ELSEVIER Journal of Monetary Economics 38 (1996) 173–193

JOURNAL OF Monetary ECONOMICS

Evaluating alternative monetary policy rules

Ray C. Fair *[,a], E. Philip Howrey[b]

[a] *Yale University, New Haven, CT 06520-8281, USA*
[b] *University of Michigan, Ann Arbor, MI 48109-1220, USA*

(Received January 1995; final version received April 1996)

Abstract

This paper examines monetary policy from an optimal control perspective. Five loss functions are minimized for each of five models, and the results are compared. The basic ('true') loss function targets inflation and unemployment. The other loss functions target, respectively, inflation alone, unemployment alone, nominal growth alone, and real growth alone. The five models are two small structural models, two VAR models, and a large structural model. A numerical procedure is presented that can handle a variety of loss functions and models.

Key words: Monetary rules; Optimal Fed behavior

JEL classification: E52

1. Introduction

There is a large and growing literature on determining the best monetary policy rule for the Federal Reserve.[1] The general approach in this literature is to 1) choose a policy instrument – usually the monetary base, M1, M2, or a short-term interest rate, 2) choose a target variable – such as nominal GDP – and a rule, and 3) examine, using a model of the economy, how the economy would have behaved under the rule versus how it in fact behaved. The rule typically expresses the policy instrument as a function of the deviation of the

* Corresponding author.

We are indebted to Al Klevorick, James Stock, and a referee for helpful comments.

[1] Some of the more recent studies are Clark (1994), Croushore and Stark (1994), Feldstein and Stock (1993), Hall and Mankiw (1993), Judd and Motley (1993), and Thornton (1995). Widely cited earlier studies include McCallum (1988), Taylor (1985), and Tobin (1980). Taylor (1985, Fn. 1, p. 61) cites much of the literature prior to 1985.

PII S0304-3932(96)01279-2

target variable from its target value. Usually the actual values of the variances of key macroeconomic variables like the real growth rate and the rate of inflation are compared to the values of the variances that would have occurred had the rule been followed. The aim is to find the rule (including the choice of the target variable) that gives (in some sense) the best overall performance of the economy.

The general question of determining the best monetary policy rule of the Fed is an obvious one to examine using optimal control techniques, but with the exception of Feldstein and Stock (1993), the literature has not been concerned with solving optimal control problems and deriving optimal monetary policy rules. In this paper we examine monetary policy from an optimal control perspective. Our general procedure is as follows.

First, we choose a particular loss function and assume that everyone (including the Fed) agrees that this is the loss function whose expected value the Fed should minimize. Call this the 'true' loss function. Second, we choose a policy instrument and use this instrument to minimize the expected value of the true loss function. Third, we use the instrument to minimize the expected value of other, simpler loss functions (such as one that targets only nominal GDP). Finally, we compare the different outcomes to see how close the minimization of the expected value of the simpler loss functions comes to the minimization of the expected value of the true loss function. We also compare the outcomes to the actual, historical outcome to see how much better the economic performance would have been had the Fed behaved by minimizing the expected value of the particular loss function in question. We use five different models for these results: two small structural models, two VAR models, and a large structural model. The policy instrument for one of the small structural models and for one of the VAR models is the money supply (M2), and the policy instrument for the other three models is the three-month Treasury bill rate. The VAR models are essentially two versions of the same model, one with an equation for the money supply and one with an equation for the bill rate.

Our results allow us to examine how much is lost by targeting, for example, only nominal GDP. If the results of minimizing the expected value of the loss function that targets only nominal GDP are close to the results of minimizing the expected value of the true loss function, then the recommendation some studies have made that only nominal GDP should be targeted may be worth considering on grounds of simplicity. If, on the other hand, the results are not close, then the recommendation should probably be rejected. The same holds for targeting only real GDP, only inflation, and only unemployment.

The control period is 1962:1–1993:2 (126 quarters). The target value for each of the various target variables in the loss functions is taken to be the mean of the actual values of the variable over the control period. The focus in this paper, as in much of the literature, is on variances and not means, and having the target values be the mean values is a way of eliminating mean effects.

The loss functions and control problems are discussed in Section 2, and the procedure for solving the control problems is explained in Section 3. The models are then discussed in Sections 4, 5, and 6, and the results are presented in Section 7. Finally, Section 8 discusses how stochastic simulation might be used in future work, and Section 9 concludes with some other suggestions for future research.

Before proceeding, it should be noted that the optimal control approach is not without criticism. Some have argued that monetary policy rules should be robust in the sense of working well across a variety of models. This may not be true of a rule derived by minimizing the expected value of a loss function using a particular model, since the rule is obviously model-specific. The robustness property would be a good one if the models in question were all equally likely of being the best approximation of the economy. If, however, all models but, say, one were poor approximations, one would want to find a rule that worked well on the good model regardless of how it worked on the other models.

