Federal Reserve Bank of Minneapolis Quarterly Review

Winter 1980

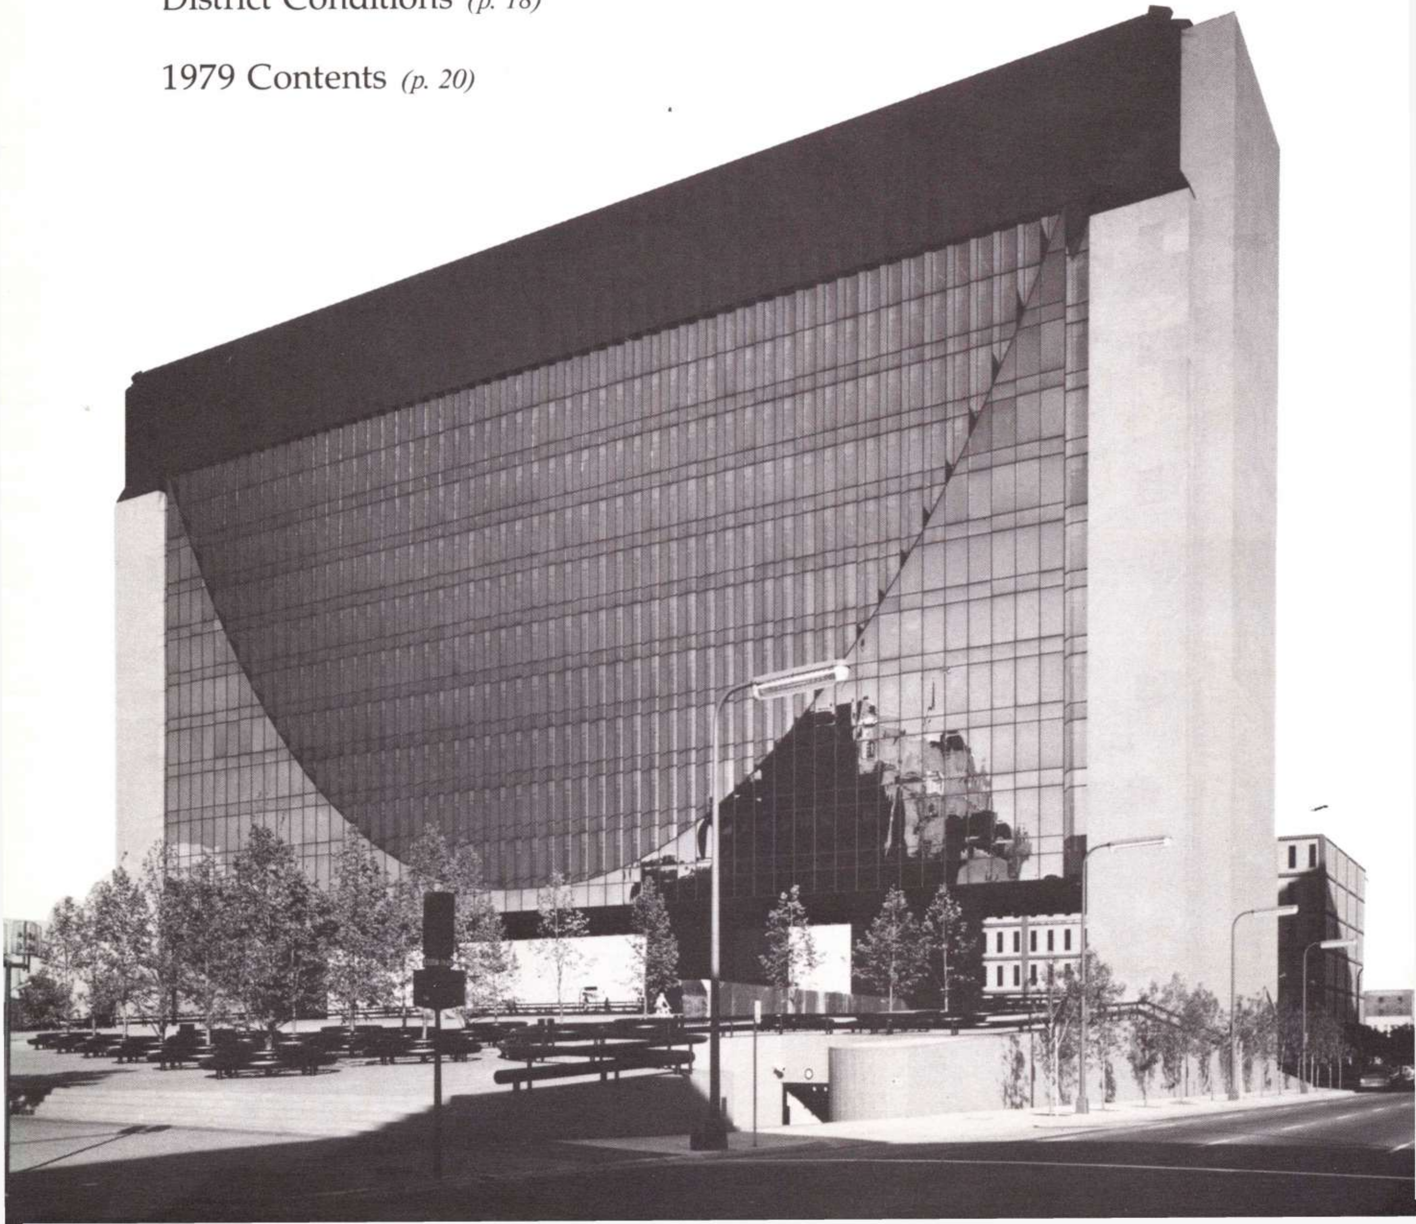

Federal Reserve Bank of Minneapolis
Quarterly Review Vol. 4, No. 1

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 Arthur J. Rolnick, Kathleen S. Rolfe, and Alan Struthers, Jr. Graphic design by Phil Swenson and charts drawn by Mary K. Steffenhagen, Graphic Services Department.

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

Estimating the Effects of the Oil-Price Shock

Preston J. Miller, *Assistant Vice President*
Thomas M. Supel, *Senior Economist*
Thomas H. Turner, *Economist*

Research Department
Federal Reserve Bank of Minneapolis

The analysis and forecast we presented in "The U.S. Economy in 1980" (earlier in this issue) were based on a methodology that took into account how consumer and business demands, production, wages and prices, and government economic policies responded to the oil-price shock. In this paper we describe this methodology in more technical detail. Specifically, we describe how we measured the impact of the oil-price shock on the economy in 1979 and on the outlook for 1980.

Our investigation resulted in three main observations. First, the oil-price shock was a major cause of the 1979 forecast errors. Second, actual monetary and fiscal policies in 1979 were tighter than historical relationships would imply, thus contributing to forecast errors in 1979 and to more pessimistic output forecasts for 1980. Third, some fundamental economic relationships underlying the structure of our model appear to have changed during the course of 1979.

The experiment we performed consisted of making and comparing three forecasts (see Table 1). We first used data through the fourth quarter of 1978 to generate a forecast, referred to as F1, for major economic variables over 1979 and 1980. We then used these data and a measure of actual oil prices through the first three quarters of 1979 to generate a second forecast of major variables, F2, a forecast we might have made at the end of 1978 had we known exactly what oil prices would be in 1979. Comparing the two forecasts allowed us to identify the effects of the oil-price shock—the differences between forecast and actual oil prices—on the outlook for the economy over each quarter of 1979 and 1980. Finally, we generated forecast F3, our model's best forecast for the fourth quarter of 1979 and all of 1980, by reestimating the model using complete information through the third quarter of 1979 (the most recent information then available) plus fourth-quarter information on interest rates.

