Federal Reserve Bank of Minneapolis
Research Department Staff Report 82

February 1983

OPTIMAL CONTROL OF THE MONEY SUPPLY

Robert B. Litterman

Federal Reserve Bank of Minneapolis

ABSTRACT

Using optimal control theory and a vector autoregressive representation of the relationship between money and interest rates, one can derive a feedback control procedure which defines the best possible tradeoff between money supply fluctuations and interest rate volatility and which could be used to reduce both from their current levels.

The debate over the proper conduct of U.S. monetary policy has intensified since October 1979 when the Federal Reserve focused its attention on reducing inflation by controlling the rate of growth of the money supply. Although most observers have given the Fed credit for reducing the trend growth of money, many have criticized it for having increased the short-run variability of money growth rates and the volatility of interest rates. The Fed is currently searching for procedures which will guarantee control over the trend growth of money and at the same time reduce the short-run fluctuations in both money and interest rates. In this paper I use optimal control theory and a time series representation for money and interest rates to derive a feedback control procedure which defines the best possible tradeoff between money supply fluctuations and interest rate volatility and which could be used to reduce both from their current levels. I first review the control theory framework, then describe the use of a time series model to represent the dynamic behavior of the system, then present the application to short-run control of the money supply, and finally address the key issue of structural stability.

OPTIMAL CONTROL THEORY

Optimal control theory is a well-developed set of mathematical tools used primarily by engineers to solve problems involving a dynamical system which responds to exogenous inputs and is subject to shocks. These tools are used here to generate a rule for targeting interest rates in order to optimally balance the competing goals of controlling the supply of money and reducing the volatility of interest rates.

In its usual form, optimal linear control theory specifies an algorithm for setting one or more inputs in order to minimize a quadratic loss function. This result, and others cited below, can be found in standard control theory texts, such as Kwakernaak and Sivan 1972, Chow 1975, and Kendrick 1981.

The textbook application of control theory to monetary policy assumes that the Fed can control either money or interest rates perfectly. (See, for example, Sargent 1979a, Sargent and Wallace 1975, and Kalchbrenner and Tinsley 1976.) The usual questions at issue are which variable the Fed should control and how it should set that variable so as to achieve a full employment, stable price path for the economy.

The standard approach attempts to derive an optimal monetary policy; it ignores the important tradeoff between the degree of monetary control and interest rate volatility. I focus on this narrower issue. I take the money target path as given, but rather than taking as given the ability of the Fed to hit its money supply target, I investigate the Fed's short-run problem of attempting to control the money supply. I assume the Fed is using open market operations to try to keep seasonally adjusted M1 as close to a long-run growth path as is feasible, on average. Open market operations, by increasing or decreasing the supply of reserves, cause the federal funds rate to go down or up, respectively. These movements in the federal funds rate cause banks and other economic agents to adjust their portfolios, leading to predictable movements in the stock of money.

I do not attempt to model the open market operations directly. Instead, I focus on the levels of the federal funds rate which emerge each week and their effects on subsequent movements in money. In the control procedure modeled here, the Fed receives, at the end of the week, the latest figures for M1 (data for the week ending two weeks earlier) and decides on a new desired level for the funds rate for the following week. Other procedures and timing relationships can easily be modeled in a similar way. In fact, I will later discuss the application of control theory with a funds rate target, a procedure in which the Fed basically sets its targeted funds rate, and with a reserves target, a procedure in which the Fed supplies reserves consistent with a chosen funds rate but does not offset shocks which may cause significant deviations within a given week.

I assume that the Fed knows the dynamic response pattern of money and interest rates and uses this knowledge to set the funds rate so that the money supply will stay near its target path. In the next section I will address the question of how to estimate the necessary response patterns. Because the money supply is subject to random disturbances, the best the Fed can do is to cause the expected value of money to be on target each week. However, in order to achieve this level of accuracy with respect to money, the Fed may have to make large changes in the funds rate each week. The required changes may easily increase over time, leading to explosive oscillations in interest rates. This is the instrument instability problem suggested by Holbrook (1972). In fact, the Fed does not try to bring the expected value of the money supply onto its target path each week. Rather, it recognizes a short-run tradeoff between reducing expected deviations of money from its target path and reducing fluctuations in interest rates. (For a recent discussion of this issue, see Radecki 1982.) In order to investigate the nature of that tradeoff, I specify a loss function which has terms penalizing both money supply deviations from target and volatility of interest rates. These two objectives are assumed to capture the most important tradeoff in the current Fed operating procedures. However, the loss function could easily be generalized to include additional goals. It might be desirable, for example, to avoid large interventions in the market, in which case one could include a term representing a cost associated with the size of the control itself.

Optimal control is most often expressed in the context of a first-order difference equation in the state vector. Let x_t be an nx1 state vector, u_t be the control, and w_t be an nx1 vector of disturbances. The laws of motion of the system are given by

(1) $$x_t = A\, x_t + B\, u_t + w_t$$

where A is an nxn matrix and B is an nx1 vector. In order to fit the monetary control problem into this framework, x_t includes current and lagged values of M1, m_t; the federal funds rate, r_t; possibly other informational variables; and a monetary target, m^*_t The Fed-controlled shock to the funds rate is u_t. The matrix A includes two or more rows of estimated coefficients which define how M1, the funds rate, and possibly other variables evolve through time. All but one of the other rows of A identify as their values in the previous state lags of m, r, and possibly other variables. The final row defines the target money supply path.

The quadratic loss function is defined to be

$$L = E\left(\sum_{s=0}^{\infty} \beta^s\left\{\left(m_{t+s} - m^*_{t+s}\right)^2 + \lambda \sum_{k=1}^{q}\left[\left(r_{t+s} - r_{t+s-k}\right)^2/k\right]\right\}\right). \tag{2}$$

The cost associated with money deviations from target is balanced with interest rate volatility, measured as a weighted sum of expected squared changes in the federal funds rate over time. Different relative costs associated with deviations from the money target path and interest rate volatility can be represented by different values of λ. More terms in the sum measuring interest rate volatility--that is, larger values of q--will lead to a smoother funds rate path. For example, a high λ, with q equal to 1, will avoid whipsawing the market--large movements in the funds rate in a given week--while still allowing significant movements over a period of time as short as, say, two or three months. A q of 12, however, will dampen considerably these longer swings as well, leaving only smooth changes in the funds rate over time.

This form for the loss function is only one of many possibilities. I choose it primarily because of its simplicity; the higher is q, the more it will respond to--that is, penalize--low-frequency variations in interest rates. A more sophisticated loss function in the linear-quadratic class could be constructed by making loss proportional to the square of particular linear combinations of expected future interest rates, the linear combinations being chosen specifically to respond to certain bands of frequencies of interest rate movements.

The loss function (2) also includes a discount factor β, which allows the loss function to give relatively less weight to future losses than to current losses. For the purposes of this paper, there is no reason to discount future losses, and the discount factor is taken to be 1. Although the expected loss is not finite when the discount factor is 1, there is a well-defined feedback rule which is the limit as β goes to 1 of rules associated with β's less than 1 which do generate finite expected losses. In fact, it may not be particularly desirable to have a finite expected loss; this requires a discount factor less than 1, which is myopic in the sense that in a steady state the average stream of losses will be larger than need be. That occurs because the feedback rule does not look far enough ahead. For example, if movements in interest rates affect money with a lag, and if future losses are heavily discounted, then interest rates are not likely to be moved in any given period, and the average loss each period will become very large since money will deviate far from its target.

Given the environment described in (1) and the loss function (2), optimal control theory answers the following question: What is the linear feedback rule for choosing u_t which, on the basis of current information, minimizes the expected future loss? The solution is a feedback matrix, F, and a rule

$$(3) \qquad u_t = -F\, x_{t-1}$$

which determines u_t as a linear function of the past state and is optimal in the sense that this choice of F generates a smaller expected loss than any other choice. Specifying different values of λ will lead to different optimal rules. Given these different rules, calculating the tradeoff between the degree of money control and interest rate fluctuations is straightforward.

TIME SERIES ANALYSIS

A special problem is encountered when optimal control theory is applied to economic systems. A key element in the optimal control framework, knowledge of the laws of motion of the system, is either missing completely or known only with a large degree of uncertainty. Engineering texts on optimal control spend little time considering this problem because engineers can usually do controlled experiments in order to directly measure the response functions to whatever degree of accuracy they need. Such experiments are impossible in economic systems, though. Instead, economists have come to rely on the laws of motion embedded in econometric models.

