Federal Reserve Bank of Minneapolis
Research Department Staff Report 109/JV

The Quantitative Significance of the Lucas Critique*

Preston J. Miller

Federal Reserve Bank of Minneapolis

William T. Roberds

Federal Reserve Bank of Atlanta

ABSTRACT

Doan, Litterman, and Sims (DLS) have suggested using conditional forecasts to do policy analysis with Bayesian vector autoregression (BVAR) models. Their method seems to violate the Lucas critique, which implies that coefficients of a BVAR model will change when there is a change in policy rules. In this paper we attempt to determine whether the Lucas critique is important quantitatively in a BVAR macro model we construct. We find evidence following two candidate policy rule changes of significant coefficient instability and of a determination in the performance of the DLS method.

*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.

1. INTRODUCTION

The Lucas critique (Lucas 1976) changed the way economists do policy analysis. The critique suggested that such analysis was flawed, as it then was conducted with macroeconometric models. Economists responded to the critique by developing new approaches.

One new approach is described in papers by Litterman (1984), Sims (1982, 1986), and Doan, Litterman, and Sims (1984). These authors propose doing policy analysis by having Bayesian vector autoregressions (BVARs) generate forecasts which are conditional on future paths of policy variables. While Doan, Litterman, and Sims (henceforth, DLS) are aware of the Lucas critique, and even acknowledge its logical validity, they question its quantitative significance. In this paper we attempt to determine whether DLS are right.

The Lucas critique maintains that the coefficients of a macroeconometric model will change when there is a change in the rule which determines policy actions based on the state of the economy. The reasoning seems unassailable. The relationships in a macroeconometric model are aggregates of individual decision rules. In a stochastic, dynamic environment, the decisions of optimizing agents depend on expectations of those future policy actions that affect their budget sets. A change in the policy rule changes individuals' expectations of future policy actions, their budget sets, and, therefore, the way they make decisions based on current information. This change in individual decision rules then translates into changes in the coefficients of aggregate relationships in macroeconometric models.

This reasoning is unassailable if one accepts the classical rational expectations definition of a policy rule change as a completely unanticipated, once–and–for–all change in the probability structure governing the evolution of policy variables or the public's information about this structure. This definition seems appropriate if policy rules changes are infrequent and largely unanticipated. However, as pointed out in Sims (1982), Sargent (1984), and Cooley, LeRoy, and Raymon (1984), this definition seems questionable if policy rule changes occur frequently or are properly anticipated by economic agents. But even in this case, the coefficients of a linear macroeconometric model will change with a change in the policy rule.

Although Doan, Litterman, and Sims don't quarrel with the logic of the Lucas critique, they do question its relevance. They argue that most policy changes that have occurred or that are usually contemplated do not resemble Lucas' once–and–for–all changes in policy rules. Rather, they argue that policy changes are like the drawings of residuals under a given policy rule. Since policy changes tend to be small deviations from existing policy, Sims argues that the Lucas critique amounts at most to small nonlinearities in a BVAR model and that these nonlinearities can be accommodated by specifying time–varying coefficients (see Sims 1982, p. 120; 1986, p. 7). Sims (1982, pp. 138–139) states that

> Policy rules in the relevant sense of that term have not changed frequently or by large amounts. The large forecast errors of recent years do not seem to be attributable mainly to shifts in the structure of predictive equations. Statistical models allowing for drift in predictive structure estimate best when the change in that structure is assumed to be slow, so that recent large predictive errors are interpreted as large random shocks to the equations, not mainly as the effect of parameter changes.

Given Sims' interpretation, the DLS conditional forecasting procedure might be expected to perform reasonably well at times of policy changes. However, if Sims' interpretation is incorrect, the DLS procedure might be expected to perform poorly at those times. That is because the estimates of coefficients used to generate forecasts conditional on future policy could be very different from the true coefficients that obtain following the policy change.

We attempt to assess the quantitative importance of the Lucas critique, and hence the validity of the DLS method, at two recent times when macro policies changed: the fall 1979 change in the FOMC's operating procedures and the fall 1981 introduction of

Reagan's budget policies. We begin by describing the small BVAR model we constructed to analyze these policy changes. We next examine some empirical evidence on the magnitude of the policy changes to determine whether they seem more like changes in policy rules or like random drawings under given policy rules. While our evidence is not decisive, it does suggest that these two policy episodes are good candidates for rule changes. We then look for empirical evidence of changes in the model's coefficients following the policy changes. We uncover some strong evidence of coefficient instability—the strongest associated with the monetary policy change. We finally check for changes in the performance of the DLS method at the times of these policy changes. We find that associated with the coefficient instability there is a deterioration in the usefulness of the DLS method.

Our findings suggest that the Lucas critique is quantitatively important, and, thus, the DLS method is not very accurate in predicting the effects of changes in policy rules. This result seems robust to variations in our model and analysis. But, does that mean the DLS method should be abandoned? In our conclusion we argue that it should not be. We are unaware of an empirical method which outperforms it, and our analysis indicates a potential for improving it.

2. THE BVAR MODEL

Our BVAR forecasting model was constructed using the methodology described in Sims (1982) and Doan, Litterman, and Sims (1984).

2.1 Data Series Chosen

For our study we chose six postwar (1948:2–1986:4) U.S. quarterly macroeconomic series:

(i) real GNP (logarithm of)

(ii) inflation (GNP deflator)

(iii) three–month Treasury–bill rate

(iv) trade–weighted value of the dollar

(v) ratio of monetary base to public debt

(vi) ratio of deficit net–of–interest to nominal GNP.

The ratio series (v) and (vi) are taken as indicators of policy: the monetary base–to–debt ratio indicates current monetary policy and the deficit–to–GNP ratio indicates current fiscal policy. Our choice of policy variables was motivated by the theoretical analyses of Wallace (1984) and Miller and Wallace (1985). Our choice also was influenced by stationarity considerations, the predictive ability of the whole model, and the responses of the nonpolicy variables to policy shocks.

A detailed description of the data series is given in Miller and Roberds (1987), and plots of the policy series are displayed in Figures 1A and 1B.

2.2 Specification Search

We considered a number of different specifications for our BVAR forecasting model. Following Sims (1982) and DLS, we allowed for a small amount of explicit time variation in the parameters of the model. (The amount of variation is not crucial: all our reported calculations would have yielded similar results with fixed–coefficient models. We used the time–varying parameter feature primarily because it led to smaller forecast errors.) Our estimation methodology was the usual Bayesian one of specifying initial first and second moments to offset the overparameterization inherent in vector autoregression models. That is, the models we considered took the form

$$y(t) = A_t(L)y(t-1) + c(t) + u(t) \qquad (1)$$

where $y(t)$ is the time t observation on the vector of six macroeconomic variables; $A_t(L)$ is the time–varying autoregressive polynomial in the lag operator L; $c(t)$ is a time–varying constant term; and $u(t)$ is a white noise error term. Following standard practice for quarterly models, we set the lag length of $A_t(L)$ at six lags.

Combining $A_t(L)$ and $c(t)$ into a common coefficient matrix $B(t)$, we assumed that a typical row $b_i(t)$ of $B(t)$ (i.e., the coefficients of a typical equation i in model (1)) follows a random walk specification

$$b_i(t) = b_i(t-1) + e_i(t) \tag{2}$$

where $e_i(t)$ is a white noise error term assumed to be independent of $u(t)$. For a given prior mean value of $b_i(t)$, ${}_{-1}b_i(t)$, a known initial covariance matrix of $b_i(t)$, ${}_{-1}\Sigma(t)$, a known covariance matrix W_i of the coefficient shocks $e_i(t)$, and a known value of the variance of each component of $u(t)$, the use of the Kalman filter allows for recursive calculation of the linear projection of the coefficient vectors $b_i(t)$ on information available at time t. Since estimation of the second moment matrices and the initial mean values needed for Kalman filtering is not computationally practical in this application (due to the large dimension of the Kalman state vector $b(t)$; see the Appendix), the essence of Sims' forecasting methodology involves using heuristic techniques that yield some informed guesses about the value of these moments. Roughly speaking, this methodology involves a search for scaling factors for these moments in order to maximize unconditional forecast accuracy. An unfortunate aspect of this methodology is that it does not yield a joint posterior distribution for the model parameters across equations, a fact which makes formal inference rather difficult. A more detailed discussion of this and other features of the BVAR technique is given in the Appendix.

Our methods for choosing the model's scaling factors closely followed those suggested in DLS, although we adopted some simplifications. Details of our specification search are relegated to the Appendix. One significant feature of our final specification is that it allows only a small amount of time variation in the parameters. The covariance matrix of coefficient shocks W was taken to be 10^{-6} times the initial covariance matrix of the coefficients $\Sigma(-1)$. A similar amount of time variation is reported in DLS.

2.3 Out-of-Sample Forecasts

The forecasting performance of the model over the period 1966:1–1986:4 is summarized in Table 1. Also given in Table 1 are forecast performance statistics for a fixed-coefficient BVAR model, as well as for a benchmark system of AR(4) fixed-coefficient univariate models, which were estimated by ordinary least squares. As can be seen from Table 1, the DLS methodology was generally successful in delivering a forecasting model with desirable out-of-sample forecasting properties. Although the fixed-coefficient BVAR was able to generate significant improvements over the benchmark forecasts at shorter horizons, its performance deteriorated over longer horizons. The greater flexibility of the DLS-type specification then resulted in large improvements over the fixed-coefficient model at longer horizons, as well as some smaller gains over the short term. In addition, the DLS approach was generally able to outperform a random walk forecast at most horizons, as evidenced by the Theil U-statistics reported in the table.

3. STRUCTURAL INSTABILITY FOLLOWING POLICY CHANGES

3.1 Evidence of Policy Changes

It seems clear that monetary policy changed in the fall of 1979, when the Federal Reserve began following a new operating procedure. And it also seems clear that budget policy changed in the fall of 1981, when the first Reagan budget was implemented. What is not clear, however, is whether either policy change would be considered large in a statistical sense or whether either change would be better characterized as a shift in a policy rule or as a random disturbance under a given policy rule. To address these issues, we informally examine the behavior of the monetary and budget policy indicators before and after the respective policy changes.

Simple plots of the monetary and fiscal policy indicators reveal somewhat unusual behavior at the times of the corresponding policy changes (see Figures 1A and 1B). In late 1979, our monetary policy indicator (the ratio of the monetary base to federal debt) leveled and then turned down sharply after a period of substantial increase. Near the start of 1982, our fiscal policy indicator (the ratio of the deficit net-of-interest to GNP) began a sustained climb from around 0 to a range of between 2 and 4%. While the plots of the policy indicators prove nothing, they at least suggest that the changes in policy may have been significant, and they provide a reason to look at additional measures of policy change.

We examined three additional measures: (1) one-step-ahead forecast errors of the policy indicators, (2) estimates of the coefficients on the first own lag in the policy indicator equations, and (3) comparisons of unconditional dynamic forecasts of the policy indicators to the actual outcomes. Measures (1) and (2) didn't provide much evidence of policy change, but measure (3) did. Together the measures suggest that the changes in policy—if any occurred—were cumulative and affected more than the first-order properties of the indicators. Measure (3) suggests that the policy changes may have been significant and may be better characterized as rule changes, since there is no tendency for the forecasts to converge to the actual outcomes.

Our examination of the first measure revealed that one-step-ahead forecast errors followed a little different pattern after the monetary policy change but not after the fiscal policy change (see Figures 2A and 2B). Errors in forecasting the monetary policy indicator in the few years following 1979 seem slightly larger on average with slightly more negative serial correlation than they were prior to 1979. Errors in forecasting the fiscal policy indicator in the few years following 1981 seem to be in the same pattern as in the years immediately preceding 1981.

Our second measure showed no evidence of significant policy changes. The coefficients on the first own lag do not seem to move unusually after either policy change (see Figures 3A and 3B).

The third measure showed different results. Dynamic unconditional forecasts of the policy indicators made at the time of the corresponding policy changes reveal large and persistent errors (see Figures 4 and 5). The size of the errors suggests that the policy changes may have been significant, and the persistence of the errors suggests they may have been in the nature of rule changes. Note that the likelihood of persistent forecast errors clearly increases with the "persistence" of the series being forecast. Since the VAR model estimated above at any given time probably contains numerous roots close to the unit circle, persistence of the model's forecast errors is not entirely unexpected. Still, we were somewhat surprised at the degree of persistence in the errors of some of the forecasts depicted in Figures 8 and 9.

If we interpret the equation for the deficit–to–GNP ratio to be the fiscal policy "rule," then it is possible that the persistent differences in the unconditional forecast from actual values plotted in Figure 5 are not due to a change in the rule (i.e., to coefficient changes) but instead are due to the errors the model made in forecasting nonpolicy variables. For instance, a larger deficit than predicted may not be due to an unforeseen tax cut but instead may represent a typical policy response to an unforeseen recession. This explanation is advanced by Barro (1986) in an econometric study of U.S. fiscal policy.

In order to determine better whether the errors plotted in Figure 5 are due to a policy change, we also computed a dynamic forecast of the deficit–to–GNP ratio conditional on the actual values of the nonpolicy variables. We found the same general pattern of persistent errors as for the unconditional forecasts, with smaller errors at short horizons but larger errors at longer horizons (see Figure 5). We interpret this as evidence suggesting a rule change.

Our finding of a change in the fiscal policy rule in the fall of 1981 contrasts with Barro's (1986) finding of no change. Since Barro's econometric methodology and data differ from ours, it is difficult to directly compare our finding to his. However, there is some evidence (Modigliani 1986) that Barro's results are distorted by his model's large forecast errors for inflation over this time period.

In summary, anecdotal evidence, simple plots, and dynamic forecasts suggest that there may have been significant changes in the monetary policy rule in 1979 and in the fiscal policy rule in 1981. The evidence is far from decisive, however. One–step–ahead forecast errors and estimates of coefficients on first own lags show fairly little in the way of unusual patterns following the policy changes.

3.2 Evidence of Coefficient Instability

If the changes in policies in 1979 and 1981 were significant changes in rules, then the Lucas critique implies that we should find evidence of coefficient instability in the model's nonpolicy equations following the changes. We examined first a number of informal measures of coefficient changes and then a more formal Box–Tiao statistic.

We began by examining one–step–ahead forecast errors (see Figures 6A–6D). If there were significant changes in coefficients, we would expect to find some large errors. However, if we found large errors, they could just be the result of large residuals. So this evidence is of the necessary, but not sufficient, variety. The errors for real GNP and inflation do not seem extraordinary following either policy change. The errors for the three–month T–bill rate seem unusual following both policy changes. After both policy changes, the errors become quite large, frequently exceeding two standard errors of forecast. The errors following the monetary policy change tend to be somewhat larger than those following the fiscal policy change. Finally, the errors for the dollar, like those for the interest rate, are large after each policy change, although somewhat larger following the monetary policy change. As a whole, these figures seem to indicate that if we were to find evidence of coefficient instability, it would be with respect to the financial variables: the T–bill rate and the value of the dollar.

We next examined the estimates of the coefficients of the nonpolicy variables on the first own lags (see Figures 7A–7H). The estimate for the coefficient on real GNP seems to settle down in about 1975, and it shows few signs of instability following the policy changes. The estimate for the coefficient on inflation also doesn't show much instability immediately following the policy changes, although it does start rising sharply beginning in 1984, and changes in the estimate seem to follow a lower–order process after

1979. In contrast to the coefficients on real GNP and inflation, the estimate of the coefficient on the first own lag in the T–bill equation does show signs of instability following the policy changes. The revision in the estimate following the monetary policy change is the largest and sharpest. It suggests that any model with time–varying coefficients that adjust smoothly would not predict well the effects of the monetary policy change on the interest rate. The estimate of the coefficient on the dollar also shows signs of instability, but most of the instability doesn't occur until 1985, well after both policy changes.

In order to check for changes in coefficients other than just those on first own lags, we examined some measures which are functions of all the coefficients. We examined decompositions of variance before and after policy changes and forecasts made at the time of the policy changes. In particular, we constructed a series of forecast error decompositions based on the model's estimated coefficients and the empirically generated variance–covariance matrix of one–step–ahead forecast errors. These decompositions were constructed for the dates 1979:3, 1980:3, 1981:4, and 1982:4, and they reveal some seemingly important shifts. The shifts from 1981:4 to 1982:4 tend to be smaller than they were between the earlier two dates, implying a less fundamental (or perhaps better anticipated) change in the model's dynamics following the fiscal policy change than following the monetary policy change. (Our method for constructing these decompositions of variance and their values are reported in Miller–Roberds 1987.)

We also considered two forecasts for each policy change and compared them to the actual outcomes. We looked at unconditional forecasts and forecasts with coefficient estimates conditioned on data up to the policy change, but with the relevant policy variable set at its actual future values. The conditional forecasts were generated according to the DLS procedure (pp. 61–71), the empirically generated variance–covariance matrix of one–step–ahead forecast errors. (This procedure lets us make use of "structural" assumptions by introducing additional constraints that structural innovations in some variables be required to be zero. We experimented with such constraints, but concluded that such "identifications" are unlikely to improve the accuracy of the model's forecasts.)

The conditional forecasts are the most likely outcomes according to our model, given the initial coefficient estimates, the future path of the relevant policy variable, and the assumption of no future parameter change (see Figures 8A–8D and Figures 9A–9D). If the policy changes were important, we would expect to see large differences between unconditional forecasts and the actual outcomes, which we do see, especially for the T–bill rate and the dollar. If the policy changes were like drawings of residuals which implied no instability in coefficients, we would expect the conditional forecasts to largely close the gaps between the unconditional forecasts and the actual outcomes. The conditional forecasts do close much of the gaps for real GNP and inflation but little of the gaps for the T–bill rate and dollar. These figures suggest that forecast errors following the policy changes tend to be large and that they would still have been large for financial variables if the actual future values of the relevant policy variable had been known at the times of the policy changes.

Our Fed policy results seem consistent with those of Blanchard (1984) who found a change in the behavior of interest rates but not in the behavior of output and prices. They also seem consistent with those of Christiano (1986) for his trend–stationary BVAR, which corresponds most closely to our model. They are not consistent, however, with those for his difference–stationary BVAR.

3.3 A More Formal Measure of Forecast Accuracy

The exercises in the preceding section are necessarily informal in nature. The DLS methodology does not yield a joint distribution of the model's coefficients that can be used to construct traditional tests of model stability. Nonetheless, we found a statistic proposed by Box and Tiao (1976) to be a useful indicator of the model's forecast accuracy over various periods. Define this statistic as

$w(i,k,t) = \epsilon(i,k,t)' D^{-1}(i,k,t) \epsilon(i,k,t)$, where

i indexes the set of variables being forecast,

k is the forecast horizon,

t is the date of the forecast,

$\epsilon(\cdot)$ is a vector of forecast errors from the k–step ahead forecast of variables i made at time t, and

$D(\cdot)$ is the variance–covariance matrix of ϵ for the time t estimate of the VAR, assuming no future parameter variation and assuming that the time t VAR is known with certainty.

For a stationary Gaussian model, as the sample size grows $w(\cdot)$ should be asymptotically distributed as $\chi^2(\dim \epsilon)$ and the standardized score

$$z(i,k,t) = \frac{w(\cdot) - \dim \epsilon(\cdot)}{[2 \dim \epsilon(\cdot)]^{1/2}}$$

should be approximately normally distributed as N(0,1) for large dim ϵ. For the VAR model of our paper, we thought it highly unlikely that this approximation would be very close to the actual distribution of z. Instead, we computed series of these scores for different groups of variables to give an indication of the comparative accuracy of the model forecasts over different time periods.

Figure 10 displays the z series of three sets of forecasts at a forecast horizon of eight quarters over the period 1966:1–1986:4. (In Figure 10, the forecast horizon was shortened for the last seven quarters of the sample. We also plotted the z's at a horizon of four quarters, and the picture was basically the same.) The first series (UNSCORB) is the series computed using the model's unconditional forecasts. The second series (COND1B) is for forecasts conditional on the deficit/GNP ratio, while the third (COND2B) series is for forecasts conditional on the base/debt ratio. The scores for the conditional forecasts were computed using only the forecast errors for the variables not conditioned on. Hence, dim ϵ (= number of variables $\times$ k) was set equal to $5{\times}8 = 40$, while for the unconditional forecasts dim ϵ was set equal to $6{\times}8 = 48$. The conditional forecast scores were computed using the unconditional covariance matrices for the variables forecast, rather than conditional covariance matrices that would be the correct choice for the Box–Tiao test. This was done in order to facilitate comparison with the model's unconditional forecasts. However, the z series computed using the conditional covariance matrices differed little from those depicted here.

A number of interesting observations can be drawn from Figure 10. First, the accuracy of all three forecasts deteriorates following the oil price shocks of the early 1970s. However, conditioning on future policy variables only seems to help appreciably in the case of the base/debt ratio. The accuracy of the forecasts also deteriorates following the 1979 monetary policy change. One might guess that much of this forecast accuracy deterioration derives from errors in forecasting interest rates over this period. This is confirmed in Section 4 when we condition forecasts on interest rates and find much less deterioration. In the case of the 1981 budget policy change, the accuracy of all forecasts gradually deteriorates after the policy change. This gradual change seems consistent with the popular notion that what changed is that deficits did not decline as the recovery progressed.

Our overall conclusion after inspection of Figure 10 and similar graphs was that conditioning on policy variables, as proposed by DLS, did not generally reduce the relative uncertainty (i.e., relative to the average level of uncertainty associated with this type of forecast and relative to a situation in which optimal forecasts would be linear and time invariant) associated with forecasts made at times of policy changes.

3.4 A Closer Examination of the DLS Procedure

A casual reading of the discussion above might convey the impression that conditioning on future values of policy variables as suggested in DLS does not in general reduce the model's errors in forecasting nonpolicy variables, even when future values of the policy variables are known with certainty. This impression would be incorrect. Above we have attempted to show that the forecast uncertainty associated with policy change episodes is greater than the average level of uncertainty associated with the model forecasts, and that this result holds for policy–conditional forecasts as well as unconditional forecasts. In this section we document for arbitrary time periods the extent to which conditioning on the model's policy variables can be expected to reduce errors in forecasting the model's nonpolicy variables. We also consider some evidence as to how much the model's conditional forecasts could be improved by accurately predicting future variations in the model's parameters.

As a benchmark, we simulated with the standard coefficient estimates the out–of–sample forecasting performance of the DLS method over the period 1976:1–1986:4. (Note that we use out–of–sample loosely here. These forecasts take into account data revisions announced after the forecast date, as well as hyperparameter settings that use subsequent data. Thus, these forecasts are likely to be more accurate than real–time forecasts.) In our simulations, we conditioned our forecasts on the next 12 quarters of the base/debt and deficit/GNP ratios, both singly and together. Our measure of forecast accuracy was the log determinant of the covariance matrix of out–of–sample forecast errors for the other variables of the model, i.e., the nonpolicy variables.

The results of this exercise, included in columns (1)–(3) of Table 2, are generally favorable to the DLS procedure. Except at the 12–quarter horizon, adding information about future values of policy variables appears to increase forecast accuracy. (This conditional forecasting procedure reduced forecasting accuracy when we carried out the exercise over the period 1966:1 through 1986:4, but we believe that result is due to the small sample sizes available for forecasting at the beginning of the period.) Moreover, an examination of forecast error covariance matrices (not reported here) indicated the forecast error variances were reduced for each of the nonpolicy variables at all horizons but 12 quarters. While these improvements in forecast accuracy are large enough to suggest that the DLS conditional forecasting technique may be of some practical utility to forecasters, they are nonetheless disappointingly small, generally on the order of an average 5% reduction in the standard deviation of the forecast error. These numbers are consistent with population values computed with the 1976:1 parameter estimates assuming no parameter uncertainty (not reported), suggesting that the small forecast error improvements reported here are representative of what may be expected using the DLS methodology.