Table 1

Forecasts for Estimating the Effects of the Oil-Price Shock

Forecast	Quarterly Data Period	Estimation Period
F1	1957:1 through 1978:4	1958:2 through 1978:4
F2	1957:1 through 1978:4 plus actual oil prices in 1979:1, 2, 3	1958:2 through 1978:4
F3	1957:1 through 1979:3 plus actual interest rates in 1979:4	1958:2 through 1979:3

The Vector Autoregression Model

Our method employs a quarterly vector autoregression (VAR) model for the national economy which gives prominence to the role of energy prices. As Sargent (1979) points out, VAR is a statistical tool, not based on economic theory, which can be successfully employed as a forecasting device.[1] The VAR model forecasts exclusively on the basis of the historical behavior of the variables over the sample period.

[1]The restrictions on the model we used were generated using statistical techniques developed in Litterman 1979. Sargent 1979, pp. 11–13, describes the role of these sorts of restrictions in using VARs for forecasting.

The model we used contains twelve variables divided into two classes: those assumed to be central in driving other variables in the system and those responding primarily to their own pasts and to the pasts of the driving variables. The four *driving variables* in our system are:

1. the NIA federal government deficit
2. the money supply, M2
3. the 4- to 6-month commercial paper rate
4. the energy items subindex of the consumer price index (energy CPI)

The first three of these driving variables are determined to a large extent by monetary and fiscal policies. The *responding variables* include the real GNP demand components:

5. consumption (consumer expenditures for nondurables and services)
6. investment (consumer expenditures for durables, residential construction, nonresidential business fixed investment, and change in business inventories)
7. government purchases (federal, state, and local)
8. net exports

the real GNP supply variables:

9. hours worked per week in private business sector
10. compensation per hour in private business sector
11. civilian employment

and a final variable:

12. the total urban consumer price index (CPI)

Before applying the VAR model to its assigned task, we compiled error statistics to ensure that our model was reasonable. The statistics were generated by estimating the model over subperiods of the sample period and forecasting eight quarters beyond those subperiods. Table 2 presents Theil U statistics for VAR forecasts and for representative major forecasts.[2] Lower values of Theil U statistics indicate better forecasting performance than higher values. The statistics in the table aren't exactly comparable, both because the forecasts used different measures of the price level (CPI and GNP deflator) and because they were predicated on different data periods. Either of these differences would make direct comparisons of the error statistics impossible. However, we do interpret the similarities among the statistics as implying that this VAR model is a reasonable forecasting device. In fact, the VAR has better error statistics over horizons of four quarters and longer, even though the forecasts from Data Resources Incorporated (DRI) and Chase Econometrics were based on more complete information.[3]

Table 2

Theil U Statistics for VAR Forecasts (1976:1–1978:4) and Two Major Forecasts (1970:3–1979:2)

Variable	Forecast Horizon (Quarters) 1	2	3	4	5	6
Real GNP						
VAR	.685	.544	.460	.360	.341	.321
DRI	.431	.415	.447	.489	.503	.504
Chase	.382	.349	.371	.419	.459	.495
Price Level						
VAR (CPI)	.259	.268	.241	.196	.159	.161
DRI (Deflator)	.160	.191	.232	.265	.292	.315
Chase (Deflator)	.156	.159	.202	.244	.284	.314

Key: 0.0 = perfect accuracy
1.0 = results could be attained by predicting no change from current value over the period

The Initial Forecast: F1

Initially, the VAR was estimated over the entire period from the second quarter of 1958 to the fourth quarter of 1978 and was used to forecast variables through the end of 1980. This forecast, F1, can be thought of as an "original best guess" that might have been made at the end of 1978 if we had assumed that the economic relationships established in the past would hold up in the future.[4] Forecast F1 predicted that for the fourth

[2]The Theil U statistic is, for a given forecast horizon, the square root of the ratio of the sum of squared forecast errors to the sum of squared actual changes. Thus, the Theil U is a dimensionless statistic. In McNees 1979, only aggregate Theil U statistics, summarizing forecasts over a range of horizons, are given. The Theil U statistics for particular horizons were kindly provided to us by McNees.

[3]The forecasts from Chase and DRI are based on late-quarter data; when their forecasts based on early-quarter data are used, our VAR looks even better.