Unfortunately, econometric models have a rather poor record as forecasters of the response of the economy to changes in policy. For example, when a key econometric relationship, the Phillips curve, was identified in the 1960s, many economists claimed it could be used as the basis for attempting to trade off higher inflation for lower unemployment. (See, for example, Tobin 1972.) After a decade of high inflation along with high unemployment, few would suggest such an approach today. The rational expectations critique of standard econometric models has provided a reasonable explanation of why those models failed, and many economists have developed a cautious, if not skeptical, attitude toward the use of control theory as a result.

At the same time that this dissatisfaction with traditional econometric models has been emerging, a number of economists--including Christopher Sims, Thomas Sargent, and staff at the Federal Reserve Bank of Minneapolis--have been developing alternative time series methods of forecasting economic variables. (Examples are in Sargent and Sims 1977, Anderson 1979, Sargent 1979b, Sims 1980, and Litterman 1981.) Not all economists would feel comfortable applying these models to the control framework, but Sims 1982 has recently made a strong defense of a time series approach to policy analysis. I take this approach, viewing Fed policy as essentially the choice of shocks to the interest rate equation in the context of a vector autoregressive representation.

I begin by constructing a vector autoregression with M1, the federal funds rate, and other variables in order to represent the laws of motion of the money market. In constructing this representation, I keep as a primary goal the desire to optimally forecast the movements of M1. For this reason I pay particular attention to a statistic measuring the out-of-sample forecasting performance of different models. I also follow the Bayesian procedures suggested by Litterman (1981) for forecasting with vector autoregressions.

I have searched through a variety of different variables, looking for those which help to predict weekly movements in seasonally adjusted M1. The federal funds rate clearly stands out as the most important. This is followed at a considerable distance by the level of commercial and industrial loans, Standard and Poor's index of 500 stocks, nonborrowed reserves, borrowed reserves, and total reserves. The <u>Business Week</u> index, a composite measure of real activity published by McGraw-Hill, shows no explanatory power. Measures of stock market volume and the discount rate do not help either.

These results are based on experiments using systems with different sets of variables to forecast M1. All systems are estimated using the same Bayesian prior, which is described in detail below. The results of some of these tests are given in Table 1.

Table 1
Forecast Performance for M1 and the Federal Funds Rate

Included variables	M1 prediction error ($billion)	Funds rate prediction error (basis points)
One-variable systems:		
M1	1.1845	
Funds rate		55.309
Two-variable systems:		
M1 and funds rate	1.1302	51.974
M1 and discount rate	1.1580	
M1 and commercial & industrial loans	1.1656	
M1 and Standard & Poor's index	1.1688	
M1 and borrowed reserves	1.1776	
M1 and total reserves	1.1804	
M1 and *Business Week* index	1.9356	
M1 and New York Stock Exchange volume	1.2023	
Three-variable systems:		
M1, funds rate, and		
Commercial & industrial loans	1.1129	52.079
Standard & Poor's index	1.1149	51.728
Borrowed reserves	1.1230	52.567
Nonborrowed reserves	1.1247	52.144
Total reserves	1.1262	51.619
Business Week index	1.1354	52.208
Discount rate	1.1388	50.023
New York Stock Exchange volume	1.1414	51.312

The prediction error in Table 1 is an out-of-sample statistic. It is based on residuals calculated by dropping each observation from the sample, one at a time, and using the estimator so obtained to generate the residual for that observation. The out-of-sample statistic is designed to distinguish variables which improve the fit only in-sample from those which actually explain out-of-sample movements. The data are weekly observations from 1976:1 through 1982:12.

Not only does the federal funds rate significantly improve the prediction of M1, but it also explains a dramatically larger share of the variation of M1 at longer horizons than any of the other variables considered. In the sets of three-variable systems in Table 1, the percentage of the one-year-ahead forecast variance explained by innovations in the funds rate varies between 49 and 73 percent. The largest share received by any of the other variables considered is 10 percent, and in several cases the share received is less than 1 percent.

For the purpose of short-run monetary control, the important aspect of the estimated model is the dynamic response of money to changes in the federal funds rate. I have found that response to be much stronger, and more stable across time periods, than the response of money to any of the other variables considered. The response to a change in the federal funds rate is also very insensitive to the addition of other variables into the money equation. Based on these results, I proceed with a bivariate autoregression using only M1 and the federal funds rate. All of the subsequent analysis could be generalized to include other variables in the state vector, but the results would probably not be materially affected.

The use of a time series representation as the basis for the dynamical structure of a control exercise is a departure from the standard econometric approach. I have decided not to estimate a structural model in this study because doing so would have greatly increased the cost and complexity of the exercise and it probably would not have led to improved estimates. In fact, as stressed by Sims (1980), the usual identifying restrictions of a structural model are likely to be false, and their application probably leads to misspecification and therefore bias in the estimation of the crucial response function. Given the strength of the evidence in the data, as seen in the lack of sensitivity to alternative specifications, the results from using a reasonable structural model would presumably be similar to those I obtain with a time series representation. However, the risks of biasing results from imposing false restrictions and inappropriate specification of dynamic structures appear to outweigh the expected benefits from a possible reduction in the variance of the estimates. Even if it would not improve the estimates, of course, one might prefer a structural model if it would be more likely to remain valid in the face of interventions. Unfortunately, construction of such an invariant structural model is an extremely difficult task. Moreover, the degree of inadequacy of the time series representation is not obvious. This issue is addressed at greater length in the last section of this paper.

IMPROVING SHORT-RUN CONTROL OF THE MONEY SUPPLY

Current Federal Reserve operating procedures do not include optimal control techniques, even though the Fed appears to be trying to solve a problem of the type which optimal control theory is designed to handle. Therefore, the solutions the Fed obtains by its current procedures may be suboptimal. It is possible, in fact, to estimate the tradeoff frontier which measures the obtainable combinations of interest rate volatility and expected deviations from monetary targets and, therefore, to measure the degree to which a change to an optimal control policy would be likely to improve operating characteristics. The tradeoffs which emerge suggest that the Fed could achieve a considerable smoothing of interest rates with little or no loss in terms of money supply control. There does not, however, appear to be much room for reducing the average size of money deviations from target. Moreover, such reductions would require large fluctuations in interest rates.

In order to understand these tradeoffs, it is first necessary to motivate my model of short-run monetary control. There are some obvious differences between my earlier discussion of that model, in which the funds rate is the control, and the usual discussion of current Fed operating procedures, which stresses reserves targets. Those differences, however, may be more apparent than real. Under current Fed policy there is an implicit role for the funds rate, and that role is the same as the one it plays in my optimal control procedure.

My interpretation of current Fed policy is based on the descriptions of operating procedures published in recent issues of the Federal Reserve Bank of New York (FRBNY) Quarterly Review and in the 1981 Federal Reserve Board of Governors staff study, New Monetary Control Procedures. In these descriptions (Davis 1979, Davis 1979-80, FRBNY 1981), the causal chain which connects changes in the nonborrowed reserve path to changes in money holdings is clearly through the level of borrowings, which affects the federal funds rate and causes banks and the public to make portfolio decisions which return money holdings to the desired level. Whether the focus is directly on the funds rate or on reserves targets, the fundamental link between open market operations and their effect on the money supply is through their effect on the funds rate.

Although there is a good deal of uncertainty over what causes money to respond to changes in the funds rate, there is general agreement that an important role is played by banks, which respond rapidly to changes in the price of reserve credit. In the first several weeks after a change in the funds rate, it is banks' decisions to make or refuse commercial loans and to buy or sell assets which transmit changes in the funds rate to changes in deposits and money. Banks do not respond directly to the level of reserves in the system, but rather to the current and expected future prices for reserves. At the margin, when deciding whether or not to make a loan, a bank compares the risk-adjusted rate of return on that loan with its alternative return from supplying those funds to the federal funds market. If other things don't change, when the funds rate goes up, for example, the level of bank loans and deposits, and hence the money supply, will go down.

The Fed has not kept secret the fact that it will occasionally modify the nonborrowed reserves path in order to affect the speed of adjustment. But any consideration of how much to adjust the nonborrowed reserves path must face the following issues: What is the effect of a change in borrowings on the funds rate? What is the response of money to changes in the funds rate, and at what level should the funds rate be targeted in order to generate the desired path for money? Thus, unless the Fed does not care about interest rate volatility, the use of the nonborrowed reserves targeting procedure does not eliminate the need to solve this optimal control problem. The monetarist policy prescription, which suggests fixing the supply of reserves no matter what happens to the money supply, in this context, amounts to a choice of loss function which cares only about hitting money targets and associates no cost with interest rate fluctuations.

The Federal Open Market Committee (FOMC) and the New York Trading Desk have long recognized that there is a tradeoff between the rapidity of reduction of short-run deviations in money and volatility of interest rates. The Committee has often expressed concern with "the possibility of whipsawing the markets and ultimately destabilizing money growth and interest rates." In a recent Fed staff study, Tinsley et al. (1981) found that there exists ". . . a well-behaved trade-off between the volatility of deviations of M1A from long-run targets and the volatility of short-term interest rates under current and alternative operating procedures that may be exploited by short-run policy."