We then asked how much of the remaining forecast error can be explained by predictable changes in the model's coefficients. Ideally, the answer to this question would involve simulation of the true conditional forecasts of the model, as opposed to the simulations just reported, which assume no parameter change in the future. Unfortunately these true conditional forecasts are, in general, quite difficult and expensive to compute. As an alternative, we considered a simple procedure which attempts to account for parameter variation in a somewhat ad hoc way. This procedure consisted of doing a DLS–type conditional forecast, then applying the Kalman smoothing algorithm using these forecasts as data. (On Kalman smoothing see, e.g., Goodwin and Sin 1984.) A second set of DLS–type conditional forecasts was then constructed using the smoothed coefficients.

We also computed a third set of forecasts conditional on policy variables. In this set of forecasts, we again applied the smoothing algorithm to the model's coefficients, only this time we used the actual data on all of the model's series to do smoothing. DLS–type conditional forecasts were then constructed using the smoothed coefficients. Although this is also clearly an ad hoc procedure with many problems of interpretation, we used these forecasts to get a rough measure of how well a forecaster might do if, in addition to the future values of some policy variables, the future values of shocks to the model's coefficients were known with perfect accuracy.

Table 2 compares the results of simulating the two heuristic procedures described above over the last ten years of data, together with the results of unconditional and conditional procedures reported earlier. As before, the forecasts are conditioned on 12 future quarters of data on both the monetary base–to–debt and deficit–to–GNP ratios. Column (4) of Table 2 reveals that for our BVAR model, the sequential forecasting procedure generally performs slightly worse out–of–sample than the DLS procedure. However, the fact that it does not perform substantially worse than the DLS procedure suggests that gains in forecasting accuracy might be possible if more sophisticated procedures were used to account for the influence of future policy variables on parameter values. That message is reinforced by the "ideal" column (5) of Table 2, which unequivocally points to coefficient variation as an important source of error in conditional forecasting. Of course, if the random walk parameter specification of equation (2) is approximately correct, then this variation would be unforecastable in an unconditional sense. However, the potential for improvement in conditional forecasting due to improved forecasts of coefficient changes remains largely unexplored.

Table 2 demonstrates that there is considerable room for improvement in the performance of policy–conditional forecasts over the "reduced form" DLS conditional forecasting procedure. But the realization of such improvements would depend on being able to accurately infer future changes in the forecasting model coefficients resulting from future policy changes. Of course this has always been the purported advantage of structural over reduced form modeling methodologies. But Table 2 does not tell us which if any structural models would be successful at delivering better conditional forecasts than DLS. The table also leaves open the possibility that more sophisticated reduced form approaches might also yield better policy–conditional forecasts.

4. ROBUSTNESS OF FINDINGS

Although our model specification and interpretation of results can be defended, they nevertheless are somewhat arbitrary. Readers have questioned whether our findings would hold up under some different specifications or considerations. Specifically, they questioned the robustness of our findings if we substituted for our monetary policy indicator, allowed for heteroskedasticity in the equations for financial variables, or compared the deterioration in performances of the DLS method following policy changes to what occurs in business downturns. While our findings seem robust to the first two modifications, we cannot distinguish clearly whether the deterioration in the DLS method is due to policy changes or to changes in the dynamics of the economy over the business cycle.

4.1 Alternative Monetary Policy Indicator

One concern about the dynamic forecasting exercises shown in Figures 8 and 9 is that these forecasts focus on incorrect measures of policy. To address this concern we recalculated the dynamic conditional forecasts, this time conditioning our forecasts on future values of T–bill rates over the relevant forecast horizon. One rationale for considering such forecasts is that the monetary policy instrument over the period might have been closer to an interest rate than to a quantity variable. A second rationale is that to the extent interest rate movements reflect anticipations of future policy actions, interest rates may capture the effects of a policy change even if the model itself does not contain an accurate measure of policy.

The results of these exercises were mixed. For the Fed policy change experiment, forecasts conditional on the path of T–bill rates were generally less accurate than those conditional on the path of the base/debt ratio. For example, Figure 11 compares the actual path of real GNP over the years 1979–82, the 1979:3 forecast conditional on the base/debt ratio (CONDL1), and the 1979:3 forecast conditional on T–bill rates (CONDL2). The T–bill conditional forecast completely misses the 1980 downturn, and predicts continued growth until the second quarter of 1981. However, the forecasts conditional on T–bill rates were generally more accurate than the forecasts conditional on the deficit/GNP ratio in predicting the effects of the Reagan tax cut. Figure 12, for example, shows that the 1981:4 interest rate—conditional forecast (CONDL2) comes closer to predicting the 1982 downturn than does the 1981:4

deficit/GNP ratio—conditional forecast (CONDL1). We also computed z series for forecasts conditioned on the T–bill rate and compared the series to those displayed in Figure 10. Conditioning on interest rates worsened the relative accuracy of forecasts following the oil price shocks of the early 1970s. Conditioning on interest rates did improve forecasts following the 1979 monetary policy change, which suggests that the deterioration in forecasts displayed in Figure 10 at this time is largely due to errors in interest rates. As was the case for the other forecasts, the accuracy of the interest rate–conditional forecast deteriorates gradually after the 1981 budget policy change. We concluded that there was no clearcut advantage of interest rate–conditional forecasts over forecasts conditioned on our original choices for policy variables.

4.2 Increased Volatility of Financial Markets

The average size of the model's forecast errors for the dollar clearly increases after the collapse of the Bretton Woods system in the early 1970s, leading some readers to question whether our results are influenced by the inability of the model to adequately correct for this change. To address this concern, we reestimated the time–varying coefficient BVAR model using a different set of initial second moments (see the Appendix for the original specification). In this new specification, the initial second moments for the time–varying model were fit using estimates from a fixed–coefficient BVAR model estimated from 1973:1 through 1986:4 (instead of 1949:3 through 1965:4). Changing the data used to calculate the initial second moments allows for the error variances of the model's equations for T–bill rates and the dollar to be scaled up to reflect the increased volatility of the financial markets over this period. Except for this difference in sample period, the new specification was identical to the original one described in the Appendix.

The new model's one–step–ahead unconditional forecast errors are shown for T–bill rates and the dollar in Figures 13A and 13B (comparable to 6C and 6D for the original model). The unconditional forecasting performance of the new model is summarized in the last column of Table 1. This table shows that the overall performance of the new specification is roughly comparable to that of the original specification. The new model forecasts better at longer horizons, and it does a better job of forecasting inflation, T–bills, and the dollar. But the model performs worse at shorter horizons and does a worse job of forecasting GNP and the two policy variables.

To gauge the performance of the new specification of the model in conditional forecasting, we replicated for the new specification the Box–Tiao scores depicted in Figure 10 and the scores for forecasts conditional on T–bill rates (COND3B). These replications are shown in Figure 14. Inspection of Figure 14 reveals that the measured accuracy of the conditional forecasts for this specification follow almost exactly the same pattern as the Box–Tiao scores for the original model specification. The one major exception is for forecasts that are conditional on T–bill rates during and after the period of the first oil shock. In the original model specification, the forecasts conditional on interest rates perform much worse than the unconditional or other conditional forecasts over this period. Under the new specification the performance of the forecasts conditioned on interest rates during the first oil shock episode is no worse than the unconditional (UNSCORB) forecasts or the forecasts conditioned on the deficit/GNP ratio (COND1B), but is still not as good as that of the forecasts conditioned on the base/debt ratio (COND2B).

On balance, we found that rescaling the initial moments of the time–varying coefficient BVAR resulted in relatively small, if any gains in the forecast accuracy of the model. Our original specification seems preferable due to the fact that it does not make explicit use of data from the forecast evaluation period (i.e., 1966–1986) in the calibration of the prior distributions of its parameters.

Another possibility for improving the forecasts of the model would be to allow for time variation in the variances of the equation error terms, as in Sims (1988). Since such variation was not allowed for in the original DLS methodology, we relegate these and other potential modifications to future research.

4.3 Forecast Deterioration Due to Changed Dynamics

Another concern about the dynamic forecasts shown in Figures 8 and 9 is that the lackluster performance of the model's forecasts over these episodes may not have resulted from policy changes associated with these episodes, but instead from changes in the economy's dynamics associated with the onset of recessions in 1980 and 1981. To address this concern, we simulated the model's dynamic forecasts for the onset of the recession associated with the 1974–75 oil shock episode. Figures 15A–F present the results of this exercise. In addition to showing the actual data over 1973–76, Figures 15A–F depict the model's unconditional forecasts as of 1973:4, and the model's 1973:4 conditional forecasts, conditional on the next 12 quarters of the base/debt ratio (CONDL) and the next 12 quarters of T–bill rates (CONDL2). On the one hand, overall accuracy of these forecasts seems roughly on par with the forecasts depicted in Figures 8 and 9, suggesting that the relative inaccuracy of these forecasts may be due to changes in dynamics over the business cycle rather than to changes in policy. On the other hand, Figures 15E and 15F suggest that both monetary and fiscal policy deviated considerably from their expected paths over this episode, lending support to the idea that policy changes are primarily responsible for the deterioration in the model's forecast accuracy. We suspect that a clean distinction between these two sources of forecast error will not easily obtain in a world where both the monetary and fiscal authorities are committed to following countercyclical policies.

5. SUMMARY AND CONCLUSION

In this paper we have considered the historical accuracy of the Doan–Litterman–Sims methodology for policy evaluation. We have done this by using Bayesian methods to fit an unconstrained VAR model to a small number of macroeconomic time series and then using this model to simulate the historical performance of the conditional forecasting technique described in Doan, Litterman, and Sims (1984). Our simulations suggest that over the last ten years conditioning on the actual future paths of policy variables would have in fact led to improvements in the accuracy of postsample forecasts for nonpolicy variables. The magnitude of these improvements is small, however—typically less than 5%. In particular, our consideration of two historical episodes associated with policy changes suggests that the performance of the DLS conditional forecasting technique is not particularly strong over such episodes. Our feeling is that this weakness reflects the fact that the DLS conditional forecasting technique ignores potential future variation in the model coefficients, which can be quite large during such policy change episodes.

Our overall conclusion is that the measurement exercises performed in this paper do not provide any evidence that the Lucas critique is invalid in an empirical sense. Instead, these exercises suggest that the DLS methodology for computing conditional forecasts is of limited use in predicting the effects of policy changes.

Our negative conclusion about DLS leaves unanswered two important questions:

1. Should the DLS method of policy analysis be rejected?

2. Should future research abandon DLS or attempt to improve it?

Our answers to these questions are more positive about DLS than our paper might suggest.

We are unaware of an empirical method for predicting the effects of policy changes that outperforms the DLS method. So we would view an argument that the DLS method should be rejected as an argument that quantitative policy analysis also should be rejected. We are not ready to make that argument.

Given the shortcomings of the DLS method which showed up in our analysis, new work on alternative methods of quantitative policy analysis would be welcome. But our analysis also showed that much could be gained with the DLS approach, if it were possible to predict changes in coefficients following policy changes. Sims (1988) has proposed altering the probability structure of BVARs as one way to put this possibility into practice. It also may be possible to predict coefficient changes by building more structure in the model with respect to policy rules and expectations conditioned on those rules (as in Miller–Roberds 1990).

ACKNOWLEDGEMENTS

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

APPENDIX: SPECIFICATION OF THE RANDOM COEFFICIENT VAR MODEL AND THE PRIOR DISTRIBUTION OF THE PARAMETERS

A.1 Summary of Model Specification

Write component i of equation (1) as

$$y_i(t) = \Sigma_{j=1}^{6}\Sigma_{\ell=1}^{6} a_{ij\ell}(t)y_j(t-\ell) + c_i(t) + u_i(t) \tag{A.1}$$

$$u_i(t) \sim \text{IIDN}(0,\sigma_i^2).$$

Defining $b_i(t) \equiv [a_{i11}(t), \cdots a_{i66}(t), c_i(t)]$, we assumed that $b_i(t)$ follows the law

$$b_i(t) = b_i(t-1) + e_i(t);\ e_i(t) \sim \text{IIDN}(0,W_i) \tag{A.2}$$

and $e_i(t) \perp u_i(s)$, for all t and s. Given an initial value of the mean ${}_{-1}\hat{b}_i(-1)$ and covariance matrix ${}_{-1}\Sigma_i(-1)$ for the initial Gaussian distribution of $b_i(-1)$, the DLS methodology applies the Kalman filter to obtain sequential estimates (conditional means) ${}_t\hat{b}_i(t)$ of $b_i(t)$. To be more specific, the Kalman filter is applied separately to each equation i, taking (A.1) as the observer equation and (A.2) as the state equation for the filter. Some inherent limitations of this technique are:

(1) Straightforward estimation of the model parameters is impractical (perhaps less so now than in 1984) due to the large dimension of the state vector $b_i(t)$, e.g., 37 (= 6 lags times 6 variables plus a constant term) for the model used in this paper. In the DLS paper, this limitation is circumvented by using heuristics to choose ${}_{-1}\hat{b}_i(-1)$, ${}_{-1}\Sigma_i(-1)$, W_i, and σ_i^2. This approach differs from the common strategy of econometric models, which is to adopt restrictions on the b_i's in order to reduce the dimension of the state (parameter) space. The DLS approach is motivated by the critique of econometric models in earlier works by Sims (1980,1982).

(2) Given the DLS methodology, in which no exact restrictions are placed on the model's coefficients, dimensionality considerations also rule out system Kalman filtering. Thus, no joint (cross–equation) posterior distribution of the model coefficients is produced by this technique without the introduction of additional assumptions regarding the covariances of parameters of different equations. The most obvious such assumption would be that the model's coefficients are independent across equations. This assumption seems unrealistic.

(3) Denoting the variables on the RHS of (A.1) as x(t–1), the one–step–ahead forecast error associated with a time t–1 forecast of $y_i(t)$ is $\nu_i(t) \equiv [y_i(t) - x(t-1)_{t-1}\hat{b}_i(t-1)]$. This error is equal to $u_i(t) + x(t-1)[_{t-1}\hat{b}_i(t-1) - b_i(t-1)]$, which is generally not Gaussian white noise unless $b_i(t)$ has been estimated without error. As noted by DLS (p. 69), forming linear conditional projections that use the sample covariance matrix of such errors is generally not a procedure that can be rigorously defended. The contradiction inherent in this aspect of the DLS methodology underscores the informal nature of their approach.

A.2. Specification of the Priors

The initial coefficient vectors $b_i(-1)$ were assumed to have prior mean given by

$$a_{ii1}(-1) = 1;\ a_{ij\ell}(-1) = 0 \text{ otherwise};\ c_i(-1) = 0. \tag{A.3}$$

The prior covariance matrices $_{-1}\Sigma_i(-1)$ were constructed by first estimating a fixed–coefficient version of the model. In this version of the model, the matrices W_i are set to zero. In this case, the variances of the equation error terms, i.e., the σ_i^2 terms, cancel out of the Kalman filter calculations and need not be specified a priori. The $_{-1}\Sigma_i$ matrices were assumed to be, except for the constant term, diagonal with the square root of the typical element being given by

$$S(i,j,\ell) = g \cdot f(i,j) \cdot s(j) / [s(i) \cdot \ell] \tag{A.4}$$

where s(i) represents the standard error of the estimate from a univariate autoregression of variable i. The "overall tightness" parameter g was initially set equal to 0.2, while the matrix of weights f(i,j) was initially taken to be given by f(i,i) = 1, and f(i,j) = 0.3 for i $\neq$ j, which are close to the default values.given by the RATS program (see Doan 1988). Experimentation with this version of the BVAR model yielded relatively poor out-of-sample forecast performance. Experimentation with the weights f(i,j) led to the specification for the weighting matrix [f(i,j)] given by

$$\begin{array}{cc} \begin{array}{cccccc} (1) & (2) & (3) & (4) & (5) & (6) \end{array} & \\ \left[\begin{array}{cccccc} 1.0 & 0.3 & 0.5 & 0.3 & 0.3 & 0.3 \\ 0.3 & 1.0 & 0.3 & 0.3 & 0.3 & 0.3 \\ 0.2 & 0.2 & 1.0 & 0.2 & 0.2 & 0.2 \\ 0.3 & 0.3 & 0.3 & 1.0 & 0.3 & 0.3 \\ 0.3 & 0.3 & 0.3 & 0.3 & 1.0 & 0.3 \\ 0.4 & 0.2 & 0.2 & 0.2 & 0.2 & 1.0 \end{array}\right] & \begin{array}{c} (1) \\ (2) \\ (3) \\ (4) \\ (5) \\ (6) \end{array} \end{array} .$$

The order of the variables in the matrix above is the same as given in the text. The prior on the constant term was taken as "flat" or uninformative. Forecasting statistics for this model are given in the "Fixed–Coefficient BVAR" column of Table 1.

Our time–varying coefficient BVAR model assumes the same prior mean for the $b_i(-1)$ as the fixed–coefficient model. To specify the prior covariance matrices for the $b_i(-1)$, we began by taking the 1965:4 estimates of the the covariance matrices from the fixed–coefficient model, scaled up by a constant factor π_1. That is, letting ${}_{1965:4}\Sigma_i$ represent the 1965:4 estimate of the covariance matrix of the coefficients of equation i of the fixed–coefficient model, we initially took ${}_{-1}\Sigma_i(-1)$ to be equal to $\pi_1 \cdot {}_{1965:4}\Sigma_i$.

At this point, priors on sums of coefficients were introduced using dummy observations of the form

$$S_{ij}\theta_{ij} = v_{ij} \tag{A.5}$$

where S_{ij} is equal to $\{\pi_2 s(j)/s(i)\}[1...1]$ or $\{\pi_3 s(j)/s(i)\}[1...1]$; θ_{ij} is the subvector of coefficients on variable j in $b_i(-1)$; and $v_{ij} \sim N(\delta_{ij},1)$. If v_{ij} had zero variance for all i and j, this would imply a model in first differences. Experimentation with out–of–sample forecasts revealed that assigning a relatively large weight (π_2) to this dummy observation for GNP and the two policy variables, and a relatively low weight (π_3) for inflation and the two financial market variables yielded the best results. Letting ${}_{-1}\Sigma_{ij}(-1)$ be the covariance matrix of θ_{ij} implied by the appropriate submatrix of ${}_{-1}\Sigma_i(-1)$, conditional on (A.5) the covariance matrix of θ_{ij} is given by [cf. DLS, p.14]

$$ {}_{-1}\Sigma^*_{ij}(-1) \equiv {}_{-1}\Sigma_{ij}(-1) - \{[{}_{-1}\Sigma_{ij}(-1)]S'_{ij}S_{ij}[{}_{-1}\Sigma_{ij}(-1)]/(1+S_{ij}[{}_{-1}\Sigma_{ij}(-1)]S'_{ij})\}. \tag{A.6} $$

The specification of the prior second moments for the model was completed by setting the covariance matrices of coefficient shocks $e_i(t)$ proportional to the prior covariance matrix of the coefficients, i.e., by setting $W_i = \pi_4 \cdot {}_{-1}\Sigma^*_i$. The variances of the equation error terms were set equal to the corresponding values from the univariate models; i.e., we set $\sigma_i^2 = (s_i^*)^2$, where s_i^* represents the standard error of the estimate from equation i of the fixed–coefficient model. Forecasting experiments with the vector π of scale factors ("hyperparameters") led us to the specification $\pi = [\ 1\ 10^3\ 1\ 10^{-6}]$. This specification comes close to a first differences specification for real GNP and the policy variables and assigns a relatively minor role to time variation in the model parameters.

A.3 Impact of the Priors

Some idea of the "tightness" of the priors used in BVAR technique can be obtained by comparing the relative sizes of the covariance matrices for the coefficients of three VAR models: an unrestricted model, the fixed–coefficient model described above, and the time–varying coefficient model. The accompanying table compares the log determinants of the 1986:4 (end–of–sample) covariance matrices of the coefficients of the three models. Column (1) of the table gives the scaled difference between the log determinants for the unrestricted VAR vs. the fixed–coefficient BVAR, and column (2) gives the scaled difference between the log determinants of the fixed– vs. the time–varying coefficient BVARs. The scaling factor is 50/37 so as to give an approximate average "percentage" change in the posterior standard deviations of the coefficients in equation i (i.e., 100% divided by 2 to get standard deviations, then divided by the dimension of the covariance matrix). A negative number represents a reduction in uncertainty; e.g., an entry of –100 means that on average the posterior standard deviation of the coefficients of equation i in the less restrictive model is roughly $e \doteq 2.7$ times as large as the posterior standard deviations in the more restrictive model. The shrinkage obtained using the time–varying model as opposed to the unrestricted model can be obtained by summing the two columns of the table. Imposing an exact linear restriction on the coefficients of any equation would yield a reduction of $-\infty$, due to the singularity of the restricted covariance matrix.

From the first column of the table, it can be seen that the use of standard "Minnesota" priors in the fixed–coefficient model results in posterior standard deviations of the model coefficients that are roughly one–third to one–fourth as small as those from an unrestricted VAR. The shrinkage is fairly uniform across equations. By contrast, the priors employed in the time–varying model are somewhat less restrictive than those employed in the fixed–coefficient model in the case of inflation, somewhat more restrictive in the case of interest rates, the dollar, and the base/debt ratio, and about equally restrictive in the case of real GNP and the deficit/GNP ratio. On the whole, the values in the table suggest that the posterior distribution of the model's coefficients implied by our final specification (i.e., the time–varying coefficient model) is not inappropriately dominated by our prior distribution for the model coefficients. The influence of our priors is no greater for three of the model's equations than for a textbook "Minnesota" type of prior. The equations where the influence of our priors is greatest are for the two financial market variables plus the monetary policy variable, i.e., variables for which many economists would generally expect optimal forecasts to be very close to those given by a random walk model.

References

Barro, R. J. (1986), "U.S. Deficits Since World War I," Scandinavian Journal of Economics, 88, 195–222.

Blanchard, O. J. (1984), "The Lucas Critique and the Volcker Deflation," American Economic Association Papers and Proceedings, 74, 211–215.

Box, G. E. P., and Tiao, G. C. (1976), "Comparison of Forecast and Actuality," Applied Statistics, 25, 195–200.

Christiano, L. J. (1986), "Money and the U.S. Economy in the 1980s: A Break From the Past?" Federal Reserve Bank of Minneapolis Quarterly Review, 10 (Summer), 2–13.

Cooley, T. F.; LeRoy, S. F.; and Raymon, N. (1984), "Econometric Policy Evaluation: Note," American Economic Review, 74, 467–470.

Doan, T. A. (1988), User's Manual: RATS (Regression Analysis of Time Series) (Version 3.00), Evanston, ILL: VAR Econometrics.

Doan, T.; Litterman, R.; and Sims, C. (1984), "Forecasting and Conditional Projection Using Realistic Prior Distributions, Econometric Reviews 3, 1–100.

Goodwin, G. C., and Sin, K. S. (1984), Adaptive Filtering Prediction and Control, Englewood Cliffs, NJ: Prentice–Hall.

Litterman, R. B. (1984), "Forecasting and Policy Analysis With Bayesian Vector Autoregression Models," Federal Reserve Bank of Minneapolis Quarterly Review, 8 (Fall), 30–41.

Lucas, R. E., Jr. (1976), "Econometric Policy Evaluation: A Critique," in: The Phillips Curve and Labor Markets, ed. K. Brunner and A. H. Meltzer, Carnegie–Rochester Conference Series on Public Policy (vol. 1), Amsterdam: North–Holland, pp. 19–46.

Miller, P. J., and Roberds, W. (1987), "The Quantitative Significance of the Lucas Critique" (unabridged), Research Department Staff Report 109, Federal Reserve Bank of Minneapolis.

_______________ (1990), "How Little We Know About the Effects of Budget Policy," Research Department Staff Report 120, Federal Reserve Bank of Minneapolis.

Miller, P. J., and Wallace, N. (1985), "International Coordination of Macroeconomic Policies: A Welfare Analysis, Federal Reserve Bank of Minneapolis Quarterly Review, 9 (Spring), 14–32.

Modigliani, F. (1986), "Comment on R. J. Barro, U.S. Deficits Since World War I," Scandinavian Journal of Economics, 88, 223–234.

Sargent, T. J. (1984), "Autoregressions, Expectations, and Advice, American Economic Review, 74, 408–415.

Sims, C. A. (1980), "Macroeconomics and Reality," Econometrica, 48, 1–48.