The literature to date has worked with very small models, and it does not seem likely that any of them is a good approximation. The large structural model considered in this paper has been extensively tested, and it seems (to us) to be much more likely to be a good approximation than are the other four models considered here and similar models in the literature. We thus put much more weight in this paper on the results using the large model, and we are not concerned with how the (implicit) rule for the large model might work for the small models. The main reason for including the results for the small models in this paper is to have a reference point to the previous literature, which, as just noted, has only been concerned with small models.

Another criticism of the optimal control approach is that it is time-inconsistent if expectations are rational. This is, of course, a problem for any rule, not just a rule derived from the solution of an optimal control problem. Some precommitment technology is needed to avoid this problem, and this is true whether the rule is optimal or whether it is chosen in some other way.

Finally, the work in this paper is based on the assumption that the parameters of the models do not change as policy rules change. Regarding the large model, it is argued in Fair (1994) that if changing policy rules cause important parameter changes, the model should not do well in the various tests that were performed on it. The model in fact does reasonably well in the tests.

2. The loss functions and control problems

Let X_t denote the log of nominal GDP, Y_t the log of real GDP, and P_t the log of the GDP price index, where $X_t = P_t + Y_t$ and where the t subscript refers to period t.[2] Denote the growth rates as: $x_t = 400(X_t - X_{t-1})$, $y_t = 400(Y_t - Y_{t-1})$,

[2] All the data used in this paper are described in Fair (1994).

and $p_t = 400(P_t - P_{t-1})$. The differences are multiplied by 400 to put the growth rates at annual rates in percentage points. Let U_t denote the level of the unemployment rate, and let x^*, y^*, p^*, and U^* denote the target values of x_t, y_t, p_t, and U_t, respectively, where, as noted above, the target values are the means of the historical values over the control period. Finally, let L_{jt} be the loss in period t associated with variable j deviating from its target, where j is either x, y, p, or U. L_{jt} is postulated to be

$$L_{jt} = (j_t - j^*)^2. \tag{1}$$

As noted in the Introduction, we consider two policy instruments in this paper, M2 and the three-month Treasury bill rate. Let M_t denote the log of M2, and let m_t denote its growth rate: $m_t = 400(M_t - M_{t-1})$. Let R_t denote the level of the three-month Treasury bill rate. In the following analysis we will use a measure of the cost associated with changing the policy instrument. The cost in period t, denoted D_{zt}, where z is either m or R, is postulated to be

$$D_{zt} = (z_t - z_{t-1})^2. \tag{2}$$

When m is the policy instrument, the 'true' loss in period t is postulated to be

$$H_t = 0.5L_{pt} + 0.5L_{Ut} + \alpha D_{mt} + \frac{0.1}{m_t^- + 2.001} + \frac{0.1}{21.001 - m_t^+}, \tag{3}$$

where m_t^- is equal to m_t if m_t is greater than -2.0 and to -2.0 otherwise and m_t^+ is equal to m_t if m_t is less than 21.0 and to 21.0 otherwise. When R is the policy instrument, the 'true' loss in period t is postulated to be

$$H_t = 0.5L_{pt} + 0.5L_{Ut} + \alpha D_{Rt} + \frac{0.1}{R_t^- - 1.999} + \frac{0.1}{16.001 - R_t^+}, \tag{4}$$

where R_t^- is equal to R_t if R_t is greater than 2.0 and to 2.0 otherwise and R_t^+ is equal to R_t if R_t is less than 16.0 and to 16.0 otherwise.

The third, fourth, and fifth terms in Eqs. (3) and (4) require some explanation. Consider first the third term, and assume that the control period of interest is 1 through T. For each control problem solved in this paper we have chosen α so that the value of $\sum_{t=1}^{T} D_{zt}$ using the optimal values of z is close to the value using the historical values of z (where z is either m or R). We have thus constrained the problem so that the period-to-period variation in the policy instrument as measured by $\sum_{t=1}^{T} D_{zt}$ is roughly the same in the optimal solution as it is historically. We are taking the actual value of $\sum_{t=1}^{T} D_{zt}$ as the maximum amount of variation in the policy instrument that the Fed allows, and we are thus constraining each solution to have roughly no more variation than this. Without this type of constraint, the solutions sometimes have huge period-to-period changes in the policy instrument, and it is not sensible to think that the

Fed would ever behave in this extreme way. The Fed undoubtedly assigns costs to changing policy instruments, which is what our treatment is doing.