The Tinsley et al. study involved simulations of the Board's monthly money market model. Its conclusions are similar to those reached here, although their approach differs in that they did not adopt an explicit control-theoretic framework, nor did they try to model the week-to-week dynamics of the money market. Pindyck and Roberts (1974) reached a similar conclusion in an earlier study which used optimal control but did not penalize interest rate volatility directly.

The optimal control approach to monetary control outlined above is an attempt to formalize the Fed's operating procedures and the implicit loss function which trades off short-run control for interest rate smoothness. Applying time series techniques to estimate the dynamics of the M1, federal funds process formalizes the Committee's attention to the lags inherent in the system. The Committee appears to be, in effect, attempting to solve this same problem; but without the benefit of optimal control theory and time series analysis, its solution may be suboptimal.

Because of the controversy surrounding the question of whether the Fed can or should peg interest rates, it is important to address this issue. What the optimal control procedure produces is a suggested level for the funds rate at a given point in time. The level next week will depend on the information observed this week. In contrast to Pindyck and Roberts (1974) and other control schemes which target interest rates in the future, this proposal includes no target for future interest rates. In the control scheme described here, if money deviates from its target, the funds rate will eventually adjust as much as is necessary to bring money back to its desired path.

It is only within the shortest time interval--that is, within a given week--that the funds rate is held fixed, and even this degree of interest rate control is not necessary. It would be possible to use the feedback rule defined here under a reserves targeting procedure that would not be much different from current procedures. Today, the FOMC picks target ranges for the funds rate and money growth rates, which the Federal Reserve Board and the Desk translate into reserves path targets. Under an optimal control approach, the Board and the Desk could compute reserves targets on a week-by-week basis, targets consistent with the funds rate given by the feedback rule. As long as the Fed is willing to cause the federal funds rate to move as needed to control the money supply, the difference between a funds and a reserves targeting procedure is not sharp.

The time series model that drives the analysis to follow is a bivariate autoregressive representation for seasonally-adjusted M1 and the funds rate. Twelve lags of each variable and a constant term are included in each equation. The model is estimated using weekly data from 1976:1 through 1982:40. The M1 data is logged and detrended. The estimation procedure is Theil's (1971) mixed estimation procedure, applied equation by equation--that is, ordinary least squares with the data sets augmented to include a set of observations representing a Bayesian prior of the type described by Litterman (1981).

The estimation is carried out using Doan and Litterman's (1981) regression analysis of time series program. Using their notation, the prior is a symmetric random walk with parameter 1. (Each variable in each equation is treated symmetrically; the coefficient on the own first lag has a mean of 1, and all other coefficients have a mean of 0.) The lag decay is harmonic with parameter 2. (The prior for the coefficient on lag j is centered around 0 with a standard error $1/j^2$ times the standard error on the first lag.) The overall tightness is 0.5. (The standard deviation of the prior distribution for the first lag of the dependent variable is 0.5.) The prior standard deviations of variables other than the dependent variable in each equation are scaled by the standard errors of univariate equations in order to take account of the different units of the variables. The prior permits the specification of a loosely parameterized model which retains good out-of-sample forecasting properties.

The coefficient estimates from this procedure can be viewed as an approximation of the posterior mean using this prior. These estimates are given in Table 2. It is not very enlightening to analyze the autoregressive representation directly, however, so I also present the moving average, or impulse response function, representation in Figures 1a and 1b.

Table 2
Coefficient Estimates

Table 2a: EQUATION (1) M1

OBSERVATIONS 318

R**2	0.855065
SSR	27.357168
SEE	0.293769
DURBIN-WATSON	1.971286

Label	Lag	Coefficient	Standard error	T-Statistic	Significant level
M1	1	.9436641	.0542227	17.40348	.0000000
	2	.0182534	.0634398	.28772	.7735550
	3	-.0198099	.0438006	-.45227	.6510709
	4	.0330111	.0285976	1.15433	.2483644
	5	-.0024560	.0195830	-.12541	.9001944
	6	-.0125440	.0140685	-.89163	.3725870
	7	-.0037547	.0105214	-.35686	.7211930
	8	-.0005963	.0081397	-.07326	.9415947
	9	-.0015183	.0064762	-.23444	.8146391
	10	-.0028283	.0052736	-.53632	.5917345
	11	-.0009250	.0043760	-.21138	.8325879
	12	-.0005145	.0036869	-.13955	.8890106
Funds rate	1	-.0294392	.0302283	-.97389	.3301089
	2	-.0219705	.0362646	-.60584	.5446208
	3	.0172461	.0234549	.73528	.4621638
	4	.0139132	.0150226	.92615	.3543671
	5	.0045868	.0101752	.45078	.6521476
	6	.0016376	.0072733	.22515	.8218554
	7	.0020154	.0054335	.37093	.7106879
	8	.0024023	.0042058	.57119	.5678693
	9	.0022535	.0033489	.67291	.5010028
	10	.0007649	.0027281	.28039	.7791749
	11	.0007203	.0026327	.31826	.7502875
	12	.0005434	.0019075	.28487	.7757411
Constant		.0548648	.0428758	1.27962	.2006780

Table2b: EQUATION (2) Funds rate

OBSERVATIONS	318
R**2	0.988180
SSR	75.916125
SEE	0.489370
DURBIN-WATSON	1.948948

Label	Lag	Coefficient	Standard error	T-Statistic	Significant level
M1	1	-.1750059	.0903259	-1.93749	.0526851
	2	.5034224	.1056801	4.76364	.0000019
	3	-.0250717	.0729645	-.34361	.7311353
	4	.0162133	.0476389	.34033	.7336023
	5	-.0357166	.0326221	-1.09485	.2735784
	6	-.0111387	.0234357	-.47528	.6345822
	7	-.0074525	.0175269	-.42520	.6706879
	8	-.0041482	.0135594	-.30592	.7596595
	9	-.0046949	.0107884	-.43518	.6634263
	10	-.0045208	.0087849	-.51461	.6068227
	11	-.0027660	.0072897	-.37943	.7043640
	12	-.0019039	.0061417	-.30999	.7565663
Funds rate	1	1.0718520	.0503554	21.28573	.0000000
	2	-.0871561	.0604108	-1.44272	.1490983
	3	-.0107947	.0390720	-.27627	.7823344
	4	.0031161	.0250252	.12452	.9009018
	5	-.0023759	.0169502	-.14017	.8885234
	6	.0026674	.0121161	.22015	.8257515
	7	.0000772	.0090513	.00853	.9931906
	8	.0015541	.0070062	.22181	.8244562
	9	.0020104	.0055787	.36037	.7185675
	10	.0011576	.0045445	.25472	.7989355
	11	.0012712	.0037702	.33718	.7359777
	12	.0011879	.0031777	.37383	.7085268
Constant		.1960725	.0714239	2.74519	.0060475

