FEDERAL RESERVE BANK OF PHILADELPHIA

Ten Independence Mall
Philadelphia, Pennsylvania 19106-1574
(215) 574-6428, www.phil.frb.org

Working Papers

Research Department

WORKING PAPER NO. 98-12

A BAYESIAN VECTOR ERROR CORRECTIONS MODEL OF THE U.S. ECONOMY

Tom Stark
Federal Reserve Bank of Philadelphia

June 1998

WORKING PAPER NO. 98-12

A BAYESIAN VECTOR ERROR CORRECTIONS MODEL OF THE U.S. ECONOMY

Tom Stark

Federal Reserve Bank of Philadelphia

June 1998

Abstract

This paper presents a small-scale macroeconometric time-series model that can be used to generate short-term forecasts for U.S. output, inflation, and the rate of unemployment. Drawing on both the Bayesian VAR and vector error corrections (VEC) literature, I specify the baseline model as a Bayesian VEC. I document the model's forecasting ability over various periods, examine its impulse responses, and consider several reasonable alternative specifications. Based on a root-mean-square-error criterion, the baseline model works best, and I conclude that this model holds promise as a workhorse forecasting tool.

1. Introduction

Many analysts argue that successful monetary policy requires monetary policymakers to be forward looking. According to this view, the Fed should consider forecasts for key macroeconomic variables before deciding the appropriate stance of monetary policy. While this view may be unobjectionable from a theoretical perspective, it is complicated in practice by the fact that there are many different structural models of the economy, each of which may yield a forecast that points Fed policymakers in a different direction. It is well known that one source of difference among these competing models is the set of identifying assumptions that model builders impose to reflect their *a priori* views about the relationships among economic variables. Bayesian techniques can be used to incorporate many different views on the economy into one model and may help to reduce the problem that occurs when policymakers are confronted with conflicting forecasts.

This paper introduces a small-scale macroeconometric time-series model that can be used to generate short-term forecasts for U.S. output, inflation, and the rate of unemployment, three variables that many analysts believe are important in monetary policy deliberations.
The model's specification draws on both the error corrections and Bayesian approaches to estimation. Following suggestions in LeSage (1990) and Joutz, Maddala, and Trost (1995), I specify the model in Bayesian vector error corrections form (BVEC). Although some final specification decisions remain, I demonstrate how the model can be used in its present form to generate out-of-sample, dynamic forecasts over various sample periods. For example, over the period 1988 to 1995, the forecasts appear to be as accurate as those produced by a model recently specified by researchers at the Federal Reserve Bank of St. Louis. In addition, the model seems to

track the acceleration of inflation over the second-half of the 1970s and the subsequent disinflationary experiences of the 1980s and 1990s; out-of-sample forecasts for the unemployment rate over recent years do not suffer from the persistent negative errors that one might expect from empirical models that have a strong natural rate mechanism. The preferred specification produces root-mean-square-errors of 0.61 percent and 0.72 percent (annual rate) for the one- and two-year ahead, four-quarter-average rates of inflation in the 1990s. For the four-quarter-average unemployment rate, the RMSEs are 0.40 percent and 0.85 percent, respectively, over the same period.

To aid in model specification and forecast evaluation, I examine root-mean-square- errors for four-quarter-average forecasts of output growth, inflation, and the rate of unemployment over one- and two-year-ahead horizons. This represents a departure from more traditional approaches, which typically emphasize forecast performance over a one-quarter-ahead horizon. However, the present criterion seems better suited to the perspective of a Fed policymaker, who may be more interested in the economy's forecastable momentum over the next year or so, rather than just over the next quarter. By examining four-quarter-average forecasts, I also eliminate the tradeoff that can arise when one makes specification adjustments that improve forecast performance over a one-quarter horizon but worsen such performance over the remaining quarters.

Although the core of my preferred specification consists of output, inflation, the rate of unemployment, and the federal funds rate, I find slightly improved performance when real M2, the 10-year Treasury bond rate, and the relative price of imports are added. The spread between the federal funds rate and the 10-year Treasury bond rate enters the model in error correction form. Thus, the preferred specification in this paper is a vector error corrections process in output growth, the change in the inflation rate, the change in the rates of interest on federal funds and 10-year Treasury bonds, the change in the unemployment rate, real M2 growth, the growth in real import prices, and one lag of the spread between the two interest rates. A variant on Litterman's (1986) approach is used to assign Bayesian priors on all coefficients except that on the error correction (interest-rate spread) term, which is assigned a diffuse prior, as suggested in Joutz et al. (1995). The model is estimated with Theil's (1963) mixed estimator.

As in Laurent (1988), I find that the spread enters the output growth equation with the appropriate sign and with statistical significance. The spread also enters significantly in the unemployment equation, and, when combined with lagged changes in the 10-year rate, produces a sizable reduction in the forecast errors for the unemployment rate. The inclusion of long-term rates also yields improved out-of-sample forecast performance over the 1980 and 1982 recessions. Real import prices enter significantly in the inflation and federal funds rate equations and improve the accuracy of the inflation forecasts; the inclusion of real M2 yields improved output forecasts over the 1990s. Interestingly, out-of-sample dynamic forecasts for real M2 growth over the last few years indicate that M2 may be back on track.

An examination of the model's impulse responses yields results that seem plausible.   For example, based on estimation through 1997Q4, a 0.6 percentage point increase in the federal funds rate produces lower output growth over the following 10 quarters, with the largest effect--a slowing in annualized, quarter-over-quarter rate of growth in output of 0.7 percentage points--occurring two quarters after the shock. The inflation rate begins to decline after five quarters, and the unemployment rate rises gradually by a maximum amount of 0.3 percentage points. The qualitative nature of these responses is consistent with Christiano, Eichenbaum and Evans' (1996) monetary "facts."

The paper is organized as follows. In section 2, I discuss model specification and estimation. Section 3 presents the results of various forecast evaluation experiments, while section 4 considers--and rejects--some proposed improvements to the baseline model. Section 5 concludes.

2. Model Specification and Estimation

Litterman's Bayesian Approach

Bayesian techniques have a couple of important advantages over traditional approaches to designing a macroeconomic forecasting model. Included among the traditional approaches is the structural approach, which includes large-scale IS/LM/Aggregate-Supply models, and the relatively more recent, reduced-form approach, as represented by the unrestricted vector autoregression model. First, as noted by Todd (1984), standard, unrestricted VARS often suffer from the problem of overfitting--fitting too many parameters to too few data points--and consequently, their forecasting performance has not been good. Bayesian VARs have been shown to yield an improvement on this front. Second, structural models are often specified with many identifying assumptions that reflect the model builder's personal views on how the economy works but that also exclude other reasonable views. The exclusions tend to take the form of zero restrictions--omitting certain variables from certain equations--and generally yield reduced-form representations for use in forecasting that are heavily restricted in ways that may be inconsistent with the data. The Bayesian approach, in contrast, allows an analyst to give nonzero weight to many different views of the economy in specifying a reduced-form forecasting model. Robert Litterman developed the Bayesian approach to estimating VARs in the early 1980s, and his approach continues to enjoy widespread acceptance today.[1]

In essence, the Litterman Bayesian estimator involves estimation subject to "fuzzy" or stochastic prior restrictions on a model's parameters in a manner reminiscent of Theil's (1963) mixed estimation technique. Indeed, under certain assumptions, the two estimators are identical, and Bayesian VARs are typically estimated with the Theil approach. If we write the i^{th} equation of an n-variable VAR with maximum lag length L as $Y_{it} = X_t'\beta + e_{it}$, where X_t is an nL x 1 vector of 1 through L lags of the system's n dependent variables and β is the corresponding vector of VAR coefficients, Litterman's (1986) restrictions can be written in the form

$$R\beta = r + v. \tag{1}$$

Equation (1) is supposed to represent an analyst's prior beliefs about certain linear combinations of VAR coefficients, where R is a nonstochastic matrix of restrictions, r is the expected value of those restrictions and v is an error term assumed to have a mean of zero and variance-covariance matrix Σ. The presence of the stochastic term v means that the analyst views the elements of β as random variables and thus (1) defines the analyst's Bayesian prior. Assuming normality, the prior is $R\beta \sim N(r,\Sigma)$. An alternative interpretation, derived by solving (1) for r, is that β is

[1]Recent examples include Artis and Zhang (1990) , Trehan's (1992) current-quarter model for predicting real GDP, Joutz, Maddala, and Trost (1995), Dua and Smyth (1995), and Bikker (1998).

nonstochastic and there exists some uncertainty about whether the restriction $R\beta = r$ holds exactly. This represents Theil's (1963) preferred interpretation.