The fourth and fifth terms in Eq. (3) insure that the optimal values of m will be between -2.0 and 21.0. The range of m in the control period is -1.34 to 20.34, and so the fourth and fifth terms keep m roughly within its historical range. We are taking this range to be the maximum that the Fed allows. Similarly, the fourth and fifth terms in Eq. (4) insure that the optimal values of R will be between 2.0 and 16.0. The range of R in the control period is 2.72 to 15.09, and so the fourth and fifth terms keep R roughly within its historical range.

The main feature of Eqs. (3) and (4) is that we are taking the true loss to be deviations of unemployment and inflation from their targets. The Fed is assumed in the final analysis to care only about these two variables (except for the cost of changing the policy instrument). As an approximation, this does not seem an unreasonable assumption, but it would be easy in future work to have the true loss depend on more variables. We do need, however, to make some assumption about the true loss in order to have a basis for comparison of alternative policy rules.

The other loss functions tried are one that targets only inflation [the first two terms in (3) and (4) replaced by L_{pt}], one that targets only unemployment [the first two terms in (3) and (4) replaced by L_{Ut}], one that targets only nominal growth [the first two terms in (3) and (4) replaced by L_{xt}], and one that targets only real growth [the first two terms in (3) and (4) replaced by L_{yt}].

As noted above, the control period of interest is 1 through T, where 1 is 1962:1 and T is 1993:2. However, in order not to have to assume that life ends in T, the control problem should be thought of as one of minimizing the expected value of $\sum_{t=1}^{T+n} H_t$, where n is chosen to be large enough to avoid unusual end-of-horizon effects near T. The overall control problem should thus be thought of as choosing the values of m or R that minimize the expected value of $\sum_{t=1}^{T+n} H_t$ subject to the model used.

3. The solution procedure

If a model is linear and the objective function quadratic, it is possible to derive analytically optimal feedback equations for the control variables.[3] In general, however, optimal feedback equations cannot be derived for nonlinear models or for objective functions with nonlinear constraints on the instruments, and a numerical procedure must be used. The following procedure was used for the results in this paper. It is based on a sequence of solutions of deterministic control problems.

The aim is to approximate the T values of z, denoted $z_1^*, \dots, z_T^*$, that would have been chosen had it been possible to derive analytically an optimal feedback

[3] See, for example, Chow (1981).

equation for z. Given information at the beginning of period 1, the overall optimal control problem is to minimize the expected value of $\sum_{t=1}^{T+n} H_t$ subject to the model. As noted above, although n is used here to avoid end-of-horizon problems, in the final analysis we are only interested in the values of the control variable through T.

For purposes of solving the control problems, the Fed is assumed to know the model (its structure and coefficient estimates). This is more than the Fed would have known historically because the model was developed after the first quarter of the control period and is estimated through the end of the control period. On the other hand, the Fed is assumed *not* to know the future values of any endogenous variable, any exogenous variable, or any error term when solving the control problems. The optimal path of z is something the Fed could have achieved had it had knowledge of the model.

The procedure for solving the overall control problem is as follows:

1. If the model has exogenous variables, which the large structural model considered in this paper has, add estimated autoregressive exogenous-variable equations to the model, so that the model has in effect no exogenous variables. [For the large structural model this is Model US+ in Fair (1994).] Set all the error terms to zero (their expected values), including the errors in the exogenous-variable equations, for periods 1 through k, where k is the length of the horizon for the first solution. Choose values of z for periods 1 through k that minimize $\sum_{t=1}^{k} H_t$ subject to the model with the error terms set to zero. This is just a deterministic optimal control problem, which can be solved, for example, by the method in Fair (1974).[4] Let z_1^* denote the optimal value of z for period 1 that results from this solution. The value of k should be chosen to be large enough so that making it larger has a negligible effect on z_1^*. z_1^* is a value that the Fed could have computed at the beginning of period 1 (assuming the model were known) having knowledge of the endogenous-variable values, the exogenous-variable values, and the error terms only up to, but not including, period 1.[5]

2. Drop the exogenous-variable equations, and set the error terms in the structural equations for period 1 equal to their historical values. Solve this version

[4] One could at this stage use stochastic simulation to compute the expected value of $\sum_{t=1}^{k} H_t$ for each function evaluation. [See, Fair (1994, pp. 271–272) for an explanation of this.] In most applications, however, the expected value computed via stochastic simulation is very close to the expected value computed by setting the error terms to zero and solving once. It is thus usually not worth the extra cost to use stochastic simulation to compute the expected values.

