


Federal Reserve Bank of Minneapolis

Quarterly Review

Fall 1984



*Above-Average National Growth
in 1985 and 1986*

Robert B. Litterman (p. 3)

*More Growth Ahead
for Ninth District States*

Hossain Amirizadeh
Richard M. Todd (p. 8)

*Improving Economic Forecasting
With Bayesian Vector Autoregression*

Richard M. Todd (p. 18)

*Forecasting and Policy Analysis With
Bayesian Vector Autoregression Models*

Robert B. Litterman (p. 30)

Federal Reserve Bank of Minneapolis

Quarterly Review Vol. 8, No. 4 ISSN 0271-5287

This publication primarily presents economic research aimed at improving policymaking by the Federal Reserve System and other governmental authorities.

Produced in the Research Department. Edited by Preston J. Miller, Kathleen S. Rolfe, and Inga Velde. Graphic design by Phil Swenson and typesetting by Barbara Cahlander and Terri Desormey, Graphic Services Department.

Address requests for additional copies to the Research Department, Federal Reserve Bank, Minneapolis, Minnesota 55480.

Articles may be reprinted if the source is credited and the Research Department is provided with copies of reprints.

The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Forecasting and Policy Analysis With Bayesian Vector Autoregression Models

Robert B. Litterman

Senior Economist

Research Department

Federal Reserve Bank of Minneapolis

At the Federal Reserve Bank of Minneapolis, the development of Bayesian vector autoregression (BVAR) models has been motivated primarily by a desire to have as accurate as possible forecasts of current and future economic conditions for use in the analysis of current monetary policy. (See Litterman 1984 and the paper by Todd in this issue.) Of course, all economic forecasts involve a great deal of uncertainty. BVAR models, however, can provide reliable measures of the uncertainty of their forecasts. In addition, the models can generate probabilities about the likelihood of future events. The national monthly BVAR model we've developed has also been designed to project the likely impacts of unexpected shocks to the economy or of changes in Federal Reserve policy.¹

This paper discusses how the two types of forecasts generated by our BVAR model—unconditional (or *baseline*) forecasts and conditional forecasts—influence the process of monetary policy analysis. The first type, baseline forecasts, predicts current and future economic conditions, given that the economy behaves as it has in the past. The second type, conditional forecasts, can be used to answer “what if” questions—questions that ask about the implications for the economy if a particular condition is imposed. The discussion of conditional forecasts is divided into two parts: conditional forecasts that assume no explicit policy change and those that assume an explicit policy change. (For more detailed discussion of conditional forecasting and policy analysis with BVAR models, see Doan, Litterman, and Sims 1984 and Sims 1982.)

Unconditional (or Baseline) Forecasts

A brief account of the way monetary policy is formulated reveals the importance of the baseline forecast to the process. The policy formulation process includes three

basic steps: first, policymakers reach agreement on a baseline forecast; second, they consider a set of feasible alternatives to the forecast; third, they choose the policy stance associated with the alternative that best suits the goals of monetary policy—price stability and sustainable growth. From this simplified account, we can see that preparing a baseline forecast of the likely future course of the economy is the first and probably most important step in formulating the appropriate stance of monetary policy at a given point in time.

Once prepared, the baseline forecast influences subsequent policy formulation in at least two important ways. First, it provides the context for considering alternative policy options. For example, the lower the rate of future inflation predicted in the forecast, the harder it is to support an argument that a tighter monetary policy (one with lower rates of money growth) is needed to reduce future inflation. Similarly, if the forecast suggests low unemployment in the future, then it is harder to defend the position that a more accommodative policy (one with higher rates of money growth) is necessary to reduce the future level of unemployment.²

¹The national monthly BVAR model we discuss here is the same used to generate the forecasts in the Litterman national outlook paper in this issue. A technical appendix describing the model in detail is available on request to the Research Department, Federal Reserve Bank of Minneapolis.

²When policymakers disagree on the proper stance of policy, their differences can frequently be traced not to disagreement on the goals of policy but rather to different views about the current economic outlook. For example, at the December 1983 Open Market Committee meeting, Mr. Martin (FR Board, 1984a, p. 120) dissented from the policy action taken because he felt it would “present a threat to the sustainability of the economic expansion.” In March 1984, Messrs. Gramley and Wallich (FR Board, 1984b, p. 514) dissented from the policy action taken because in their view “the strength of the economic expansion warranted more restraint now”; in contrast, Mr. Martin (p. 514) felt the policy was too tight “in light of the vulnerability of key sectors of the economy to rising interest rates.”

A second way the baseline forecast influences policy formulation is by acting as a benchmark against which alternative scenarios or policy options are compared. Some assumption about future policy will be embedded in the baseline forecast—usually an assumption of no change in current policy. Alternative assumptions are then generally considered in terms of their impacts on the baseline forecast. In this way, the forecast also plays a key role in defining the range of policy options. For if policymakers generally agree on the direction and magnitude of the effects a change in policy stance will generate, then agreement on the baseline forecast together with the effects of these changes defines the set of feasible choices. Thus, the role of the baseline forecast is crucial because a different baseline forecast will lead to a different set of feasible outcomes for the economy, and the preferred policy from this set is likely to be different from the preferred policy chosen from a different feasible set.