On the basis of the success of the pure random walk model in describing the behavior of many economic time series, Litterman developed a set of baseline prior restrictions that most analysts adopt to this day. These baseline priors help analysts assign numerical values to r, R, and Σ. Formally, these priors--known as the Minnesota priors--assume each variable in the VAR follows a random walk (possibly with drift). This means that each coefficient in the VAR has a mean of zero except the coefficient on the own first lag, which has a prior mean of unity and implies that r is a vector of zeros, except in the positions that correspond to the first own lag in each equation. A value of unity is entered in those positions. The restriction matrix R is constructed such that each individual coefficient is equated with its corresponding value in r plus an error. When values for r and R are selected in this manner, the analyst's prior can be written as $Y_{it} = Y_{it-1} + e_{it}$. When the VAR includes variables that trend, like real GDP, a drift term is added.

Litterman also provided recommendations for the elements of Σ. He assumed zero covariance between elements of β. Letting SD(i,j,l) denote the standard deviation of the analyst's prior on the coefficient attached to the l^{th} lag of the j^{th} variable in the i^{th} equation, Litterman suggested a parameterization based on a few hyperparameters according to $SD(i,i,l) = \lambda/l$ and, for $i \neq j$, $SD(i,j,l) = \theta\lambda(\sigma_i/\sigma_j)(1/l)$. This parameterization implies that the analyst is more certain about his prior beliefs--the standard deviations are smaller--as the lag, l, increases. The parameter λ measures the overall tightness of the priors and θ allows for relative differences in the tightness of beliefs about the values of the cross-lag coefficients. The term (σ_i/σ_j) adjusts for the effect on coefficient magnitudes of scale differences in the underlying data. In a practical implementation of Litterman's technique, many analysts generalize the standard deviation formulas according to

$$SD(i,j,l) = \lambda\omega_{ij}(\sigma_i/\sigma_j)(1/l)^{\delta}, \tag{2}$$

where ω_{ii} is usually, though not necessarily, unity and δ may be less than unity if the analyst prefers the standard deviations decline less rapidly as the lag length increases. Typically, analysts will calibrate the values of λ, ω_{ij}, and δ to yield optimal forecast performance, but Hamilton (1994) suggests baseline values of λ=0.20, ω_{ij}= 0.5, ω_{ii}= 1.0, and δ=1.0.

Specification of the Baseline Model

My forecasting model applies the Litterman technique to a VAR composed of the log of real, chain-weighted GDP (RGDP), the log of the GDP chain-weighted price index (PGDP), the rate on federal funds divided by 400 (RFF), the log of the real, chain-weighted price of imports (RPIM), the unemployment rate divided by 100 (U), the log of real M2 (RM2), and the rate on 10-year Treasury bonds divided by 400 (RTB). As noted, all data with the exception of the interest rates and the rate of unemployment are in logs. Real M2 is defined as nominal M2 divided by the GDP price index; real import prices are defined similarly. The data set includes observations from 1959Q1 to 1997Q4. Quarterly observations on M2, the interest rates, and the unemployment rate are quarterly averages of monthly data. All data are obtained from standard sources.

In specifying the model, I make two departures from the Litterman recommendations. First, I find that forecast performance is enhanced by <u>imposing</u> a unit root on the underlying data, including the rates of inflation and unemployment, rather than assigning--as Litterman recommends--a unit root prior with a finite (and small) variance around that prior. Thus, rather than assuming $Y_{it} = \alpha Y_{it-1} + e_{it}$ with a prior that $\alpha=1.0$, I impose $\Delta Y_{it} = e_{it}$. Webb (1995) notes that such a modification improves the forecast performance of his Bayesian VAR as well. The present model includes five lags of the first differenced data. Second, to each equation in the system, I add the error correction term RFF_{t-1}-RTB_{t-1}. Entering the spread in this fashion implies that the rates on federal funds and 10-year Treasury bonds are cointegrated with a cointegrating vector (1,-1). The Bayesian prior on this term is diffuse, as in LeSage (1990) and Joutz et al. (1995), and the inclusion of this term defines the system as a vector error correction model [Engle and Granger (1987)].[2] Assuming that the model is dynamically stable (i.e., all system eigenvalues on or inside the unit circle), the model's dynamics ensure that any long-run forecast will obey the cointegrating relation, RFF = RTB + constant.

---

[2]Unit root tests on an earlier version of the data set that extended through 1997Q2 yielded standard results: RGDP, ΔPGDP, RFF, RTB, RM2, RPIM appear to be I(1). In addition, U also appears to be I(1). Importantly, RFF-RTB appears I(0), suggesting that the rate on federal funds is cointegrated with the rate on 10-year Treasury bonds with a cointegrating vector of (1,-1). A preliminary examination for the presence of additional cointegrating relations, using the Johansen and Juselius (1990) technique, yielded inconclusive results. By adjusting both the lag length and information set, it was possible to find that RFF-ΔPGDP is I(0), suggesting that the federal funds rate and inflation are cointegrated and thus that the real federal funds rate is stationary and that a long-run Fisher effect holds. However, the need to adjust both the lag length and the information set, combined with the evidence presented in King and Watson (1997) against a long-run Fisher effect, led me to not impose such a relation on my forecasting model. Standard unit root tests on RFF-ΔPGDP are also ambiguous, although such tests do indicate conclusively that RTB-ΔPGDP is I(1). Below, I examine whether the model's forecasts are harmed by not imposing a stationary real rate.

In summary, the preferred forecasting model is a Bayesian vector error correction model with five lags and an error correction term given by RFF_{t-1}-$RTB_{t-1,}$ which I express as

$$\text{Model 1: } BVEC(5)\{\Delta RGDP, \Delta^2 PGDP, \Delta RFF, \Delta RPIM, \Delta U, \Delta RM2, \Delta RTB\}. \quad (3)$$

A typical equation in the system (3) is the equation for real GDP growth, given by

$$\begin{aligned} \Delta RGDP_t = c_1 &+ \beta_{11}(L)\Delta RGDP_{t-1} + \beta_{12}(L)\Delta^2 PGDP_{t-1} + \beta_{13}(L)\Delta RFF_{t-1} \\ &+ \beta_{14}(L)\Delta RPIM_{t-1} + \beta_{15}(L)\Delta U_{t-1} + \beta_{16}(L)\Delta RM2_{t-1} \\ &+ \beta_{17}(L)\Delta RTB_{t-1} + \rho_1(RFF_{t-1} - RTB_{t-1}) + e_{1t}\ , \end{aligned}$$

where $\beta_{ij}(L)$ is a polynomial in the lag operator, c_1 is a constant that reflects steady-state output growth and the steady-state spread, and ρ_1 is the error correction coefficient. The Bayesian prior for the mean of the coefficients in $\beta_{ij}(L)$ is zero, and a diffuse prior is employed for c_1 and ρ_1, as is true in all other equations.

Parameter estimates for this model are shown in Table 1, with estimates significant at the 10 percent level highlighted. The estimation period is 1960Q4 to 1997Q4, with observations from 1959Q1 to 1960Q3 serving as pre-sample values. All estimation uses the following values of the hyperparameters in (2): λ=0.20, δ=0.50, ω_{ii}= 1.0. The values of ω_{ij}, for $i \neq j$, vary from equation to equation and by variable within an individual equation, although the forecast performance is not much affected by these choices.

As Table 1 shows, the error correction term, denoted EC(t-1), enters significantly in the RGDP equation [p-value (not shown), 0.003], the RFF equation (p-value, 0.068), the U equation (p-value, 0.000), and in the RTB equation (p-value, 0.063).[3] Also of interest is that the unemployment rate and the relative price of imports enter significantly in the PGDP equation, and that both interest rates enter significantly in the RM2 equation. The results for the error correction term are particularly encouraging, especially given recent findings that similar financial spreads appear useful in predicting future output and inflation in other countries [Davis and Henry (1994) and Davis and Fagan (1997)].