Figures 1a and 1b

IMPUSE RESPONSE FUNCTIONS

Figure 1a. RESPONSE TO A 1 PERCENT DEVIATION IN M1

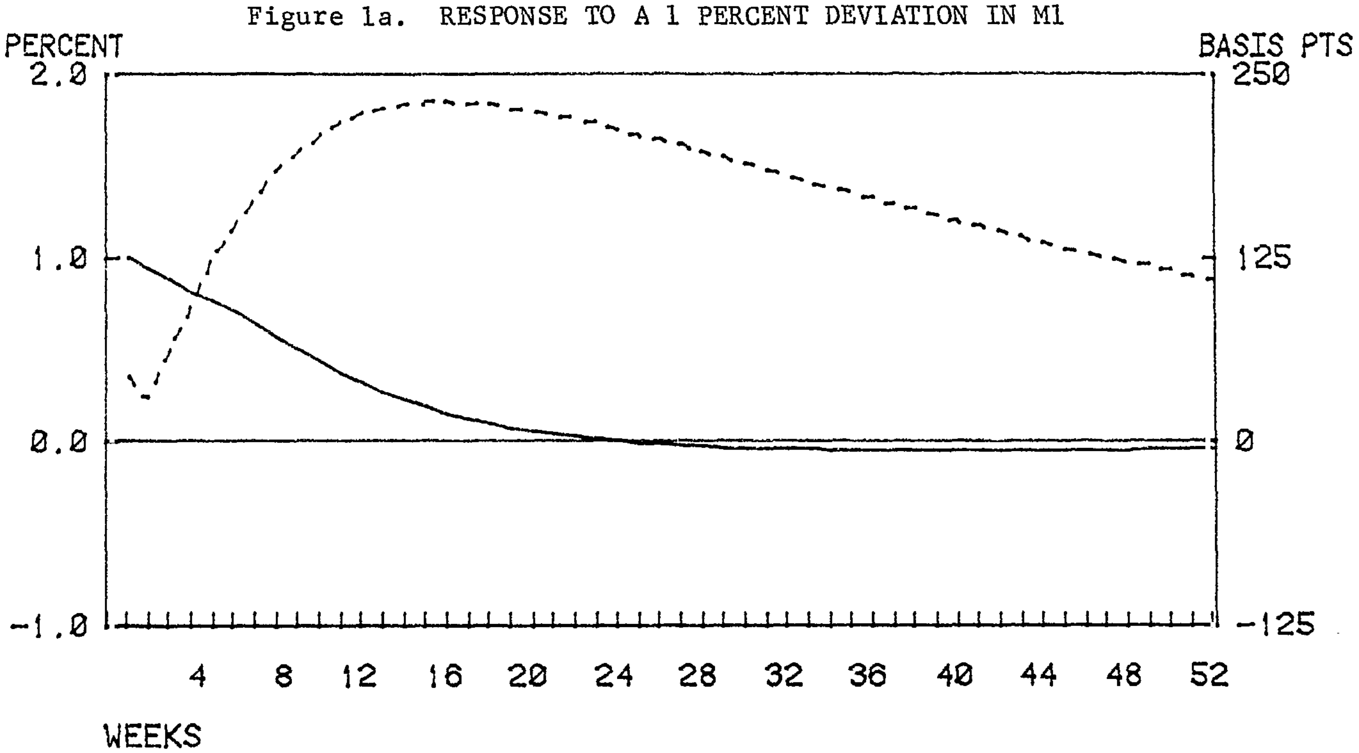

Figure 1b. RESPONSE TO A 100 BASIS POINT SHOCK IN FEDERAL FUNDS

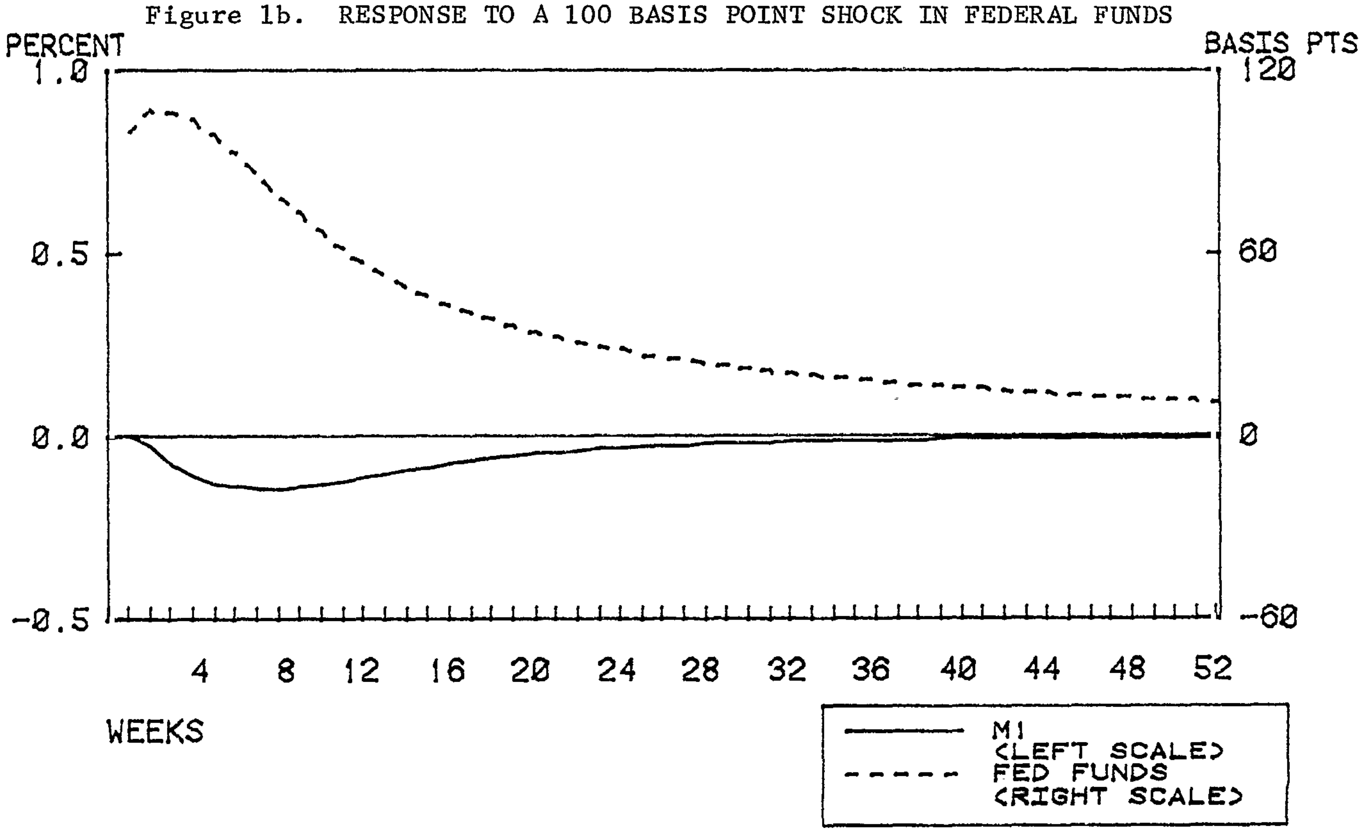

The state vector for this exercise includes 12 lags of M1, 12 lags of the federal funds rate, a constant, and a money target:

$$x_t = (m_t, m_{t-1}, \ldots, m_{t-12}, r_t, r_{t-1}, \ldots, r_{t-12}, 1, m_t^*).$$

The equation of motion is given by

$$x_t = A\, x_{t-1} + B\, u_t + w_t$$

where

$$u_t = -F\, x_{t-1}$$

defines the control. The control, u_t, is a scalar variable defined as a linear combination of the previous state vector by the feedback vector F. The vector F is generated by the solution of a matrix Riccati equation. B is a vector of zeros with a one as the 13th element, corresponding to the element r_t in the state vector. The vector w_t has zeros everywhere except in its 1st and 13th elements, which are white noise error terms with a covariance matrix equal to the estimated covariance matrix of the residuals from the post-October 1979 data. This covariance matrix is as follows:

Table 3
Covariance Matrix of Innovations

	M1	Funds rate
M1	.0860	.0382
Funds rate	.0382	.2387

Correlation .2669

The matrix A is given below:

$$
\begin{array}{ccccccccc}
a^{1,1} & a^{1,2} & \cdots & a^{1,12} & a^{1,13} & \cdots & a^{1,24} & a^{1,25} & 0. \\
1. & 0. & \cdots & 0. & 0. & \cdots & 0. & 0. & 0. \\
0. & 1. & \cdots & 0. & 0. & \cdots & 0. & 0. & 0. \\
\vdots & \vdots & & \vdots & \vdots & & \vdots & \vdots & \vdots \\
0. & 0. & \cdots & 0. & 0. & & 0. & 0. & 0. \\
a^{2,1} & a^{2,2} & \cdots & a^{2,12} & a^{2,13} & \cdots & a^{2,24} & a^{2,25} & 0. \\
0. & 0. & \cdots & 0. & 1. & \cdots & 0. & 0. & 0. \\
0. & 0. & \cdots & 0. & 0. & \cdots & 0. & 0. & 0. \\
\vdots & \vdots & & \vdots & \vdots & & \vdots & \vdots & \vdots \\
0. & 0. & \cdots & 0. & 0. & \cdots & 0. & 0. & 0. \\
0. & 0. & \cdots & 0. & 0. & \cdots & 0. & 1. & 0. \\
0. & 0. & \cdots & 0. & 0. & \cdots & 0. & 0. & g
\end{array}
$$

where g is the targeted growth rate of money (on a week-to-week basis). The $a^{i,j}$ are the coefficients from the time series model described above which determines m_t and r_t as a function of the lagged state.

This control model corresponds to a world in which the Fed at the beginning of the week picks a shock, u_t, which it does not modify as the week progresses. The model is designed to simulate a reserves targeting procedure in which the level of nonborrowed reserves to supply during the week is chosen so as to cause an optimal movement in the funds rate. Because there are unforeseen shocks during the week, given by w_t, the funds rate has a stochastic element which is not under the Fed's control.