_______________ (1982), "Policy Analysis With Econometric Models," Brookings Papers on Economic Activity, 1, 107–152.

_______________ (1986), "Are Forecasting Models Usable for Policy Analysis?" Federal Reserve Bank of Minneapolis Quarterly Review, 10 (Winter), 2–16.

_______________ (1988), "A Nine Variable Probabilistic Macroeconomic Forecasting Model," unpublished manuscript.

Wallace, N. (1984), "Some of the Choices for Monetary Policy," Federal Reserve Bank of Minneapolis Quarterly Review, 8 (Winter), 15–24.

Table 1. Theil U-Statistics[a] for
Out-of-Sample Forecasting Performance, 1966:1-1986:4

		BVAR Model		
			Random-Coefficient	
Variable	Univariate Model	Fixed-Coefficient	Original	New
1-Step Horizon				
Log Real GNP	0.843	0.836	0.807	0.824
Inflation	1.132	0.913	0.937	0.827
T-bill Rate	1.018	1.033	1.022	0.978
Dollar	0.951	0.958	0.966	0.937
Base/Debt	0.653	0.637	0.637	0.763
Deficit/GNP	1.107	1.020	1.000	1.017
Log Determinant[b]	-8.146	-9.020	-9.110	-8.400
Average Improvement[c]		(7.3%)	(0.7%)	(-5.9%)
4-Step Horizon				
Log Real GNP	0.749	0.747	0.650	0.718
Inflation	1.300	0.825	0.897	0.742
T-bill Rate	1.088	1.005	0.983	0.909
Dollar	0.999	1.066	1.076	0.970
Base/Debt	0.868	0.776	0.745	0.900
Deficit/GNP	1.153	0.999	0.891	0.984
Log Determinant	2.607	1.205	0.291	0.280
Average Improvement		(11.7%)	(7.6%)	(7.7%)
8-Step Horizon				
Log Real GNP	0.681	0.830	0.557	0.624
Inflation	1.197	0.846	0.881	0.856
T-bill Rate	1.158	0.886	0.871	0.856
Dollar	0.985	1.034	1.044	0.947
Base/Debt	1.005	0.941	0.832	0.993
Deficit/GNP	1.088	0.906	0.823	0.970
Log Determinant	6.658	5.982	3.970	3.329
Average Improvement		(5.6%)	(16.8%)	(22.1%)
12-Step Horizon				
Log Real GNP	0.595	0.955	0.459	0.515
Inflation	1.099	0.915	0.855	0.813
T-bill Rate	1.201	0.828	0.826	0.819
Dollar	0.950	0.897	0.909	0.891
Base/Debt	1.094	1.087	0.878	1.042
Deficit/GNP	1.053	0.900	0.815	0.953
Log Determinant	8.427	8.243	5.169	4.243
Average Improvement		(1.5%)	(25.6%)	(27.8%)

NOTES FOR TABLE 1

[a]Ratio of RMS forecast error to RMS forecast error of a forecast of no change.

[b]Value of log determinant of the covariance matrix of out-of-sample forecast errors at the given forecast horizon.

[c]Approximation to average percentage reduction in standard error of the forecast obtained by taking the difference in log determinants and multiplying by 8.33. (Divide by 12 to get standard errors for 6 variables and multiply by 100 to get percent.) This first improvement is of the fixed-coefficient BVAR over the univariate model; the second is of the random-coefficient BVAR over the fixed-coefficient BVAR.

Table 2. Out-of-Sample Forecasting Performance Measured by Log Determinants of the Covariance Matrix of Forecast Errors for Nonpolicy Variables (average percentage improvements shown in parentheses)[a]

		Conditional on[b]				
				Both (1) & (2)		
Step	Unconditional	Base/Debt (1)	Deficit/GNP (2)	DLS (3)	Sequential[c] (4)	Ideal[d] (5)
1976:1-1986:4						
1	-5.209	-5.404 (2.43)	-5.416 (2.58)	-5.525 (3.94)	-5.518 (3.87)	-6.981 (22.14)
4	0.05814	-0.3204 (4.73)	-0.2369 (3.69)	-0.5440 (7.58)	-0.4704 (6.61)	-3.001 (38.25)
8	2.125	1.711 (2.73)	2.038 (1.08)	1.742 (4.78)	2 1.837 (3.60)	-1.692 (47.71)
12	2.599	3.168 (-7.12)	2.726 (-1.59)	2.987 (-4.86)	2.991 (-4.90)	-1.812 (55.14)

[a]Approximate average percentage reduction in standard error of the conditional forecast over the unconditional forecast, obtained by taking the difference in log determinants and multiplying by 12.75. (Divide by 8 to get standard errors for 4 variables and multiply by 100 to get percent.) A negative number indicates a deterioration of forecasting performance relative to the unconditional forecast.

[b]Conditioned on the next 12 quarters ahead of the indicated variable, except at the end of the sample, where the conditioning set is truncated.

[c]Forecasts obtained by Kalman smoothing the model coefficients using the DLS conditional forecasts, then doing a second DLS-type conditional forecast using the smoothed coefficients.

[d]Forecasts obtained by Kalman smoothing the model coefficients using the next 12 quarters of actual data.

APPENDIX TABLE

Relative Log Determinants of Coefficient Covariance Matrices

Variable	Fixed vs. Unrestricted (1)	Time-Varying vs. Fixed (2)
Log Real GNP	−118.27	2.18
Inflation	−124.56	61.89
T-bill Rate	−147.08	-51.91
Dollar	−129.99	-54.92
Base/Debt	−124.11	-64.06
Deficit/GNP	−139.89	13.66