Given the success of the error correction term, it is of some interest to ask whether its statistical significance is unique to the particular sample. In Table 2, I report estimates of the error correction coefficient for sample periods that end in 1979Q4, 1984Q4, 1989Q4, and 1994Q4. As before, statistical significance at the 10 percent level is highlighted. The table shows that the error correction spread between the federal funds rate and the 10-year Treasury bond rate (RFF_{t-1}-RTB_{t-1}) always enters significantly in the RGDP and U equations; however, there is evidence of coefficient instability in the RGDP equation. In addition, prior to the 1990s, the spread also enters significantly in the RM2 equation; once the sample extends into the 1990s, however, the magnitude of that coefficient drops dramatically, as does its statistical significance. This may reflect ongoing restructuring of the S&L industry [Carlson and Byrne (1992)] and increased availability and use of bond and stock mutual funds [Darin and Hetzel (1994)] over the

[3]Because of the nature of the restrictions placed on the equations, the degrees of freedom calculation differs from the normal method. In this paper, degrees of freedom are calculated as T-D, where D is the number of coefficients with diffuse priors (2), and not T-K, where K is the number of regressors.

early 1990s. I conclude that despite some indication of coefficient instability, the spread is important in the system under study. Below, I provide forecasting evidence to suggest this is, indeed, the case.

3. Forecast Evaluation

Methodology

In this section, I present the results of various forecasting exercises using the model described in (3). As noted earlier, I measure forecast performance with the root-mean-square-error criterion (RMSE) and compute sequences of dynamic, out-of-sample forecasts over a one- through eight-quarter-ahead horizon. To aid in assessing the model's performance, each sequence of eight quarterly forecasts is combined into two four-quarter-average forecasts in the following manner. For variables that trend (RGDP, PGDP), I compute

$$(Y_{hat,\ T+4|T} - Y_{T})*100.0, \text{ and}$$
$$(Y_{hat,\ T+8|T} - Y_{hat\ ,T+4|T})*100.0, \qquad (4)$$

where "hat" indicates a forecast for the log level of the variable Y, and T+J|T indicates a forecast of that variable's value at T+J given information through time period T. Thus, these magnitudes measure four-quarter-average growth over the next four quarters and over the subsequent four quarters, expressed in percentage points. For variables that do not trend over time (RFF, U), I compute

$0.25*(X_{hat,\ T+1|T} + X_{hat,\ T+2|T} + X_{hat,\ T+3|T} + X_{hat,\ T+4|T})$, and

$$0.25*(X_{hat,\ T+5|T} + X_{hat,\ T+6|T} + X_{hat,\ T+7|T} + X_{hat,\ T+8|T}), \tag{5}$$

both expressed in annualized percentage points, where X is the level of the variable in question, like the rate on federal funds or the rate of unemployment. These magnitudes are four-quarter arithmetic averages. A rolling regression framework is used to compute a quarterly sequence of these four-quarter-average forecasts, with each quarter's eight quarterly forecasts based upon adding one more observation to the sample before reestimating and re-forecasting the model. Thus, the forecasts are based on parameters estimated on data that are not included in the forecast horizon. This procedure closely approximates real time forecasting, but it is not equivalent because I compute forecast errors based on the data as they now exist, rather than on the data as they would have existed around the time such forecasts would have been made.[4] Forecast errors are constructed by subtracting the forecasts in (4) - (5) from their corresponding historical values, and the RMSEs are constructed from these errors in the usual manner by squaring the errors, averaging the squared errors, and taking the square root of the average. Higher RMSEs suggest poorer forecast performance. The first sequence of eight quarterly forecasts, from which I compute the magnitudes in (4)-(5), covers the period 1975Q1 to 1976Q4 and is based on estimation through 1974Q4; the second sequence covers the period 1975Q2 to 1977Q1, based on estimation through 1975Q1, etc.

---

[4]The Federal Reserve Bank of Philadelphia is currently constructing a sequence of real time data sets--called Real Time Data Sets for Macroeconomists (RTDSM)--that will enable analysts to compute exact real time forecasts and forecast errors with models such as the one presented here.

Forecasting With the Baseline Model

Figures 1.a to 1.g plot the four-quarter-average, one-year-ahead (top panel) and two-year-ahead (bottom panel) forecasts (dashed line) and associated historical values (solid line) for the BVEC model in equation (3) using the rolling regression framework discussed above.
The forecasts begin in the mid-1970s, and the vertical boxes indicate recessions. The implied forecast errors are represented by the vertical difference between the dashed and solid lines.

Figure 1.a shows that at the one-year-ahead horizon, the model seems to capture well the rise in inflation over the late 1970s and the subsequent disinflationary experiences in the early 1980s and in the 1990s. Note that the forecasts are, at times, subject to some erratic, high frequency movements. Preliminary experimentation suggests that it is possible to smooth these movements--by adjusting the hyperparameters in (2) to yield tighter priors--but at the cost of increased bias, leaving the RMSE roughly unchanged. The extent to which such adjustments are desirable depends on the specific model user's loss function, and thus the results of such experiments are not reported.

Figure 1.b shows that the model captures quite well the movement in real GDP over the 1980 and 1982 recessions but misses badly in the 1990 recession. Since 1996, the model captures the general upward trend in output growth, although the exact strength of such growth comes as a bit of a surprise to the model.

Figure 1.c shows the model's forecasts for the rate on federal funds, expressed at a quarterly rate. At the year-ahead horizon, the model's forecasts track the general trend in the funds rate fairly well. One notable exception occurs just prior to the 1981-82 recession, when the model predicted a lower funds rate and the actual rate rose. The model also misses badly during

the recession, but this likely reflects the initial large error, noted above, combined with the serial correlation inherent in the year-ahead forecast errors. In contrast, it is interesting to note that the model correctly predicts a lower funds rate just prior to and throughout the 1990-91 recession. Since late 1995, the model has correctly predicted an unchanged rate on federal funds.

The model does a relatively poor job of predicting growth in real import prices, as shown in Figure 1.d. This reflects the fact that import prices are relatively volatile and that the model's specification is designed to predict output, inflation, and the rate of unemployment, not import prices. In future work, I plan to experiment with alternative ways to model import prices.

Figure 1.e shows that the model does a good job of tracking the unemployment rate. Of particular interest is the lack of any tendency to greatly overpredict the unemployment rate at a year-ahead horizon in recent years, although systematic, but small, year-ahead overpredictions, are evident. For the 1990s as a whole, the mean forecast error for the year-ahead, four-quarter-average unemployment rate is 0.02 percentage points.

Figure 1.f documents historically large one-year-ahead real M2 growth forecast errors in the early 1990s; more recently, however, the errors have moderated, suggesting a return of M2 to its historical relationship with its key macroeconomic determinants. Mehra (1997), using a much more sophisticated M2 equation, reaches similar conclusions.

Figure 1.g shows the model's forecasts for the rate on 10-year Treasury bonds, expressed at a quarterly rate. Not surprisingly, the results are qualitatively similar to those for the rate on federal funds: the model produces relatively large forecast errors around the 1981-82 recession and relatively small errors over the period since late 1995.

Figures 2.a to 2.c plot the implied forecast errors for inflation, real GDP growth, and the rate of unemployment. Most interesting about these figures is that the forecast errors in recent years do not seem unusually large relative to the historical record.

As an aid in assessing the model's quantitative performance, Table 3 records the root-mean-square-errors associated with the four-quarter-average forecasts, as described in equations (4)-(5), for the core variables in the model. These variables are the rate of inflation () PGDP), real GDP growth () RGDP), the unemployment rate (U), and the federal funds rate (RFF). Column two, labeled Model 1, shows the one- and two-year-ahead RMSEs produced by the preferred specification over various sample periods: the second half of the 1970s, the 1980s, the 1990s, and the total period (second half of the 1970s through the 1990s). All RMSEs are expressed in percentage points and, for) PGDP,) RGDP, and RFF, annualized. Although the table is provided mainly for use in comparing forecasts generated by the present model with those generated by alternative models, a couple of features are noteworthy. First, based on a comparison of RMSEs across variables, the unemployment rate is relatively easy to forecast, with RMSEs in the range of 0.4 to 0.6 and 0.5 to 0.9 for the one- and two-year-ahead horizons, respectively. In contrast, real GDP growth and the federal funds rate are relatively harder to forecast. Second, the 1990s' RMSEs tend to be lower than those of the 1970s and 1980s, suggesting that the 1990s have been relatively easier to forecast than the 1970s and 1980s.