Measures of Uncertainty

Unlike standard econometric models, which rely on judgmental adjustment of their forecasts, BVAR models use a statistical approach that requires no such adjusting. A major benefit of this approach is that it provides more information about our forecasts by allowing us to generate scientific measures of their uncertainty. In many respects, these measures are as important as the forecasts themselves. They can be used to produce *confidence bands*—bands which enclose a region of outcomes around the forecast and indicate the probability that actual outcomes will fall within the region. (For examples of baseline forecasts with confidence bands, see Charts 1–3 of the Litterman national outlook paper in this issue.)

The confidence bands surrounding our model's baseline forecasts are 70 percent wide: this means that actual values can be expected to fall outside these bands almost one-third of the time. Readers not familiar with the uncertainties of macroeconomic forecasting may be surprised by how wide these bands are. As measures of uncertainty, however, they give policymakers important information. An example where measures of uncertainty provide important information is the problem of forecasting state revenues. (For a detailed discussion of this problem, see Litterman and Supel 1983.) Because policymakers in state government must trade off the costs of excessive taxation with the costs of possible revenue shortfalls, quantification of the size of likely shortfalls is a key piece of information in deciding how large a reserve fund to maintain. The degree to which standard models are in-

adequate in providing information of this kind is apparent when we try to make sense out of the alternative scenario analyses which commercial forecasting services often provide as their measures of uncertainty.

Probabilities of Future Events

More than just providing measures of uncertainty, the BVAR model provides a realistic, multivariate probability description of the possible future paths of the economy. What this means is that the model can provide the answer to any question about the probabilities of future events. For example, policymakers are often interested in knowing how likely it is that a particular event, such as a recession, will occur over a given time period. All such questions can be answered by running the model through a large number of simulations and counting the proportion of times the event of interest occurs. (The results of such an experiment in describing the probability of a U.S. recession appear in the Litterman national outlook paper in this issue.)

Why can't other models be used in the same way to answer questions about the probabilities of events? Actually, they could—if we could trust their results. The key difference between BVAR models and standard ones in this respect is that BVAR models do not require judgmental adjustment in order to generate reasonable results. Just as we can't trust standard models' forecasts or their measures of uncertainty because they are not accurate unless subjectively adjusted, we can't trust the answers that such models would generate to these types of questions.

Cost is another factor that often prevents standard models from being used to answer questions about the probabilities of future events. Running a sufficiently large number of simulation experiments can be prohibitively expensive, especially on standard models. Granted, running simulation experiments on our complete monthly BVAR model tends to be expensive as well; smaller BVAR models, however, are relatively inexpensive to simulate. So for many purposes we might be willing to look just at the results from simulating a smaller model, such as the core sector of our model. For example, one such BVAR model, whose forecast accuracy has proved competitive with standard models, can be simulated using a microcomputer. (This model is described in Litterman 1984.)

Beyond the Baseline

The uses of our BVAR model discussed so far have been founded on the baseline forecasts generated by our model. While these uses certainly are an important part of the

policy analysis process, it may be necessary to go further. Often we are interested in knowing what the implications are if some condition is satisfied. Because these “what if” questions require that the stated condition be imposed, they must use conditional forecasts, rather than baseline forecasts, to generate answers. There are, however, two kinds of conditional questions, each requiring very different kinds of assumptions before the model can generate conditional forecasts that provide satisfactory answers.

Conditional Forecasts I: No Explicit Policy Assumptions

The first kind of conditional question is one which asks only about the implications of some event occurring. For example, What are the implications for the economy if stock market prices rise by 10 percent? This kind of question is relatively easy to answer, at least in the sense that it requires no explicit additional policy assumptions, merely a mathematical manipulation (known as *taking a conditional expectation*) of the probability structure inherent in the baseline forecast. In particular, if the question concerns the effect of an event today on the forecast, the answer is given directly by a certain representation of the model’s structure, called the *impulse response functions*. The set of response functions (the *moving average representation* of the model) is defined as the responses of the model to a particular set of one-time shocks. This representation provides a complete description of the dynamic structure of the model which, when depicted graphically, is much more comprehensible than the usual representation given in terms of the coefficients of the model’s equations.

The responses of several of the model’s variables to our hypothetical shock to stock prices are shown in Charts 1–4. These responses can be interpreted in two ways. First, we can think of them as representing how the system would respond if it were at rest (that is, if the current and past values of all system variables were zero) and then were subjected to a one-time shock to stock prices. Second, we can interpret the response functions as the change in a forecast that would occur as a result of a one-time surprise change in the shocked variable.³ The responses shown indicate that, as a result of the stock price shock, there is an initial surge in growth—inflation-adjusted gross national product (real GNP)—that dwindles by the end of two years. The growth surge is associated with a sustained reduction in unemployment and is also accompanied by rising inflation.

In Table 1, we show how the response functions can be added to the baseline forecast to give a direct answer to the original question. Thus, the forecast, conditional on a rise in stock prices, is just the baseline forecast plus the response functions, aggregated and expressed as quarterly growth rates.