FIGURE TITLES

FIGS. 1A-1B. POLICY SERIES

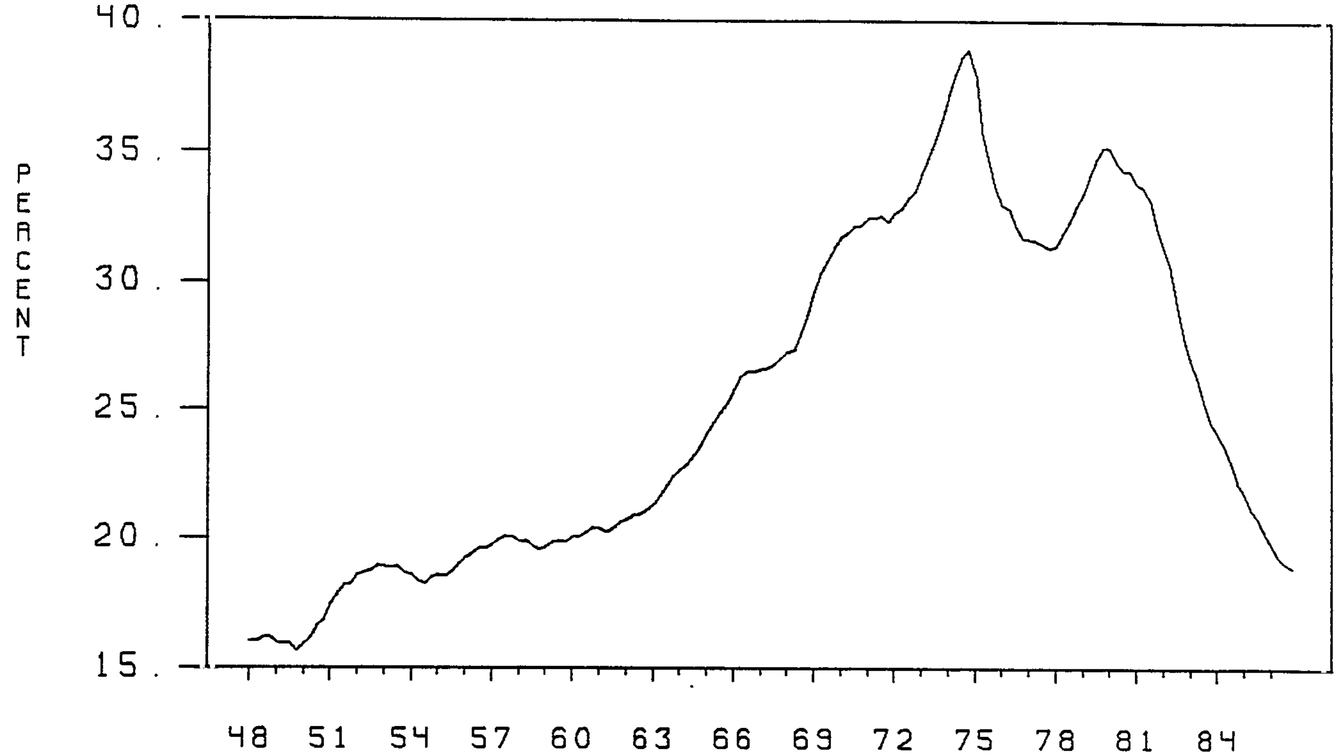

FIG. 1A. RATIO OF MONETARY BASE TO DEBT

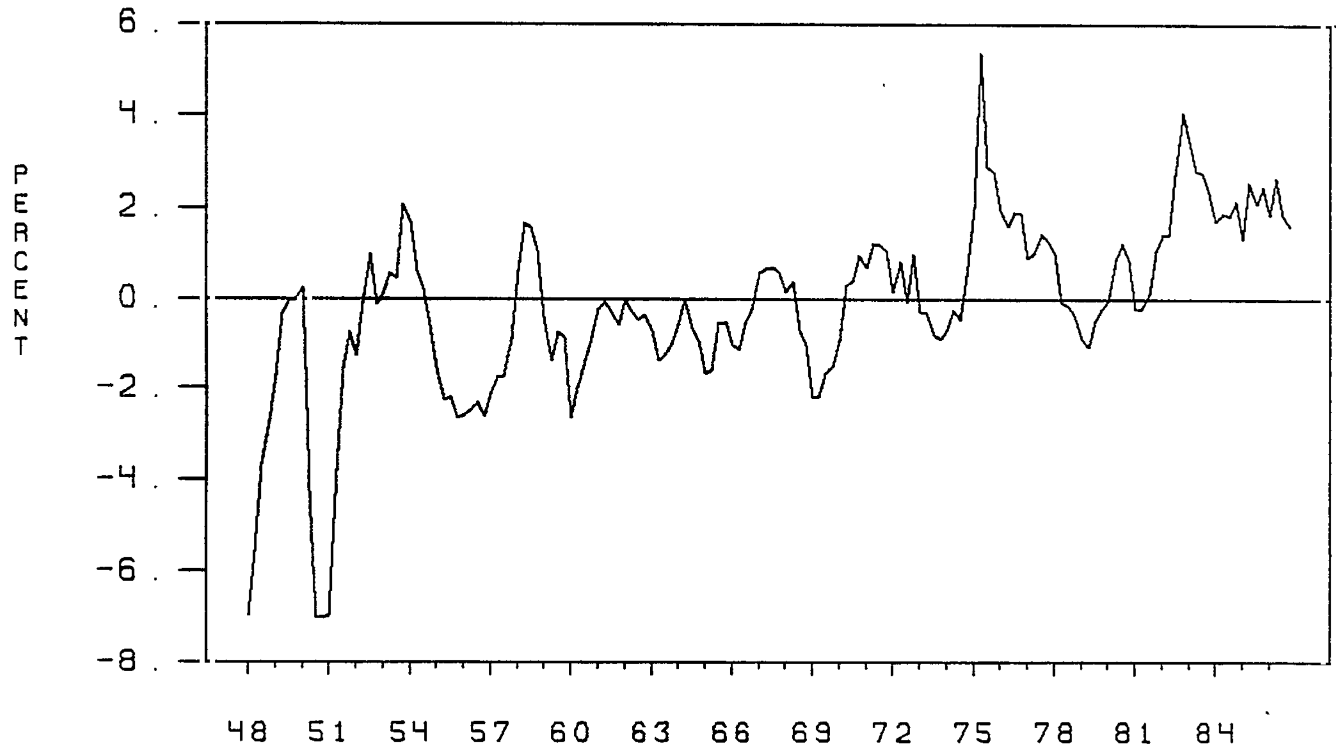

FIG. 1B. RATIO OF PRIMARY DEFICIT TO GNP

FIGS. 2A-2B. ONE-STEP-AHEAD FORECAST ERRORS

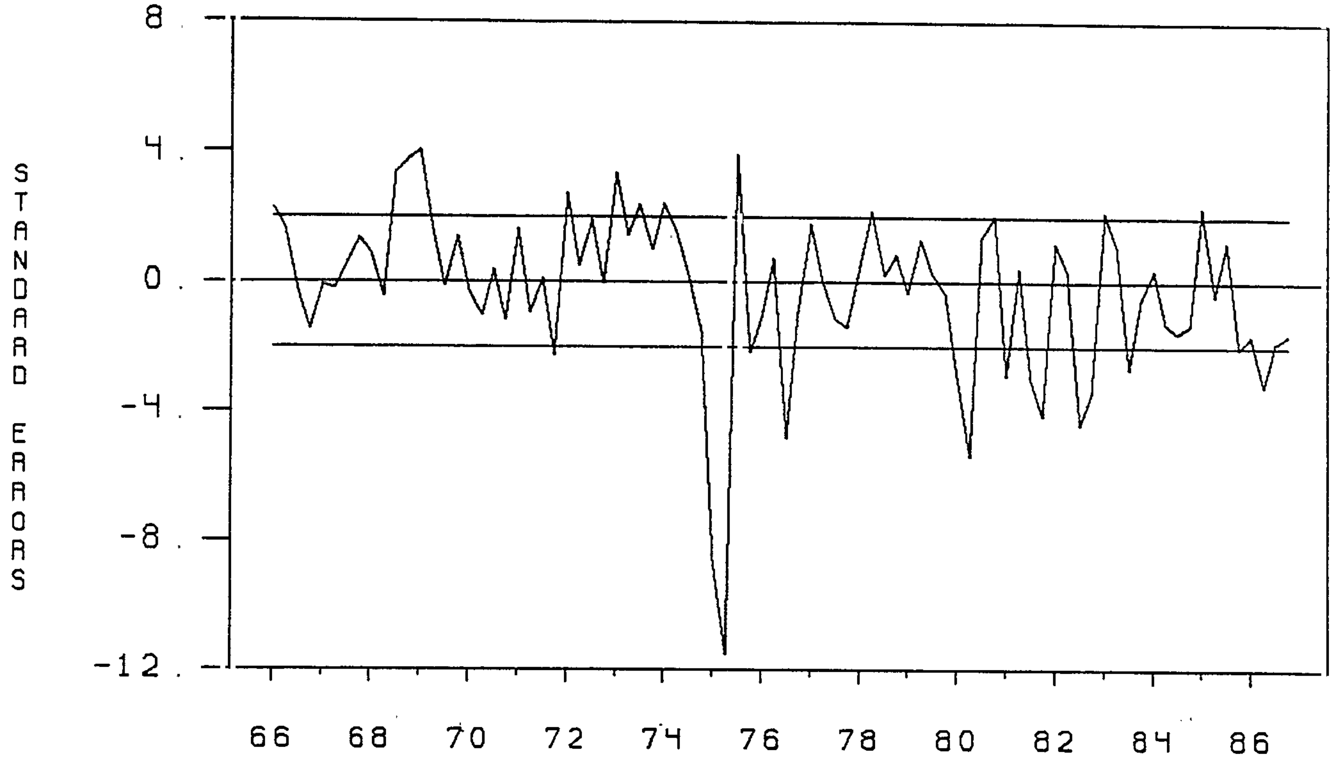

FIG. 2A. RATIO OF MONETARY BASE TO DEBT

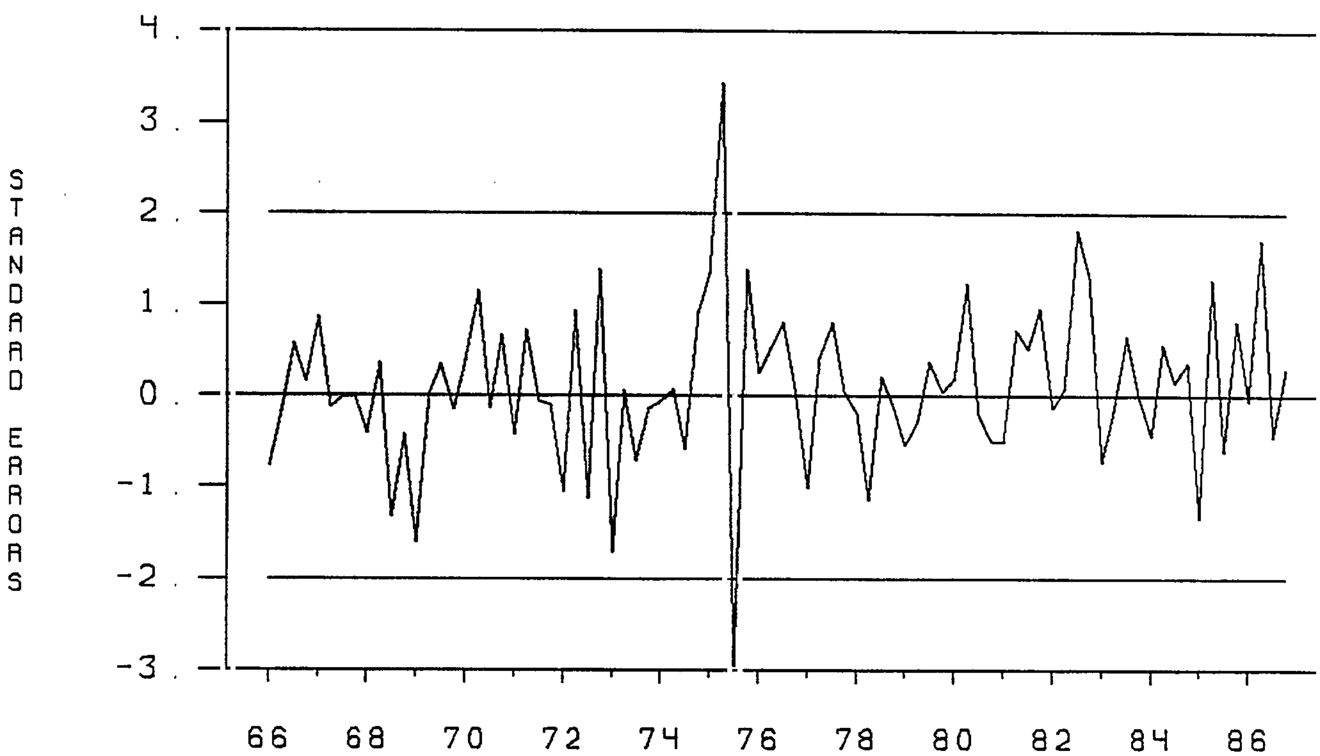

FIG. 2B. RATIO OF PRIMARY DEFICIT TO GNP

FIGS. 3A-3B. COEFFICIENT ON FIRST OWN LAG

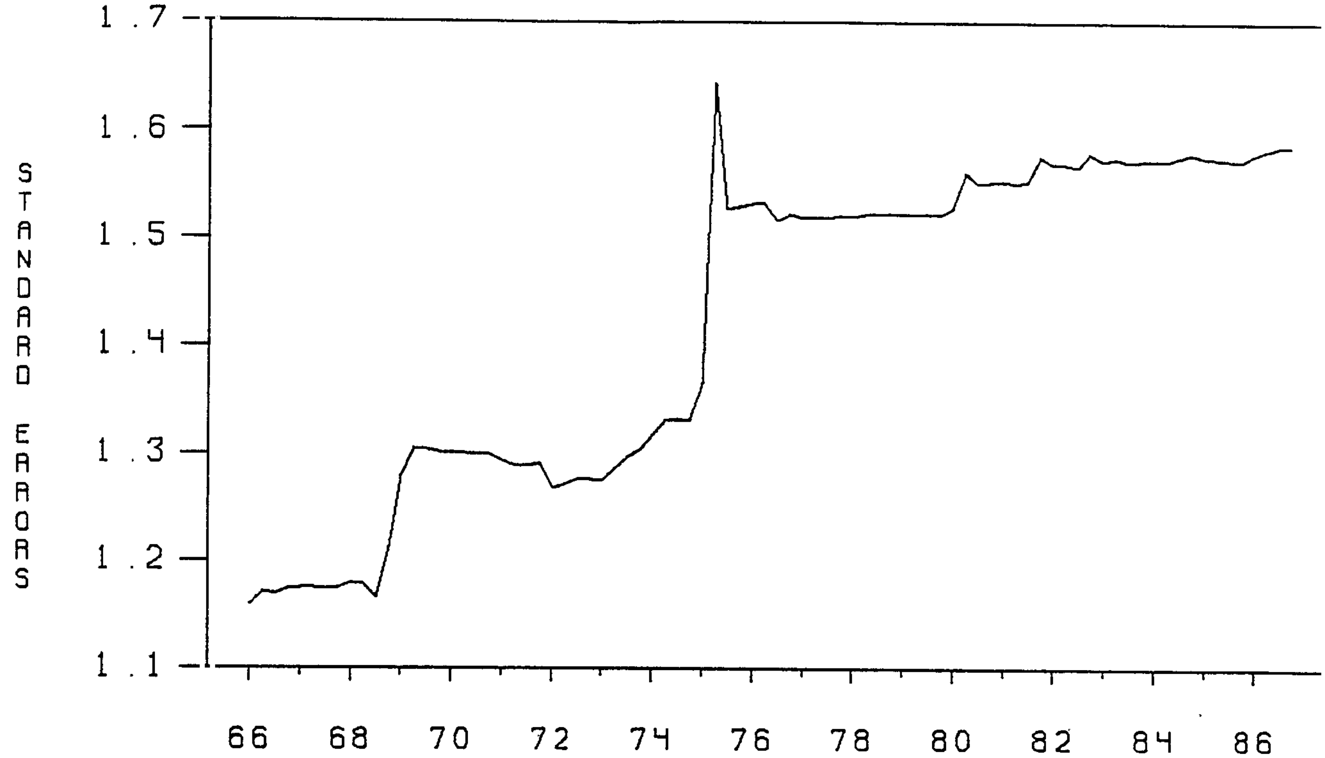

FIG 3A. RATIO OF MONETARY BASE TO DEBT

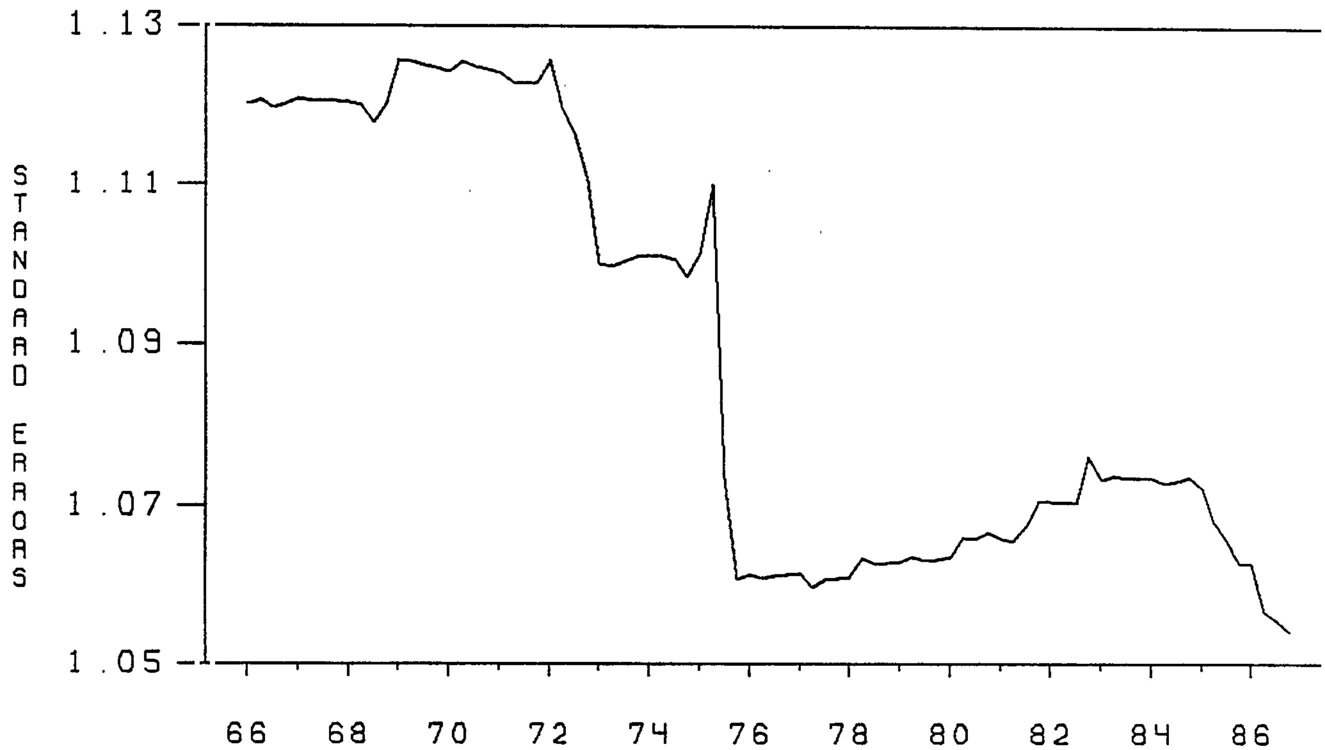

FIG 3B. RATIO OF PRIMARY DEFICIT TO GNP

FIG. 4. FED POLICY CHANGE EXPERIMENT

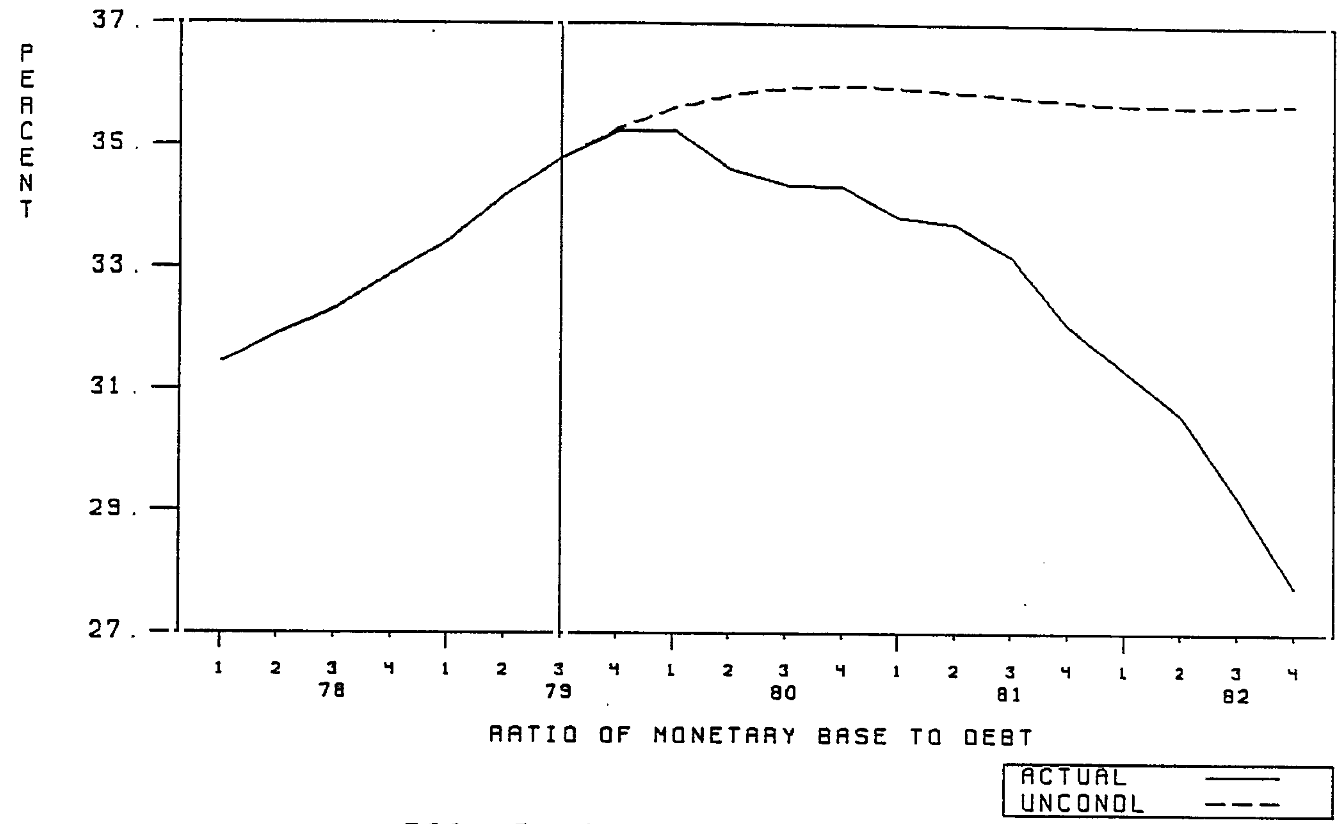