Analyzing the Role of Departures From the Litterman Specification

There is a great deal of uncertainty in theoretical and empirical macroeconomics about the number of unit roots in individual time series and about the number of unit roots or common trends in multivariate systems. From the empirical perspective, it is well known that unit root tests have low power against local alternatives; less well known is that they may even have low power against not-so-local alternatives [Rudebusch (1993)]. From the theoretical perspective, some real business cycle models suggest one unit root in real variables, like real GDP, while, as noted by Jones (1995), endogenous growth models suggest that real GDP may have two unit roots.

In principle, Litterman's Bayesian approach is designed to handle such uncertainty; however, in specifying the baseline model, I simply imposed a single unit root on all the variables and one cointegrating relation (six common trends) on the system. It is natural to ask how these choices affect the model's forecasts. This section considers whether it is inappropriate to impose the unit root constraint and the RFF/RTB cointegrating relation on the system under study. The next section considers whether an additional cointegrating relation--represented by a stationary real rate--is appropriate. Since the relative price of imports is not usually included in VARs, it is also of interest to examine whether its inclusion improves the forecasts for the model's core variables.

To address these issues in a forecasting context, I estimated and simulated the following variants to the baseline model:

Model 2: BVAR(6){RGDP,) PGDP, RFF, RPIM, U, RM2, RTB},

Model 3: BVAR(5){) RGDP,) ^{2}PGDP,) RFF,) RPIM,) U,) RM2,) RTB},

Model 4: BVEC(5){) RGDP,) ^{2}PGDP,) RFF,) U,) RM2,) RTB}.

Model 2 removes the differencing that appears in the baseline model's variables. In this variant, each variable in the system is thought to be I(1), and the Bayesian priors are set to reflect this view. Because the system, Bayesian considerations aside, is no longer in error corrections form, I label the model BVAR rather than BVEC. It should be noted, however, that because the variables appear in "levels," cointegration between RFF and RTB is not ruled out; rather, it is just not imposed. Similarly, although unit root priors are assigned in estimating the model, the associated standard deviations are not zero and thus the data are permitted to modify the prior. An additional lag is added to the model, and all hyperparameters are as set in the baseline. This model is designed to test for the importance of differencing. Model 3 expresses all variables in first-difference form, but, in contrast to the baseline model, drops the spread from the right-hand side and thus rules out cointegration between RFF and RTB. The absence of the possibility of cointegration defines this model as a BVAR, and this specification is designed to test for the importance of the error correction term, RFF-RTB, in the baseline model. The final model, model 4, is identical to the baseline model with the exception that it drops RPIM from the system and is designed to assess the importance of import prices, particularly for the inflation forecasts.

The root-mean-square-errors associated with these additional models are reported in Table 3 under the columns labeled model 2, model 3, and model 4. Model 1 is the preferred specification. As shown in the table, the results concerning the issue of differencing are fairly conclusive: regardless of the variable, forecast horizon, or sample period, failure to difference (model 2) leads to a substantial deterioration in forecast performance, as measured by a higher RMSE. Particularly noteworthy are the "total" results for inflation and real GDP growth at the two-year horizon: the inflation RMSE rises from 1.5 to 2.4 and that for real GDP growth from

2.5 to 3.0. On the basis of these results, I conclude that in this model it is important to difference the data and then apply the appropriately modified Litterman baseline priors. It is encouraging to note that these findings confirm those reported by Webb (1995) and, more recently, by Christoffersen and Diebold (1997) that proper specification of the order of integration is important.

Concerning the role of the error corrections spread between the federal funds rate and the 10-year Treasury bond rate, the results are less striking; the spread does, however, appear useful in reducing unemployment-rate forecast errors. A comparison of the results for model 1 (error corrections spread included) and model 3 (excluded) shows that the RMSE for the one-year-ahead RGDP forecasts is a bit lower, on average, when the spread is included but that the reverse is true for the two-year-ahead RMSE. No such tradeoff exists for the unemployment rate. For this variable, including the spread in the specification leads to an average reduction in the two-year-ahead RMSE of about 0.3 percentage points, from 1.1 in model 3 to 0.8 in model 1, and almost no change in the year-ahead RMSE. A reasonable conclusion is that the RFF/RTB cointegrating relation and associated error corrections spread play a useful role in the model.

The results suggest that including the price of imports does improve the inflation forecasts. A comparison of the columns labeled model 1 (import prices included) and model 4 (excluded) suggests that the improvement is on the order of 0.2 percentage points in the RMSE at both the one- and two-year horizons. On average, the other forecasts are not much affected by including import prices in the model.

A Comparison With the Anderson-Hoffman-Rasche VEC Forecasting Model

Recently, researchers at the Federal Reserve Bank of St. Louis have estimated a forecasting model that is similar in many respects to the present BVEC model. Anderson, Hoffman, and Rasche (AHR, 1998) specify a vector error corrections forecasting model consisting of real GDP, two measures of inflation (based on the CPI and the GDP implicit price deflator), real M1, and the rates on federal funds and 10-year Treasury bonds. Four cointegrating relations are imposed on their system. First, as in the present model, AHR assume that the rates on federal funds and 10-year Treasury bonds are cointegrated. In addition, the two interest rates are thought to be cointegrated with inflation, each with a cointegrating vector equal to (1,-1). This implies that each real rate of interest is a stationary variable and, therefore, that the model incorporates a long-run Fisher relation. Third, the two inflation rates are thought to be cointegrated with a cointegrating vector of (1,-1), implying the two move one-for-one in the long run. Finally, M1 velocity is thought to be cointegrated with interest rates. AHR interpret this last relation as a long-run M1 demand function.

There are several important differences between the BVEC model and the AHR VEC model. First, the AHR model imposes two additional cointegrating relations that imply: a stationary real rate (a Fisher relation) and a long-run money demand function. Second, AHR estimate their model with non-Bayesian methods. Third, the two models contain a slightly different set of variables and, hence, forecasts from the two are based on different information sets: the BVEC adds the unemployment rate but eliminates the CPI measure of inflation; AHR measure GDP prices by the GDP implicit deflator. The BVEC, in contrast, uses the official, chain-weighted version of this concept. Fourth, AHR use real M1 rather than real M2. Fifth, AHR estimate over a sample that begins in 1956Q1, three years earlier than the starting point of the data in my sample. Finally, AHR include dummy variables at various points in their analysis and model specification and exclude observations from 1979Q4 to 1981Q4 prior to estimation. The present model makes no use of dummy variables, and no observations are excluded in estimating the model.

Although some of these differences may complicate any direct comparison of forecasts from the two models, I find it useful, nevertheless, to make such comparisons--as a check on the reasonableness of the forecasts from the present model. The experiment for the comparison is as follows. Following the discussion in AHR (p. 13), I estimate the BVEC through 1987Q4 and use the resulting parameter estimates to generate a sequence of one- through eight-quarter-ahead quarterly forecasts over the period 1988Q1 to 1995Q4. Starting with an information set that includes the 1987Q4 observation, I forecast 1988Q1 to 1989Q4. The information set is then updated to include 1988Q1, and a new set of forecasts for 1988Q2 to 1990Q1 is generated, without reestimating the model. Thus, the procedure is similar to the one described earlier except there is now no reestimation prior to generating a new sequence of eight-step-ahead forecasts. There are 32 one-step-ahead forecasts, 31 two-step-ahead forecasts, etc., from which to compute RMSEs, just as in AHR.

In Table 4, I compare RMSEs for the one-, two-, four-, and eight-step-ahead forecasts from the two models for the following variables: the log level of real GDP (RGDP), the log first difference of real GDP () RGDP), the log first difference of the GDP measure of the price level () PGDP), and the two interest rates (RFF, the federal funds rate; and RTB, the rate on 10-year

Treasury bonds).[5] I adopt AHR's convention concerning the units in which to report the RMSEs, as indicated in the table. The table shows that both models produce RMSEs of roughly the same magnitude, although the BVEC's RMSEs do appear a bit smaller. Whether this reflects inherent differences in model performance, the particular forecast period or horizons analyzed, or differences in the underlying data used to estimate the models is unknown at this point; it is, thus, inappropriate to draw any firm conclusions about the relative merits of the two models.