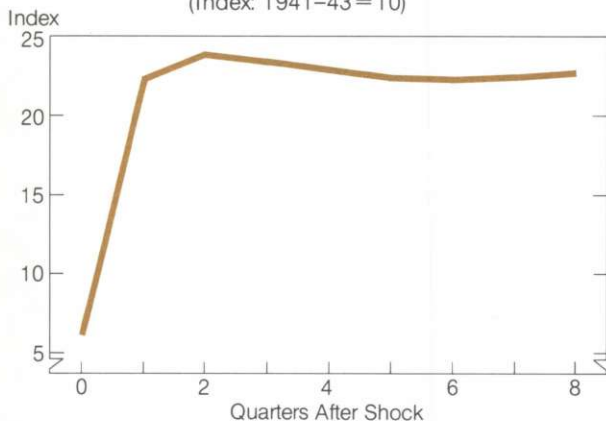
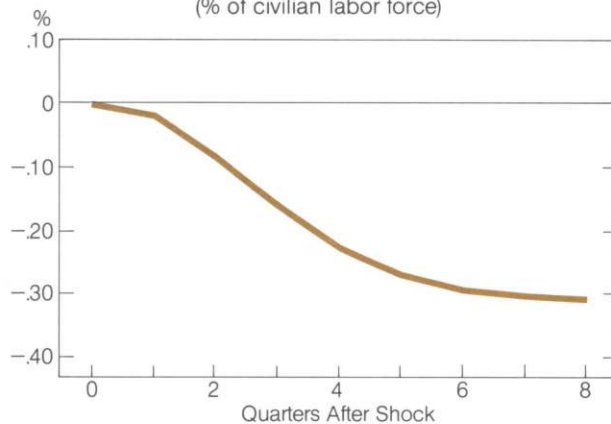
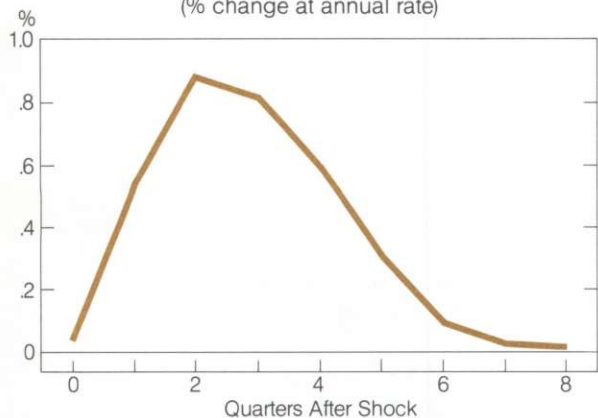
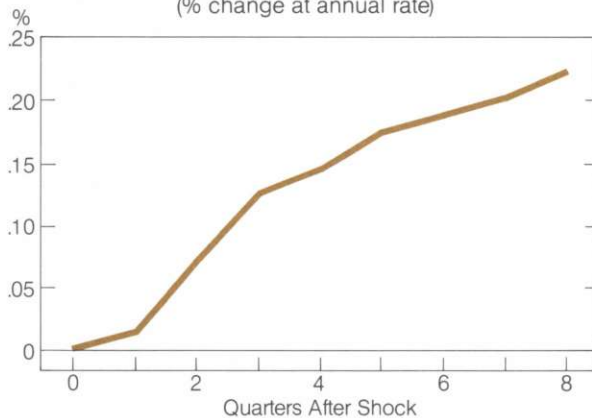
Defining a Shock

When we refer to a stock price shock, or to any other type of “what if” condition, it is important that we be extremely precise about what we mean by that condition. In particular, when we think of a shock to stock prices, do we mean to take into account those other changes in variables usually associated with stock price changes? In the responses shown in Charts 1–4, we have considered an increase in stock prices that has no effect in the current period on the values of any other variables in the model. This definition of the shock makes sense if we think of the stock price movement as being caused by something not expected to have any contemporaneous impact on the other variables in the system. Of course, such an assumption may be highly suspect, especially if stock price shocks historically have been strongly correlated with some other variables. This method of defining what we mean by a shock, however, is just one possibility.

Another way to define a shock is by choosing an ordering of the model’s variables and then including as part of the shock those components of movements in variables ordered after the shocked variable. For example, if stock prices were said to be first in an ordering, we would include as part of the stock price shock a contemporaneous movement in each remaining variable in the model. The size of these components of the shock would be based on the way unexpected movements (called *innovations*) in stock prices have varied historically in relation to contemporaneous innovations in each other variable. This definition of the shock to stock prices makes sense if we want to act as if we know nothing about what caused the

³In our BVAR model, stock prices are modeled as being very close to a *random walk*—an assumption that today’s value is the best forecast of a variable’s future values. In continuous time, the response function of such a variable looks like a step function. The response function shown here, however, follows a very different shape for two reasons. First, we use monthly averaged data. In this case, an increase in the variable implies that its value at the end of the month is most likely to be higher than the monthly averaged value; this, in turn, implies that there will be a further increase in the most likely value for the following month. Second, we take quarterly averages of the monthly averaged data. In this case, we model a shock in the third month of the quarter. Thus, the shock to stock prices shows up only to a small extent in the quarter in which it occurred and is followed by a large increase in subsequent quarters, as shown in Chart 1.