FIG. 5. REAGAN TAX CUT EXPERIMENT

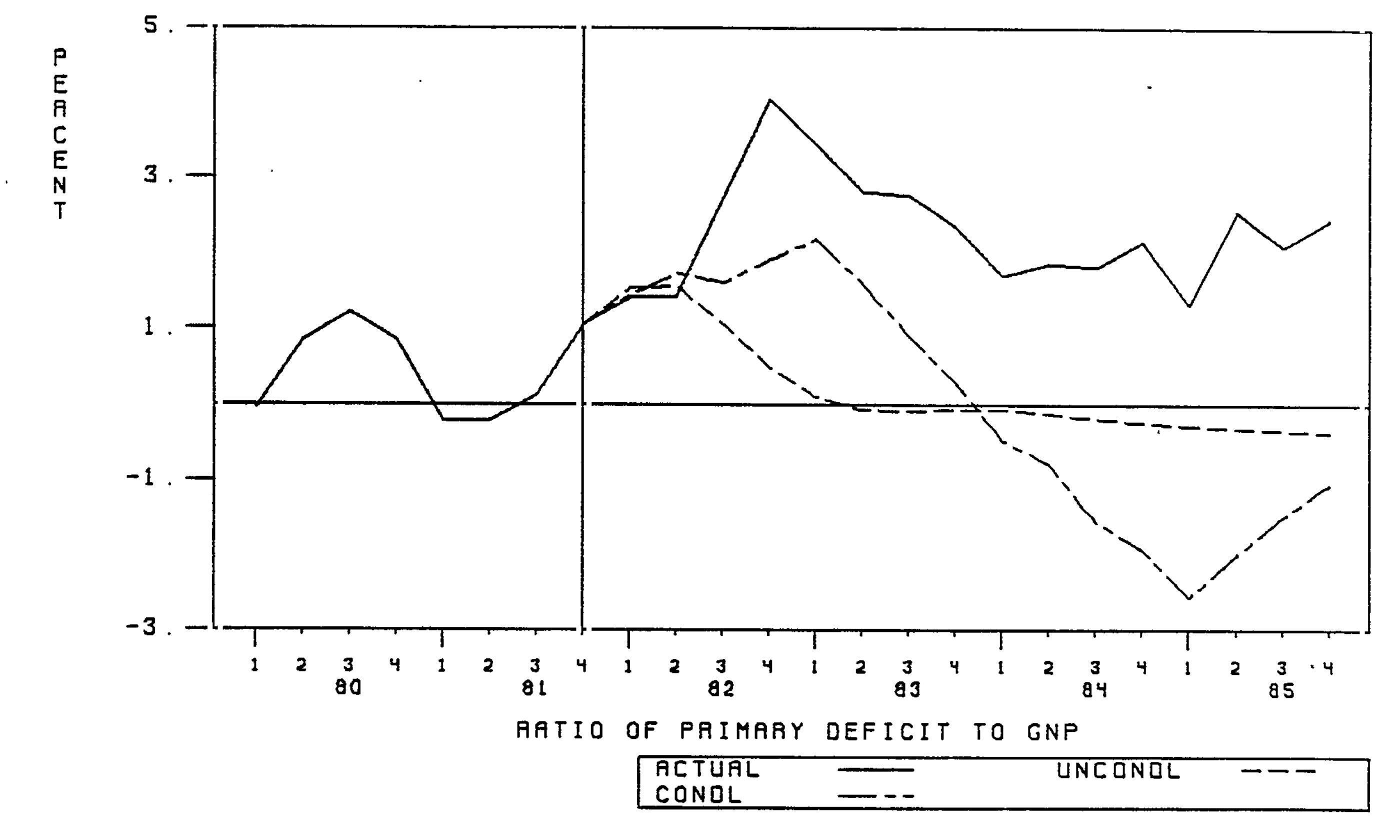

FIGS. 6A-6D. ONE-STEP-AHEAD FORECAST ERRORS

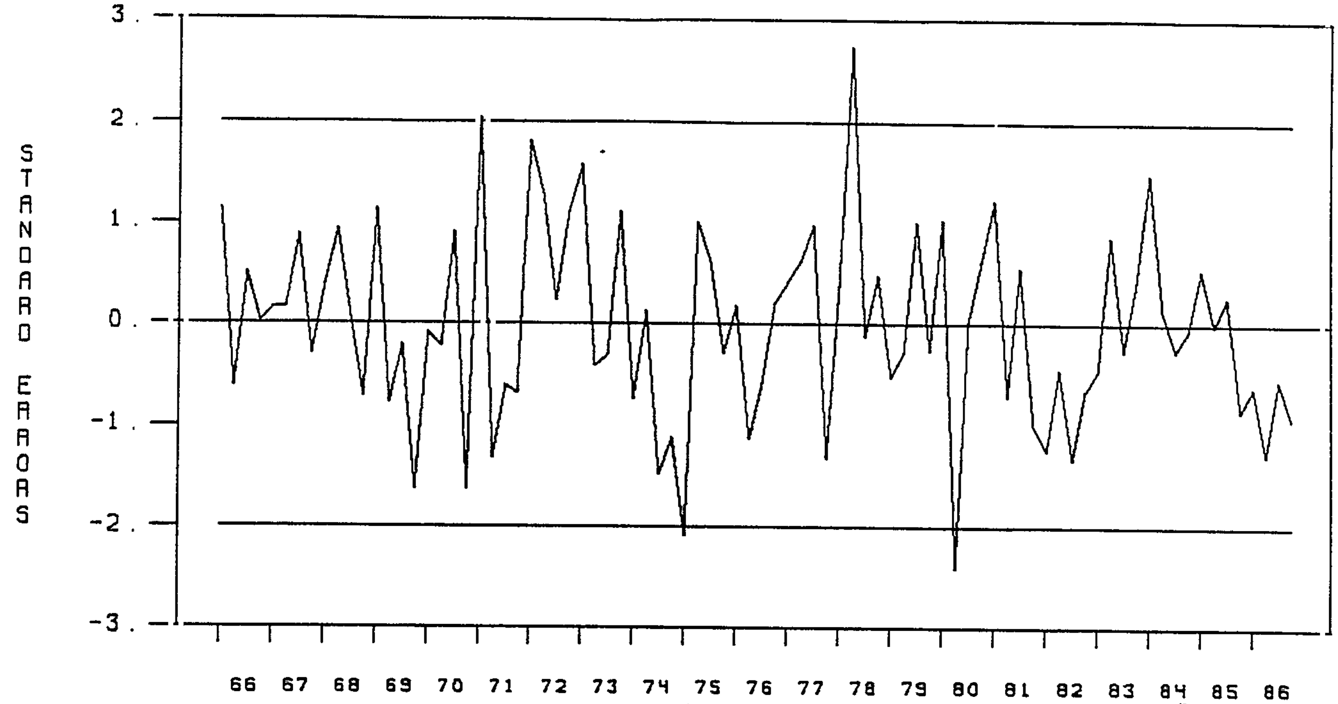

FIG. 6A. LOG OF REAL GNP

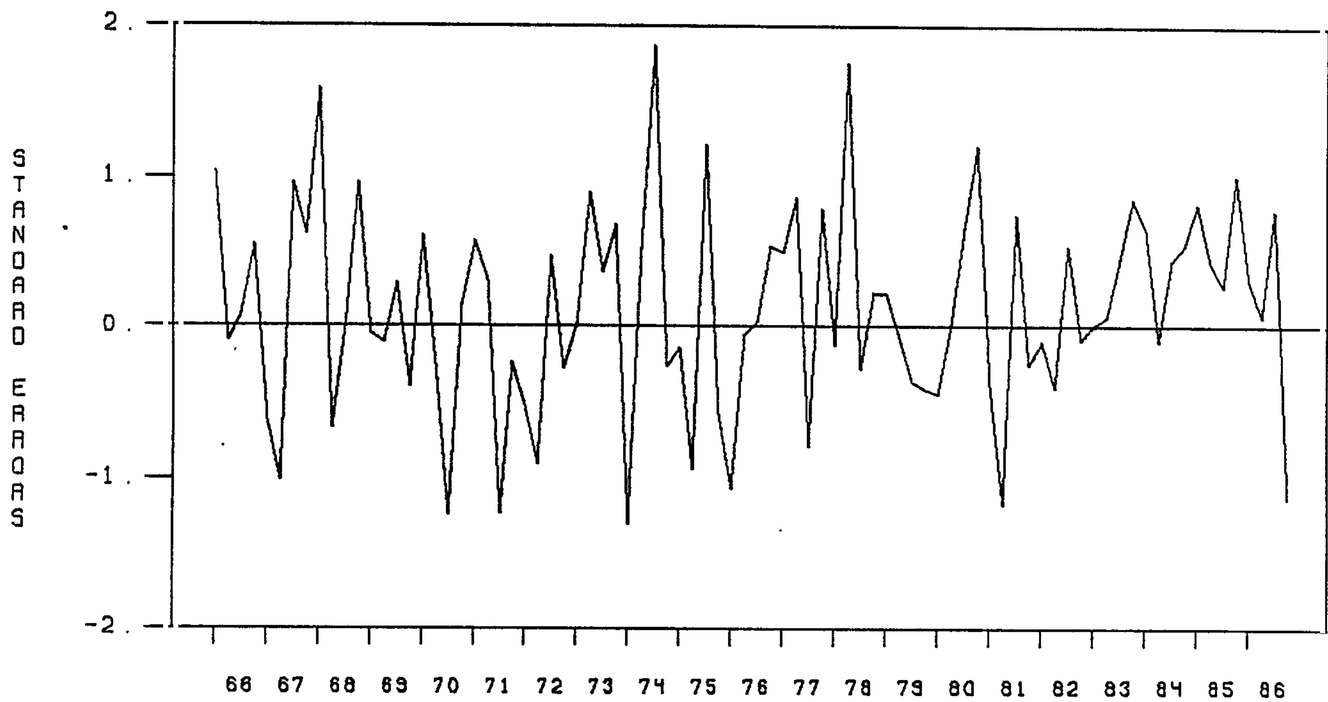

FIG. 6B. INFLATION (GNP DEFLATOR)

FIGS. 6A-6D. ONE-STEP-AHEAD FORECAST ERRORS

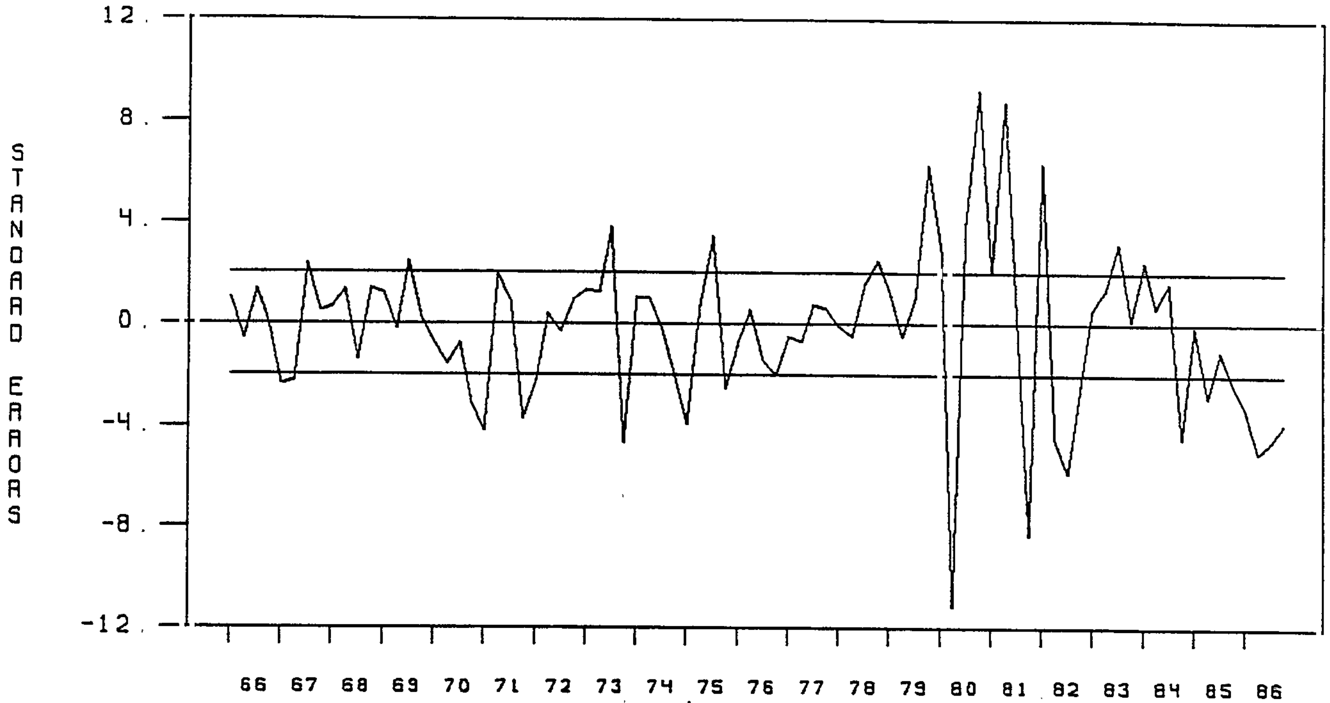

FIG. 6C. 3-MONTH TREASURY-BILL RATE

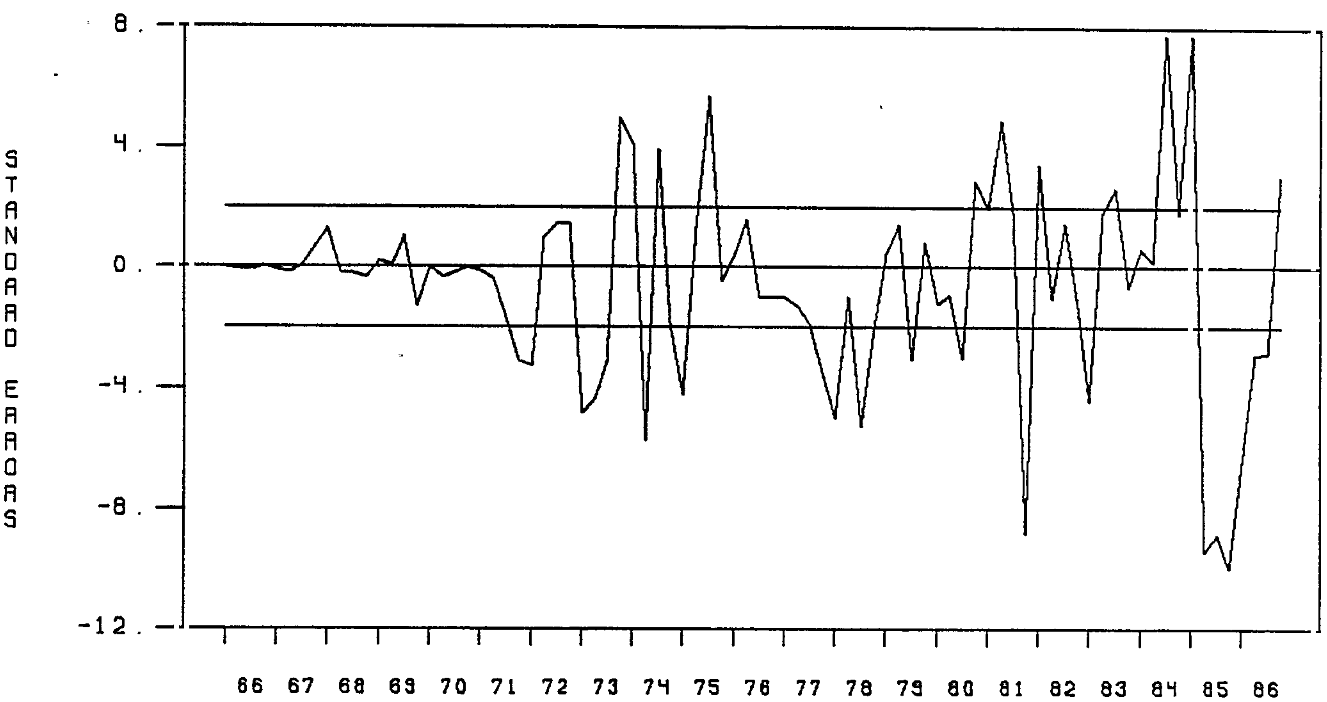

FIG. 6D. VALUE OF THE TRADE-WEIGHTED DOLLAR

FIGS. 7A. AND 7B. LOG OF REAL GNP

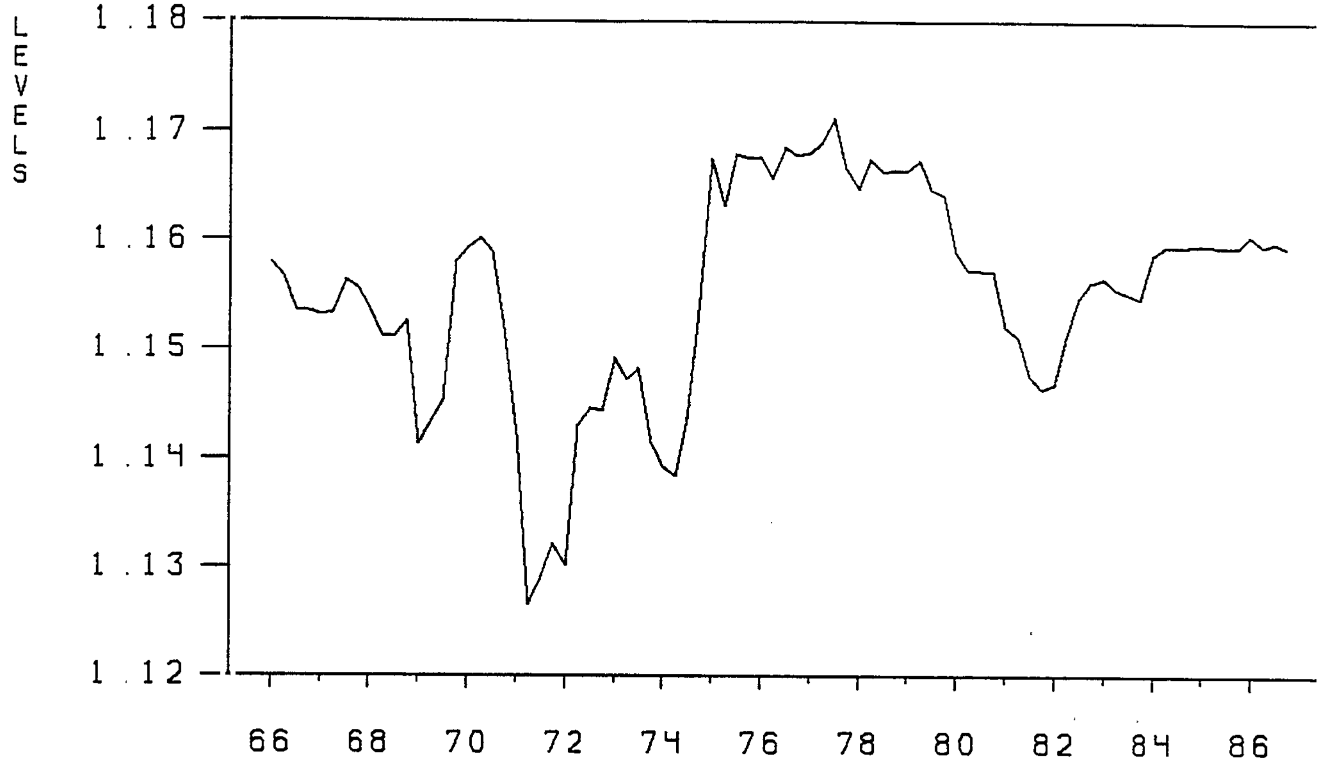

FIG. 7A. COEFFICIENT ON FIRST OWN LAG

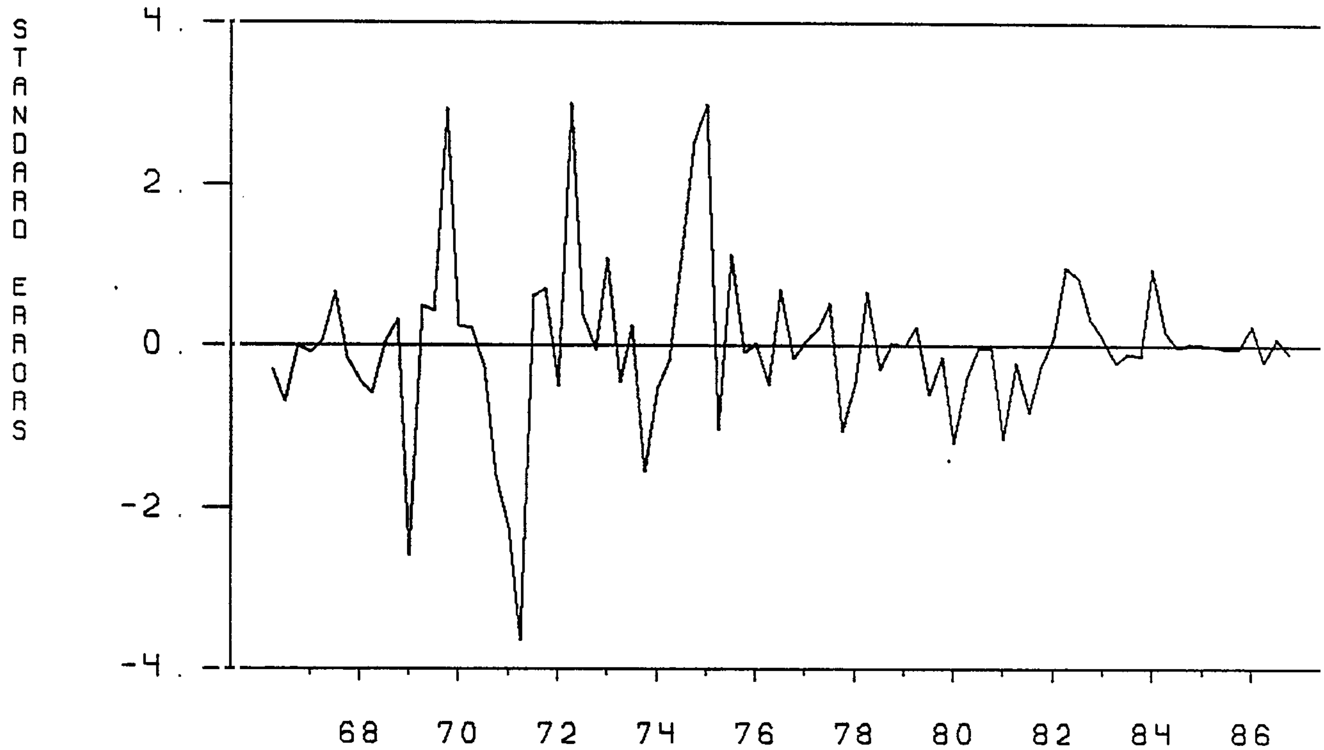

FIG. 7B. STANDARDIZED FIRST DIFFERENCES

FIGS. 7C. AND 7D. INFLATION (GNP DEFLATOR)

