month Forecast Errors

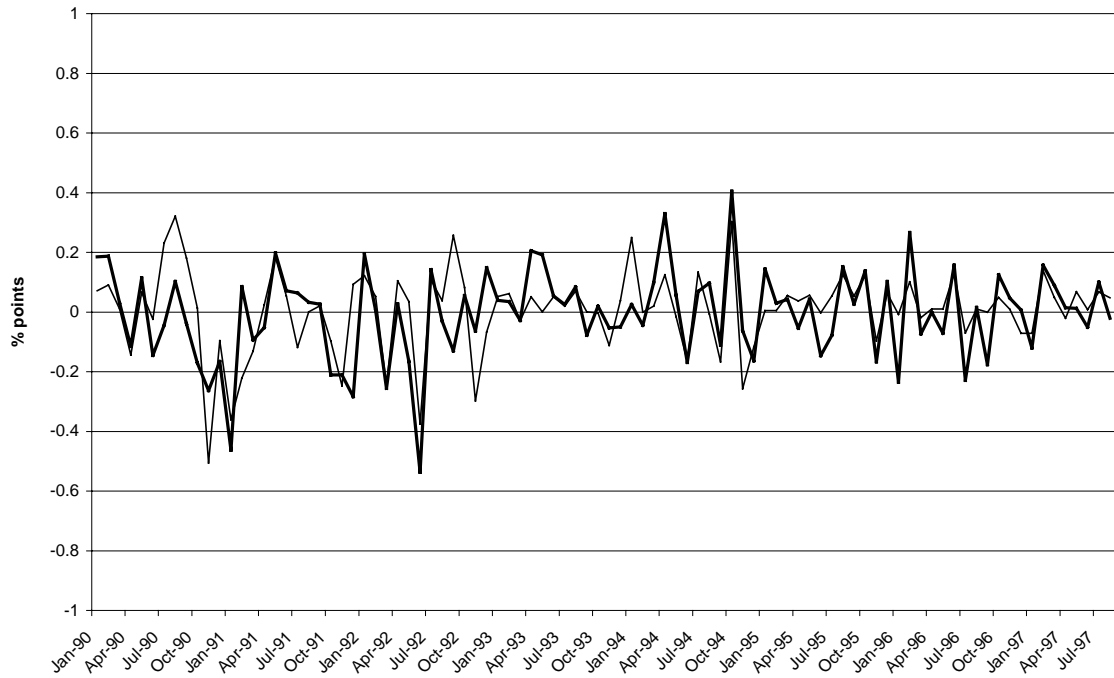


FIGURE 2

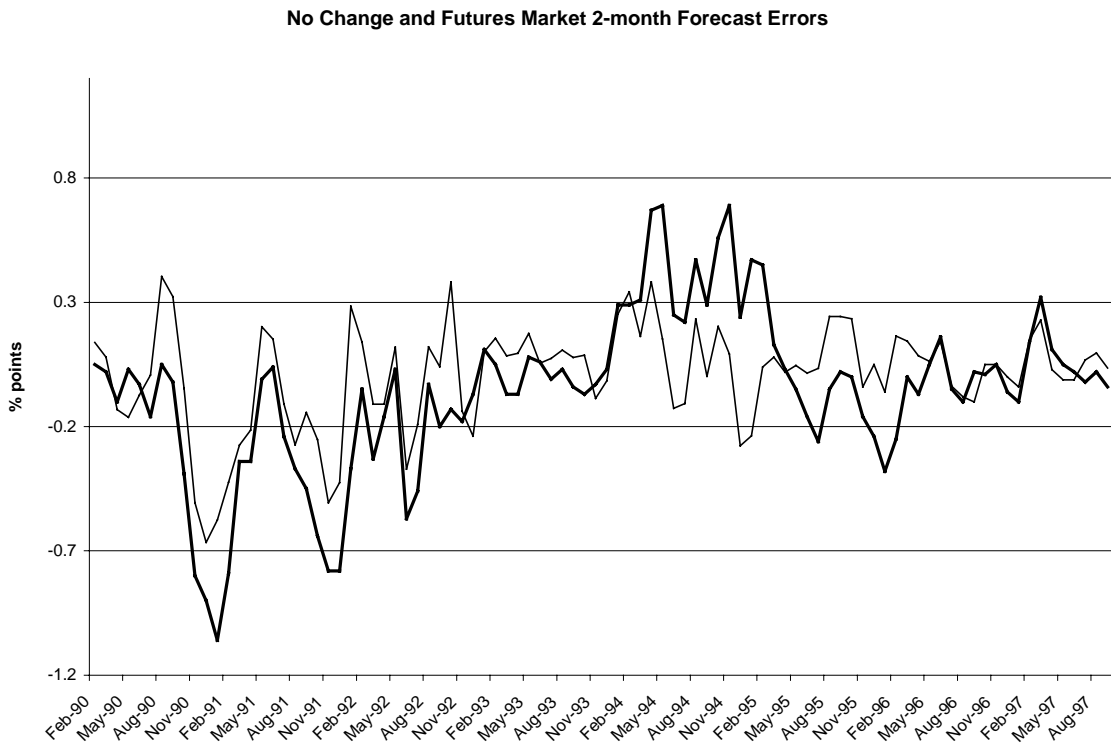
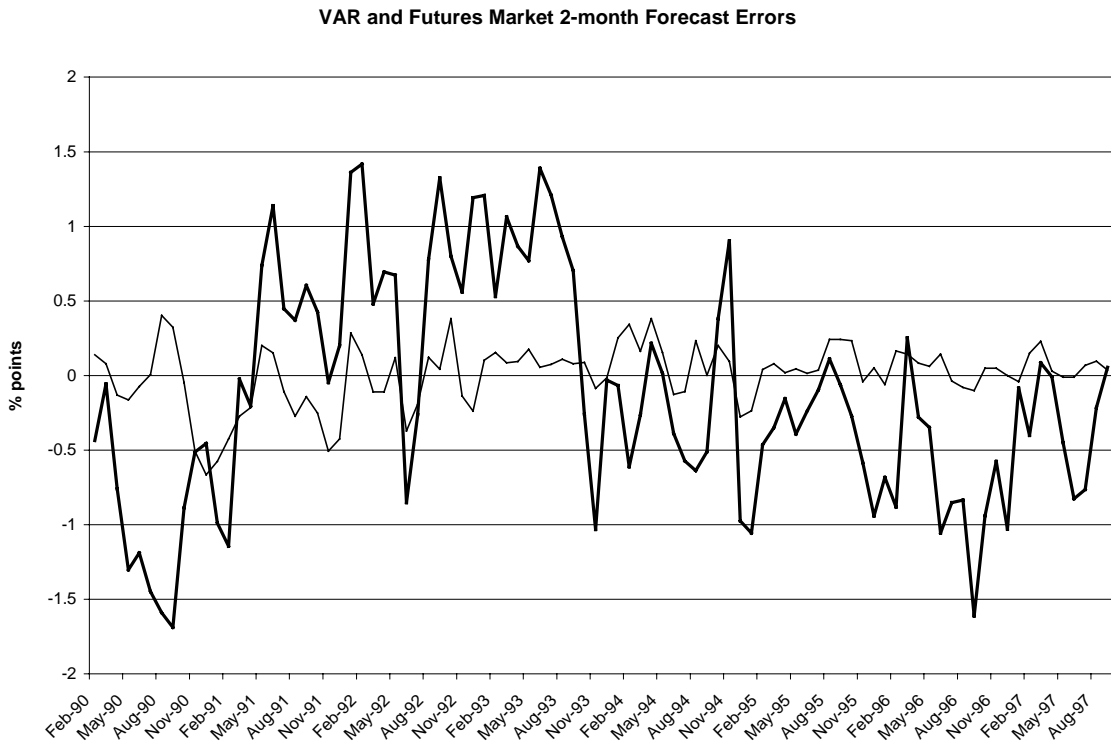