In order to model a funds rate targeting procedure, the state vector is augmented to include the next disturbance to the funds rate equation. The A matrix is augumented by a column which is zeros except for a one in the 13th row, corresponding to the funds rate equation. This inclusion allows the feedback rule to respond to the disturbance during the week in which it occurs. Responding to the disturbance is a way to model an operating procedure in which the funds rate is targeted each week and reserves are supplied or demanded by the Fed as necessary to keep the rate within a narrow band. In this approach, the only difference between a funds rate targeting procedure and a reserves targeting procedure is the Fed's ability to respond to the disturbance: under the funds rate procedure the Fed can respond, and under the reserves procedure it cannot. Thus, using this approach implies that there will always be more noise under a reserves targeting procedure. For this reason, I will focus on the control strategy using the funds rate targeting procedure defined above.

The control problem is solved here using the passive-learning stochastic control algorithm described by Kendrick (1981) for use in cases with additive and multiplicative uncertainty. That is, the feedback control takes into account both the uncertainty due to the additive error term and the uncertainty due to the fact that the coefficients in the A matrix are not known exactly. The control is passive in the sense that no attempt is made to shock the system in order to learn more about the coefficients. In this application the control solution which does take account of coefficient uncertainty in the passive sense is very similar to the solution in which the coefficients are treated as known. This result suggests that there would be very little gain available through an active-learning algorithm.

To this point, no mention has been made of the fact that the money supply is not observed contemporaneously with the funds rate. For the purpose of optimal control, there is an important separation of the problem of setting a control from the problem of observing the current state. (For a statement of this result, see Bertsekas 1976.) The implication of this result is that when one or more of the most recent observations of money are not available, the optimal strategy is to form the best linear prediction of these values of money and then to proceed as if they had been observed.

In practice, depending on the day of the week, the lag between the observation of the funds rate and M1 varies between 7 and 12 business days. I model this as a two-week lag in the weekly data. Thus, I proceed in two steps. First I form the optimal linear forecast of the most recent two weeks of money data; then I proceed as above. The forecasting exercise is conditional on the two advanced observations on the funds rate. (The optimal linear forecasting procedure in this case is described in Example 13.5 of Doan and Litterman 1981.) The astute reader will have realized that the conditional forecast depends on the reduced form, which is a function of the feedback control rule; but the feedback control rule itself is a function of the conditional forecast. Thus, with two unobserved values of money in the state vector, the problem of finding the optimal control rule requires a simultaneous solution with the problem of generating a conditional forecast. Actually, the problem is not all that serious. The method described below has worked quite well with very little additional computing expense.

The solution procedure is a simple iteration. The reduced form in the first step is derived by solving the matrix Riccati equation for a feedback control vector, F, and plugging it into the state equation:

$$(4) \qquad x_t = (A-BF)\ x_{t-1} + w_t.$$

The conditional forecast of $\hat{x}_{t-1}$, given observations on a subvector of x_{t-1}, can be written as

$$(5) \qquad \hat{x}_{t-1} = G\ x_{t-1}$$

where G is a matrix which has zeros in the columns corresponding to the unobserved components of x_{t-1}. For a given G, the reduced form is

$$(6) \qquad x_t = (A-BFG)\ x_{t-1} + w_t\ .$$

This reduced form implies a new G, and so on. Note that each iteration adds two lags to the state vector. Thus, in principle, the reduced form has an infinite autoregressive representation. In practice, within the relevant range of λ's, iterating between these two equations quickly leads to convergence of G and the reduced form transition matrix, (A-BFG). Notice that this iterative procedure does not require repeated solution of the matrix Riccati equation which determines F. In Figures 2a and 2b, I illustrate the reduced form responses of money and the funds rate to a 1 percent innovation in money for two particular values of λ. The response of the funds rate could be viewed as a Fed reaction function under an optimal control approach. Note that there

Figures 2a and 2b

RESPONSE FUNCTIONS UNDER OPTIMAL RULE

RESPONSE TO A 1 PERCENT DEVIATION IN M1

Figure 2a. LAMBDA = .2

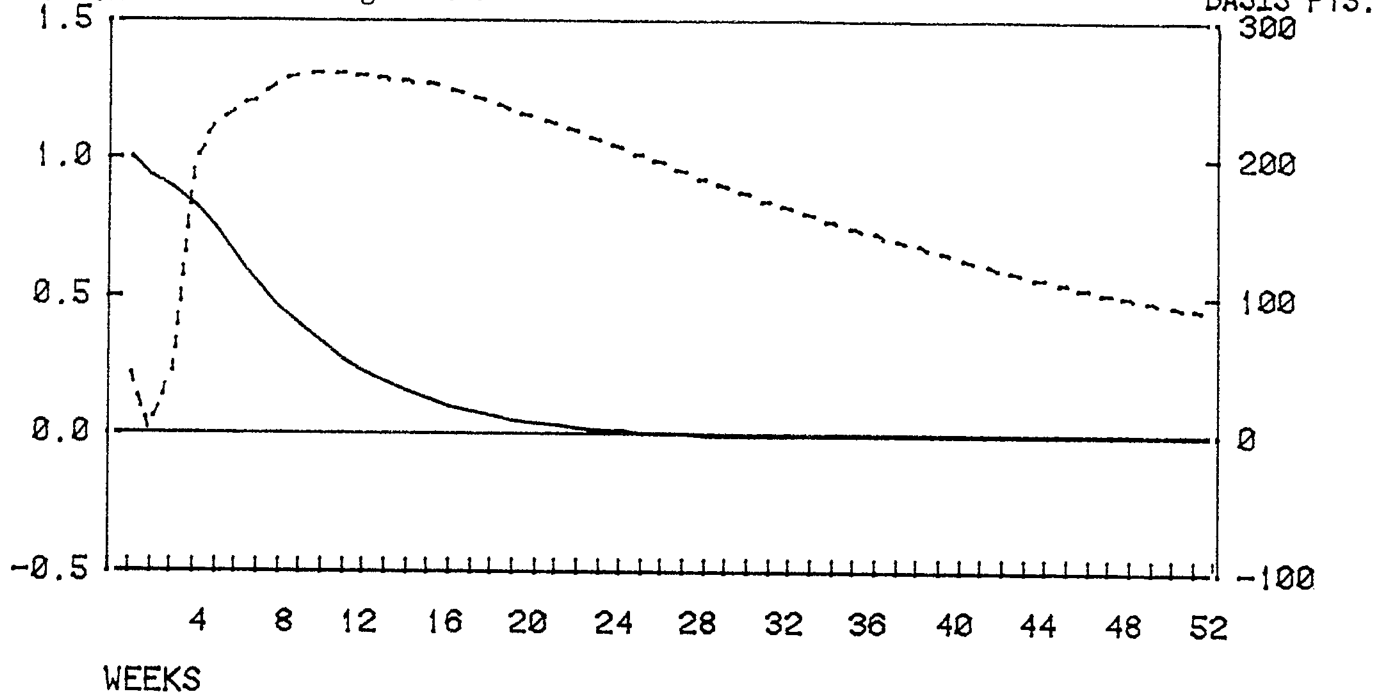

Figure 2b. LAMBDA = 6.