Charts 1-4

**Responses of Selected Variables
to a 10 Percent Rise in Stock Prices****Chart 1 Stock Prices**
(Index: 1941-43=10)**Chart 3 Unemployment Rate**
(% of civilian labor force)**Chart 2 Growth (Real GNP)**
(% change at annual rate)**Chart 4 Inflation (GNP Deflator)**
(% change at annual rate)

Sources of basic data: Standard & Poor's 500-Stock Index; U.S. Departments of Commerce and Labor

current stock price shock and nothing about the contemporaneous movements in any of the other variables. Ultimately, the right way of defining the shock to answer the original question depends on precisely what we mean by the condition "a rise in stock market prices."

A Variation

The answer to the stock price question was fairly simple because the condition only involved the effect of one

particular shock to the system. Other conditional forecasts of this type are often slightly more complicated because the condition implied in the question is more complicated. For example, we might ask, What is the implication for the economy if stock market prices rise by 10 percent over the next twelve months? In this question, the condition does not define a single shock. What we can do to answer the question, though, is to generate the forecast which satisfies the condition and is the most likely, given the probability

Table 1
Projected Effect of a 10 Percent Rise in Stock Prices*

Actual Data, 1st Quarter 1983–3rd Quarter 1984;
 Projected Values, 4th Quarter 1984–4th Quarter 1986

Period	On Stock Prices (% Change)			On Real GNP (% Change)			On Unemployment Rate (% of Civilian Labor Force)			On Inflation (GNP Deflator) (% Change)		
	Actual Data or Baseline Forecast	Response Function	Conditional Forecast	Actual Data or Baseline Forecast	Response Function	Conditional Forecast	Actual Data or Baseline Forecast	Response Function	Conditional Forecast	Actual Data or Baseline Forecast	Response Function	Conditional Forecast
Quarterly**												
1983:1	36.06%	—	—	3.31%	—	—	10.37%	—	—	5.04%	—	—
1983:2	47.56	—	—	9.38	—	—	10.10	—	—	2.62	—	—
1983:3	7.01	—	—	6.81	—	—	9.40	—	—	3.10	—	—
1983:4	0.57	—	—	5.95	—	—	8.47	—	—	4.37	—	—
1984:1	-12.38	—	—	10.09	—	—	7.87	—	—	4.42	—	—
1984:2	-11.00	—	—	7.10	—	—	7.47	—	—	3.34	—	—
1984:3	12.85	—	—	1.89	—	—	7.47	—	—	3.72	—	—
1984:4†	14.68	16.93%	31.60%	-0.95	0.04%	-0.91%	7.39	0.00%	7.39%	4.27	0.00%	4.27%
1985:1	0.98	44.37	45.35	2.38	0.54	2.92	7.39	-0.02	7.37	2.89	0.02	2.90
1985:2	2.76	3.00	5.75	3.54	0.88	4.42	7.39	-0.08	7.30	3.40	0.07	3.48
1985:3	3.61	-1.77	1.84	4.41	0.81	5.22	7.35	-0.16	7.19	3.39	0.12	3.52
1985:4	2.54	-1.02	1.51	4.48	0.60	5.07	7.35	-0.23	7.08	3.00	0.15	3.15
1986:1	3.29	-1.57	1.72	4.07	0.31	4.38	7.25	-0.27	6.98	3.30	0.17	3.47
1986:2	3.37	-0.62	2.74	3.82	0.09	3.91	7.20	-0.30	6.91	3.20	0.19	3.39
1986:3	3.33	-0.06	3.27	3.74	0.02	3.76	7.18	-0.31	6.87	3.18	0.20	3.38
1986:4	3.30	0.14	3.44	3.77	0.01	3.78	7.17	-0.31	6.87	3.32	0.22	3.54
4th Quarter Over 4th Quarter												
1984:4/1983:4	0.23	3.51	3.74	4.44	0.01	4.45	-1.07	0.00	-1.07	3.94	0.00	3.94
1985:4/1984:4	2.46	9.81	12.28	3.70	0.71	4.41	-0.09	-0.23	-0.32	3.17	0.09	3.26
1986:4/1985:4	3.32	-0.53	2.79	3.85	0.11	3.96	-0.13	-0.08	-0.22	3.25	0.20	3.45

*The conditional forecast, indicated by shading, shows the effect of the stock price shock. The forecast is determined by adding the impulse response functions to the model's baseline forecast. Due to rounding, the sum may not precisely equal its components.

**Percent changes are at annual rates.

† The hypothetical shock to stock prices occurs in December 1984 of this quarter.

Sources of basic data: Standard & Poor's 500-Stock Index; U.S. Departments of Commerce and Labor

structure defined by the model. Another way to put this— and this is the way we proceed computationally—is to ask, **What is the most likely set of innovations** (which formally means the smallest set where the measure of size, or *metric*, is defined by the covariance matrix of innovations) which, when input as shocks to the model, leads to a path that satisfies the condition? (The details of this necessary computation are given in Doan, Litterman, and Sims 1984.)