[5] If the policy instrument enters the model with a one-quarter lag, as it does in the small models examined in this paper, this must be taken into account in the solution procedure. For example, if $z_t = m_{t-1}$, then $m_1^*, \ldots, m_{k-1}^*$ are chosen to minimize $\sum_{t=1}^{k} H_t$. The value for m_0 is the historical value.

of the model for period 1 using z_1^* instead of its actual value and using the actual values of all the other exogenous variables. (This is just a deterministic simulation for period 1.) The solution values from this simulation are what the model estimates would have occurred historically in period 1 had the Fed chosen z_1^* and had the exogenous-variable values and error terms been what they were historically.

3. Repeat steps 1 and 2 for periods 2 through T. For an arbitrary period s, use the solution values of all endogenous variables for periods $s-1$ and back, as well as the values of z_{s-1}^* and back.

The solution values of the endogenous variables carried along from period to period in the above procedure are estimates of what the economy would have been like had the Fed chosen $z_1^*,\ldots,z_T^*$ and had the exogenous-variable values and error terms been what they were historically. Given knowledge of the model, these are values that the Fed could have achieved. Note that once $z_1^*,\ldots,z_T^*$ are chosen, the solution values for periods 1 through T can be generated all at once. This can be done by 1) dropping the exogenous-variable equations, 2) setting the structural error terms to their historical values, 3) using the actual values of all the exogenous variables (except z), 4) using $z_1^*,\ldots,z_T^*$, and then 5) solving the model for periods 1 through T. These are the values referred to below as 'solution values'. Note that these values are not the Fed's expected values after it has made its decision each period. The Fed's expected values are based on the use of the exogenous-variable equations and zero values of the error terms, whereas the solution values are based on the actual values of the exogenous variables (except the policy instrument) and error terms. The results of minimizing the expected value of each of the five loss functions for each of the five models are presented in Table 1 later in the article.

4. Two small structural models

Money instrument – Model S(m)

Model S(m) consists of a simple aggregate demand equation,[6] a Phillips curve, a variant of Okun's law, and the identity relating x_t, y_t, and p_t:

$$x_t = \alpha_1 + \alpha_2 x_{t-1} + \alpha_3 m_{t-1} + \varepsilon_{1t}, \tag{5}$$

$$p_t = \beta_1 + \beta_2 p_{t-1} + \beta_3 p_{t-2} + \beta_4 p_{t-3} + \beta_5 U_t + \varepsilon_{2t}, \tag{6}$$

$$U_t - U_{t-1} = \gamma_1 + \gamma_2 U_{t-1} + \gamma_3 U_{t-2} + \gamma_4 y_t + \gamma_5 y_{t-1} + \varepsilon_{3t}, \tag{7}$$

$$x_t = y_t + p_t. \tag{8}$$

[6] This is the nominal income equation that McCallum (1988) used in his work.

The estimated system, using the sample period 1960:1–1993:2, is shown below. Parameter estimates were obtained by the method of iterated three-stage least squares (3SLS) with the predetermined variables x_{t-1}, m_{t-1}, p_{t-1}, p_{t-2}, p_{t-3}, U_{t-1}, and U_{t-2} as instrumental variables for the second and third equations. (Absolute values of t-statistics are given in parentheses.)

$$\underset{(3.65)}{x_t = 3.00} + \underset{(2.38)}{0.190x_{t-1}} + \underset{(4.55)}{0.425m_{t-1}}, \tag{9}$$

$$SE = 3.62, \qquad DW = 2.06,$$

$$p_t = \underset{(3.85)}{1.65} + \underset{(5.94)}{0.479p_{t-1}} + \underset{(1.85)}{0.165p_{t-2}} + \underset{(3.93)}{0.321p_{t-3}} - \underset{(3.37)}{0.243U_t}, \tag{10}$$

$$SE = 1.19, \qquad DW = 2.00,$$

$$U_t - U_{t-1} = \underset{(5.12)}{0.355} + \underset{(4.68)}{0.335U_{t-1}} - \underset{(4.94)}{0.362U_{t-2}} - \underset{(4.11)}{0.044y_t} - \underset{(2.67)}{0.015y_{t-1}}, \tag{11}$$

$$SE = 0.19, \qquad DW = 2.06.$$

Several properties of this system should be noted. First, the estimated system is dynamically stable in the variables x, p, and U, with characteristic roots $0.9581 \pm 0.1072i$, $-0.2514 \pm 0.3139i$, 0.3779, and 0.1896. Second, the data do not reject the accelerationist form of the Phillips curve ($\beta_2 + \beta_3 + \beta_4 = 1$), with a χ^2 value of 0.59 (p-value $= 0.44$), but do reject the growth-rate form of Okun's law ($\gamma_4 = -\gamma_5$), with a χ^2 value of 5.97 (p-value $= 0.01$).