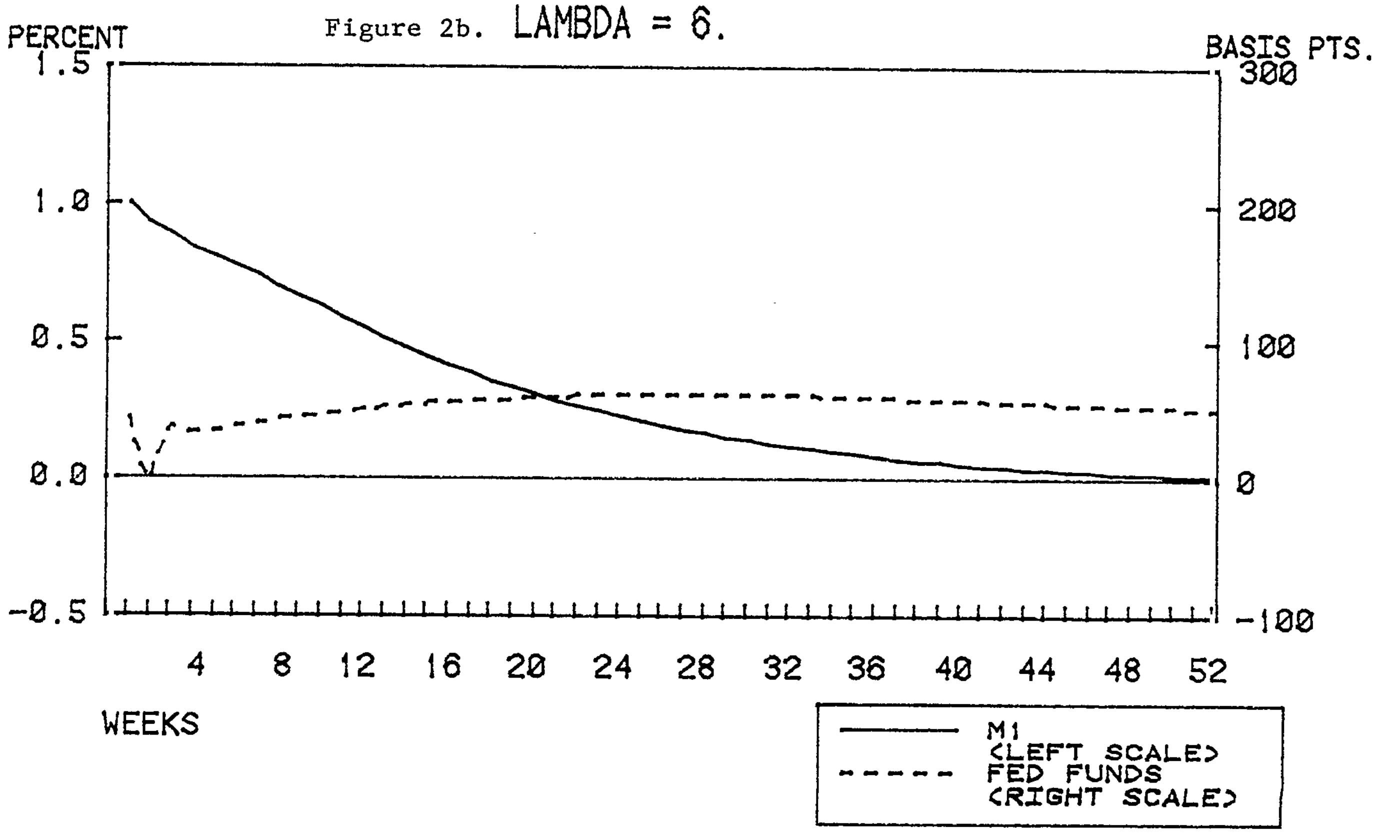

is a two-week lag in the response of interest rates to the money innovation. This is due to the delay in the observation of the money innovation. In Figure 2b, with a larger λ, more weight is given to smoothing interest rates; this causes a smaller interest rate response and a longer delay in returning money to the target path.

Once the optimal feedback rule has been calculated, taking into account the lagged observation of money, the probability laws of the controlled system are determined so measures of expected interest rate volatility and money supply deviations can be calculated. This means calculating the set of points, associated with different values of λ, which represent the best possible solutions to the problem of minimizing both money supply deviations and interest rate volatility.

To illustrate this tradeoff, I present estimates of the minimum obtainable cost combinations in Figure 3. Different relative weights attached to the goals of money control and interest rate smoothness will lead to different optimal feedback rules and thus to different points on the graph. Again, by connecting these points, I trace out the tradeoff curve, or possibility frontier, from which the Fed can choose optimal procedures.

In order to generate these points, I start with the vector autoregressive representation, which generates a set of one-step-ahead forecast errors, or shocks, for the period over which it is estimated. These shocks are then used in a simulation exercise to answer the question of how much better could the Fed have done in the post-October 1979 period, had it been following

Figure 3

POST - OCTOBER 1979 POSSIBILITY FRONTIER
TRADEOFF BETWEEN INTEREST RATE FLUCTUATIONS AND
MONEY SUPPLY DEVIATIONS FROM TARGET

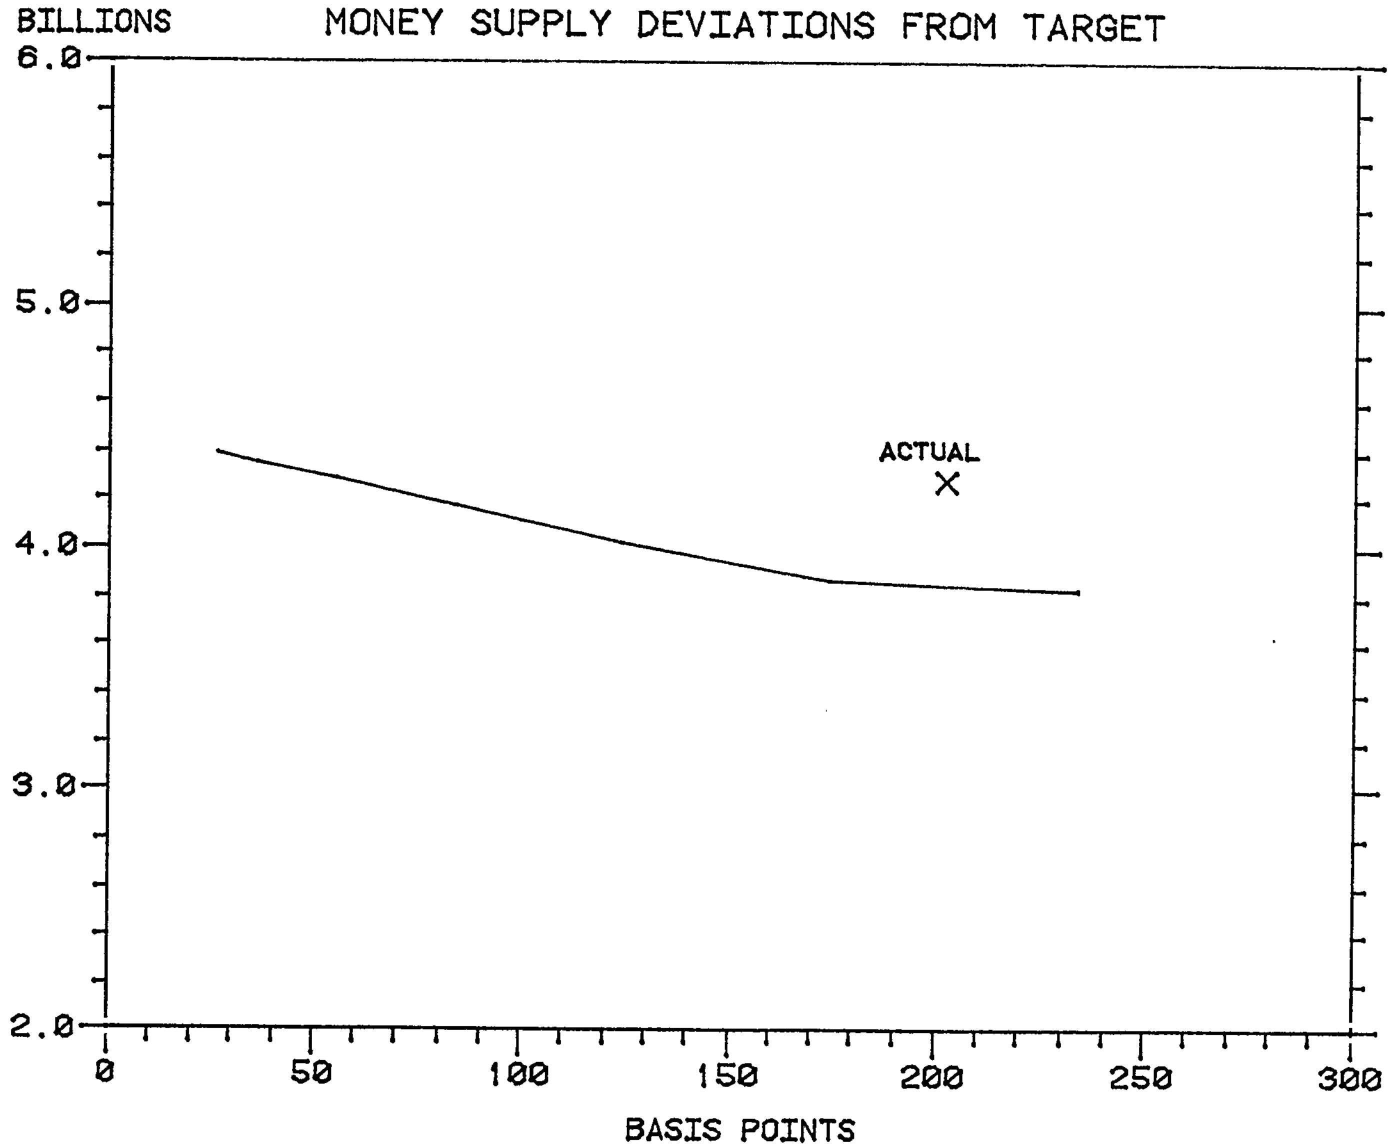

an optimal control policy. First I define a target path for this period. Since I am focusing on short-run control, I will take as the target the long-run trend fitted to the logged M1 data. For any particular values of λ, I can generate the paths the state variables would have taken assuming:

- The state evolved according to the vector autoregressive representation.
- An optimal control policy had been in force.
- The same set of shocks hit the system.