Conditional Forecasts II: With An Explicit Policy Assumption

A much more difficult problem arises when the conditional question either directly or indirectly requires an explicit assumption about policy. Suppose the question is, **What are the implications for the economy if the value of the dollar drops by 10 percent?** We could answer this question, just as we did the earlier question about stock

prices, by defining a particular shock and using the impulse response functions to generate a conditional expectation. Such an answer is apt to be much less satisfactory in this case, however, because the value of the dollar both responds to and is a concern of monetary policy.

Suppose we are asking the question because we are concerned that unless the Federal Reserve responds to a fall in the dollar, the drop will cause an increase in future inflation as imports become more expensive and domestic competitors raise their prices. If we define the depreciation in the dollar as last in an ordering of variables (so that a shock to the dollar includes no components of other variables) and then generate the conditional expectation, our answer includes not only a definition of the structure of the shock but also an implicit policy assumption: it assumes that everything in the economy, including monetary policy, responds in the most likely way. In this case,

without defining monetary policy, we still have included it in the forecast.

This conditional expectation assuming policy as usual is shown in Charts 5–10. First, we show the hypothetical shock to the value of the dollar. Next, we indicate the response of interest rates (the 3-month Treasury bill rate) to the dollar depreciation; this response shows that the dollar depreciation can be expected to lead to higher interest rates. Then we show the effect of the depreciation on four other variables. The overall effect of the dollar shock on growth (real GNP) is rather small. There is, however, stronger growth in the money supply (M1) as a result of the shock. The dollar shock has virtually no impact on stock prices because we assume that the dollar depreciation is not associated with any contemporaneous change in other variables. Stock prices are found to behave such that if the dollar shock has no immediate impact, there is not likely to be a significant delayed effect on stock prices either. Although the net effect of the dollar shock on expected real growth is very small, it does cause a large increase in the inflation (GNP deflator) forecast.

What if the original question about the effect of the dollar depreciation were motivated by a concern over what happens if the Fed does not respond as usual to the fall in the dollar? Would the conditional expectation of the effect of the dollar depreciation, assuming policy as usual, be a satisfactory answer? No, not very. The problem with the answer is that it builds in the usual policy response, whereas we're worried about what would happen if this usual policy were not followed. Thus, as with the previous question about stock prices, we have to be very careful about what we really mean with our question. One way we might make our question more precise is by asking, What is the likely effect of a 10 percent depreciation of the dollar if the Fed attempts to hold interest rates along the path generated by the baseline forecast, rather than responding as it has in the past to such a shock? (That is, What is the effect if the Fed tries to hold interest rates at the level they would have been if the dollar shock had not occurred?) Now, however, we have raised a more difficult question—a question which the BVAR model alone cannot answer. For we cannot project the effects of this new monetary policy until we add an explicit assumption about how monetary policy actions designed to hold interest rates along a given path would contemporaneously affect the other variables in the system.

Defining What We Mean By Policy

Answering a question about policy in the context of the

BVAR model thus requires us to define what we mean by policy in terms of the model's variables. As it turns out, we can answer this kind of policy question if we can define a type of shock that represents the contemporaneous impact of a one-time Fed action, such as an unexpected open market operation. Once we have defined that shock, we can then use the impulse response functions to map out the entire dynamic impact of this shock, which we will call a *Fed policy action*. We then define what we usually refer to as monetary policy as a *sequence of Fed policy actions*. Finally, by summing up the dynamic responses caused by a sequence of policy actions, we generate the response of the system due to the impact of that monetary policy.

Problems With Our Definition

We might feel that this definition of monetary policy as a sequence of Fed policy actions is problematic for a number of reasons. For one, it may not be easy to define a particular shock that represents the contemporaneous impact of a Fed policy action. We might suspect, for example, that the impact of a Fed action might depend on the circumstances in which it is taken. After all, the Fed does not directly control any of the variables that appear in our BVAR model. The Fed actually has a number of actions it can take in terms of open market operations, changing the discount rate, imposing credit controls, and so forth. It may not be realistic to assume that the contemporaneous impact of the underlying Fed actions on the observed variables in the model can always be modeled as a scalar multiple of the same linear combination of innovations. We might suspect, for example, that as the Fed takes stronger actions, the impact is not always linear in terms of the observed effects. Perhaps small actions show up only in interest rates, but large actions may have important impacts on the value of the dollar, stock prices, and other variables.

Another problem in our definition could arise because Fed actions can occur throughout a month, rather than at a particular point within a month. Even if the immediate impact of a Fed action is adequately represented as a scalar multiple of a vector of changes in variables on a daily or weekly basis, the impact of the action will generate a dynamic response; when represented in monthly averaged data, the effect will depend on at what point during the month the action took place. In general, the result will not be a scalar multiple of the same combination of components. Suppose, for example, that the Fed action in daily data is simply a change in interest rates, but that after a delay of two weeks the money supply also begins to