To gain additional perspective on how the BVEC tracks at a one-quarter horizon, Figure 3 plots the one-quarter-ahead forecasts, corresponding historical values, and the implied forecast errors (defined as actual minus forecast) for the five variables analyzed in Table 4, plus the unemployment rate. As documented earlier, panel B shows that the model's forecasts for real GDP growth miss badly in the 1990 recession; however, the inflation forecasts (panel C) track nicely and capture the apparent downward shift in trend inflation after the recession. The model also does a nice job of tracking Fed behavior (panel E) in the recession.

Dynamic Multipliers in the BVEC

An examination of a model's dynamic multipliers is an important part of any model evaluation process. Substantial doubt about a model's ability to produce reasonable forecasts exists when such a model possesses multipliers that are at odds with standard theory. Accordingly, this section considers the BVEC's response to two shocks, a monetary policy shock and a shock to the import price variable. I take the federal funds rate as the measure of monetary

[5]The RMSEs for the AHR VEC model are taken from their Table 7 on page 29.

policy and, thus, the BVEC equation for the federal funds rate is viewed as a Fed reaction function in this exercise.

Unfortunately, it is not possible to compute these multipliers--or impulse responses--from a reduced-form model, such as the BVEC. Such computations require an identified structural model. Modern time series analysis has suggested several--all relatively controversial--ways to identify a VAR. In this paper, I follow Sims (1980) in using a Cholesky decomposition of the BVEC's variance matrix of innovations to achieve statistical identification. This method of identification imposes a classical, lower triangular (or recursive) structure on the contemporaneous relationships among the variables in the model. Although many different recursive structures are possible, each yielding a potentially different set of impulse responses, I impose the following ordering to demonstrate one way in which the BVEC responds to shocks:) RGDP, $)^{2}$PGDP,) RPIM,) U,) RTB,) RFF,) RM2. This ordering corresponds closely to the baseline ordering in Christiano, Eichenbaum and Evans (1996). Because RFF, the federal funds rate, appears second to last in this ordering, the Fed is assumed to adjust policy in response to contemporaneous movements in real GDP, the price level, import prices, the unemployment rate, and the 10-year rate but not to contemporaneous movements in real M2. For the same reason, real GDP, the price level, import prices, the unemployment rate and the 10-year rate are assumed not to respond contemporaneously to a monetary policy shock.

Figures 4 and 5 show the response of the BVEC to unit shocks to the model's orthogonalized innovations in the federal funds rate and import price equations. This represents a positive 0.6 percentage point shock to the federal funds rate (Figure 4) and a positive 6.5 percentage point shock to the real price of imports (Figure 5). All percentages are expressed in annualized percentage-point form. In Figure 4, I also plot the responses for selected variables that are implied by the well known Macroeconomic Advisers, LLC (MA) structural macroeconomic model.[6] These are plotted with dashed lines.

Figure 4 shows that, on average, a positive 0.6 percentage point shock to the federal funds rate produces a path for the rate over the following four years that lies above a hypothetical no-shock baseline path (panel E). Thus, the model estimates that, on average, a monetary policy shock tends to be persistent. In general, both models yield responses for output growth (panel A), the inflation rate (panel B), and the unemployment rate (panel C) that are consistent with many macroeconomists' beliefs about the effects of a monetary policy tightening: output growth slows, the inflation rate falls, and the unemployment rate rises. Indeed, these responses accord with Christiano, Eichenbaum and Evans' (1996) monetary policy "facts." It is interesting to note that the MA model (dashed lines) produces a much larger effect on the unemployment rate, and, consequently, on the inflation rate. The BVEC inflation responses are consistent with the conventional wisdom that it takes about a year and a half before a monetary policy tightening affects the inflation rate.

Figure 5 shows the effect of a positive 6.5 percentage point (annual rate) shock to the real price of imports. On average, such a shock reduces output growth (panel A) and raises inflation

---

[6]These responses, based on the March 1998 ("803") model, were computed as follows. First, I computed and recorded the MA baseline forecasts for the variables of interest. Then, I exogenized the federal funds rate and add-factored the level of the federal funds rate over 16 quarters to produce a new path for that rate such that the quarter-by-quarter differences between the new path and the baseline path equaled the sequence of impulse responses for the level of the federal funds rate produced by the BVEC in panel E of Figure 4. I then resimulated the MA model, recorded the new forecasts, and subtracted the baseline forecasts from the new forecasts. The results of this subtraction are plotted with dashed lines.

(panel B) and the unemployment rate (panel C). The model indicates that, historically, the Fed has responded to such a disturbance by tightening policy (panel E). It is interesting to note that inflation and unemployment are estimated to move in the same direction in response to an import price shock. Given the recent resurgence of interest in discovering Phillips-curve correlations in U.S. data [e.g., King and Watson (1994) and Fuhrer (1995)], these results point to the importance of controlling for additional influences, like import prices, that could bias the estimated inflation-unemployment relationship.

4. Can the BVEC's Forecasting Performance Be Improved Upon?

In this final section, I ask whether there are any obvious ways to improve upon the forecasting performance of the BVEC given in equation (3). After respecifying the basic model in numerous ways and conducting additional forecasting experiments, I find no such obvious improvement.

Early Bayesian VAR users often included selected components of aggregate demand, such as some of the components of investment spending, in their specifications. In this section, I examine (alternatively) whether the addition of the logs of real business fixed investment (F), personal consumption expenditures (C), both personal consumption expenditures (C) and stock prices (S), or total government consumption and gross investment (G) helps to improve upon the performance of the BVEC.

The existence of a cointegrating relation between interest rates and the inflation rate, implying that the real rate is stationary and that a long-run Fisher effect exists, is an unsettled issue in macroeconometrics. As noted earlier, Anderson, Hoffman and Rasche (1998) find that such a relation is useful in their model, as do Shapiro and Watson (1988) and Gali (1992). Yet, King and Watson (1997), working with bivariate systems, find evidence against a long-run Fisher effect and King, Plosser, Stock, and Watson (1991) opt to not impose the Fisher relation on their baseline specification. My own evaluation of this literature is that the real rate is probably not stationary and, thus, there may not be a Fisher effect. Based on this evaluation, I did not impose real rate stationarity on the BVEC. The second part of this section discusses how the forecasting results are affected when real rate stationarity is imposed.[7]

As a final check on the BVEC specification, I ask how much the "Bayesian part" of the model buys in forecasting performance. If the Bayesian priors have no effect on the model's forecast performance, the model can be reestimated by conventional--and easier--methods.

Aggregate Demand Components

Many macroeconomic theories suggest that disturbances to individual components of aggregate demand can have important dynamic effects on output, inflation, and the rate of unemployment. For instance, Keynesian theories state that autonomous movements in consumption and investment expenditures can affect output with a multiplier that exceeds unity. Keynesian, monetarist [e.g., Andersen and Jordan (1968)], and real business cycle [e.g.,

---

[7]There is a third, intermediate, possibility concerning the relationship between the Fisher relation and the stationarity of the real rate of interest: The real rate may be nonstationary even when the Fisher effect holds. Such a possibility exists in any model in which a steady state increase in the rate of inflation is accompanied by a one-for-one increase in nominal interest rates and in which the model allows one or more additional shocks to have a permanent effect on the real rate. Imposing such a restriction on the present model is possible--but complicated, because it implies a zero restriction on certain elements of the vector of long-run multipliers for the nominal interest rate.

Christiano and Eichenbaum (1992) and Baxter and King (1993)] models allow for the possibility that changes in fiscal policy, including changes in government purchases, affect future output. In specifying the baseline model, I eliminated these theories from consideration by imposing zero restrictions. However, the Bayesian approach allows one to incorporate such theories into the forecasting model while still maintaining the prior that the theories do not matter for forecasting purposes.

This section reports the results of adding aggregate demand components, one by one, to the BVEC model in equation (3). In each case the additional component enters the model in log first difference form, and the forecast evaluation methodology described in section 3 is repeated. Four forecasting experiments are conducted, one each for personal consumption expenditures, personal consumption expenditures plus equity prices (measured by the log first difference of the S&P 500 index), business fixed investment, and total government consumption and gross investment.