It should be clear that a tradeoff curve defines a broad set of possible feedback rules. Each point on the curve represents a feedback rule which is optimal for a particular weighting of the Fed's goals. Deciding which rule should be chosen from this set means deciding how relatively important are the goals of money control and interest rate stability; that is beyond the scope of this analysis. What this analysis does tell the Fed is how costly, in terms of interest rate volatility, closer control of the money supply is (and vice versa).

My estimates of both the tradeoff curve, based on the data since early October 1979, and the actual costs in this period are shown in Figure 3. Note that the curve is fairly flat and the actual is fairly close to it. Together these estimates imply two main results:

- Short-run money deviations from target cannot be reduced much from recent levels without incurring large increases in interest rate volatility.

- Short-run interest rate volatility can be reduced quite a bit from recent levels without reducing the degree of money control.

Readers familiar with the behavior of money and interest rates before October 1979 may find an apparent discrepancy between that behavior and my tradeoff curve. The tradeoff curve implies that more interest rate volatility is associated with closer control of the money supply. Since October 1979, however, both money deviations from trend and interest rate volatility have increased. This discrepancy does not refute the existence of the tradeoff curve. The tradeoff penalizes money deviations from target, not money growth volatility. The use of a trend growth rate for money as a basis for computing deviations from target badly underestimates the true situation before October 1979. Although there is no exact measure of how close the Fed has come to hitting its target, all indications are that the Fed has been closer to its desired trend growth path recently than it was several years ago.

Even given the above qualifications, however, there is still a large increase in money stock volatility in recent years which my model does not account for. Some possible explanations include the increase in financial innovations, such as the nationwide introduction of checkable interest-earning accounts in January 1981; the credit controls of spring 1980; the more general tightening of policy over the whole period; and finally, the fact that seasonal effects are harder to remove from recent data than from data around which there are several years of observations.

In my analysis, these effects end up as unexplained shocks to money. Since the tradeoff curve I have estimated is based on shocks of the size experienced between late 1979 and late 1982, it does not apply to other periods. Whenever shocks to money are larger (or smaller), the tradeoff curve for that period will be higher or (lower) than the one in Figure 3.

The model I used to calculate the tradeoff curve can also demonstrate how different the weekly history of money and interest rates might have been if the Fed had chosen an optimal procedure suggested by my results. Figures 4a and 4b compare the actual history of M1 and the federal funds rate during 1980-82 with what the model suggests could have been accomplished under an optimal control procedure which weighted stabilizing interest rates more highly than actual results imply the Fed did. The comparison suggests that the funds rate could have been smoothed considerably with little or no adverse effect on money control. Clearly, though, this particular optimal control solution does not promise to reduce money deviations from their current level. Also, note that, while in this simulation the smoothed interest rate is usually lower than the actual rate, that will not generally be true. For example, in a period of falling rates, a smoothed rate will usually be higher than otherwise.