Charts 5–16

**Effect on Selected Variables
of a 10 Percent Depreciation in the Value of the Dollar***

1st Quarter–3rd Quarter 1984, Actual;
4th Quarter 1984–1986, Projected

Charts 5–10

I. No Explicit Policy Assumptions

- Actual
- Baseline Forecast
- Shocked Forecast—Policy as Usual

Charts 11–16

II. With Explicit Policy Assumption**

- Actual
- Shocked—Policy as Usual
- ▨▨▨▨▨ Shocked—Explicit Policy Assumption

Chart 5 Value of Trade-Weighted Dollar
(Index: March 1973=100)




Chart 11 Value of Trade-Weighted Dollar
(Index: March 1973=100)




Chart 6 Response of Interest Rates †
(3-Month T-Bill Rate) (%)




Chart 12 Response of Interest Rates †
(3-Month T-Bill Rate) (%)




Chart 7 Growth (Real GNP)
(% Change at Annual Rate)





Chart 13 Growth (Real GNP)
(% Change at Annual Rate)



I. No Explicit Policy Assumptions (cont)

II. With Explicit Policy Assumption (cont)**

Chart 8 Money Supply (M1)
(% Change at Annual Rate)




Chart 14 Money Supply (M1)
(% Change at Annual Rate)




Chart 9 Stock Prices (S&P 500)
(Index: 1941-43=10)




Chart 15 Stock Prices (S&P 500)
(Index: 1941-43=10)




Chart 10 Inflation (GNP Deflator)
(% Change at Annual Rate)





Chart 16 Inflation (GNP Deflator)
(% Change at Annual Rate)



*Effect of hypothetical shock occurring in December 1984

**The explicit policy assumption is that the Fed responds to the dollar depreciation by holding the interest rate to the predetermined path set by the baseline forecast.

† The response of interest rates is shown as the deviation from the baseline forecast.

Sources of basic data: Board of Governors of the Federal Reserve System, U.S. Departments of Commerce and Labor, Standard & Poor's 500-Stock Index

change. The Fed action in monthly data will then affect both components if the action occurs in the first half of the month, but only the interest rate component if it takes place in the last half.

Finally, and perhaps most importantly, we might worry about the *Lucas critique*, the criticism that the structure of the system may not be independent of the sequence of Fed policy actions taken. (See Lucas 1976.) If the behavior of agents in the economy depends on aspects of monetary policy which have remained stable over the model's historical sample, then when those aspects of policy change, agents' behavior can also be expected to change in ways that cannot be predicted by the model. The resulting change in the structure of the economy would mean that the impulse response functions, which describe the economy's response to policy, would change. Policies designed under the assumption that the responses will be the same as observed during some earlier period may therefore not generate the desired results.

On the one hand, none of these problems can be easily dismissed. They suggest just a few of the reasons why policy analysis is a very difficult task and why any policy analysis involves a large degree of uncertainty. On the other hand, there is no reason, in principle, why the policy responses predicted by our BVAR model would not be correct when the policy actions considered are similar to those observed in the historical period over which the model is estimated. Even though BVAR models do not explicitly include expectations variables, they can, for example, be expected to generate correct predictions as long as expectations are affected by policy in the same way as they have been in the past. And although it may be difficult to defend the identifying assumptions necessary in order to use BVAR models for policy analysis, it should be recognized that these problems are not unique to the BVAR modeling technique. To generate policy conclusions, all models require similar identifying assumptions, which are likely to be difficult to defend, and all are thus subject to a large degree of uncertainty. The reason we prefer to use BVAR models for policy analysis is that the most crucial assumption we make in using any model for policy analysis—and the one likeliest to be the most difficult to defend—is that the model incorporates a reasonable probability description of the economy. Because we believe the forecasting evidence suggests that BVAR models have a definite advantage in this respect, we feel more comfortable using this type of model for many kinds of policy analysis.

Still, we must keep in mind that all of our projections have large degrees of uncertainty and that the confidence bands we could generate around our conditional forecasts would be at least as wide as those discussed in the Litterman national outlook paper in this issue. Moreover, since no formal statistical procedure is used to generate our explicit policy assumptions, the confidence bands cannot take into account the additional uncertainty inherent in these assumptions.

Identifying the Effect of a Policy Action

Recognizing the uncertainties associated with any attempt to analyze policy with econometric models, we can proceed to attempt to answer the question we posed above: What would be the effect of a 10 percent depreciation of the dollar if the Fed tried to hold interest rates to a predetermined path—that is, the one generated by the baseline forecast? The first step is to define the vector of components that we will identify as the immediate effect of a policy action. For this purpose, in the context of our BVAR model, we will use a vector that has three nonzero elements; in particular, we will assume that Fed policy actions have an immediate impact on interest rates, the value of the dollar, and stock prices. By *immediate* we mean within the month that the Fed policy action takes place, recognizing that problems arise when we consider actions occurring at different times of the month. We are thus assuming that monetary policy takes at least one month to have any impact on real output, inflation, inventories, total credit, and the money supply.⁴

We can generate a rough guess of the appropriate sizes of the nonzero components by looking at how these variables have changed in recent years in response to identifiable Fed policy actions. Although it is never easy to quantify Fed policy actions, for the purposes of this exercise we have attempted to identify such actions since fall 1982. Since that time a very important indicator of Fed policy, the *federal funds rate* (the overnight rate banks pay each other for reserves), has behaved to some extent like a step function with some random fluctuation. By looking at a plot of the daily time path of the federal funds rate, we can identify approximately when changes in Fed policy have occurred. On Chart 17, we have identified twelve such

⁴Because this kind of assumption is difficult to be very confident about, it would make sense to check the robustness of the analysis to this assumption by making alternative assumptions and comparing the results. For the purpose of illustrating the method, however, we consider just this one possibility.