The mean of the vector ($x\ p\ U$) as a function of the mean of m is given by

$$\begin{pmatrix} \mu_x \\ \mu_p \\ \mu_U \end{pmatrix} = \begin{pmatrix} 3.71 \\ 0.76 \\ 6.70 \end{pmatrix} + \begin{pmatrix} 0.52 \\ 0.49 \\ -0.07 \end{pmatrix} \mu_m, \tag{12}$$

where μ denotes the mean. It follows from this that a one percentage point increase in the growth rate of the money supply will in the long run increase the growth rate of real output by 0.03 percentage points ($0.52 - 0.49$) and decrease the unemployment rate by 0.07 percentage points. The model thus exhibits modest monetary nonneutrality in the long run.

As for all the models, the results for Model S(m) are based on the solution of 126 deterministic optimal control problems for each of the five loss functions. The value of k used for the optimizations for the four small models was 16 quarters. The value of α in Eqs. (3) and (4) that was used for Model S(m) was 0.015. All the calculations for the four small models were done using the RATS program.

Interest rate instrument – Model S(R)

Model S(R) is the same as Model S(m) except that Eq. (5) is replaced by

$$y_t = \alpha_1 + \alpha_2 y_{t-1} + \alpha_3 R_{t-1} + \varepsilon_{1t}. \tag{13}$$

The same procedures were used for Model S(R) as were used for Model S(m). The predetermined variables y_{t-1}, R_{t-1}, p_{t-1}, p_{t-2}, p_{t-3}, U_{t-1}, and U_{t-2} were used as instrumental variables for Eqs. (6) and (7). The estimated system, using the sample period 1954:1–1993:2, is

$$\begin{aligned} y_t = \underset{(5.36)}{3.77} + \underset{(3.85)}{0.283}\,y_{t-1} - \underset{(2.87)}{0.288}\,R_{t-1},\\ SE = 3.83, \qquad DW = 2.09, \end{aligned} \tag{14}$$

$$\begin{aligned} p_t = \underset{(3.30)}{1.45} + \underset{(5.72)}{0.429}\,p_{t-1} + \underset{(2.50)}{0.200}\,p_{t-2} + \underset{(3.71)}{0.282}\,p_{t-3} - \underset{(2.30)}{0.174}\,U_t,\\ SE = 1.34, \qquad DW = 1.84, \end{aligned} \tag{15}$$

$$\begin{aligned} U_t - U_{t-1} = \underset{(7.08)}{0.495} + \underset{(5.50)}{0.332}\,U_{t-1} - \underset{(5.80)}{0.368}\,U_{t-2} - \underset{(6.76)}{0.084}\,y_t - \underset{(1.09)}{0.007}\,y_{t-1},\\ SE = 0.25, \qquad DW = 2.18. \end{aligned} \tag{16}$$

The characteristic roots are 0.9510, 0.9416, $-0.2612 \pm 0.4783i$, 0.3910, and 0.2834, which indicate that the model is stable. The accelerationist form of the Phillips curve is rejected at the 0.10 but not the 0.05 level, with a χ^2 value of 3.41 (p-value = 0.06). The growth-rate form of Okun's law is rejected, with a χ^2 value of 7.45 (p-value = 0.01). The mean of the vector (y p U) as a function of the mean of R is given by

$$\begin{pmatrix} \mu_y \\ \mu_p \\ \mu_U \end{pmatrix} = \begin{pmatrix} 5.26 \\ 15.56 \\ 0.38 \end{pmatrix} + \begin{pmatrix} -0.40 \\ -2.01 \\ 1.03 \end{pmatrix} \mu_R\,. \tag{17}$$

The value of α used was 0.250.

5. Two VAR models

Money instrument – Model V(m)

Model V(m) consists of four equations, one each for x, p, U, and R, where the right-hand-side variables consist of a constant and three lags each of x, p, U, R, and m. The same sample periods and procedures were used here as were used above except that the estimation technique was ordinary least squares instead of 3SLS. The value of α used was 0.015.

Interest rate instrument – Model V(R)

Model V(R) is the same as Model V(m) except that the equation for R is replaced by an equation for m. The right-hand-side variables are the same, and the same sample periods and procedures were used. The value of α used was 0.250.

6. A large structural model – Model L(R)

Model L(R) is presented in Fair (1994). It is quarterly and consists of 30 stochastic equations and 101 identities. It was estimated by two-stage least squares over the 1954:1–1993:2 period. The following is a brief discussion of the properties of the model that are relevant for present purposes.[7] Remember that the model is based on the assumption that expectations are not rational and the assumption that the parameters of the structural equations do not change when policy rules change.