In each case, the additional variable's lags enter on the right-hand side of each equation of the BVEC (3), and an additional equation explaining that variable is added to the model. The Bayesian priors on the coefficients attached to the original variables in each of the original equations are unchanged. Priors on the coefficients attached to the additional variable and in the additional equation vary on a case-by-case basis but were chosen in a manner that most macroeconomists would find unobjectionable. The general principle employed was to use very loose priors (an T of 0.9) on the additional variable's coefficients in the output equation, thus giving the data a good deal of flexibility to detect an influence running from the additional variable to output. In general, the additional variable's coefficients were assigned relatively tight priors (for instance, an T– 0.5) in all other equations, with the exception of those on the coefficients attached to the own lags in the own equation (T=1.0). Priors on the coefficients in the additional (own) equation were set in a similarly reasonable manner. For example, in the own equation for consumption expenditures, I assigned relatively loose priors to the coefficients on lagged output and interest rates, thus allowing the data some flexibility to detect a relationship running from output and interest rates to consumption. The priors in the stock-price equation were set relatively tight to reflect the idea that the S&P 500 index is, with a high degree of confidence, very close to a random walk with drift and that no variables in the information set Granger-cause stock prices. Of course, different priors could change the results, but preliminary experimentation suggests the results are robust.

Root-mean-square-errors for the preferred baseline model (Model 1) and the respecified models are displayed in Table 5 for the total period analyzed, the mid-1970s through 1997Q4.[8] Table 5 shows that the addition of various components of aggregate demand to the preferred model yields RMSEs for the core variables that are virtually identical to those generated in the baseline model. For example, the RMSEs for four-quarter-average output growth at one- and two-year horizons are about 1.82 percent and 2.50 percent, respectively, regardless of the modification tested. For the four-quarter-average inflation rate, the RMSEs are about 0.93 percent and 1.52 percent, again, regardless of the modification. These results suggest little reason to increase the size of the model for the purpose of increasing the accuracy of the forecasts for the core variables.

---

[8]The subperiod results (mid-1970s, 1980s, and 1990s) are similar to those for the total period presented here.

Of course, analysts may have an interest in the BVEC's ability to forecast some of the components of aggregate demand. The results in Table 5 suggest that one could add these components to the model without changing the model's essential structure. Using the model that includes personal consumption expenditures as an example, Figure 6 shows the sequence of four-quarter-average, one- and two-year-ahead forecasts (dashed line) for growth in personal consumption expenditures and the corresponding historical values (solid line).

As was the case with the real GDP-growth forecasts (Figure 1.b), the consumption-growth forecasts track fairly well at the one-year horizon over the 1980s' recessions but miss badly in the 1990 recession. Since 1994, there is an unusually close correspondence between forecasted and actual consumption growth at the one-year horizon.

Real Rate Stationarity

To examine the effect of not imposing a stationary real rate on the BVEC, I respecified the model in (3) to include one lag of the level of $RFF - \Delta PGDP$. That is, an additional error corrections term was added to yield a modified version of the model. A typical equation in the new system is given by

$$
\begin{aligned}
\Delta RGDP_t = c_1 &+ \beta_{11}(L)\Delta RGDP_{t-1} + \beta_{12}(L)\Delta^2 PGDP_{t-1} + \beta_{13}(L)\Delta RFF_{t-1} \\
&+ \beta_{14}(L)\Delta RPIM_{t-1} + \beta_{15}(L)\Delta U_{t-1} + \beta_{16}(L)\Delta RM2_{t-1} \\
&+ \beta_{17}(L)\Delta RTB_{t-1} + \rho_1(RFF_{t-1} - RTB_{t-1}) + \rho_2(RFF_{t-1} - \Delta PGDP_{t-1}) + e_{1t}\,,
\end{aligned}
$$

where $\rho_2(RFF_{t-1} - \Delta PGDP_{t-1})$ is the new, error corrections term.

The results of forecasting this model, presented in the third column of Table 6 for the period mid-1970s to 1997Q4, indicate that adding a stationary real rate produces a higher inflation RMSE, particularly at the two-year horizon, where the RMSE rises from 1.52 percent in the preferred model to 2.36 percent in the model that includes the real rate restriction. The RMSEs for the other variables, however, appear roughly unaffected. An examination of the subsample results (not reported) indicates that the increase in inflation RMSE is attributable to the new model's relatively poor forecast performance in the 1980s; inspection of a time series plot of the two-year-ahead inflation forecasts (not shown) points to relatively large overpredictions in the early 1980s as the culprit.[9]

Bayesian Priors

The last column of Table 6 shows how forecast performance is affected by removing the Bayesian priors. To generate the RMSEs in this column, I reestimated the preferred model using OLS. The results confirm the motivation provided at the beginning of this paper--that Bayesian priors can help to improve the forecast performance of VARs. A comparison of columns two and four indicates that the RMSEs are higher--in some cases, much higher--without the priors. Concentrating on the one-year-ahead results, the RMSEs rise from 0.93 percent to 1.35 percent for inflation, 1.82 percent to 2.60 percent for output growth, and from 0.52 percent to 0.72

---

[9]Following up on some recent research by Mehra (1996), I also examined forecasts from a model that imposes real rate stationarity by adding one lag of RFF-) $PGDP_4/4$, where) $PGDP_4$ is the four-quarter-average inflation rate, defined by $PGDP_t$ - $PGDP_{t-4}$, and RFF is divided by 400. (Mehra used a three-year average of past inflation to define his "real rate.") This adjustment improves the inflation-forecast performance of the stationary real rate specification, but such performance remains worse than that of the preferred specification.

percent for the rate of unemployment. The subsample results (not reported) indicate that the improvement provided by the Bayesian approach extends across all samples but that the relative differences are most pronounced in the mid-1970s and 1980s.

A reasonable assessment of the results in this section is that the baseline BVEC fares well against some obvious alternative specifications. Based on the RMSE criterion, the addition to the BVEC of several aggregate demand components yields no improvement in performance for any of the core variables. Imposing an additional cointegrating relation in the form of a stationary real rate of interest yields substantially worse inflation forecasts, and eliminating the Bayesian part of the specification leads to an all-around deterioration in performance.

5. Summary

This paper presents a small-scale macroeconometric time series model that is useful for generating short-term forecasts for U.S. real GDP, inflation, and the rate of unemployment. The model--which is specified in the Bayesian vector error corrections framework that has been suggested by others--includes real GDP, the GDP price index, the rate of unemployment, the rates on federal funds and 10-year Treasury securities, real M2, and the relative price of imports.

The out-of-sample, dynamic forecasting performance of the model is documented over various sample periods starting with the mid-1970s and extending through 1997Q4. The baseline model is then subject to several modifications in an attempt to discover whether its forecast performance can be improved upon. I find the baseline specification works best.

Although further testing on the baseline is planned, I believe the basic model holds promise as a workhorse, short-term forecasting tool.