[4]The crucial importance of this assumption is emphasized in Sargent 1979, pp. 8 and 13, and in Lucas and Sargent 1979.

quarter of 1979 over the fourth quarter of 1978, real GNP would grow by 1.9 percent and inflation, as measured by the CPI, would be 8.7 percent.[5]

The Effects of the Oil-Price Shock in 1979: F2

Changes in the price of oil are captured in this model through movements in the energy CPI. In order to determine the effects of surprise increases in oil prices on our model's forecast, we treated the quarterly differences between the VAR forecast of the energy CPI and its actual value in the first three quarters of 1979 as a series of shocks or innovations. By using the estimated coefficients in the VAR and information in the covariances of the residuals from the model, we were able to estimate the likely effects of the energy-price shocks on the future of all variables in our model. The procedure used in this analysis is discussed in the Appendix.

The revisions to F1 resulting from this series of shocks in energy prices create a new forecast, F2. The difference between F1 and F2, then, is our estimate of the effects of the oil-price shock.

According to the comparison of F2 and F1 in the top half of Table 3, the oil-price shock over 1979 reduced growth in real output by 1.7 percentage points and increased the inflation rate by 2.9 percentage points. In addition, the federal deficit in calendar year 1979 tightened significantly in response to the oil-price shock, declining by $14.3 billion.

The Oil-Price Shock and the 1979 Forecast Errors: F3

Forecast errors are the differences between actual and predicted values. Our forecast errors for 1979, then, are the differences between the F3 projections (our best estimate of actual values) and the F1 projections (our initial predicted values). The F3 forecast was obtained from a reestimated VAR using all available information through the first three quarters of 1979 plus an estimate of the actual fourth-quarter level of the 4- to 6-month commercial paper rate, based on weekly observations through the first eleven weeks of the quarter.[6] The

[5]This VAR forecast is a bit more pessimistic than that in our winter 1979 *Quarterly Review*. One reason is that in this issue Supel interpreted the fall 1978 policy actions to be an understood change in the policy rule, while the VAR interpreted them to be a shock, that is, random drawings under the existing rule. An understood change in policy rule can be expected to have a larger price effect and a smaller output effect than a policy shock.

[6]The data available when these forecasts were made were the first revisions to the third-quarter national income accounts data.

Table 3

Forecast Errors and the Impact of the Oil-Price Shock in 1979

Forecast for 1979	Percentage Change Fourth Quarter Over Previous Fourth Quarter				Annual Average	
	Real GNP	CPI	Energy CPI	M2	NIA Federal Deficit ($ Bil.)	4–6 Mo. Coml. Paper Rate (%)
F2	0.2	11.6	35.9	9.4	−27.9	8.6
F1	1.9	8.7	10.3	8.6	−42.2	9.7
Effect of Oil-Price Shock in 1979 (F2 − F1)	−1.7	+2.9	+25.6	+0.8	+14.3	−1.1
F3	0.5	12.6	42.3	8.0	−13.1	11.0
F1	1.9	8.7	10.3	8.6	−42.2	9.7
Forecast Error in 1979 (F3 − F1)	−1.4	+3.9	+32.0	−0.6	+29.1	+1.3

difference between the forecast commercial paper rate and the observed rate, which was 2.5 percentage points, was imposed on this reestimated VAR as a fourth-quarter shock independent of any other shocks to the model.

The comparisons in Table 3 support the notion that 1979 forecast errors were largely attributable to the oil-price shock. This shock accounts for a significant part of the forecast errors for growth in both real output and the CPI during 1979.

However, when the effect of the oil-price shock and the forecast errors are compared, there appear to be other important shocks in 1979. Among these was the tightening of the policy-related variables in F3 relative to F2: growth in the money supply was slower while the interest rate was higher, and the deficit was smaller.

The Effects of the Oil-Price Shock on the 1980 Forecast

In 1980, as in 1979, the oil-price shock had an impact on our forecast for real GNP and CPI. Our estimates of the effects of the oil-price shock on the outlook for 1980 were obtained by comparing the F2 and F1 forecasts for that year.