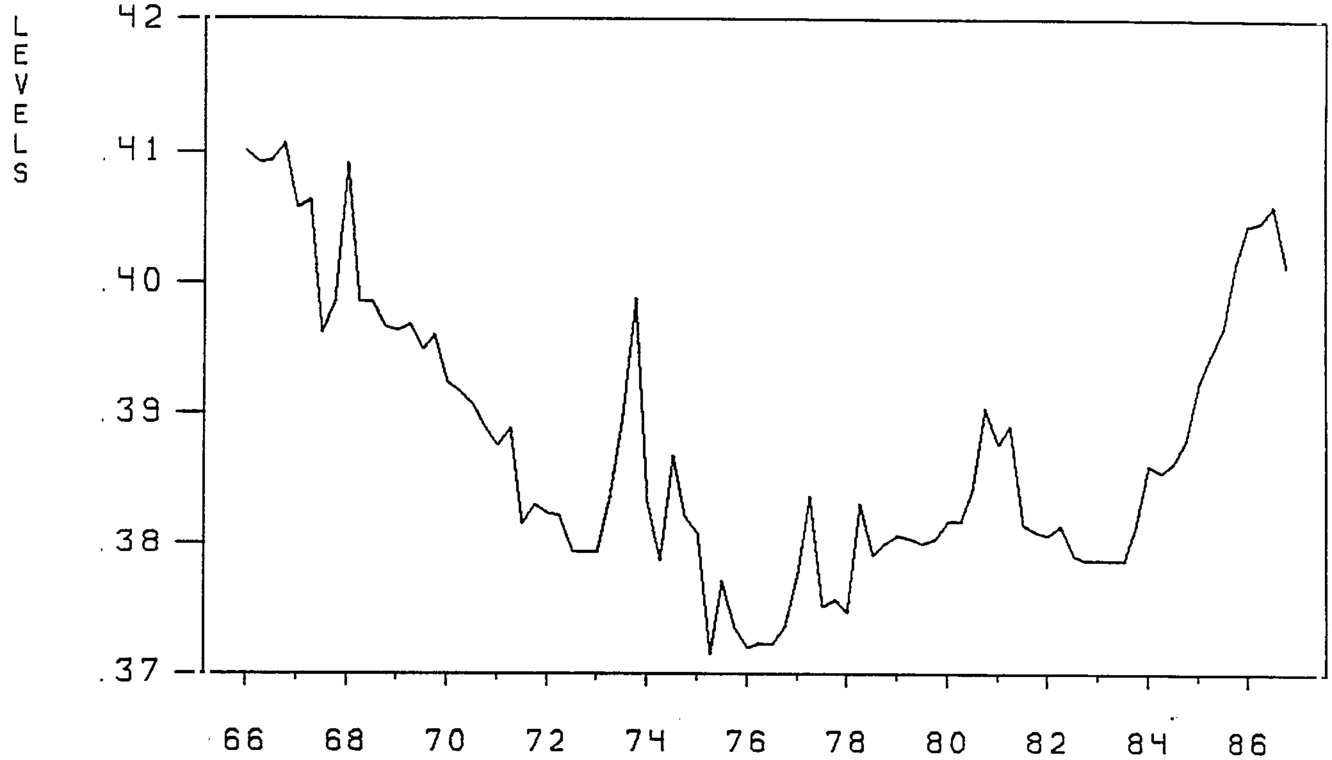

FIG. 7C. COEFFICIENT ON FIRST OWN LAG

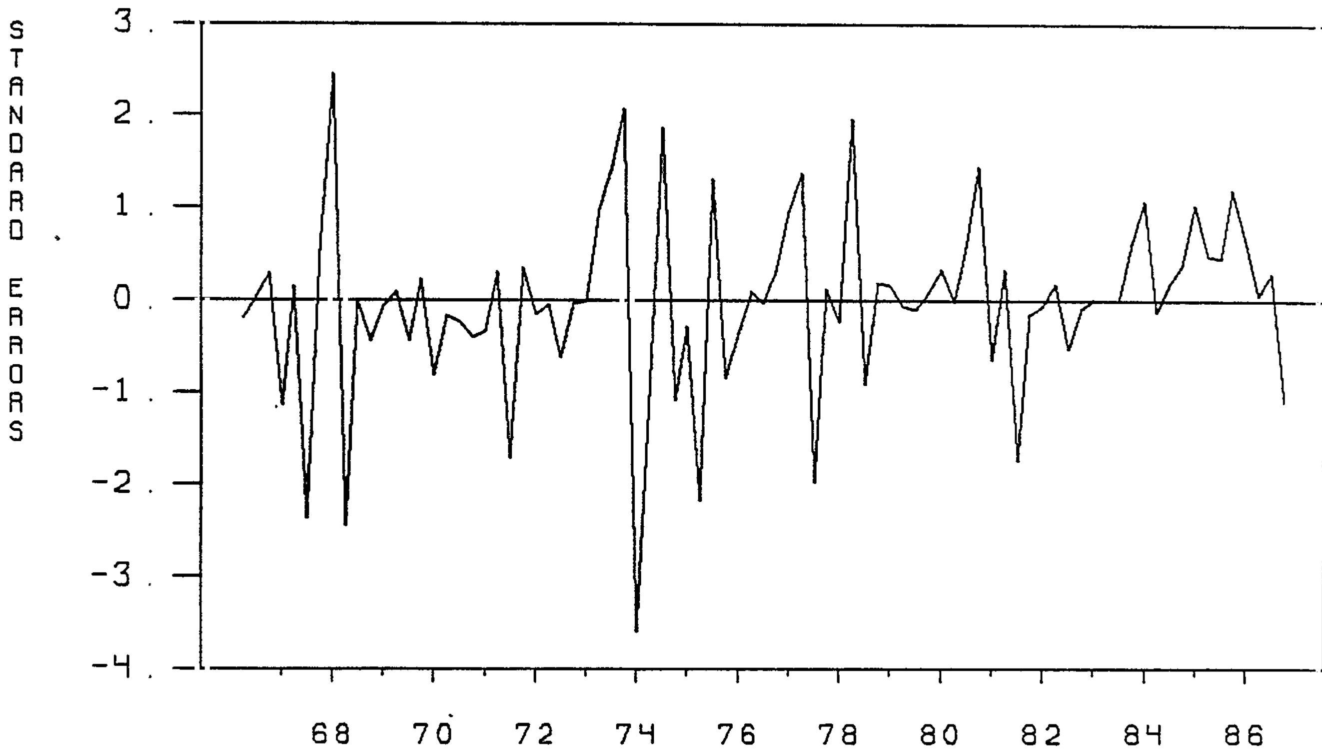

FIG. 7D. STANDARDIZED FIRST DIFFERENCES

FIGS. 7E. AND 7F. 3-MONTH TREASURY BILL RATE

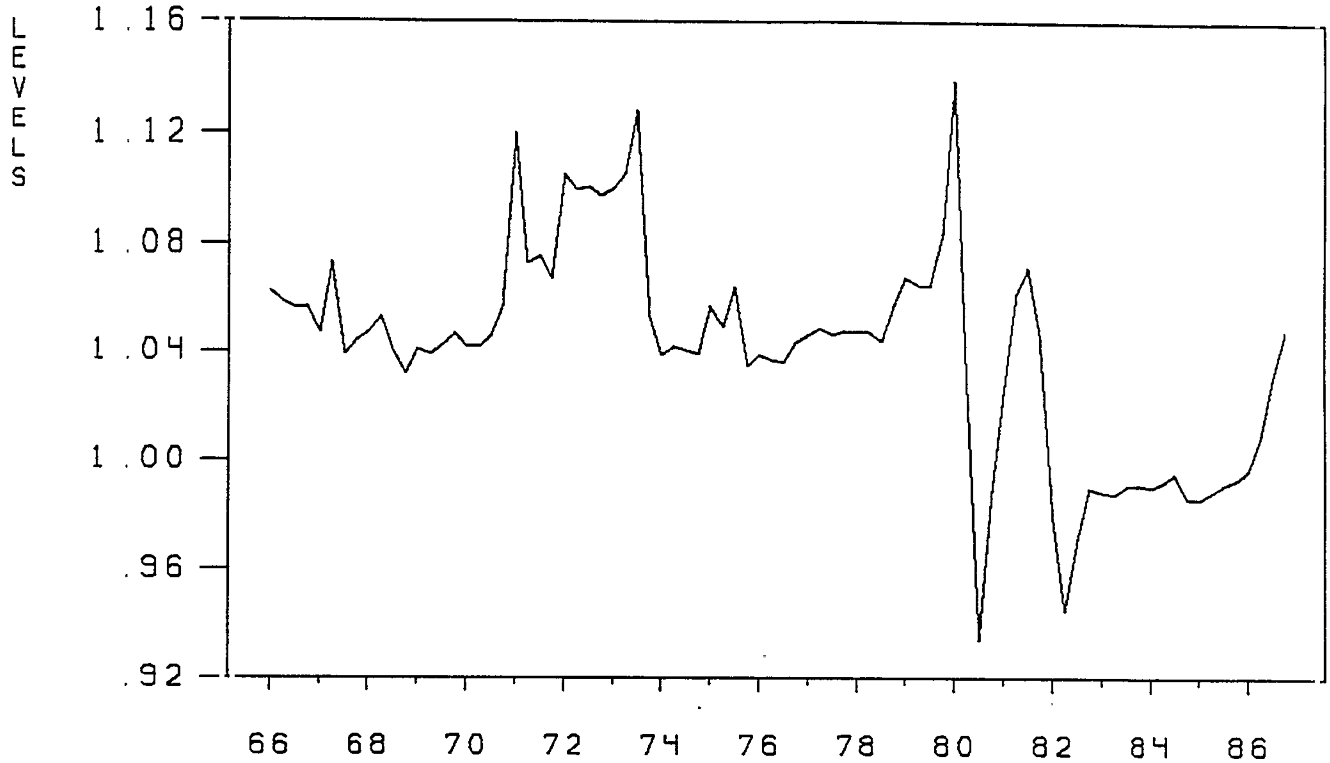

FIG. 7E. COEFFICIENT ON FIRST OWN LAG

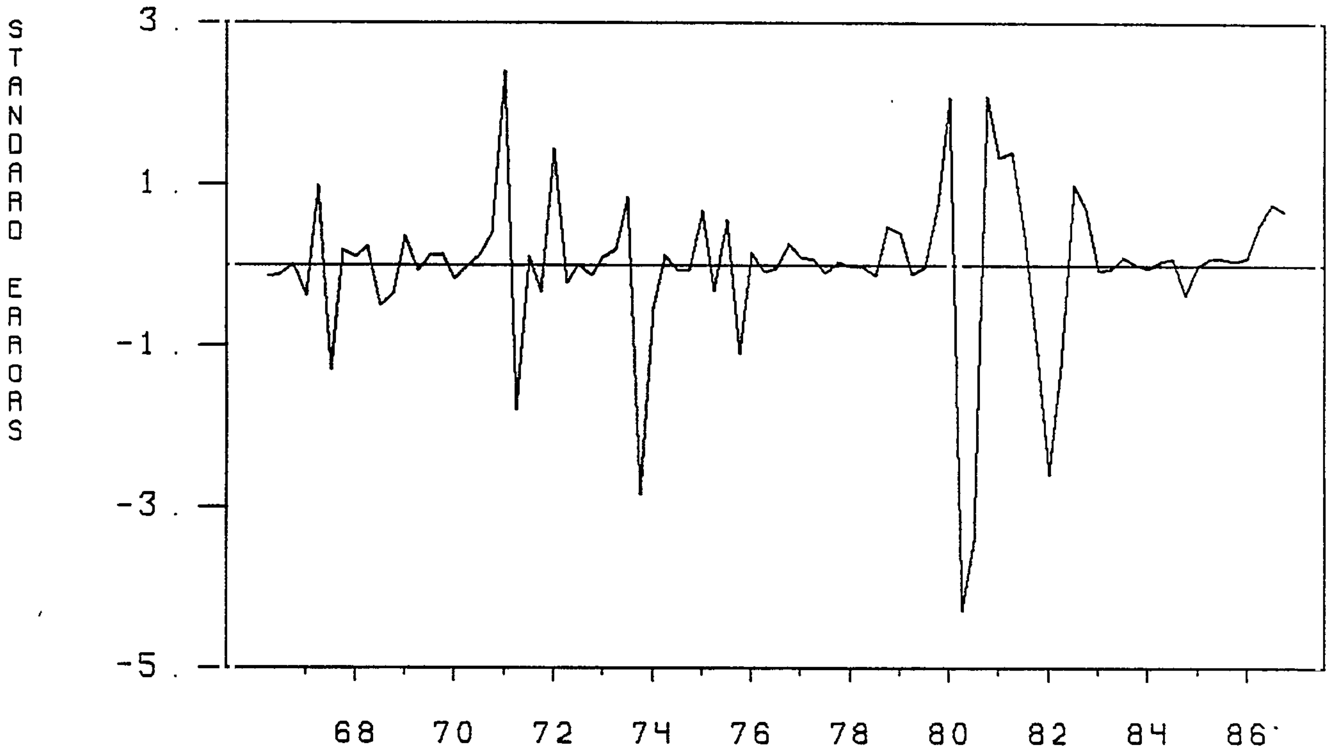

FIG. 7F. STANDARDIZED FIRST DIFFERENCES

FIGS. 7G. AND 7H. VALUE OF THE TRADE-WEIGHTED DOLLAR

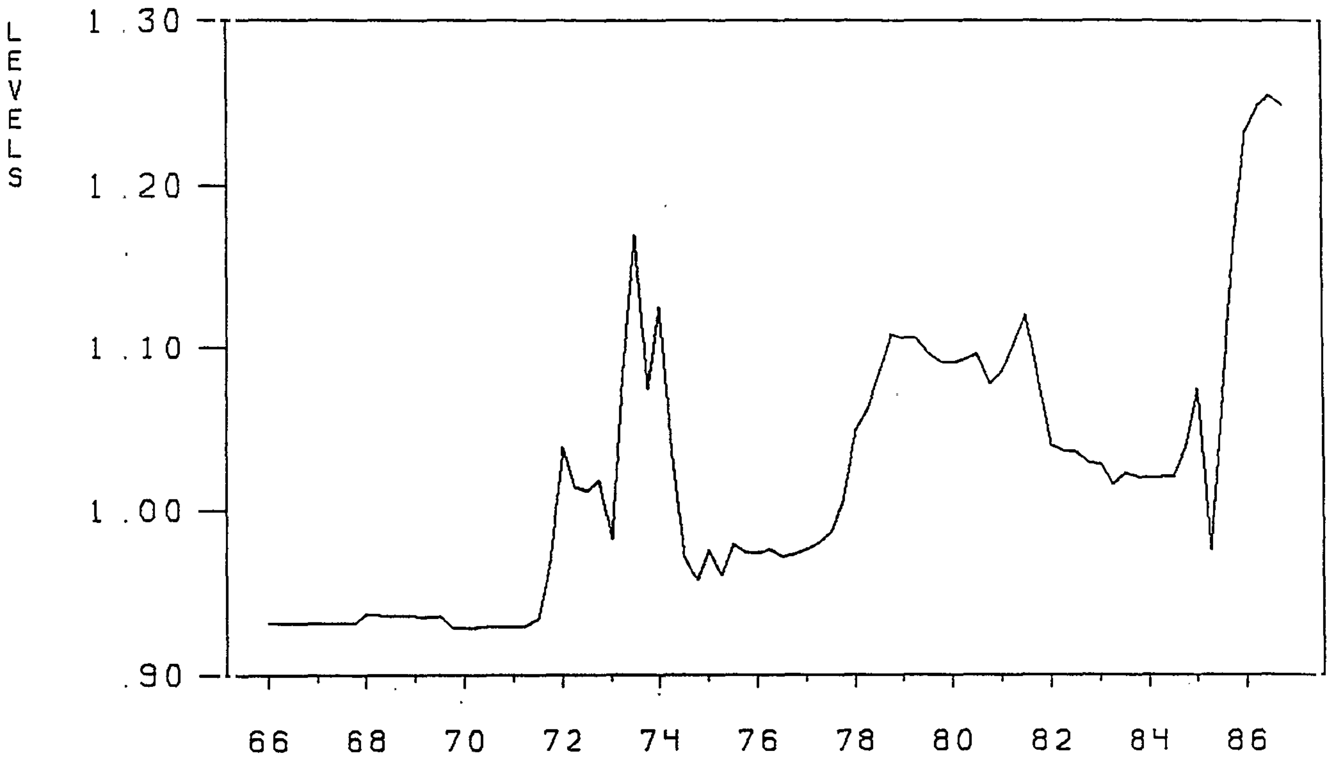

FIG. 7G. COEFFICIENT ON FIRST OWN LAG

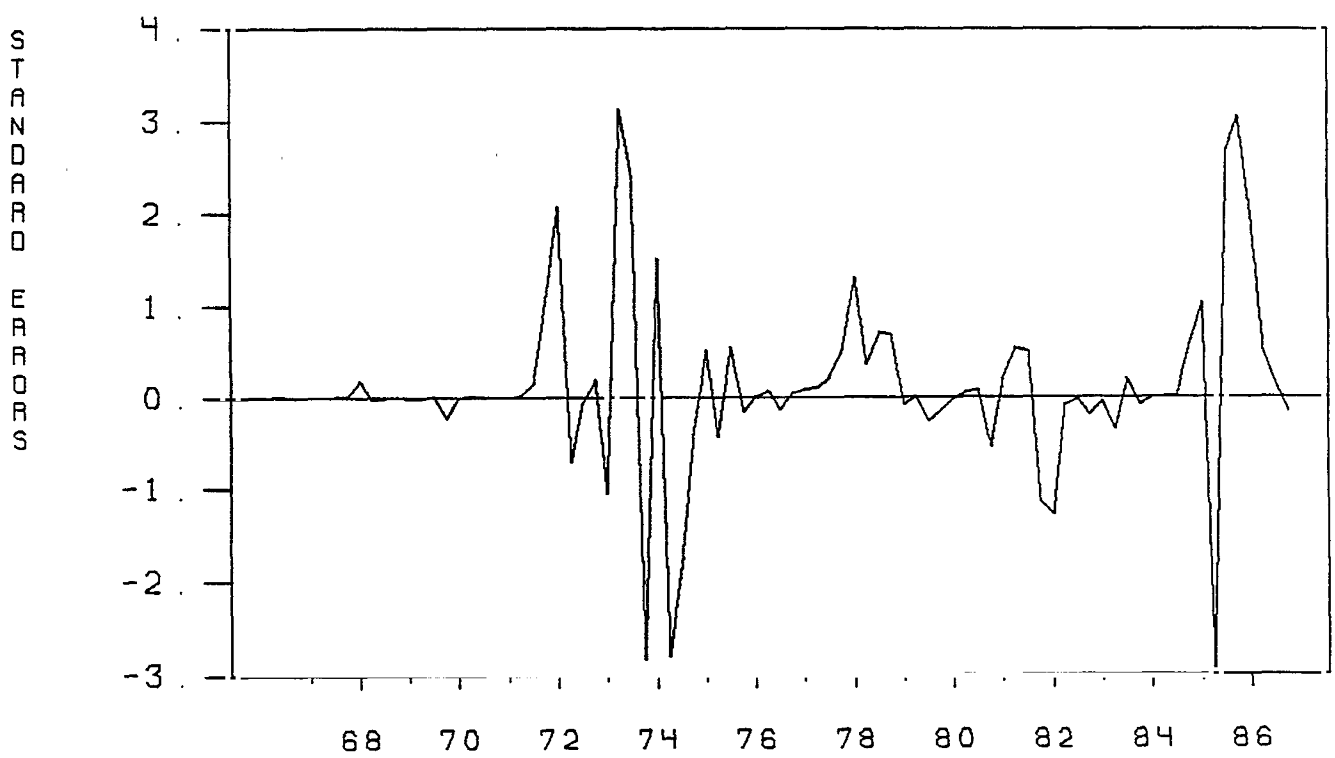

FIG. 7H. STANDARDIZED FIRST DIFFERENCES

FIGS. 8A.-8D. FED POLICY CHANGE EXPERIMENT

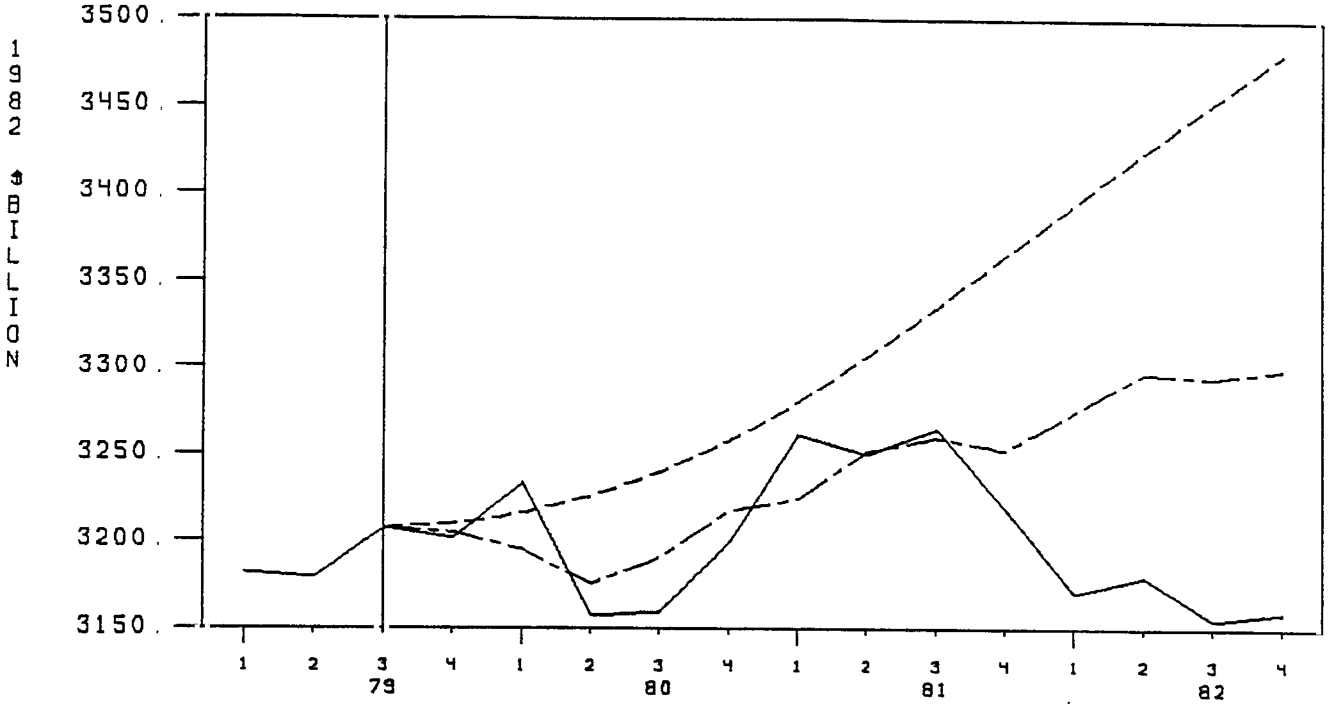

FIG. 8A. REAL GNP

ACTUAL ——— UNCOND - - -
CONDL —- -

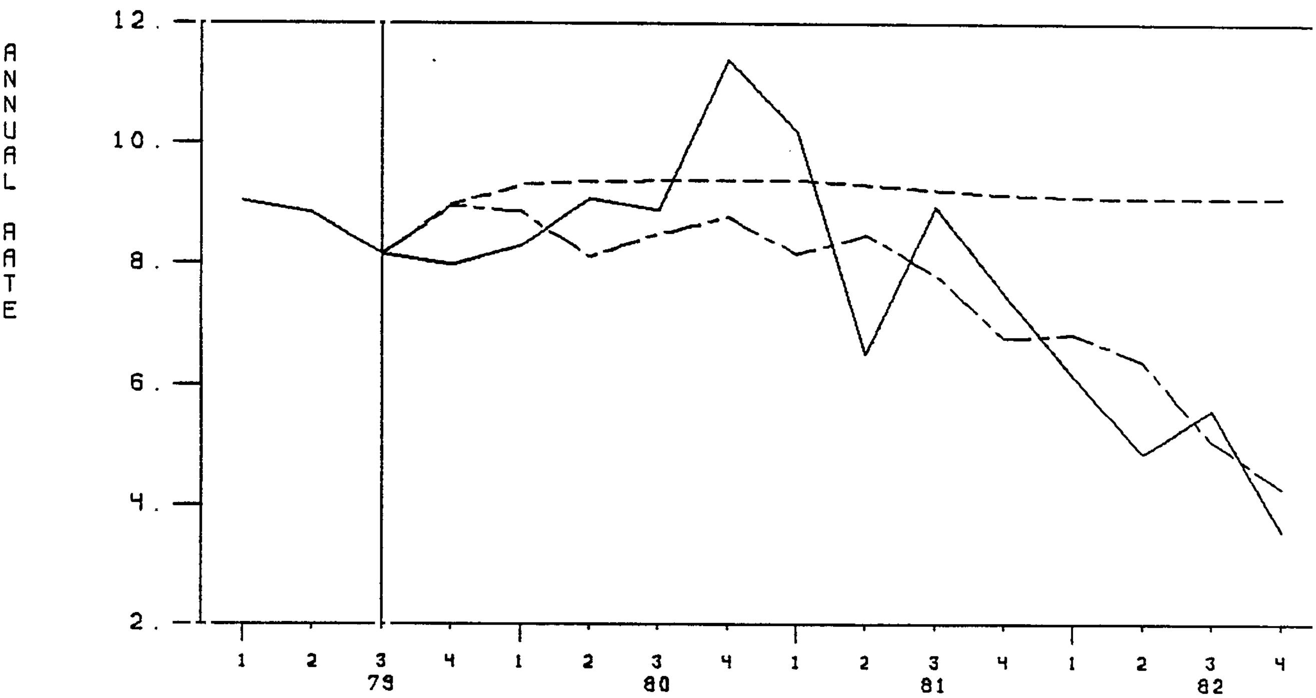

FIG. 8B. INFLATION (GNP DEFLATOR)

ACTUAL ——— UNCOND - - -
CONDL —- -

FIGS. 8A.-8D. FED POLICY CHANGE EXPERIMENT

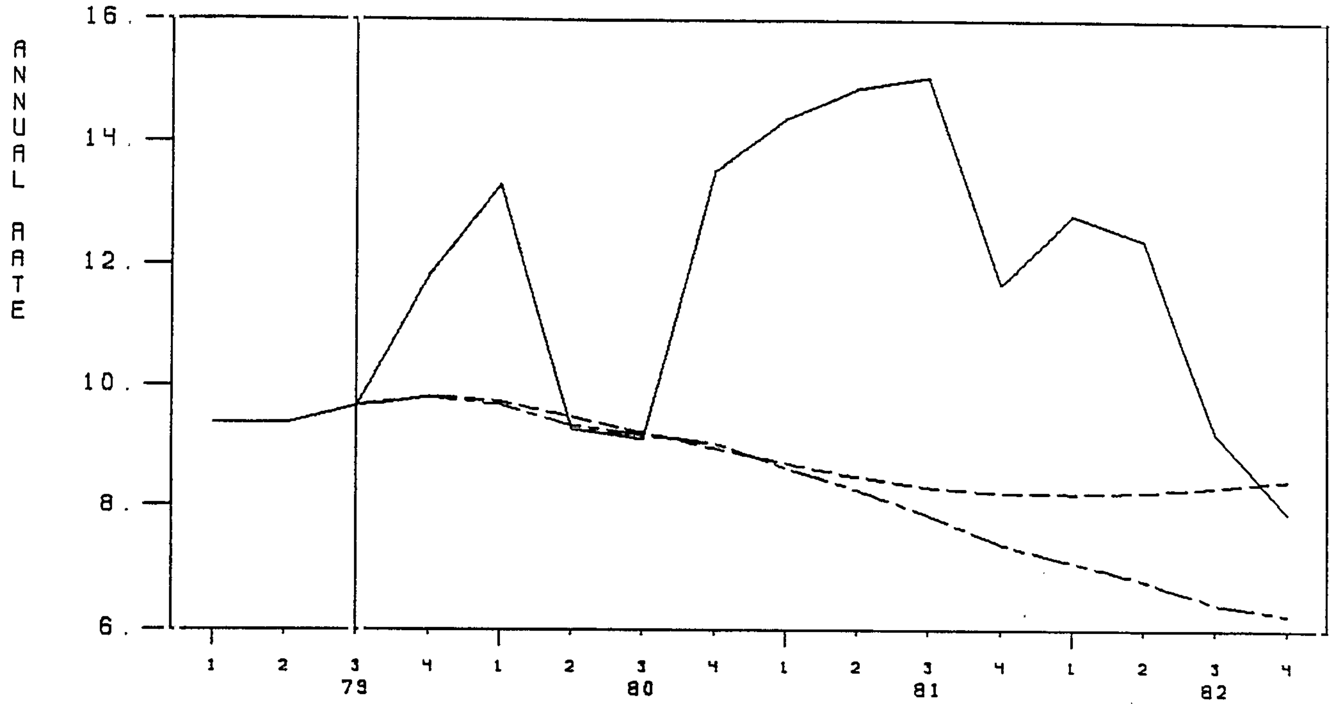

FIG. 8C. 3-MONTH TREASURY-BILL RATE

ACTUAL ——— UNCOND – – –
CONDL —– – –

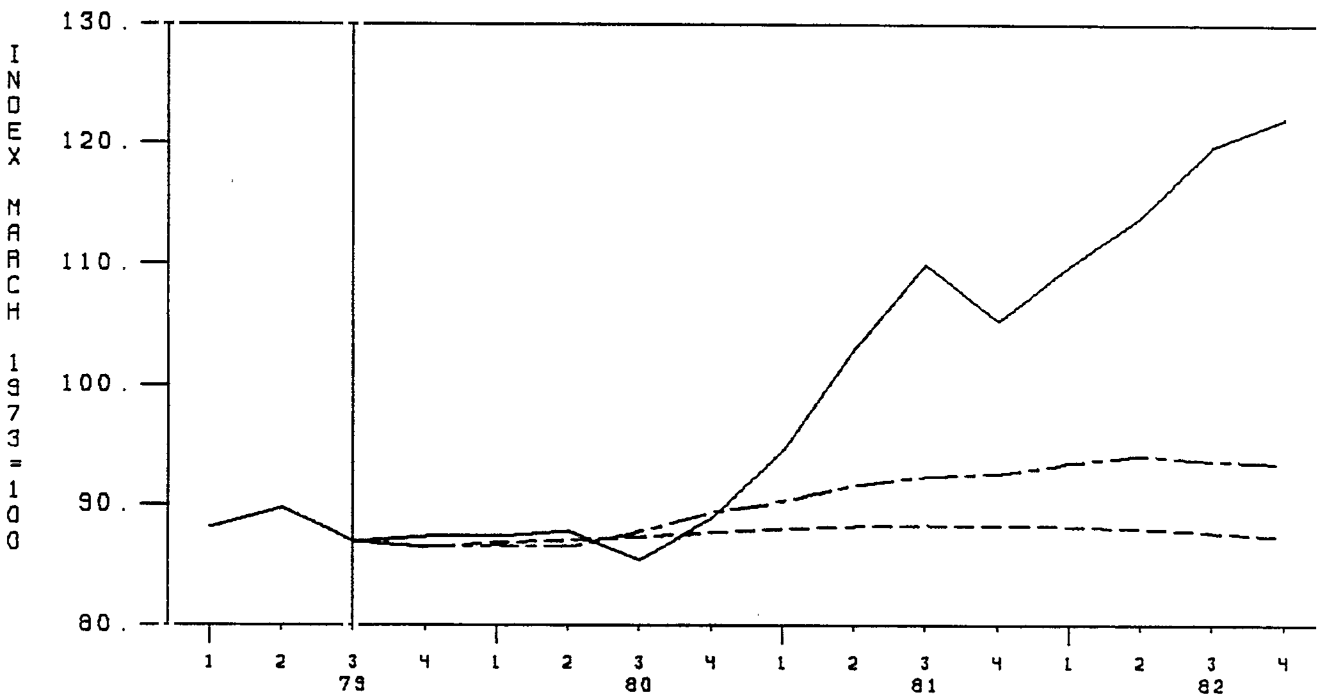

FIG. 8D. VALUE OF THE TRADE-WEIGHTED DOLLAR

ACTUAL ——— UNCOND – – –
CONDL —– – –

FIGS. 9A.-9D. REAGAN TAX CUT EXPERIMENT

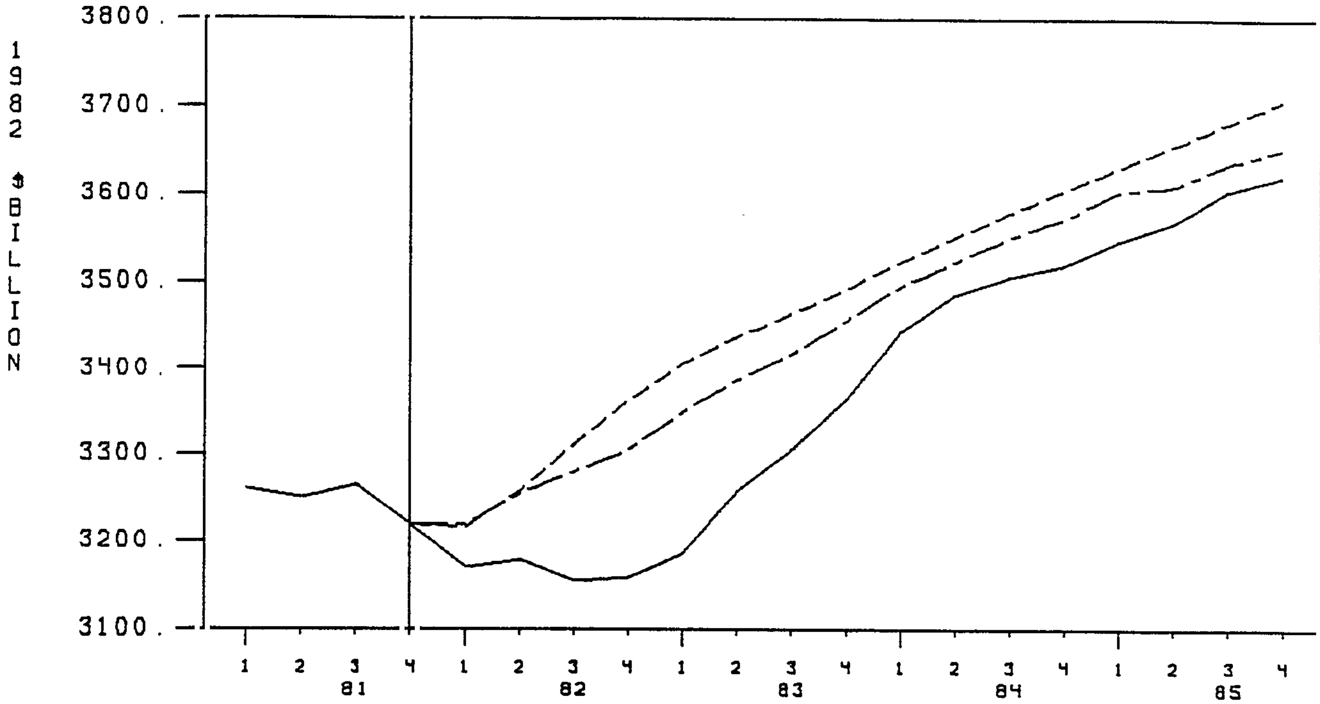

FIG. 9A. .REAL GNP

ACTUAL ——— UNCOND – – –
CONDL —-–

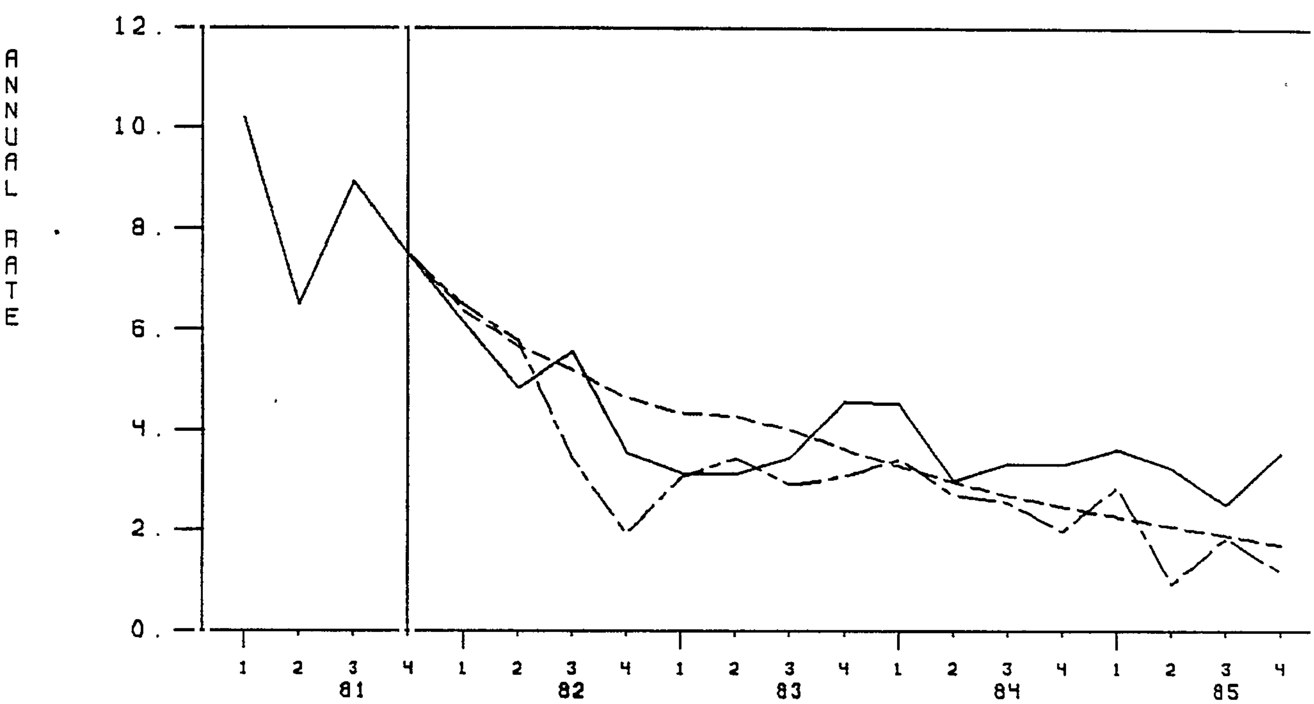

FIG. 9B. INFLATION (GNP DEFLATOR)