There are six sectors in the model: household, firm, financial, foreign, state and local government, and federal government. All the flows of funds among these sectors and all balance-sheet constraints are accounted for in the model. This is done by linking the national income and product accounts and the flow of funds accounts. The sum of the savings across the six sectors is zero; some sector's expense is some other sector's revenue. Also, one sector's financial asset is some other sector's financial liability. Accounting for all flows of funds means that there is an explicit Fed open market operations variable in the model, the amount of federal government securities outstanding. This is the main 'tool' of the Fed. The other two tools are the reserve requirement rate and the discount rate.

The basic version of the model includes an interest rate reaction function of the Fed, which is estimated. This is the monetary policy rule that the Fed is estimated to have followed over the period. The three-month Treasury bill rate is on the left-hand side of this equation and variables that are assumed to affect Fed behavior are on the right-hand side. In this version of the model the money supply is endogenous; the Fed through open market operations each quarter achieves the bill rate implied by this equation.

Two term structure equations in the model link the bill rate to two long-term interest rates, a bond rate and a mortgage rate, where the long rates are a function of the current and lagged bill rates. The bill rate thus affects long-term rates through these term structure equations.

There are four household expenditure equations in the model, explaining 1) consumption of services, 2) consumption of nondurables, 3) consumption of durables, and 4) residential investment. There is also an import equation. Interest rates appear as explanatory variables in these equations – the bill rate in the services equation and the mortgage rate in the others. (The bill rate is taken as a proxy for short-term rates in general, and the mortgage rate is taken as a proxy for long-term rates in general.) The coefficient estimates of the interest rate

[7] A detailed description of the model and of the tests that were performed on it is in Fair (1994).

variables are all negative (and significant except in the import equation). There is also a nonresidential fixed investment equation in the model, and the bond rate appears in this equation with a negative and significant coefficient estimate. Therefore, through this part of the specification a decrease in interest rates increases household expenditures and nonresidential fixed investment and increases imports. The increase in imports is not large enough to offset the other increases, so that through this part of the specification there is a net increase in the demand for domestic goods in the model when interest rates fall.

The bond rate also appears in a stock price equation, where a decrease in the bond rate has a positive effect on stock prices. The value of stocks is part of household wealth, and household wealth (lagged once) appears as an explanatory variable in two of the four household expenditure equations with positive estimated coefficients. A decrease in interest rates thus has a positive effect on household expenditures through stock prices.

Disposable income appears in the household expenditure equations and in the import equation. The net effect of an increase in disposable income on the demand for domestic goods is positive – the positive effect from the expenditure equations outweighs the negative effect from the import equation. Since interest receipts of households are part of income, a decrease in interest rates (and thus interest receipts) has a negative effect on income and thus on demand through this part of the specification. This negative effect is not large enough to offset completely the positive effects discussed above, and the total effect on output from a fall in interest rates is positive.

The following is a summary of how a decrease in the bill rate affects the economy according to the model:

1. Long-term rates fall over time through the term structure equations.
2. Stock prices rise, which is a rise in household wealth.
3. Interest payments of the federal government fall, which leads to a fall in interest receipts of the household sector and thus, other things being equal, household income.
4. The fall in interest rates and rise in wealth have a positive effect on expenditures on domestic goods, and the fall in income from the fall in interest receipts has a negative effect. The net effect is positive.

The price level is a nonlinear function of, among other things, the difference between potential output and actual output, a measure of demand pressure. As output approaches 4 percent above potential, the price level approaches infinity. The price of imports also has an important effect on the domestic price level.

The unemployment rate is determined as one minus the ratio of employment to the labor force. There are three labor force participation equations for the household sector, one for each of three age–sex groups, that determine the labor force, and there is a demand for labor equation for the firm sector that determines most of employment.

When the optimal control problems were solved, the estimated interest rate reaction function of the Fed was dropped from the model, and the interest rate determined by this equation (the three-month Treasury bill rate) was taken to be the control variable, i.e., the policy instrument. This means that the open market operations variable of the Fed, the amount of government securities outstanding, is endogenous; its value each quarter is whatever is needed to achieve the desired value of the bill rate.

The solution of the 126 deterministic optimal control problems for each loss function took about 12 hours on a Pentium 90 computer for Model L(R). To save computer time, the value of k used for the optimizations was taken to be 8 quarters instead of 16 quarters used for the small models. Some experimentation suggested that a value of k of 8 was large enough for the addition of more quarters to have a fairly small effect on the optimal value of R for the first quarter. The value of α used was 0.500. All the calculations for the model were done using the Fair–Parke program.