Table 4 suggests that over 1980 the oil-price shock by itself would not have produced either negative real growth or significantly greater inflation. Predicted real GNP growth is lower in 1980 by 1.2 percentage points because of this shock, but it is still significantly positive. Also because of the oil-price shock, inflation in the CPI is projected to increase during 1980 by only about 0.5 percentage points. This result is partly due to the model's prediction that oil prices would stabilize in the year following the price shock, showing only a 1.6 percent increase in the energy CPI over 1980. Based on past experience, this prediction seems reasonable: oil prices did not increase at all in 1975 following their quadrupling in 1974.

The Oil-Price Shock and the 1980 Forecast Revisions

Our latest forecast for 1980, F3, calls for a very small decline in real GNP and for continued double-digit inflation. It is thus the combination of 1979 oil-price increases together with the monetary and fiscal policy shocks of 1979 that leads to our pessimistic projections for 1980.

The oil-price shock explains a surprisingly small

Table 4

Forecast Revisions and the Impact of the Oil-Price Shock in 1980

Forecast for 1980	Percentage Change Fourth Quarter Over Previous Fourth Quarter: Real GNP	CPI	Energy CPI	M2	Annual Average: NIA Federal Deficit ($ Bil.)	4–6 Mo. Coml. Paper Rate (%)
F2	2.1	8.4	1.6	8.0	−81.2	6.1
F1	3.3	7.9	9.6	9.4	−71.8	8.0
Effect of Oil-Price Shock in 1979 (F2 − F1)	−1.2	+0.5	−8.0	−1.4	−9.4	−1.9
F3	−0.1	11.4	22.1	7.6	−54.5	10.0
F1	3.3	7.9	9.6	9.4	−71.8	8.0
Forecast Revisions* (F3 − F1)	−3.4	+3.5	+12.5	−1.8	+17.3	+2.0

*Combined effects of 1979 data and reestimation of the model portion of total 1980 forecast revisions. The difference between F3 and F1 is the total revision in our forecast for 1980, based on new 1979 information and reestimated economic relationships. Table 4 shows that the effect of reestimation and of all the new information including the oil-price shock (F3 – F1), clearly, is much larger than the effect of the oil-price shock alone (F2 – F1) for virtually all variables. While the total revision in real GNP growth is –3.4 percentage points, the oil-price shock alone can explain a decline of only 1.2 percentage points. And while the total revision in the CPI rate of inflation is 3.5 percentage points, the oil-price shock alone can explain an increase of only 0.5 percentage points. In the case of most other variables, the portion of the forecast revisions explained directly by the oil-price shock is not large.

The most obvious effect of new information and reestimation is that oil prices are predicted to continue their rapid rise throughout 1980. Table 4 shows, for instance, that the 1979 oil-price shock by itself (F2 – F1) is projected to lower the rate of increase in energy prices by about 8 percentage points through 1980, but the oil-price shock and all other new information (F3 – F1), in contrast, are projected to speed the rate of increase in energy prices by over 12 percentage points.[7]

Did fundamental economic relationships change?

The conclusions of our study so far—that the oil-price shock contributed greatly to the 1979 forecast errors and that the oil-price shock and policy shocks together shaped the outlook for 1980—appear consistent with each other. Nevertheless, there are reasons for exercising caution in drawing these conclusions, for another one of our findings casts some doubt on them.

The validity of our results and conclusions so far depends on the reasonableness of the three forecasts. A VAR generates reasonable forecasts as long as there are no important breaks from the past in economic relationships. But if, for example, the fall 1978 policy actions constituted an important break in economic policy, then all three forecasts, built on the assumption that old relationships still hold, may be questionable.

In fact, the surprisingly large difference between the F2 and the F3 forecasts for 1980 suggests that a break in past economic relationships is possible. Furthermore, a comparison of the original estimated VAR (which gave us F1 and F2) and the reestimated VAR (which resulted in F3) also suggests that there were significant changes in the estimated coefficients in the model. The F2 projections were derived by treating the oil-price shock simply as three observations deviating randomly from the forecast mean of the original estimated model. If these observations were not random shocks but rather were realizations of fundamental changes in economic processes, then the coefficients in the equations of the model would show significant changes. And if the coefficients change significantly, of course, then structural change cannot be easily dismissed.