ACTUAL ——— UNCOND – – –
CONDL —-–

FIGS. 9A.-9D. REAGAN TAX CUT EXPERIMENT

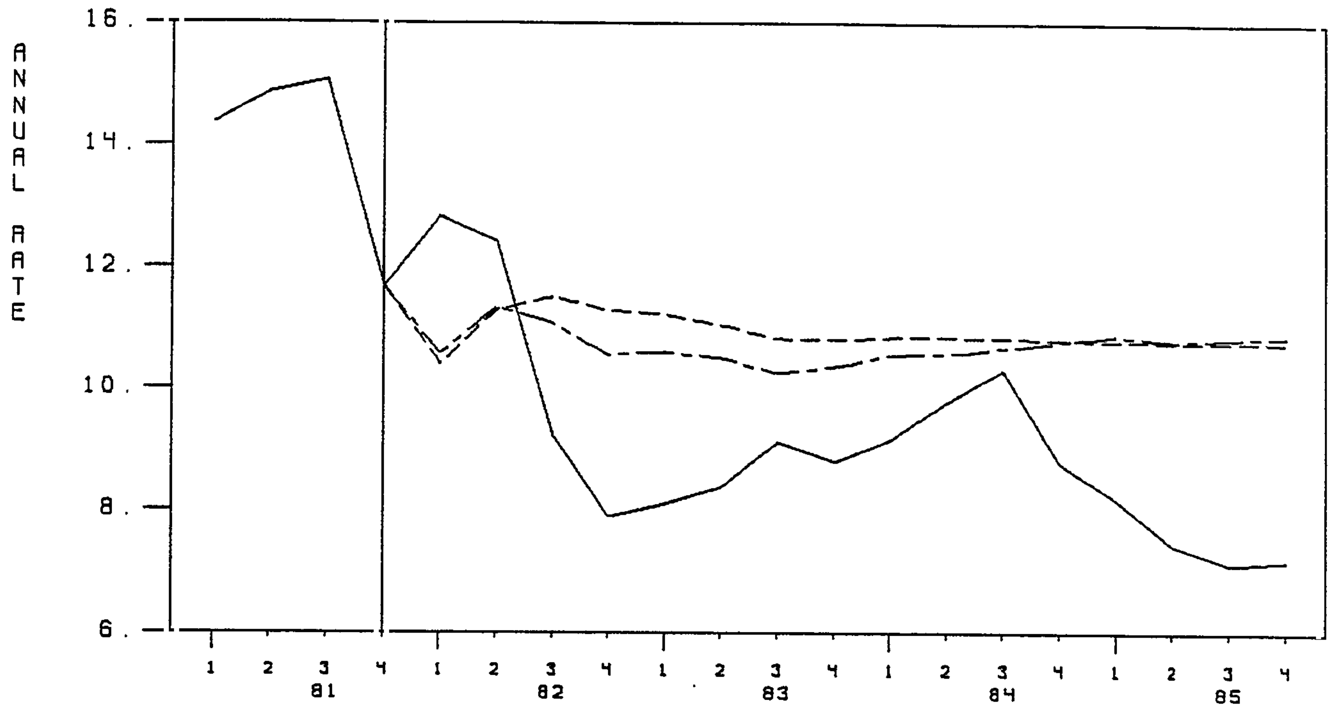

FIG. 9C. 3-MONTH TREASURY-BILL RATE

ACTUAL	———	UNCOND	— — —
CONDL	——--		

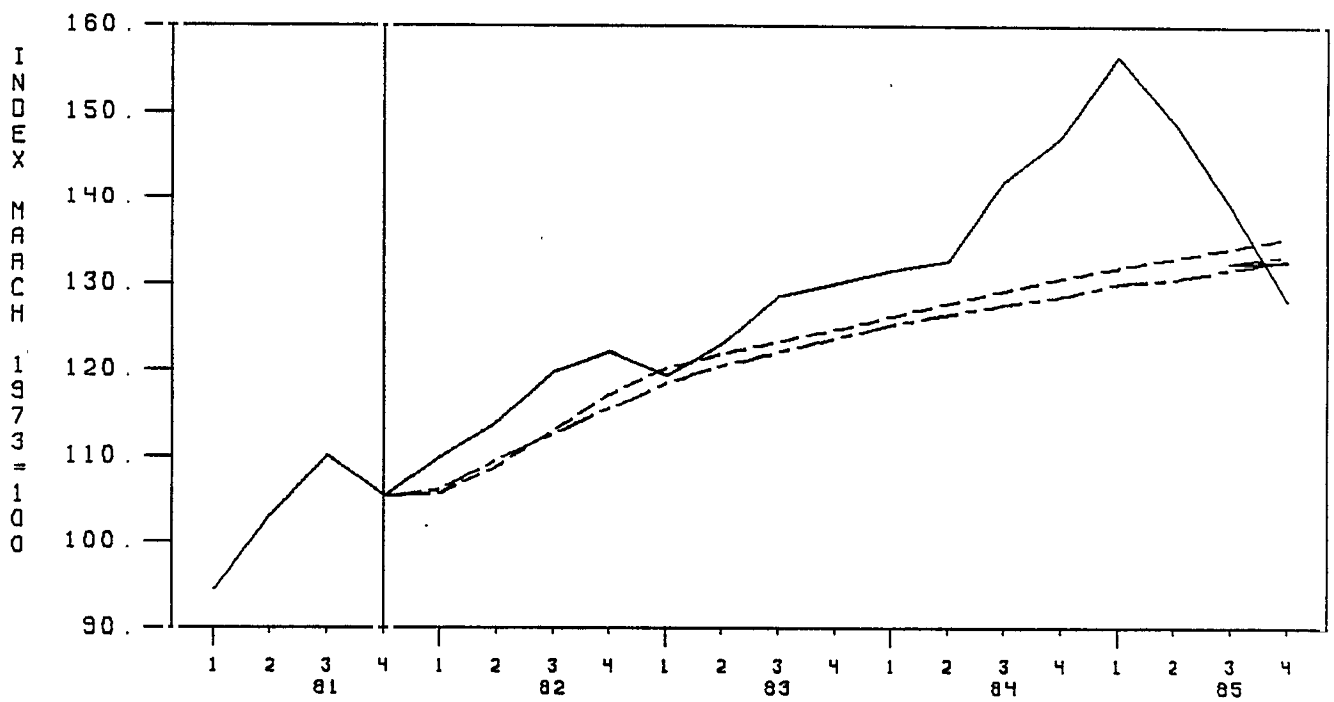

FIG. 9D. VALUE OF THE TRADE-WEIGHTED DOLLAR

ACTUAL	———	UNCOND	— — —
CONDL	——--		

FIG. 10 STANDARDIZED BOX-TIAO SCORES

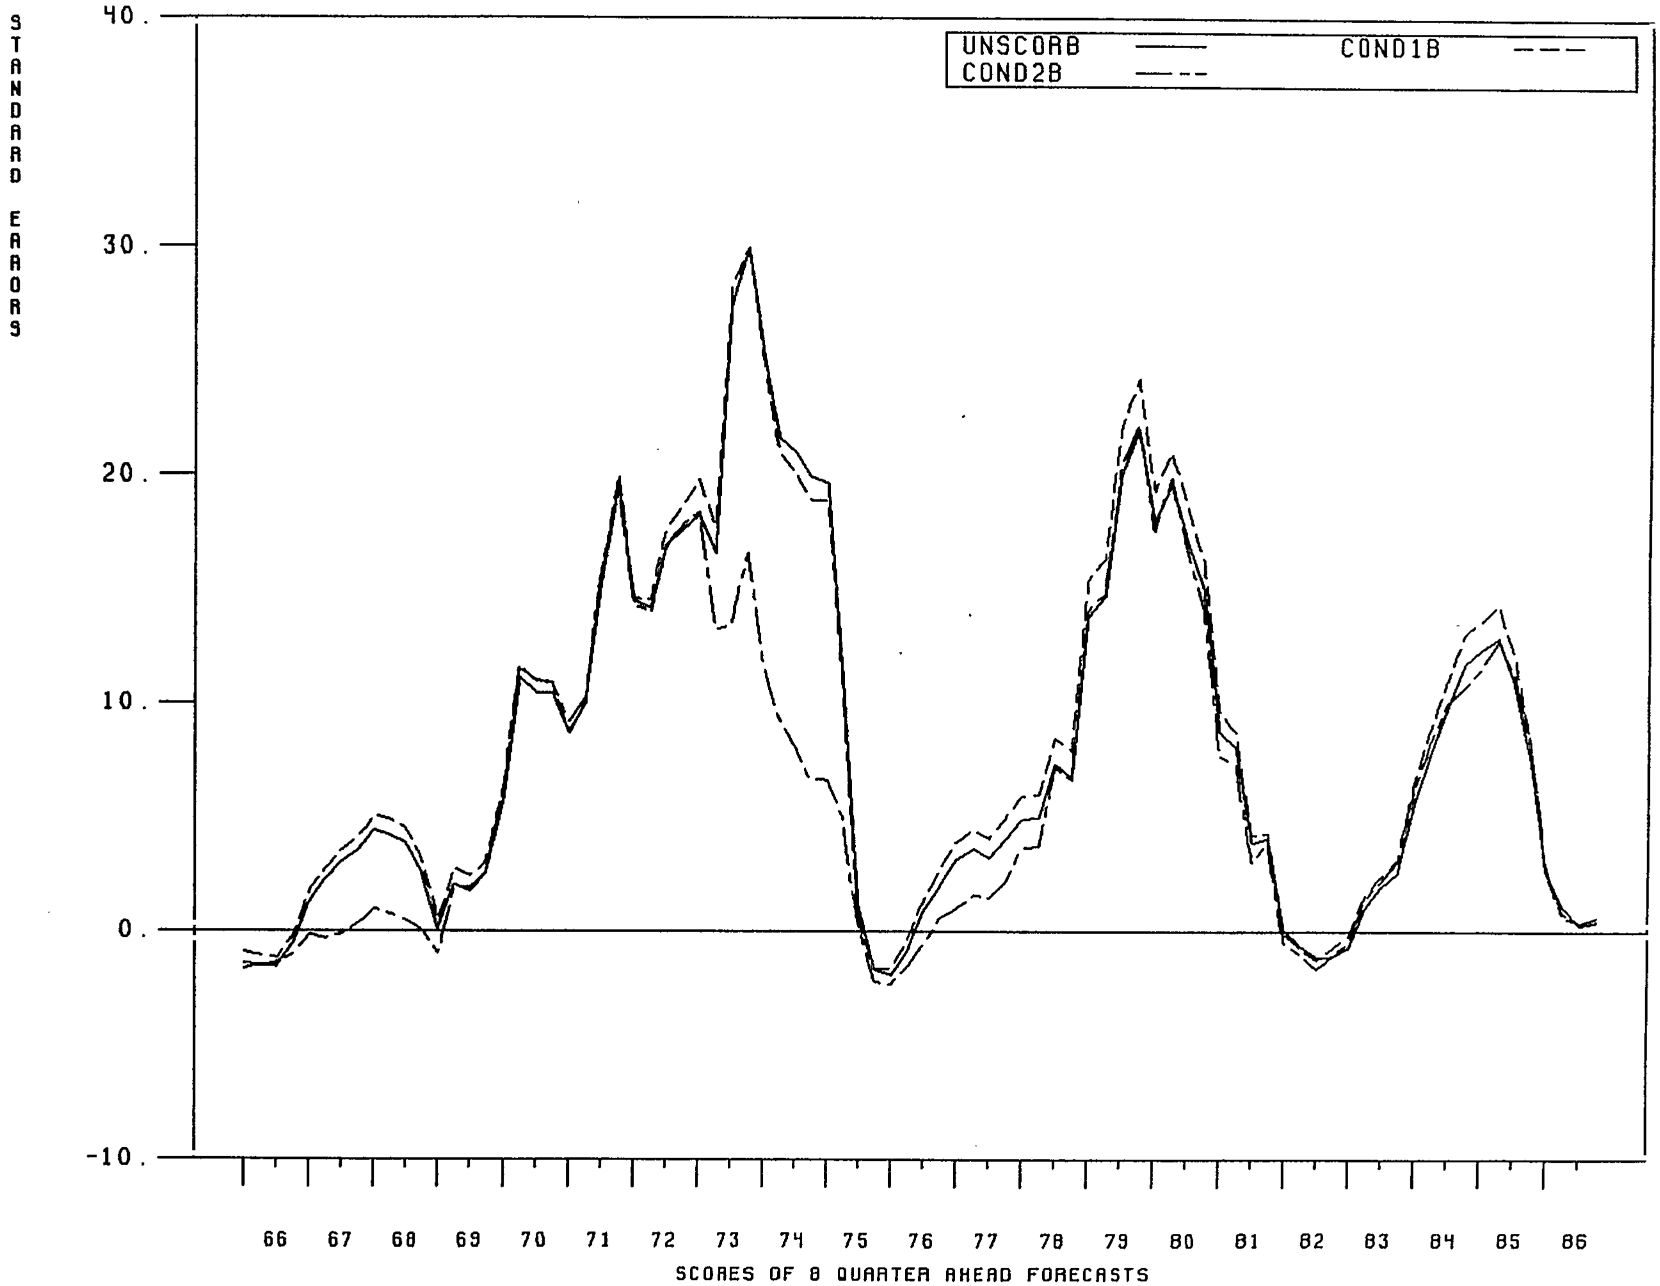

FIG. 11. FED POLICY CHANGE EXPERIMENT

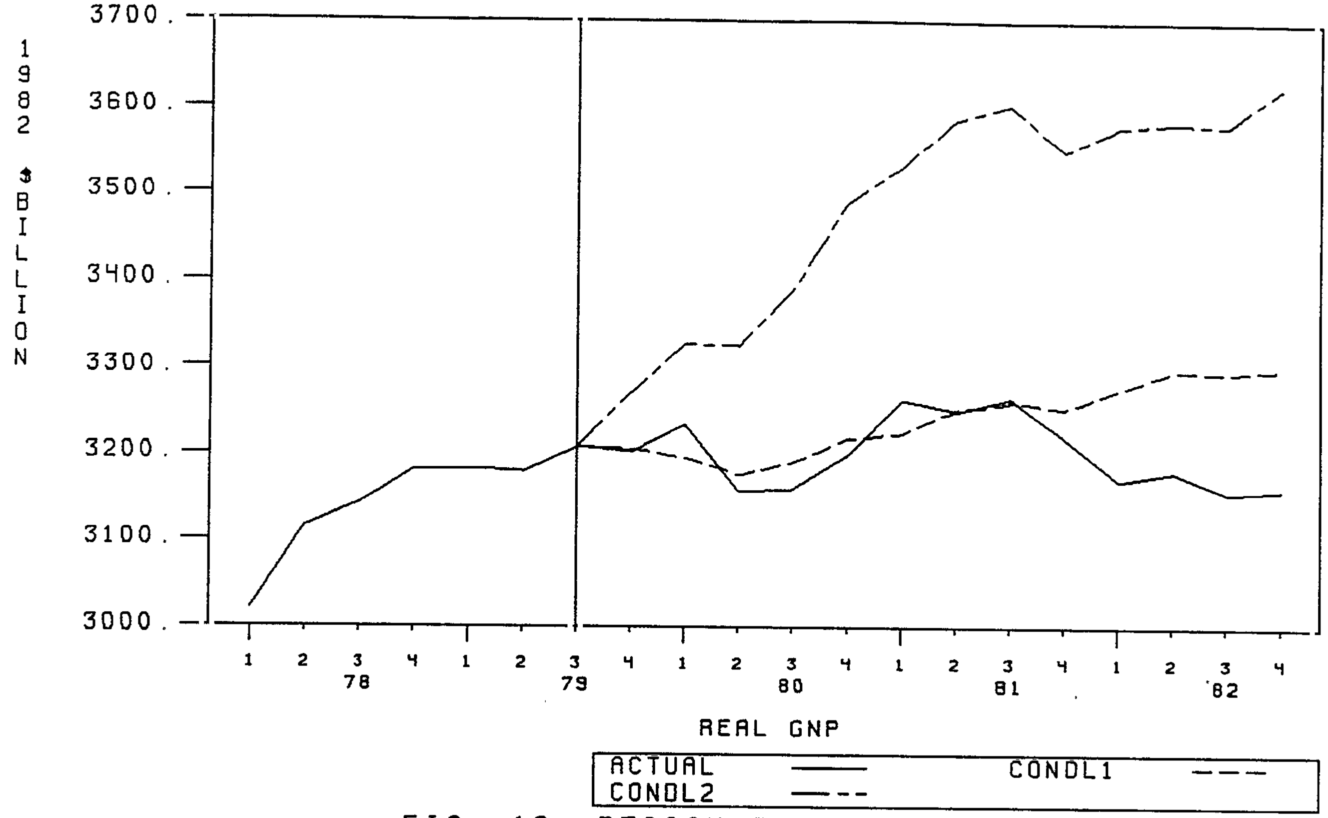

FIG. 12. REAGAN TAX CUT EXPERIMENT

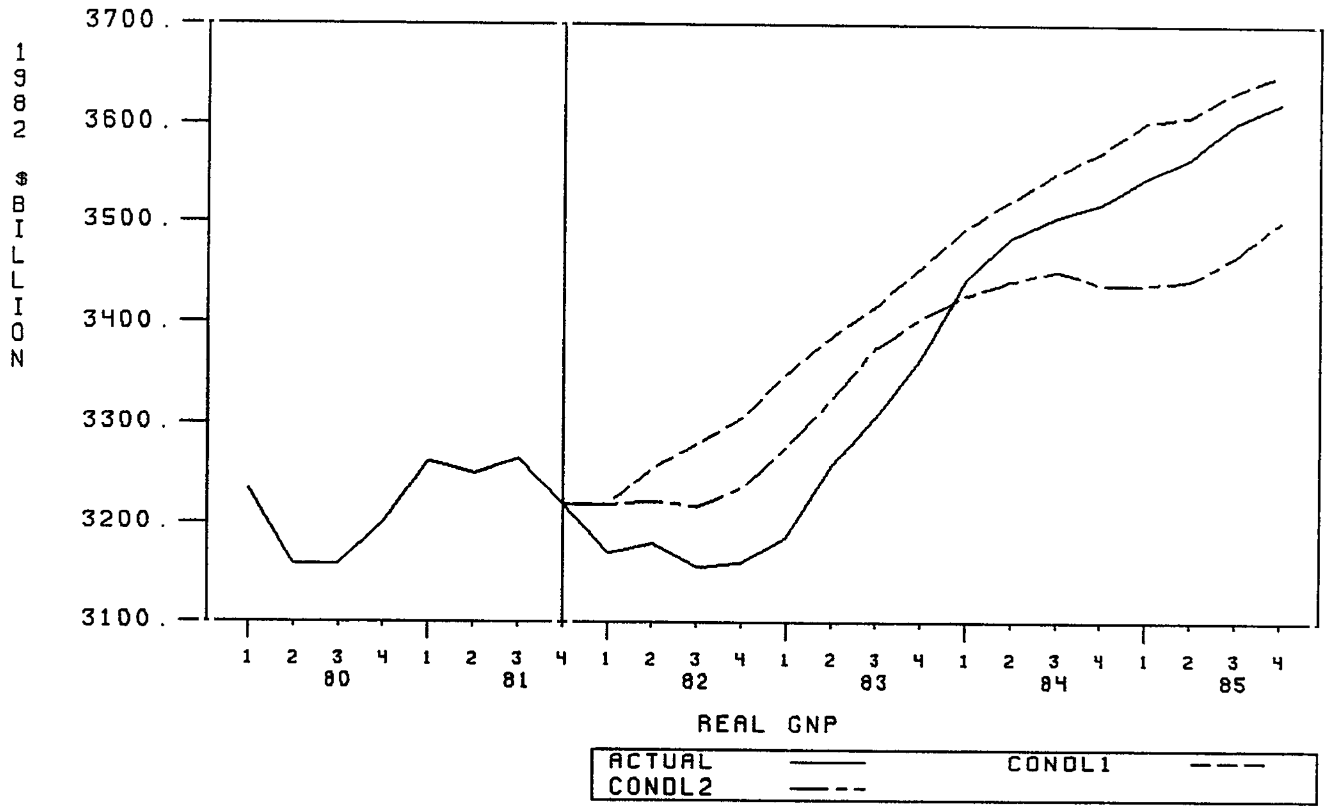

FIGS. 13A & 13B. ONE-STEP-AHEAD FORECAST ERRORS

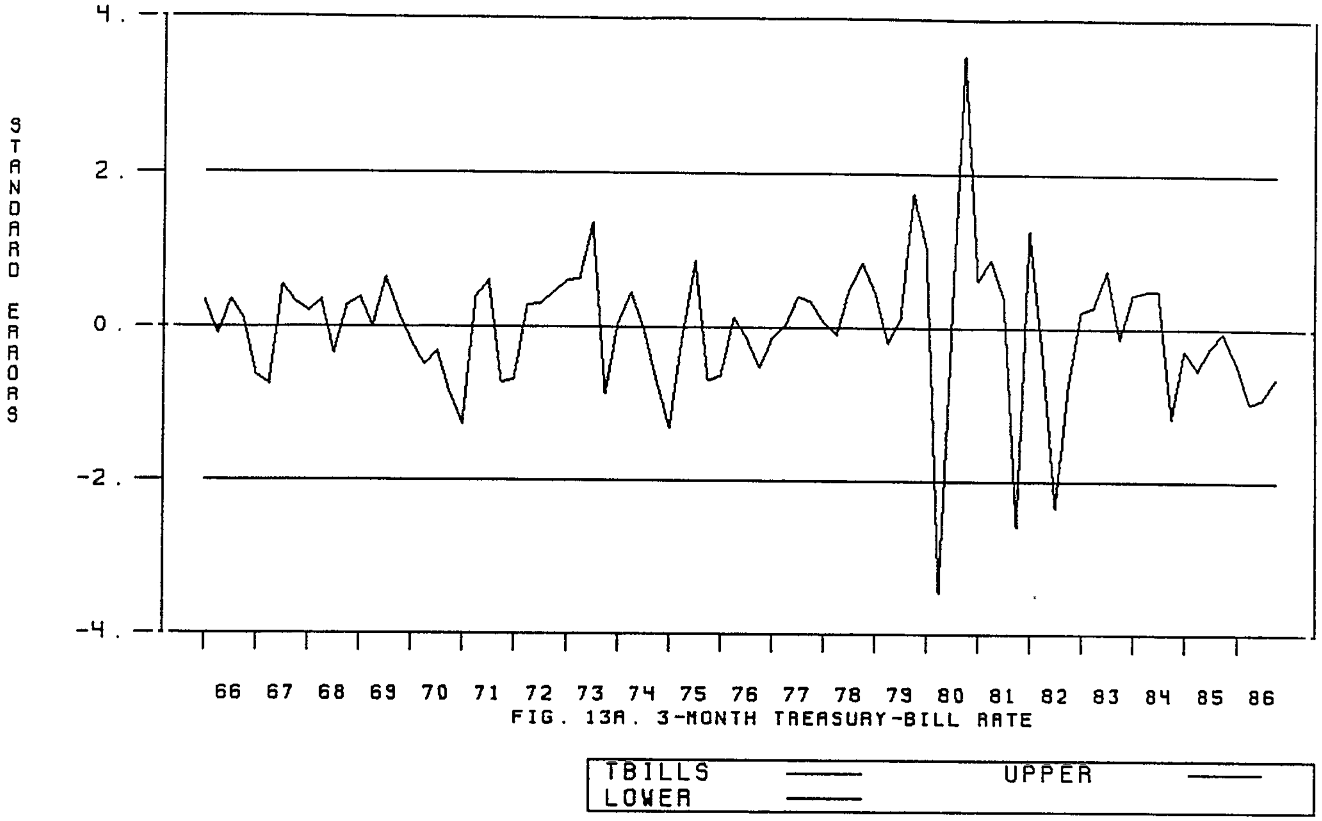

FIG. 13A. 3-MONTH TREASURY-BILL RATE

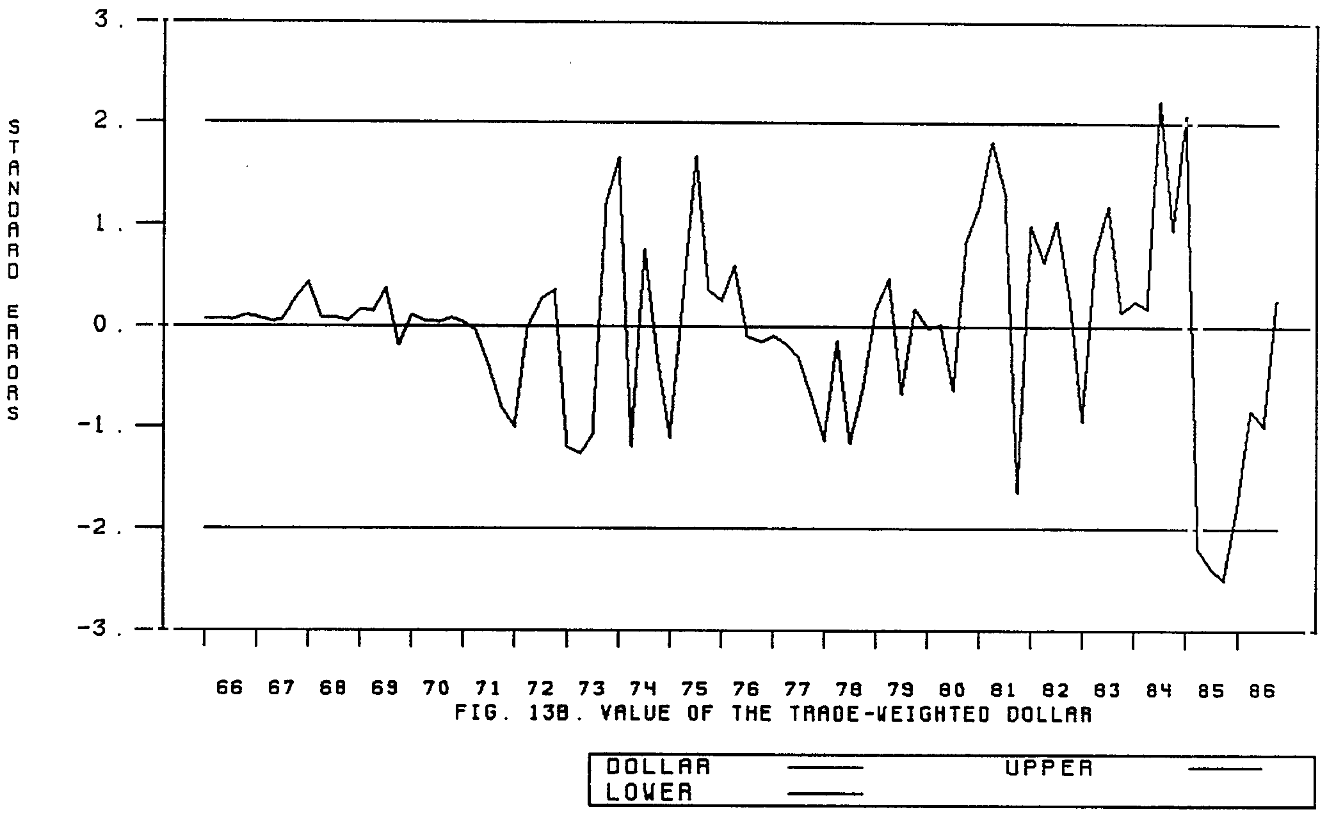