7. Discussion of the results

The results are presented in Table 1. Consider first the values in the 'Actual' row versus the values obtained by minimizing the expected value of the true loss function, which are in the first row for each model. Remember that the true loss function weights inflation and unemployment equally. When the expected value of this loss function is minimized for Model L(R), the unemployment loss (Q_U) falls from 1.59 to 0.92, but the inflation loss (Q_p) actually rises (from 2.41 to 2.62). The other four models have both inflation loss and unemployment loss falling. Model S(R) is closest to Model L(R) in that the fall in the inflation loss is small relative to the fall in the unemployment loss.

Consider next how the policy of targeting only nominal growth (x) does compared to the policy of minimizing the expected value of the true loss function.[8] For Model L(R) the increase in $Q_{p,U}$ when only nominal growth is targeted is 16.3 percent (from 1.96 to 2.28). The other percentage increases are 12.6 for Model S(m), 3.8 for Model S(R), 10.8 for Model V(m), and 13.0 for Model V(R). All models except S(R) thus show a noticeable increase in loss when only nominal growth is targeted, with the largest increase for L(R).

[8] One aspect of our results of targeting only nominal growth can be compared to a result of Feldstein and Stock (1993). Feldstein and Stock dealt with a VAR model similar to Model V(m) (although the unemployment rate is not part of their model), and they obtained a reduction of a little over 20 percent in the variability of nominal growth when they minimized the expected value of a loss function that targeted only nominal growth. For Model V(m) the reduction in Q_x when only nominal growth is targeted is 13.8 percent in Table 1 (from 3.91 to 3.37).

The policy of targeting only inflation (p) works well for the four small models in that for each model the value of $Q_{p,U}$ when only p is targeted is fairly close to the value when the true loss function is minimized. This is not true, however, for L(R). Conversely, the policy of targeting only unemployment (U) works well for L(R) in that the value of $Q_{p,U}$ when only U is targeted is fairly close to the value when the true loss function is minimized, but this is not true for the four small models. The policy of targeting only real growth (y) does not work well for any model.

There is one odd result in Table 1 that should be noted before proceeding further. When only U is targeted using Model V(m), Q_U actually increases (from 1.59 to 1.65). In this model m has little affect on U, and when only U is targeted, the value of m tends to go close to one of the bounds implied by the penalty terms in the loss function and to stay there for many periods. The penalty terms are in effect dominating the results.[9] The fact that m has little affect on U in

Table 1
Results for the five models and five loss functions

	$Q_{p,U}$	Q_p	Q_U	Q_x	Q_y	C_m	C_R	α^*
Actual	2.04	2.41	1.59	3.91	3.93	2.72	0.89	—
Model S(m)								
True loss funct.	1.67	2.03	1.21	4.22	4.49	2.70	—	0.015
Target only p	1.76	1.85	1.66	4.66	5.35	3.27	—	0.015
Target only U	2.36	3.27	0.67	4.78	3.95	2.32	—	0.015
Target only x	1.88	2.06	1.67	3.50	3.62	0.34	—	0.015
Target only y	2.68	3.68	0.91	4.64	3.39	1.07	—	0.015
Model S(R)								
True loss funct.	1.83	2.34	1.09	4.07	4.02	—	0.86	0.250
Target only p	1.84	2.10	1.53	3.94	4.22	—	1.08	0.250
Target only U	2.03	2.74	0.88	4.34	3.86	—	0.77	0.250
Target only x	1.90	2.05	1.73	3.79	4.08	—	0.95	0.250
Target only y	2.59	3.31	1.57	4.66	3.65	—	0.25	0.250
Model V(m)								
True loss funct.	1.48	1.70	1.21	3.91	4.25	2.43	—	0.015
Target only p	1.49	1.67	1.29	4.12	4.59	2.92	—	0.015
Target only U	2.36	2.90	1.65	4.06	3.58	0.74	—	0.015
Target only x	1.64	1.85	1.41	3.37	3.64	2.53	—	0.015
Target only y	2.76	3.52	1.68	4.19	3.39	2.88	—	0.015

[9] When the penalty terms were removed from the objective function, the $S(m)$ optimal control solutions with U only and y only as the targets were explosive. All of the other small-model optimal control solutions were stable when the penalty terms were removed from the objective function, although the instrument exhibited extremely unrealistic variations in some cases.