Table 1. Estimated Coefficients, Model 1, 1960Q4 to 1997Q4

	) RGDP	**) ^{2}PGDP**	**) RFF**	**) RPIM**	**) U**	**) RM2**	**) RTB**
) RGDP(t-1)	0.0109	-0.0109	0.0269	0.0982	-0.0331	0.0828	0.0095
) RGDP(t-2)	0.0417	0.0002	0.0253	0.0721	**-0.0391**	0.0547	0.0025
) RGDP(t-3)	-0.0462	0.0125	-0.0019	-0.0356	-0.0170	0.0082	-0.0004
) RGDP(t-4)	0.0263	0.0061	-0.0143	0.1133	-0.0020	-0.0373	-0.0005
) RGDP(t-5)	-0.0346	-0.0096	**0.0264**	0.0801	0.0133	-0.0153	-0.0024
) ^{2}PGDP(t-1)	0.0152	**-0.2771**	-0.0052	0.0564	0.0043	-0.0136	-0.0265
) ^{2}PGDP(t-2)	0.0000	**-0.2370**	**0.1288**	0.0061	0.0002	0.0023	0.0173
) ^{2}PGDP(t-3)	0.0005	-0.0606	0.0476	0.0149	0.0002	-0.0020	0.0051
) ^{2}PGDP(t-4)	0.0013	0.0055	-0.0033	-0.0142	-0.0008	-0.0005	-0.0016
) ^{2}PGDP(t-5)	0.0008	-0.0238	-0.0023	0.0012	0.0001	0.0035	0.0033
) RFF(t-1)	**0.5563**	0.0263	0.0699	0.0228	**-0.1817**	**-0.3784**	-0.0209
) RFF(t-2)	-0.3733	-0.0139	**-0.1435**	0.0132	0.0283	-0.0951	-0.0644
) RFF(t-3)	0.2845	-0.0140	0.0755	-0.0018	-0.0445	-0.0728	0.0155
) RFF(t-4)	-0.0602	0.0086	0.0363	0.0061	-0.0084	0.1378	0.0190
) RFF(t-5)	-0.0657	0.0293	**0.1592**	-0.0050	-0.0161	**-0.3152**	0.0461
) RPIM(t-1)	-0.0177	**0.0321**	**0.0303**	**0.4998**	-0.0055	-0.0044	0.0055
) RPIM(t-2)	0.0009	**0.0191**	-0.0047	-0.0357	0.0036	-0.0014	-0.0005
) RPIM(t-3)	0.0003	0.0057	-0.0014	**0.1222**	0.0058	-0.0006	0.0012
) RPIM(t-4)	0.0083	-0.0047	-0.0024	0.0151	0.0031	0.0001	0.0012
) RPIM(t-5)	0.0248	-0.0052	0.0075	-0.0879	-0.0094	0.0003	0.0042
) U(t-1)	-0.2772	**-0.2383**	**-0.1848**	-0.0503	**0.3565**	**0.4178**	-0.0334
) U(t-2)	0.1341	0.1027	0.0177	-0.0227	-0.0398	0.0258	-0.0215

) U(t-3)	0.0980	-0.0168	-0.0111	0.0057	-0.0350	-0.0049	-0.0080
) U(t-4)	0.1005	**-0.0952**	0.0237	-0.0014	-0.0644	0.0563	0.0055

Table 1

(cont.)

	) RGDP	**) ^{2}PGDP**	**) RFF**	**) RPIM**	**) U**	**) RM2**	**) RTB**
) U(t-5)	-0.0889	-0.0375	0.0265	-0.0014	0.0543	0.0396	0.0107
) RM2(t-1)	0.1159	0.0328	0.0067	-0.0031	-0.0191	**0.5200**	0.0007
) RM2(t-2)	**0.1572**	0.0094	0.0027	0.0068	-0.0163	0.0415	0.0004
) RM2(t-3)	0.0515	0.0009	0.0056	0.0038	-0.0098	0.0961	-0.0000
) RM2(t-4)	0.0147	-0.0092	-0.0037	0.0068	0.0043	-0.0008	-0.0002
) RM2(t-5)	0.0649	-0.0051	-0.0039	0.0031	-0.0039	0.0643	-0.0000
) RTB(t-1)	-0.0573	-0.1022	**0.2989**	0.0152	0.2377	**-1.8497**	**0.1572**
) RTB(t-2)	-0.2264	-0.0034	**-0.3683**	-0.0080	0.1098	**0.7751**	-0.0686
) RTB(t-3)	0.5700	0.0135	**0.2215**	-0.0209	-0.0579	0.1415	0.0603
) RTB(t-4)	-0.0851	-0.0204	0.0431	-0.0148	0.1325	0.1223	-0.0458
) RTB(t-5)	-0.4228	-0.0224	-0.0322	-0.0140	**0.2117**	0.3500	**-0.1327**
const.	**0.0040**	-0.0001	**-0.0007**	-0.0018	**0.0013**	0.0011	0.0001
EC(t-1)	**-0.6116**	0.0337	**-0.1058**	0.5930	**0.2258**	-0.0600	**0.0618**
R^2	0.43	0.44	0.50	0.44	0.61	0.71	0.25
F	0.0070	0.0021	0.0019	0.0163	0.0021	0.0052	0.0011

Table 2. Estimated Error Correction Coefficients, Various Sample Endpoints

	) RGDP	**) ^{2}PGDP**	**) RFF**	**) RPIM**	**) U**	**) RM2**	**) RTB**
1979Q4	**-1.0707**	0.1225	-0.1116	**1.3431**	**0.3289**	**-0.8979**	0.0464
1984Q4	**-0.9172**	0.0418	-0.1117	0.6233	**0.2497**	**-0.7418**	**0.0694**
1989Q4	**-0.8957**	0.0490	-0.1087	0.5952	**0.2504**	**-0.5048**	0.0590
1994Q4	**-0.6556**	0.0306	**-0.1031**	0.6351	**0.2359**	-0.0405	0.0470
1997Q4*	**-0.6116**	0.0337	**-0.1058**	0.5930	**0.2258**	-0.0600	**0.0618**

*: This row is also reported in Table 1.

Table 3. RMSEs for Four-Quarter Averages in Annualized Percentage Points

	Model 1	**Model 2**	**Model 3**	**Model 4**
) PGDP (inflation)				
1-year ahead				
2nd-half of 1970s	1.331	1.293	1.381	1.487
1980s	0.939	1.444	0.864	1.239
1990s	0.612	0.703	0.590	0.593
Total	0.932	1.198	0.908	1.113
2-year ahead				
2nd-half of 1970s	1.848	2.320	2.158	1.816
1980s	1.845	2.988	1.737	2.138
1990s	0.720	1.375	0.750	0.695
Total	1.523	2.395	1.531	1.684
) RGDP (real GDP growth)				
1-year ahead				
2nd-half of 1970s	1.580	1.977	1.842	1.705
1980s	1.880	2.579	2.150	1.802
1990s	1.864	2.729	1.786	1.917
Total	1.820	2.532	1.968	1.826
2-year ahead				
2nd-half of 1970s	2.742	4.937	1.839	2.697
1980s	2.701	2.652	2.552	2.647
1990s	2.054	2.446	1.818	2.024
Total	2.484	3.048	2.197	2.440

Table 3 (cont.)

	Model 1	**Model 2**	**Model 3**	**Model 4**
U (unemployment rate)				
1-year ahead				
2nd-half of 1970s	0.588	0.846	0.722	0.552
1980s	0.572	0.761	0.607	0.547
1990s	0.394	0.569	0.471	0.387
Total	0.518	0.716	0.587	0.496
2-year ahead				
2nd-half of 1970s	0.536	1.077	1.206	0.495
1980s	0.841	1.391	1.021	0.801
1990s	0.850	1.472	1.180	0.843
Total	0.805	1.380	1.113	0.779
RFF (federal funds rate)				
1-year ahead				
2nd-half of 1970s	1.382	2.309	1.489	1.296
1980s	2.313	2.875	2.295	2.354
1990s	0.734	1.010	0.831	0.719
Total	1.721	2.258	1.743	1.731
2-year ahead				
2nd-half of 1970s	2.150	3.963	2.496	2.084
1980s	4.195	5.794	4.284	4.320
1990s	1.925	2.914	2.264	1.966
Total	3.223	4.626	3.394	3.302

Model 1 (preferred): BVEC(5){) RGDP,) ^{2}PGDP,) RFF,) RPIM,) U,) RM2,) RTB}
Model 2 (levels) : BVAR(6){RGDP,) PGDP, RFF, RPIM, U, RM2, RTB}
Model 3 (no EC) : BVAR(5){) RGDP,) ^{2}PGDP,) RFF,) RPIM,) U,) RM2,) RTB}
Model 4 (no RPIM) : BVEC(5){) RGDP,) ^{2}PGDP,) RFF,) U,) RM2,) RTB}

Table 4. RMSE Comparison With AHR Model, Units Indicated, 88Q1 to 95Q4

	AHR VEC	**BVEC (Model 1)**
RGDP (percent of actual)		
1-step	0.008	0.007
2-steps	0.016	0.012
4-steps	0.035	0.023
8-steps	0.065	0.043
) RGDP (percent, qtrly rate)		
1-step	0.008	0.007
2-steps	0.009	0.007
4-steps	0.012	0.007
8-steps	0.011	0.008
) PGDP (pct. pts., ann. rate)		
1-step	1.59	0.76
2-steps	1.43	0.99
4-steps	1.29	0.86
8-steps	1.33	0.85
RFF (pct. pts., ann. rate)		
1-step	1.14	0.51
2-steps	2.09	0.86
4-steps	2.36	1.35
8-steps	2.02	2.49
RTB (pct. pts., ann. rate)		
1-step	0.44	0.45
2-steps	0.75	0.70
4-steps	1.32	1.00
8-steps	1.87	1.21