FIG. 13B. VALUE OF THE TRADE-WEIGHTED DOLLAR

FIG. 14 STANDARDIZED BOX-TIAO SCORES FOR RESPECIFIED MODEL

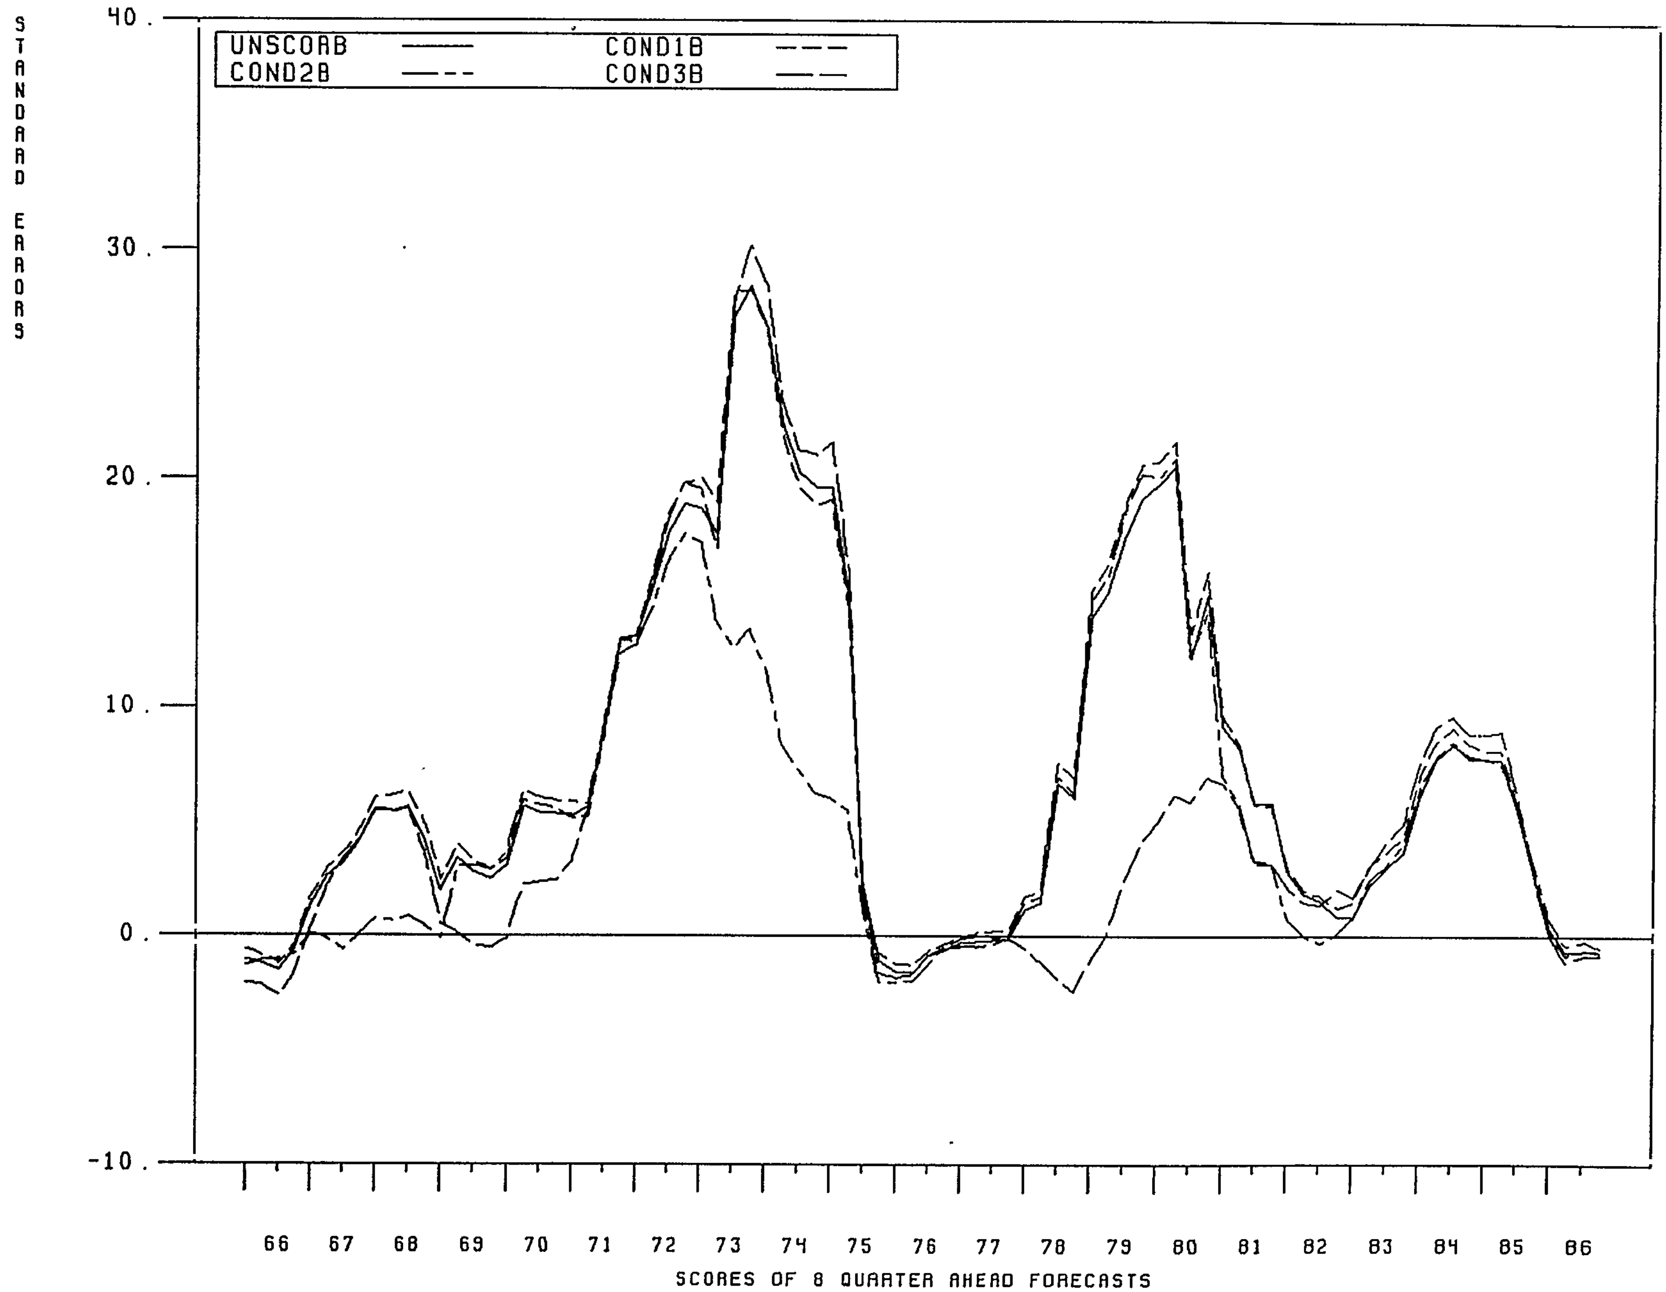

FIGS. 15A.-15F. OIL SHOCK EXPERIMENT

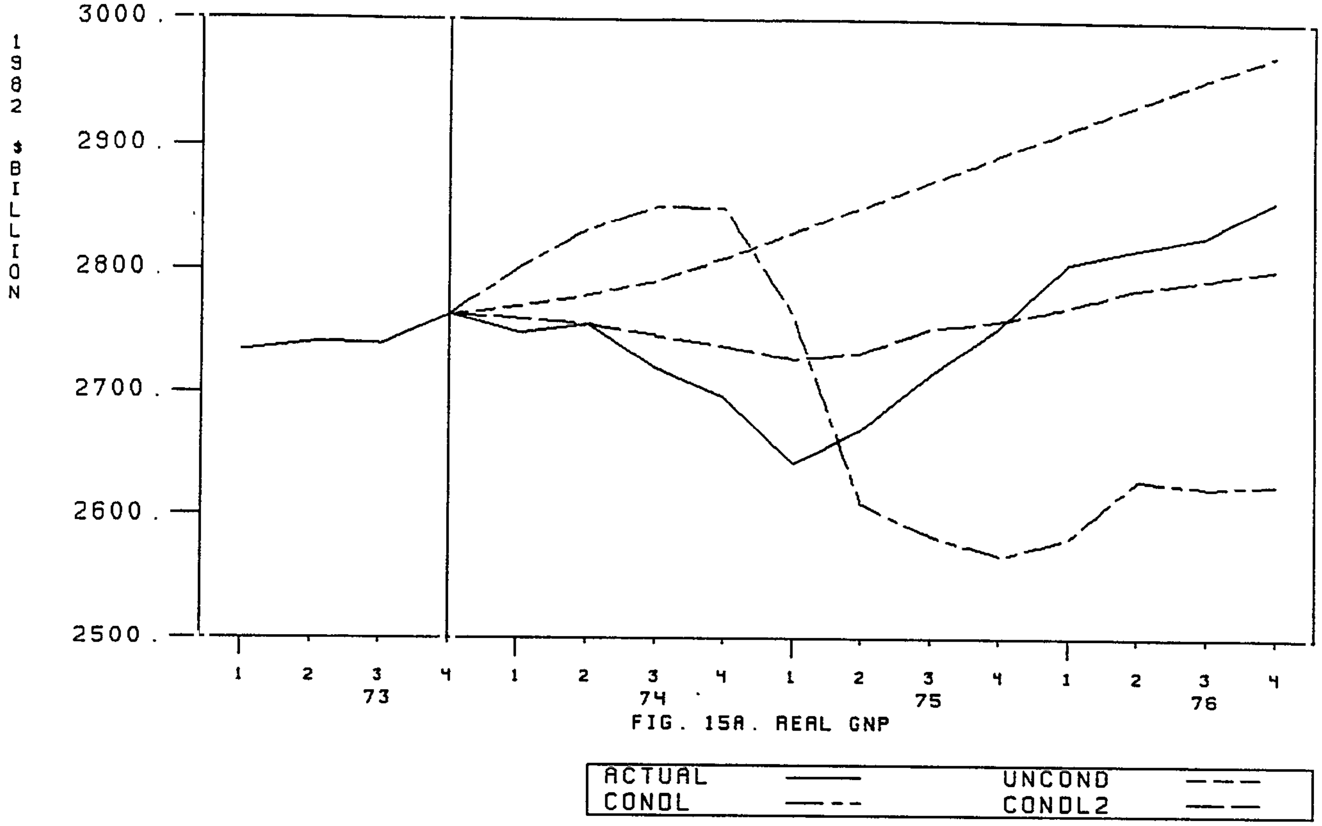

FIG. 15A. REAL GNP

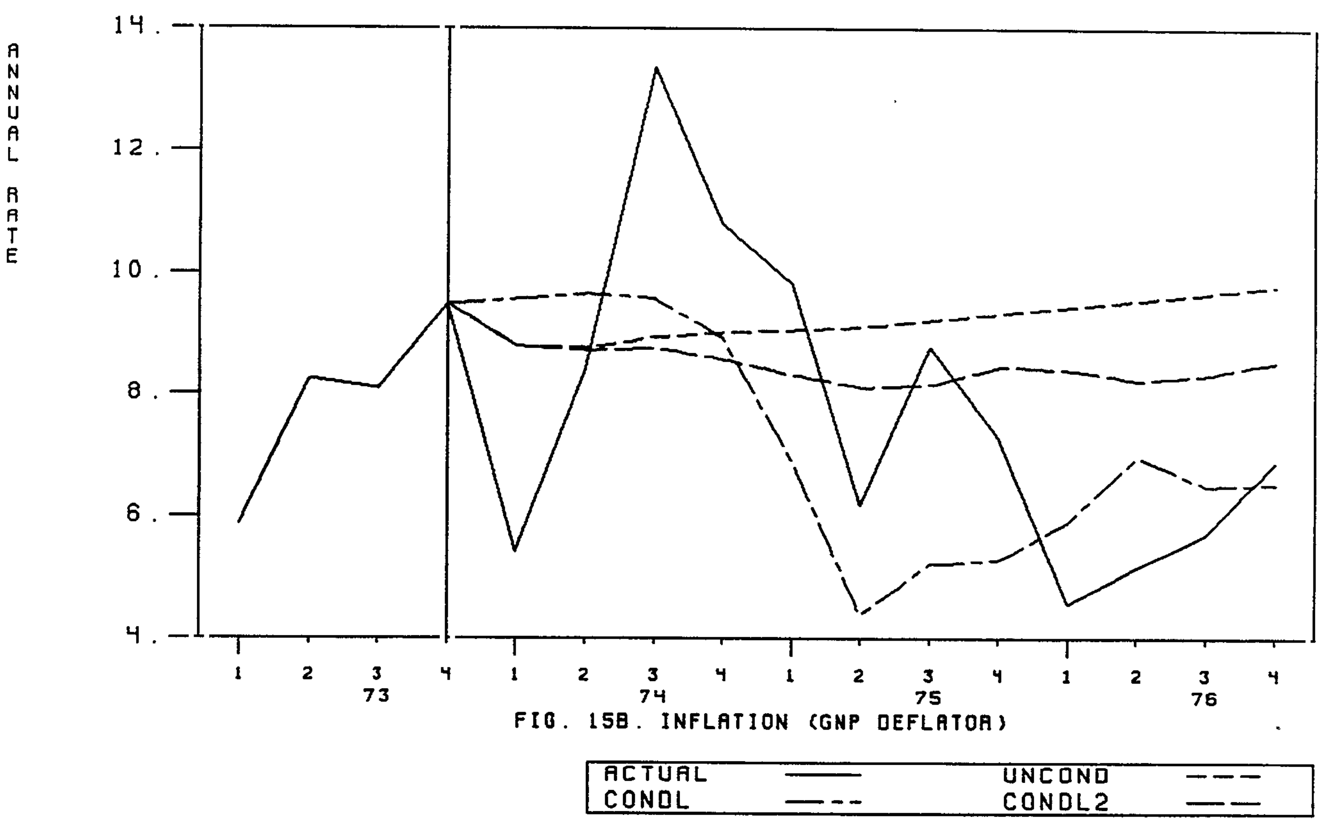

FIG. 15B. INFLATION (GNP DEFLATOR)

FIGS. 15A.-15F. OIL SHOCK EXPERIMENT

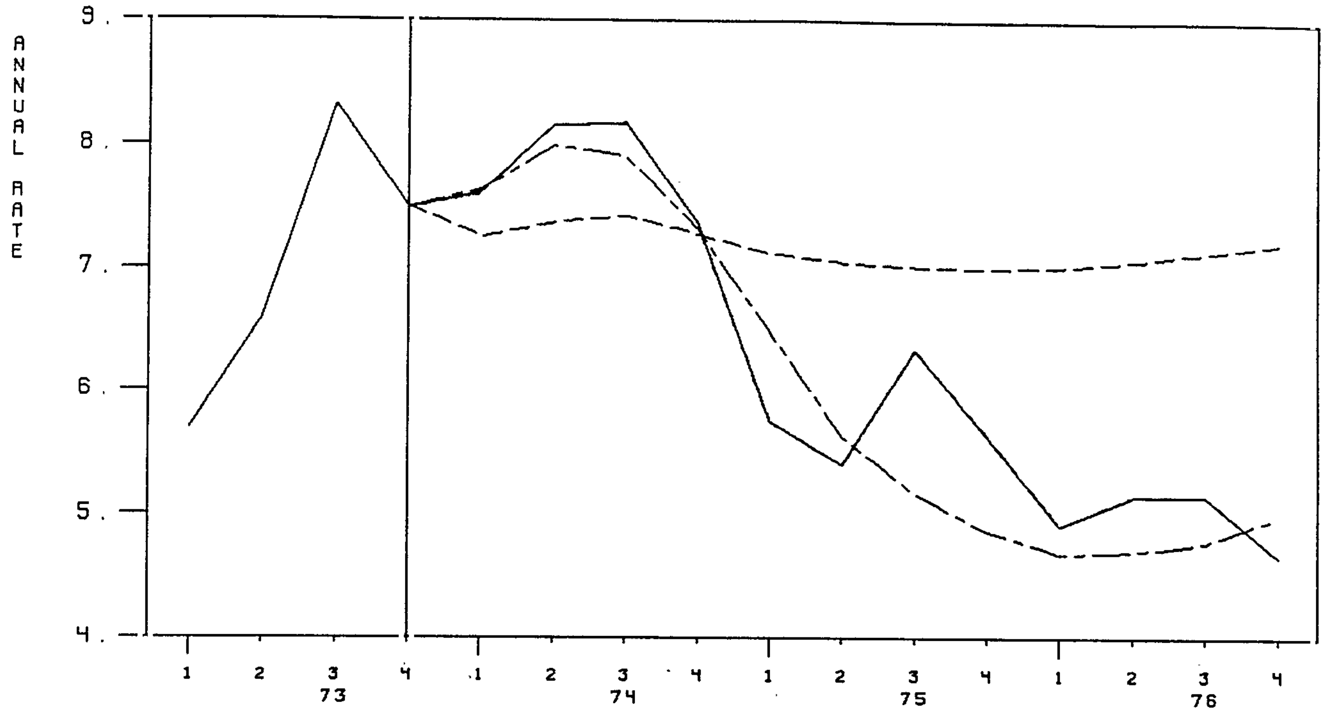

FIG. 15C. 3-MONTH TREASURY-BILL RATE

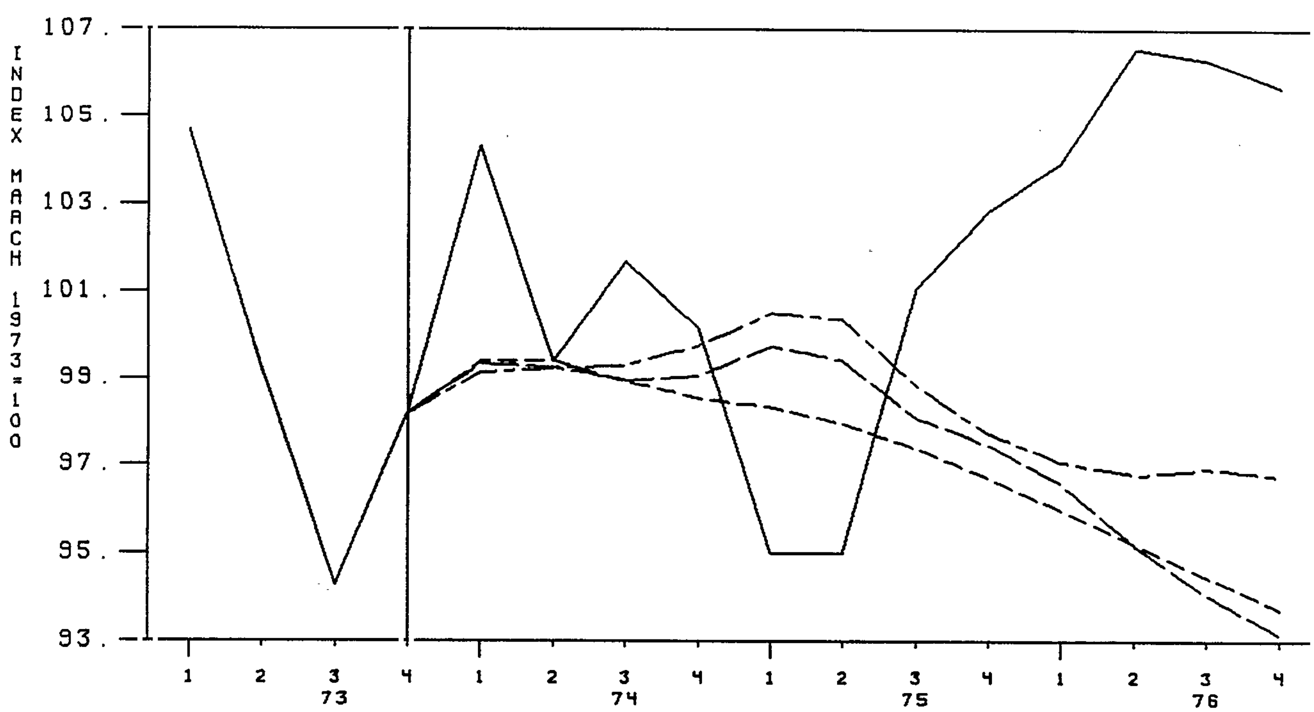

FIG. 15D. VALUE OF THE TRADE-WEIGHTED DOLLAR

FIGS. 15A.-15F. OIL SHOCK EXPERIMENT

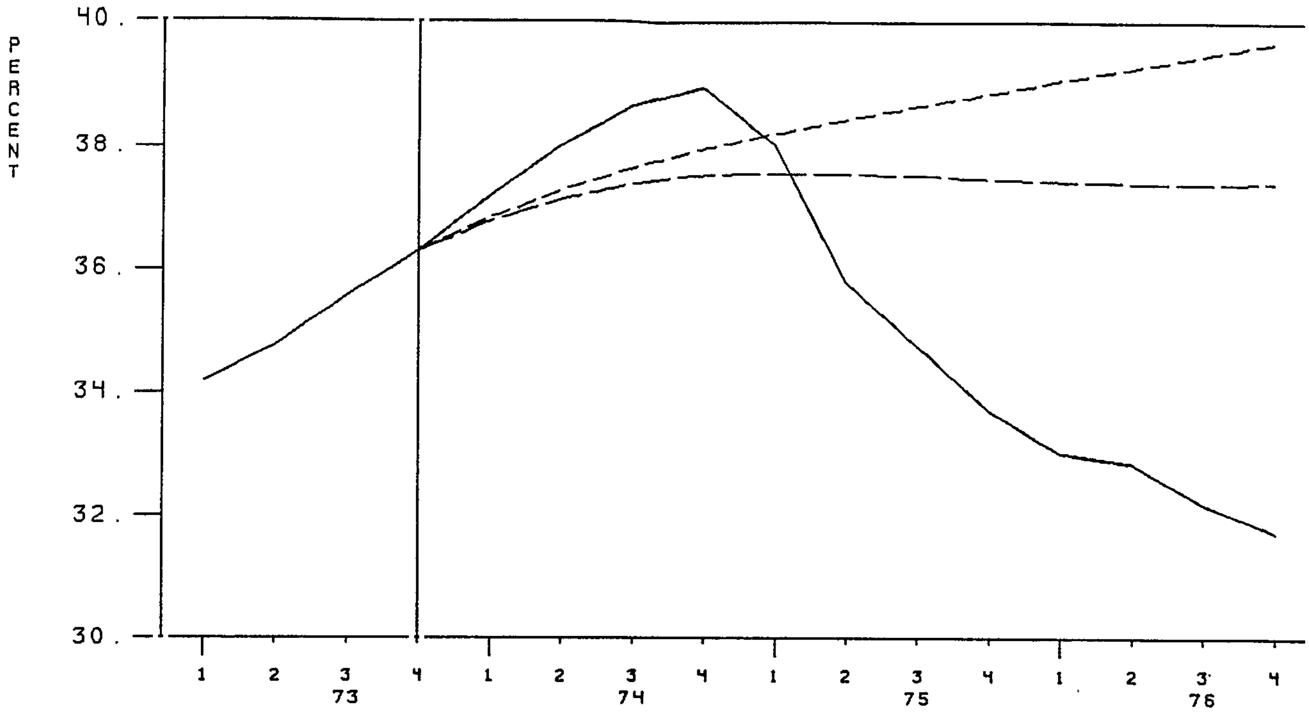

FIG. 15E. RATIO OF MONETARY BASE TO DEBT

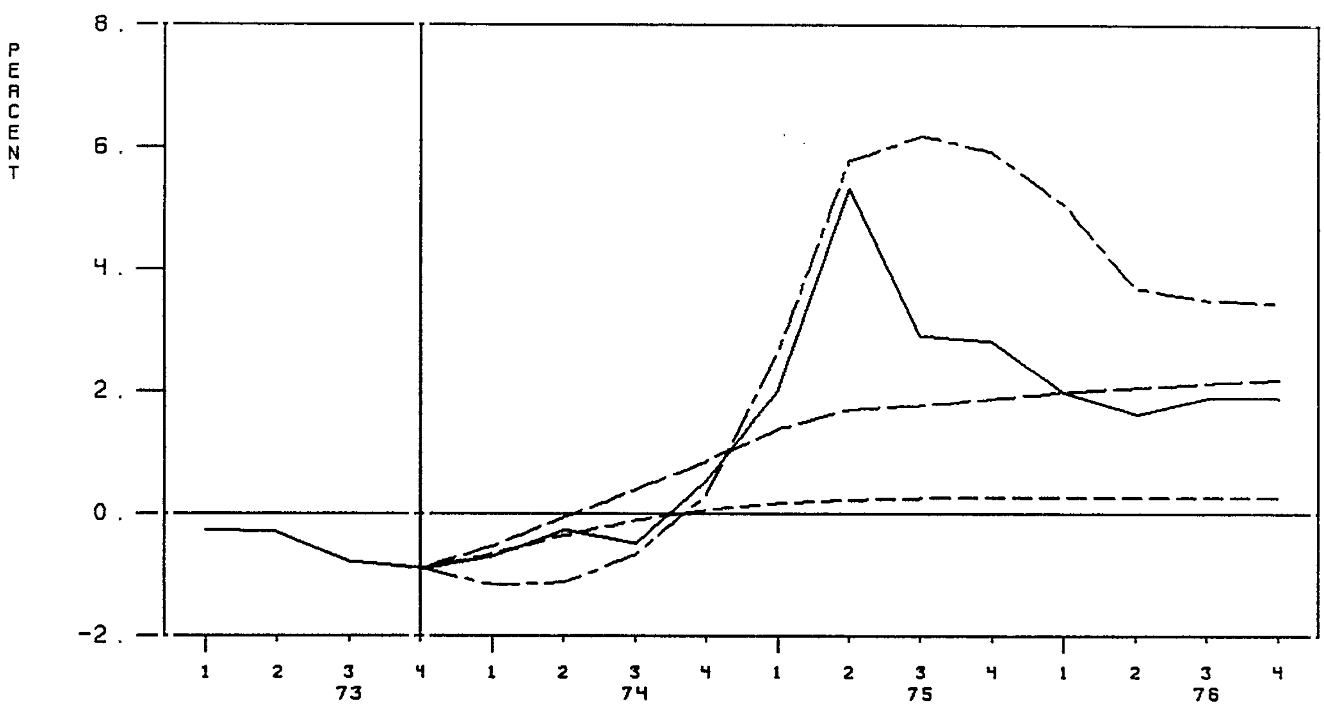

FIG. 15F. RATIO OF PRIMARY DEFICIT TO GNP