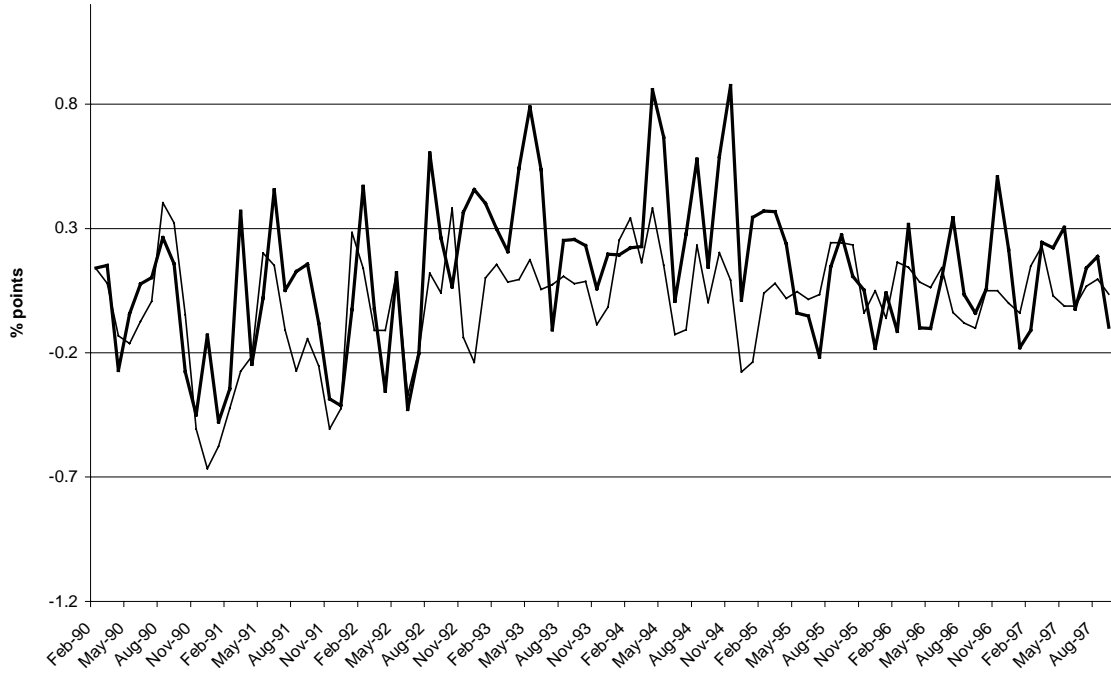


FIGURE 2 CONTINUED

DVAR-SIC and Futures Market 2-month Forecast Errors



BVAR and Futures Market 2-month Forecast Errors