Figures 4a and 4b

SIMULATIONS COMPARING ACTUAL WITH CONTROLLED FEDERAL FUNDS TARGETING PROCEDURE

Figure 4a. M1 DEVIATIONS FROM TARGET

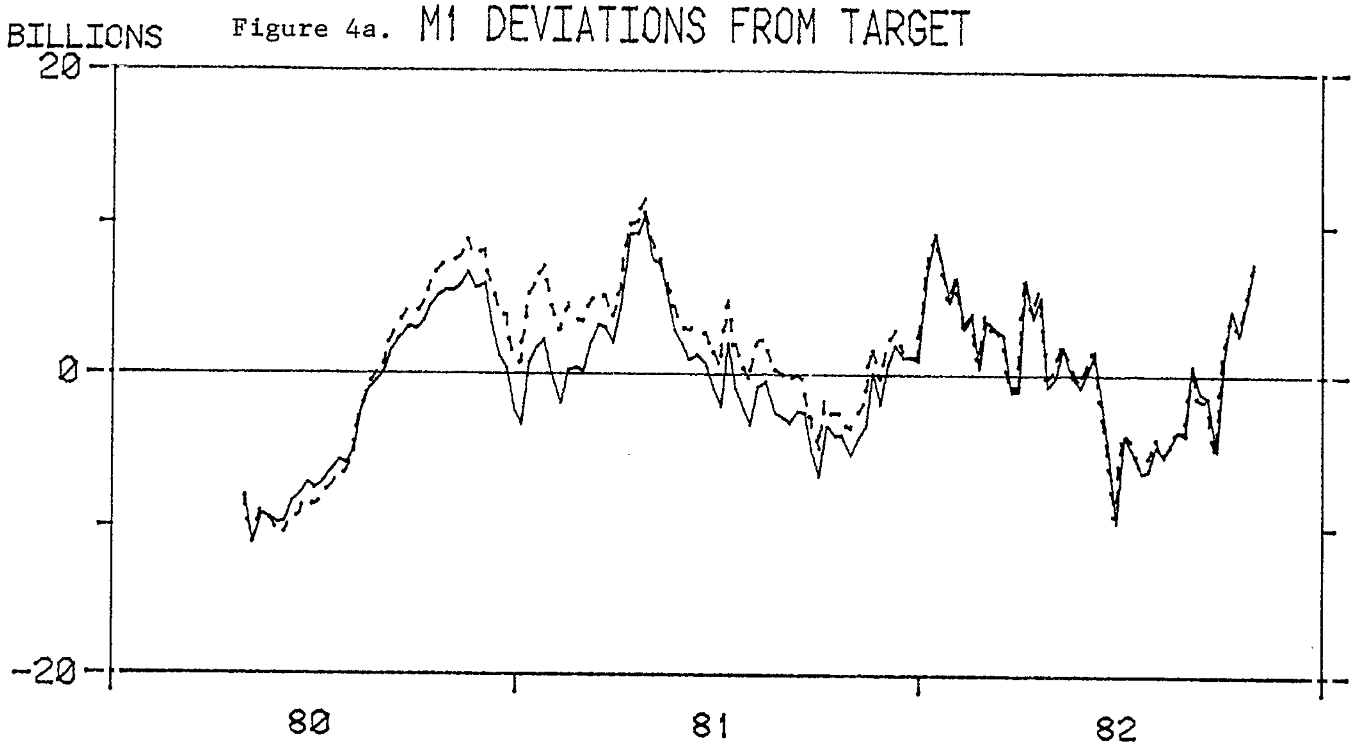

Figure 4b. FEDERAL FUNDS RATE

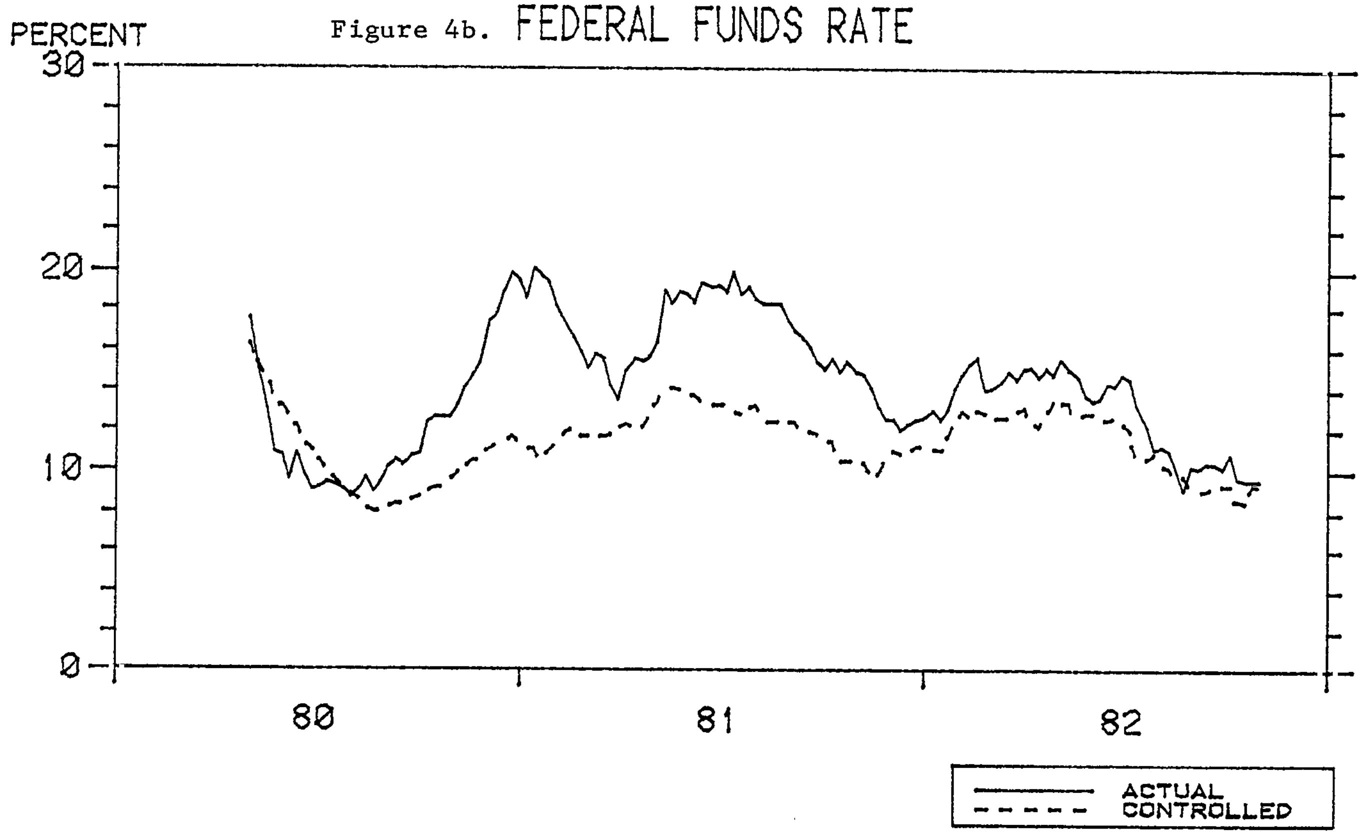

EVIDENCE ON STRUCTURAL STABILITY

There is no guarantee that changes in the operating procedures of the Fed would leave unaffected the important dynamics of the money market on which this procedure depends. There is evidence which suggests, however, that the impact would not be large.

A key assumption of the above exercise is that the dynamics of the money market variables would not change too much as a result of the adoption of an optimal control rule. Whether or not this is likely to be true is a key question; it is, after all, the focus of the rational expectations criticism of traditional econometric exercises of this type. According to the rational expectations argument, changes in the policy rule of the government will lead to changes in the actions of agents in the economy, and the new dynamic behavior of the economy is likely to be far different from the old behavior. (For a forceful exposition of this viewpoint, see Lucas 1976.)

At an abstract level, it is impossible to defend a time series representation against this criticism. The best one can do is to question, for a particular relationship, whether the effects of a given intervention will be important. Here the relationship of interest is the response of money holdings to movements in interest rates, and the intervention is the adoption of an optimal control strategy. One relevant issue is how stable that relationship has been to similar interventions, if there have been any, in the past.

With respect to the money market, we are now in the fortunate circumstances of having one bit of empirical evidence which may help resolve this issue. In October 1979 the Federal Reserve made a change in operating procedures which arguably was a more striking change than would be the adoption of the optimal control techniques proposed here. If the dynamics of the system were not affected too much by the recent change, then there is good reason to hope that they would not be too sensitive to the change proposed here.

Unfortunately, testing for structural change can be a tricky proposition. For example, it is obvious from the data that something changed in October 1979. The standard errors of innovations in M1 and the funds rate are many times larger after that date. The question of interest, however, is whether there is evidence that the response function of M1 to a shock in the funds rate changed. Based on visual inspection of the response functions presented above and a statistical test described here, there is no reason to believe that the response of money changed significantly when the Fed changed its operating procedures.

The test is as follows. One-step-ahead forecasts of money are made separately based on the data before and after the change. The forecasts are made out-of-sample, in a sense to be made precise below. If there has been a significant change in structure, then making forecasts using the full sample should lead to larger errors in both subsamples. In fact, using 12 lags, the forecasts of money in the first half of the sample improve only marginally after dropping the second half, and the forecasts of money in the second half improve considerably using estimates based on the full sample. Using two lags, the forecasts based on the full sample are better in each subsample than the forecasts based on the subsample alone.

The out-of-sample nature of the test is that for each period, the forecast of money for that period is based on an estimator using all observations in the relevant sample except that period's observation. The reason for this procedure is that if the test is done in-sample, then the subsample estimates must fit better. One version of the standard Chow test for structural stability is based on the asymptotic distribution of the size of this in-sample improvement. (See, for example, Sims 1980.) Asymptotically, my residuals and test statistic will have the same distribution. The fact that there is little or no improvement in the two subsamples means that the change in structure, if it occurred at all, was not large.

Table 4
Stability Test Results

Period	Prediction error based on subsample ($billion)	Prediction error based on full sample ($billion)
Estimated using 2 lags		
1976:13 to 1979:40	0.58473	0.57186
1980:1 to 1982:12	1.65326	1.64932
Estimated using 12 lags		
1976:13 to 1979:40	0.59865	0.60660
1980:1 to 1982:12	1.73695	1.60863

These empirical results suggest that the change in money response would not be large if the Fed were to modify its behavior by adopting an optimal control strategy.

CONCLUSION

This application of optimal control theory and time series analysis has identified an important tradeoff between degrees of short-run monetary control and interest rate volatility. Two principal conclusions emerge:

- Application of optimal control theory would likely improve Federal Reserve operating procedures.
- Interest rate volatility can be reduced considerably from current levels without adversely affecting the degree of monetary control achieved.

REFERENCES

Anderson, Paul A. 1979. Help for the regional economic forecaster: Vector autoregression. Federal Reserve Bank of Minneapolis Quarterly Review 3 (Summer): 2-7.

Bertsekas, Dimitri P. 1976. Dynamic programming and stochastic control. New York: Academic Press.

Chow, Gregory C. 1975. Analysis and control of dynamic economic systems. New York: Wiley.

Davis, Richard G. 1979. Broad credit measures as targets for monetary policy. Federal Reserve Bank of New York Quarterly Review 4 (Summer): 13-22.

__________. 1979-80. The monetary base as an intermediate target for monetary policy. Federal Reserve Bank of New York Quarterly Review 4 (Winter): 1-10.

Doan, Thomas A., and Litterman, Robert B. 1981. User's manual RATS version 4.1. Processed. Minneapolis, Minn.: VAR Econometrics.

Federal Reserve Bank of New York (FRBNY). 1981. Monetary policy and open market operations in 1980. Federal Reserve Bank of New York Quarterly Review 6 (Summer): 56-75.

Holbrook, Robert S. 1972. Optimal economic policy and the problem of instrument instability. American Economic Review 62 (March): 57-65.

Kalchbrenner, J. H., and Tinsley, P.A. 1976. On the use of feedback control in the design of aggregate monetary policy. American Economic Review 66 (May): 349-55.

Kendrick, David. 1981. Stochastic control for economic models. New York: McGraw-Hill.

Kwakernaak, Huibert, and Sivan, Raphael. 1972. Linear optimal control systems. New York: Wiley-Interscience.

Litterman, Robert B. 1981. A Bayesian technique for forecasting with vector autoregressions. Working Paper. Massachusetts Institute of Technology.

Lucas, Robert 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 1: 19-46. Amsterdam: North Holland.

Pindyck, Robert S., and Roberts, Steven M. 1974. Optimal policies for monetary control. Annals of Economic and Social Measurement 3 (January): 207-37.

Radecki, Lawrence. 1982. Short run monetary control: An analysis of some possible dangers. Federal Reserve Bank of New York Quarterly Review 7 (Spring): 1-10

Sargent, Thomas J. 1979a. Macroeconomic Theory. New York: Academic Press.

__________. 1979b. Estimating vector autoregressions using methods not based on explicit economic theories. Federal Reserve Bank of Minneapolis Quarterly Review 3 (Summer): 8-15.

Sargent, Thomas J., and Sims, Christopher A. 1977. Business cycle modeling without pretending to have too much a priori economic theory. In New methods in business cycle research: proceedings from a conference, ed. C. A. Sims, 45-109. Minneapolis: Federal Reserve Bank of Minneapolis.

Sargent, Thomas J., and Wallace, Neil. 1975. "Rational" expectations, the optimal monetary instrument, and the optimal money supply rule. Journal of Political Economy 83 (April): 241-54.

Sims, Christopher A. 1980. Macroeconomics and reality. Econometrics 48 (January): 1-48.

__________. 1982. Policy analysis with econometric models. In Brookings Papers on Economic Activity, ed. William C. Brainard and George L. Perry, 1: 107-52. Washington, D.C.: Brookings Institution.

Theil, Henri. 1971. Principles of econometrics. New York: Wiley.

Tinsley, Peter A.; von zur Muehlen, Peter; Trepeta, Warren; and Fries, Gerhard. 1981. Money market impacts of alternative operating procedures. In New monetary control procedures, vol. 2. Federal Reserve Staff Study. Washington, D.C.: Federal Reserve Board of Governors.

Tobin, James. 1972. Inflation and unemployment. American Economic Review 62 (March): 1-18.