Our sample size makes it impossible to perform a rigorous test for structural change. However, the way variables in the reestimated VAR responded to shocks in the energy CPI did change very clearly over an eight-quarter horizon from the original estimates of the model. This suggests that a structural change is at least a possibility.

Although the differences in the F2 and F3 forecasts for 1980 suggest the possibility of a structural change, they do not provide clear-cut evidence. In going from F2 to F3, we not only reestimated the coefficients of the equations in the model, but also added information on other variables in addition to the energy CPI. The differences between the two forecasts, therefore, could be due either to different coefficients or to different information sets.

However, if the differences in the forecasts were mainly due to the reestimated coefficients, then it is likely that some structural change has occurred. In order to isolate the effects of the reestimated coefficients, we must disentangle them from the effects due solely to new data. To do this, we conditioned forecasts from both estimated versions of the model on the same information set. We used data through the third quarter of 1979 to generate forecasts for the fourth quarter of 1979 and for all of 1980. In neither case did we impose the interest rate information from the fourth quarter of 1979. Since the information conditioning the forecasts is now the same, any differences must be caused by the new coefficients. These two forecasts are compared in

[7]One reason for this difference is that the F2 forecast weights most heavily the oil-price hikes of 1974, while the reestimated model weights the 1979 rises most heavily. Such significant differences in forecast paths for energy prices in 1980 suggest that there were important differences in the oil-price shocks of 1974 and 1979. In fact, in 1974 oil prices quadrupled in the first quarter and then remained stable, while in 1979 oil prices escalated in each quarter.

Table 5

The Effects of Reestimation on the 1980 Outlook

Forecast for 1980	Percentage Change Fourth Quarter Over Previous Fourth Quarter: Real GNP	CPI	Energy CPI	M2	Annual Average: NIA Federal Deficit ($ Bil.)	4–6 Mo. Coml. Paper Rate (%)
Reestimated VAR	0.3	11.3	20.7	8.0	–53.5	8.2
Original VAR	1.2	8.0	1.8	8.1	–84.5	7.0
Difference = Effects of Reestimation	–0.9	+3.3	+18.9	–0.1	+31.0	+1.2

Table 5.[8]

The effects of reestimation on the 1980 forecasts further suggest that 1979 may have been characterized by some changed economic relationships. Real output and inflation projections again show sizable changes, and the reestimated model predicts continued rapid increases in energy prices through 1980, while the original model shows virtually no change. Reestimation also results in a sizable projected tightening in the 1980 federal budget and in an increase in the average annual interest rate. While the evidence suggests that 1979 may have been characterized by some fundamental changes in relationships, it does not pinpoint the source of such changes. Either a new energy environment or new economic policies could be the source of any changes.

While these observations do not provide rigorous or conclusive proof that any sort of structural change in economic relationships occurred, they are strongly suggestive and should not be dismissed lightly. Should such changes have actually occurred, forecasts from any model based on macroeconomic relationships, including ours, would have to be interpreted with caution.

Summary

Our VAR model indicates that 1979 forecast errors were largely attributable to unexpected increases in oil prices during the year. Moreover, monetary and fiscal policies turned out to be tighter than historical relationships led us to believe they would be. A reestimated VAR, incorporating data through the first three quarters of 1979, suggests that attempts to forecast may have been thwarted by some changes in fundamental economic relationships. If such changes have occurred, then the problems we encountered in forecasting, given a policy change, and in predicting responses to a shock, given a change in economic relationships, would have been encountered no matter what method of inquiry was used.

[8]The differences in projected growth over the four quarters of 1979 were relatively small, since only the fourth quarter was forecast.

Appendix
The VAR and Impulse Response Functions

To see how our experiment works, we will characterize the VAR in the following way. (A more rigorous discussion may be found in Litterman 1979 or in Sargent 1979.) Let z_t be a vector (12×1, in this case) of all variables in the model at time t. The M^{th}-order vector autoregression ($M = 5$, in our case) for the z_t process is

$$(1) \qquad z_t = A_1 z_{t-1} + \ldots + A_M z_{t-M} + e_t$$

where $A_1, \ldots, A_M$ are (12×12) matrices, and where the (12×1) vector e_t is independent of $z_{t-1}, \ldots, z_{t-M}$. What relation (1) says is that the value of each component of z_t depends on the history of itself and of every other variable in the model plus a contemporaneous random component.