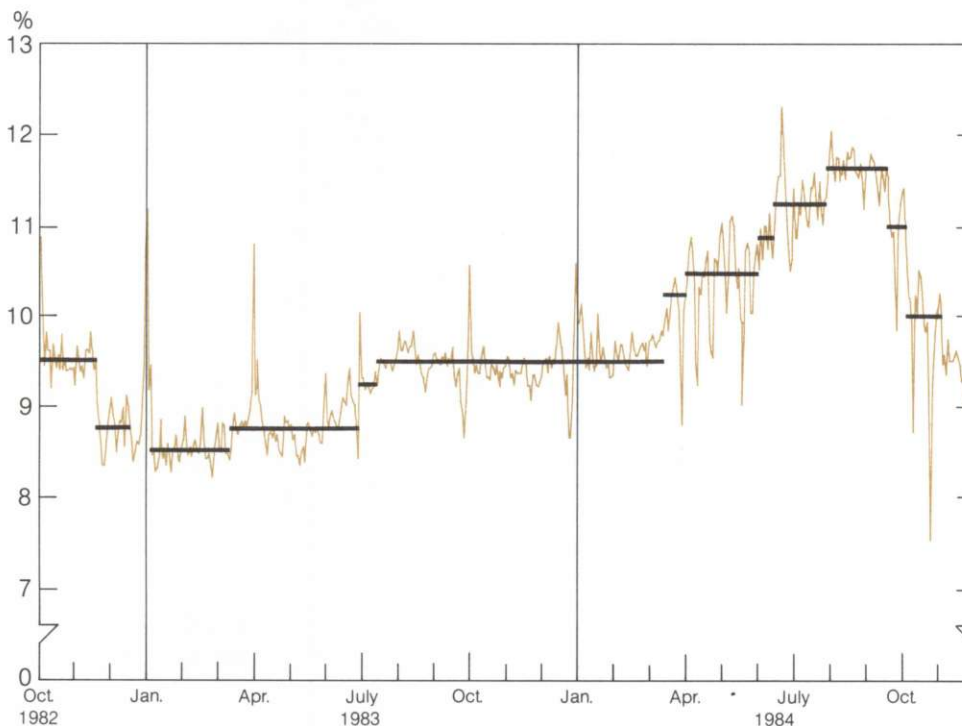
changes over this period.⁵ Once the dates of these changes are identified, we measure the average level of the funds rate before and after these dates and then measure the average levels of the nonzero elements (the 3-month Treasury bill rate, the value of the dollar, and stock prices).⁶ We then regress the measured changes in each of these variables on the given changes in the funds rate to generate estimates of the vector of effects of policy actions. This procedure generates an estimate that for each 1 percentage point increase in the federal funds rate due to a

⁵The dates of these changes are based solely on a subjective interpretation of the data: no attempt has been made to incorporate information on official actions of the Federal Open Market Committee or the Federal Reserve Board into their calculation. Some of the changes seem to take place instantaneously, others over a period of a few days. These policy changes are also much easier to identify after the fact than while they are occurring. The funds rate clearly exhibits a considerable amount of random fluctuation; only after the rate has fluctuated around a new level for several days does it become apparent that a policy change has taken place.

⁶We measure these nonzero elements during a week centered on the last day of the old level and during a week centered two weeks after the first day of the new level. We assume that a policy change becomes apparent, and its effects on these financial variables occur, during this period.

Chart 17
Daily Federal Funds Rate With Steps*

October 1, 1982–November 26, 1984



*The federal funds rate is the overnight rate that banks charge each other for reserves. The steps indicate Fed policy changes, based on the author's subjective interpretation of the data. In delineating these steps, no attempt has been made to incorporate information on official actions of the Federal Open Market Committee or the Federal Reserve Board.

Source of basic data: Board of Governors of the Federal Reserve System

policy action, we observe a 0.4 percentage point increase in the 3-month Treasury bill rate, a 1.4 percent increase in the value of the dollar, and a 1.1 percent decrease in the stock price index.⁷

Specifying a Sequence of Policy Actions

Having identified the contemporaneous effect of a single Fed policy action in the model's monthly variables, we next have to specify the sequence of policy actions we wish to use in this exercise. The question we have posed (phrased in terms of holding interest rates along a predetermined path following a dollar depreciation) is made operational by defining the policy action, in this context, to be a shock each period with a Treasury-bill component large enough to return the interest rate to the level generated in the baseline forecast. The sequence of Fed actions is then determined by shocking the system with a depreciation of the dollar and then following this shock by the sequence of Fed actions that causes the level of the Treasury bill rate to return to the baseline forecast. The policy actions will have contemporaneous impacts on the value of the dollar and stock prices, as estimated above, but will have delayed effects on all the other variables in the system.