Table 5. RMSEs for Variable Additions to the BVEC, Four-Quarter Averages in Annualized Percentage Points, Mid-1970s to 1997Q4

	Model 1 (preferred)	**Add:** **) C**	**Add:** **) C&) S**	**Add:** **) F**	**Add:** **) G**
) PGDP (inflation)					
1-year ahead	0.932	0.928	0.928	0.932	0.933
2-year ahead	1.523	1.518	1.522	1.519	1.524
) RGDP (real GDP growth)					
1-year ahead	1.820	1.821	1.820	1.860	1.826
2-year ahead	2.484	2.499	2.484	2.501	2.489
U (unemployment rate)					
1-year ahead	0.518	0.516	0.511	0.521	0.519
2-year ahead	0.805	0.802	0.798	0.813	0.811
RFF (federal funds rate)					
1-year ahead	1.721	1.719	1.710	1.719	1.718
2-year ahead	3.223	3.218	3.206	3.215	3.216

Table 6. RMSEs for Real-Rate Stationarity and Non-Bayesian Variants to the Baseline Model, Four-Quarter Averages in Percentage Points, Mid-1970s to 1997Q4

	Model 1 (preferred)	**Stationary Real Rate**	**No Bayesian Priors**
) PGDP (inflation)			
1-year ahead	0.932	1.242	1.349
2-year ahead	1.523	2.359	2.378
) RGDP (real GDP growth)			
1-year ahead	1.820	1.776	2.597
2-year ahead	2.484	2.546	5.069
U (unemployment rate)			
1-year ahead	0.518	0.556	0.723
2-year ahead	0.805	0.908	1.129
RFF (federal funds rate)			
1-year ahead	1.721	1.725	2.796
2-year ahead	3.223	3.051	4.437

References

Andersen, Leonall C. and Jerry L. Jordan. "Monetary and Fiscal Actions: A Test of Their Relative Importance in Economic Stabilization," Federal Reserve Bank of St. Louis *Review* 50 (November 1968), pp. 11 - 24.

Anderson, Richard G., Dennis L. Hoffman, and Robert H. Rasche. "A Vector Error Correction Forecasting Model of the U.S. Economy," Federal Reserve Bank of St. Louis, Unpublished Manuscript, (April, 1998).

Artis, M.J. and Zhang, W. "BVAR Forecasts for the G-7," *International Journal of Forecasting* 6 (1990), pp. 349 - 362.

Baxter, Marianne and Robert G. King. "Fiscal Policy in General Equilibrium," *American Economic Review* 83 (June 1993), pp. 315 - 334.

Bikker, J.A. "Inflation Forecasting for Aggregates of the EU-7 and EU-14 with Bayesian VAR Models," *Journal of Forecasting* 17 (1998), pp. 147 - 165.

Carlson, John B. and Susan M. Byrne. "Recent Behavior of Velocity: Alternative Measures of Money," Federal Reserve Bank of Cleveland *Economic Review* 28 (Quarter 1, 1992), pp. 2 - 10.

Christiano, Lawrence J. and Martin Eichenbaum. "Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations," *American Economic Review* 82 (June 1992), pp. 430 - 450.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. "Identification and the Effects of Monetary Policy Shocks," in *Financial Factors in Economic Stabilization and Growth*, edited by Mario I. Blejer, Zvi Eckstein, Zvi Hercowitz, and Leonardo Leiderman, Cambridge University Press, 1996.

Christoffersen, Peter F. and Francis X. Diebold. "Cointegration and Long-Horizon Forecasting," Federal Reserve Bank of Philadelphia Working Paper 97-14, September 1997.

Darin, Robert and Robert L. Hetzel. "A Shift-Adjusted M2 Indicator for Monetary Policy," Federal Reserve Bank of Richmond *Economic Quarterly* 80 (Summer 1994), pp. 25 - 47.

Davis, E.P. and S.G.B. Henry. "The Use of Financial Spreads as Indicator Variables: Evidence for the United Kingdom and Germany," IMF *Staff Papers* 41 (September 1994), pp. 517 - 525.

Davis, E. Phillip and Gabriel Fagan. "Are Financial Spreads Useful Indicators of Future Inflation and Output Growth in EU Countries?" *Journal of Applied Econometrics* 12 (1997), pp. 701 - 714.

Dua, Pami and David J. Smyth. "Forecasting US Home Sales Using BVAR Models and Survey Data on Households' Buying Attitudes for Homes," *Journal of Forecasting* 14 (1995), pp.217 - 227.

Engle, Robert F. and C.W.J. Granger. "Co-Integration and Error Correction: Representation, Estimation, and Testing," *Econometrica* 55 (March 1987), pp. 251 - 276.

Fuhrer, Jeffrey C. "The Phillips Curve Is Alive and Well," Federal Reserve Bank of Boston *New England Economic Review* (March/April 1995), pp. 41 - 56.

Gali, Jordi. "How Well Does the IS-LM Model Fit Postwar U.S. Data?" *Quarterly Journal of Economics* CVII (May 1992), pp. 709 - 738.

Hamilton, James D. *Time Series Analysis*, Princeton, New Jersey: Princeton University Press, 1994.

Johansen, Soren and Katarina Juselius. "Maximum Likelihood Estimation and Inference on Cointegration--With Applications to the Demand for Money," *Oxford Bulletin of Economics and Statistics* 52 (1990), pp. 169 - 210.

Jones, Charles I. "Time Series Tests of Endogenous Growth Models," *Quarterly Journal of Economics* CX (May 1995), pp. 495 - 525.

Joutz, Frederick L., G.S. Maddala, and Robert P. Trost. "An Integrated Bayesian Vector Autoregression and Error Correction Model for Forecasting Electricity Consumption and Prices," *Journal of Forecasting* 14 (1995), pp.287 - 310.

King, Robert G., Charles I. Plosser, James H. Stock, and Mark W. Watson. "Stochastic Trends and Economic Fluctuations," *American Economic Review* 81 (September 1991), pp. 819 - 840.

King, Robert G. and Mark W. Watson. "The Post-War U.S. Phillips Curve: A Revisionist Econometric History," *Carnegie-Rochester Conference Series on Public Policy* 41 (1994), pp. 157 - 219.

King, Robert G. and Mark W. Watson. "Testing Long-Run Neutrality," Federal Reserve Bank of Richmond *Economic Quarterly* 83 (1997), pp. 69 - 101.

Laurent, Robert D. "An Interest Rate-Based Indicator of Monetary Policy," Federal Reserve Bank of Chicago *Economic Perspectives* 12 (January/February 1988), pp. 3-14.

LeSage, James P. "A Comparison of the Forecasting Ability of ECM and VAR Models," *Review of Economics and Statistics* (November 1990), pp. 664 - 671.

Litterman, Robert B. "Forecasting With Bayesian Vector Autoregressions--Five Years of Experience," *Journal of Business & Economic Statistics* 4 (January 1986), pp.25 - 38.

Mehra, Yash P. "Monetary Policy and Long-Term Interest Rates," Federal Reserve Bank of Richmond *Economic Quarterly* 82 (Summer 1996), pp. 27 - 49.

Mehra, Yash P. "A Review of the Recent Behavior of M2 Demand," Federal Reserve Bank of Richmond *Economic Quarterly* 83 (Summer 1997), pp. 27 - 43.

Rudebusch, Glenn D. "The Uncertain Unit Root in Real GDP," *American Economic Review* 83 (March 1993), pp. 264 - 272.

Shapiro, Matthew D. and Mark W. Watson. "Sources of Business Cycle Fluctuations," NBER *Macroeconomics Annual* (1988), pp. 111 - 148.

Sims, Christopher A. "Macroeconomics and Reality," *Econometrica* 48 (January 1980), pp. 1 - 48.

Theil, H. "On the Use of Incomplete Prior Information in Regression Analysis," *Journal of The American Statistical Association* 58 (June 1963), pp. 401 - 414.

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

Trehan, Bharat. "Predicting Contemporaneous Output," Federal Reserve Bank of San Francisco *Economic Review* 2 (1992), pp. 3 - 12.

Webb, Roy, H. "Forecasts of Inflation from VAR Models," *Journal of Forecasting* 14 (1995), pp. 267 - 285.