If we solve the system of difference equations (1), we arrive at a vector moving average representation

$$(2) \qquad z_t = \sum_{i=0}^{\infty} C_i e_{t-i}, \qquad C_0 = I$$

where each C_i is (12×12) in our model. Now (2) says that each component of vector z at time t depends on a random shock in that component at time t plus the effects of random disturbances in all other variables in previous periods. From (1) we know that the random vector e_t is simply the forecast error from a linear least squares prediction of z_t based on its own history. Thus, we may use (2) to trace out the likely effects of innovations, or unexpected shocks, in any one variable in the model on the future of all variables in the model, assuming that the responses remain the same as they were in the past. When used in this way, the sequence of matrices C_1, C_2, . . . in the moving average representation is called an *impulse response function*. (See Sims 1980 or Litterman 1979 for a more complete discussion of this analytical device.)

The steps outlined above work fine as long as the contemporaneous innovations in different variables are independent. But when we want to characterize the response of the model to innovations taking into account nonzero cross covariances, we must normalize the model in some way. The method employed here is to orthogonalize the errors according to a given ordering of the variables. Thus, a lower triangular matrix B_0 is found which will transform the vector e_t into a vector u_t having zero cross covariances,

$$u_t = B_0 e_t.$$

Premultiplying all the terms in (1) by B_0 leads to the representation

$$(3) \qquad B_0 z_t = B_0 A_1 z_{t-1} + \ldots + B_0 A_M z_{t-M} + B_0 e_t$$

$$= D_1 z_{t-1} + \ldots + D_M z_{t-M} + u_t$$

where $D_j = B_0 A_j$, $j \geqslant 1$. Since B_0 is lower triangular, (3) says that the first component of z_t, z_{1t}, depends on the history of all variables in the model plus a contemporaneous disturbance term. The second component z_{2t} depends on the contemporaneous value of z_{1t}, the history of all variables in the model, and a contemporaneous random term. The final component will depend not only on the history of, but also on the contemporaneous values of, all other variables in the model, as well as on its own history. Thus, the use of (3) and of the solution to it for determining effects of innovations is equivalent to reestimating the VAR allowing contemporaneous values of some variables to affect values of other variables in a manner reflected in the order of the variables and in the triangular characteristics of B_0.

In our experiment, the energy CPI was the first variable, z_{1t}, in the orthogonalization order. Thus, it was not affected by contemporaneous values of any other variable in the model. Following, in order, were the commercial paper rate, the deficit, M2, consumption, government purchases, investment, net exports, employment, hours, compensation per hour, and the

CPI. The first step of the experiment was to take the difference between the actual and forecast energy price index in 1979:1 as the innovation u_{1t}. After the effects on all variables of this innovation were taken into account, the difference between the new 1979:2 forecast of the energy price index and the actual value was imposed as a second shock. The procedure was repeated for the third-quarter innovation. The result was that we imposed successive innovations in the energy price index of 1.2, 5.3, and 5.8 standard deviations from the forecast mean in the first three quarters of 1979.

References

Litterman, Robert B. 1979. Techniques of forecasting using vector autoregressions. Research Department Working Paper 115. Federal Reserve Bank of Minneapolis, Minnesota.

Lucas, Robert E., Jr., and Sargent, Thomas J. 1979. After Keynesian macroeconomics. *Federal Reserve Bank of Minneapolis Quarterly Review* 3 (Spring): 1–16.

McNees, Stephen K. 1979. The forecasting record for the 1970s. *New England Economic Review* (September/October): 33–53.

Sargent, Thomas J. 1979. Estimating vector autoregressions using methods not based on explicit economic theories. *Federal Reserve Bank of Minneapolis Quarterly Review* 3 (Summer): 8–15.

Sims, Christopher A. 1980. Macroeconomics and reality. *Econometrica* 48 (January): 1–48.