The Results

The results of this exercise in modeling policy are shown in Charts 11–16. We can compare the results of this sequence of policy actions with the results in the earlier conditional forecast without an explicit policy assumption, seen previously in Charts 5–10. Since the conditional forecast with an implicit assumption of policy as usual predicts the interest rate would rise, we know that the effect of the sequence of policy actions is to lower interest rates. We also see that the sequence of policy actions leads to more rapid money growth than otherwise would have occurred. We can thus characterize the policy of holding interest rates along the path predetermined in the baseline forecast as being a more accommodative monetary policy than normally would have been expected. This more accommodative policy is modeled as a sequence of surprise changes in interest rates. We can see on Chart 12 that the policy action in the quarter following the shock is a lowering of interest rates.⁸ The lowering leads contemporaneously to higher stock prices as well as to a slight further decrease in the value of the dollar. The more accommodative policy response also leads to higher growth (real GNP), but only at the cost of a further increase in likely inflation rates, especially after a one-year lag. It is not hard

to predict that an unanticipated accommodation in monetary policy leads to more real growth and inflation, but the model allows us to quantify how much inflation is likely to occur and when it will appear—a quantification crucial in evaluating the tradeoffs involved in setting policy.

Notice that the sequence of policy actions and the response functions do not, of themselves, define the monetary policy in this exercise. What they do define is the necessary change in policy, relative to the implicit policy incorporated into the conditional forecast. The exercise thus illustrates how it is not necessary to disentangle current policy from the economic structure in order to generate policy conclusions. (Such identification, which is a part of traditional policy analysis, would require additional assumptions.)

Variations

We could, of course, ask more complicated policy-related questions such as, What policy-induced path for interest rates (or money growth) would be expected to lead to some other particular inflation result? In general, there would be more than one such path. In order to generate a unique sequence of policy actions consistent with a given degree of inflation, we would need some additional constraints. For example, we could ask, What would be the smallest (where smallest might be defined in terms of a minimum sum of squares metric) such sequence of policy actions?

In other cases, we may not want to ask about the effects of a particular policy. Instead we may want to ask, What is the optimal policy relative to a given loss function? (An example of this kind of analysis is described in Litterman 1982.) The technical details of solving an optimization problem make such an exercise appear much more complicated, but the basic procedure is really quite similar to our policy modeling exercise. In effect we use well-known mathematical techniques to find a function which at each point in time defines the policy action expected to minimize the loss function.

Summary

The formulation of monetary policy requires the ability to continually weigh and choose from among various policy

⁷While the magnitude of these responses is fairly sensitive to the two-week interval—the responses tend to increase for at least four weeks after the change in funds rate—the ratios of the three variables in the model are fairly stable. For the purpose of this exercise, only these ratios matter.

⁸In subsequent periods, however, the policy action may be either to lower or raise interest rates, depending on whether the forecast of rates at each period (in the absence of a policy action) is above or below the baseline target path.

alternatives. For this purpose it is important to have an econometric model that can generate accurate forecasts, project the range of likely outcomes, answer "what if" questions, and predict the effects of alternative policies. We have illustrated here how our BVAR model can be used to perform each of these tasks.

Although other types of models have traditionally been used to perform these tasks, they generally require significant judgmental adjustment in order to generate reasonable results. As a result, such models have not inspired much trust. The main virtue of the BVAR model is its ability to generate a realistic probability description of the future course of the economy—a description which does not require any adjustment. This virtue of the BVAR model makes it, we think, a useful tool for both forecasting and policy analysis.

References

- Doan, Thomas; Litterman, Robert; and Sims, Christopher. 1984. Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews* 3 (1): 1-100.
- Federal Reserve Board of Governors (FR Board). 1984a. Record of policy actions of the Federal Open Market Committee. *Federal Reserve Bulletin* 70 (February): 115-20.
- _____. 1984b. Record of policy actions of the Federal Open Market Committee. *Federal Reserve Bulletin* 70 (June): 509-14.
- Litterman, Robert B. 1982. Optimal control of the money supply. *Federal Reserve Bank of Minneapolis Quarterly Review* 6 (Fall): 1-9.
- _____. 1984. Forecasting with Bayesian vector autoregressions—Four years of experience. Paper presented at the American Statistical Association's Annual Meeting, sponsored by the Business and Economic Statistics Section and held at Philadelphia, Penn., August 12-14. Also, forthcoming in proceedings volume and as Research Department Staff Report, Federal Reserve Bank of Minneapolis.
- Litterman, Robert B., and Supel, Thomas M. 1983. Using vector autoregressions to measure the uncertainty in Minnesota's revenue forecasts. *Federal Reserve Bank of Minneapolis Quarterly Review* 7 (Spring): 10-22.
- Lucas, Robert E., Jr. 1976. Econometric policy evaluation: A critique. In *The Phillips curve and labor markets*, ed. Karl Brunner and Allan H. Meltzer. Carnegie-Rochester Conference Series on Public Policy 1: 19-46. Amsterdam: North-Holland.
- Sims, Christopher A. 1982. Policy analysis with econometric models. *Brookings Papers on Economic Activity* 1: 107-52.