Table 1 (continued)

	$Q_{p,U}$	Q_p	Q_U	Q_x	Q_y	C_m	C_R	α^*
Model V(R)								
True loss funct.	1.61	1.97	1.15	3.77	3.92	—	0.82	0.250
Target only p	1.71	1.88	1.53	3.92	4.35	—	1.10	0.250
Target only U	2.02	2.76	0.74	4.19	3.61	—	0.61	0.250
Target only x	1.82	1.85	1.78	3.40	3.69	—	0.75	0.250
Target only y	1.94	2.47	1.19	3.68	3.26	—	0.70	0.250
Model L(R)								
True loss funct.	1.96	2.62	0.92	4.08	3.68	—	0.54	0.500
Target only p	2.27	2.34	2.19	3.93	3.93	—	0.41	0.500
Target only U	2.03	2.72	0.92	4.32	3.74	—	0.75	0.500
Target only x	2.28	2.41	2.15	3.80	3.73	—	0.93	0.500
Target only y	2.32	2.87	1.60	4.32	3.67	—	0.82	0.500

$Q_q = \sqrt{(1/T)\sum_{t=1}^{T}(\hat{q}_t - q^*)^2}$, where $q = p, U, x$, or y.

$Q_{p,U} = \sqrt{(1/T)[0.5\sum_{t=1}^{T}(\hat{p}_t - p^*)^2 + 0.5\sum_{t=1}^{T}(\hat{U}_t - U^*)^2]}$.

(The hatted values are the solution values using the optimal values of the policy instrument, the actual values of the exogenous variables, and the historical residuals.)

$C_z = \sqrt{(1/T)\sum_{t=1}^{T} D_{zt}}$, where $z = m$ or R.

α^* = value of α used in Eq. (3) or (4).

x = 400×change in the log of nominal GDP.

y = 400×change in the log of real GDP.

p = 400×change in the log of the GDP price index.

U = level of the unemployment rate (percentage points).

m = 400×change in the log of the money supply.

R = level of the short-term interest rate (percentage points).

The values in the 'Actual' row are obtained using the actual values of q instead of the solution values.

Control period is 1962:1–1993:2.

Model V(m) can also be seen by comparing the results of minimizing the true loss function to those of targeting only p. These two sets of results are very close, which means that when the true loss function is being minimized, the only action is in controlling p, not U. Because of this feature of Model V(m), little weight should be attached to the U results for this model.

Continuing with the results in Table 1, it is possible to use them to examine the trade-off between unemployment variability (Q_U) and inflation variability (Q_p). This trade-off is what Taylor (1979) calls a 'second-order' Phillips curve. Table 1 provides three points on this curve for each model. The weights on U and p are 0.5 and 0.5 for the true loss function, 1.0 and 0.0 when only U is targeted, and 0.0 and 1.0 when only p is targeted. Table 2 presents for each model except V(m) the changes in Q_U and Q_p in going from zero weight on

Table 2
Changes in Q_U and Q_p from Table 1
$\Delta Q_U = Q_U$ (target only U) $- Q_U$ (target only p), $\Delta Q_p = Q_p$ (target only U) $- Q_p$ (target only p).

Model	ΔQ_U	ΔQ_p
Model S(m)	−0.99	1.42
Model S(R)	−0.65	0.64
Model V(R)	−0.79	0.88
Model L(R)	−1.27	0.38

inflation (targeting only U) to zero weight on unemployment (targeting only p). (Model V(m) is excluded for the reasons discussed above.) For Model L(R) there is a 1.27 drop in unemployment variability at a cost of a 0.38 rise in inflation variability, a trade-off of 3.3 to 1 (1.27/0.38). The smallest trade-off is for Model S(m) of 0.7 to 1 (0.99/1.42). The trade-offs for the other two models are about 1 to 1.

As mentioned in the Introduction, we are inclined to put more weight on the results for Model L(R) because it is likely to be more accurate. One of the main differences between L(R) and the small models concerns the trade-off between unemployment variability and inflation variability. The trade-off is much higher for L(R) – 3.3 to 1 – than for the small models – about 1 to 1. Two other differences are that targeting inflation works fairly well for the small models, but not for L(R), and targeting unemployment works fairly well for L(R), but not for the small models. Two similarities are that targeting nominal growth does not work very well (except for Model S(R)) and targeting real growth does not work well. This first similarity suggests that the widely held view in the literature that targeting nominal growth is a good idea may not be right.

Further details

The figures in Table 1 hide many of the details of the results, and it is of interest to consider a few of these details. Fig. 1 shows plots of the actual and solution (hatted) values of p and U for Model S(R), along with the actual and optimal values of R. The optimal and solution values are from the results of minimizing the expected value of the true loss function. Fig. 2 shows the same plots for Model L(R). For both models the optima correspond to higher unemployment rates in the 1960s and lower unemployment rates in the 1980s. To achieve this, the bill rates are higher in the 1960s and generally lower in the 1980s. For L(R) there is little difference between the optimal and actual inflation rates, whereas for S(R) the optimal inflation rates are generally higher than the actual rates in the 1980s. These plots are, of course, consistent with the result in Table 1 that unemployment variability is lowered more than inflation variability

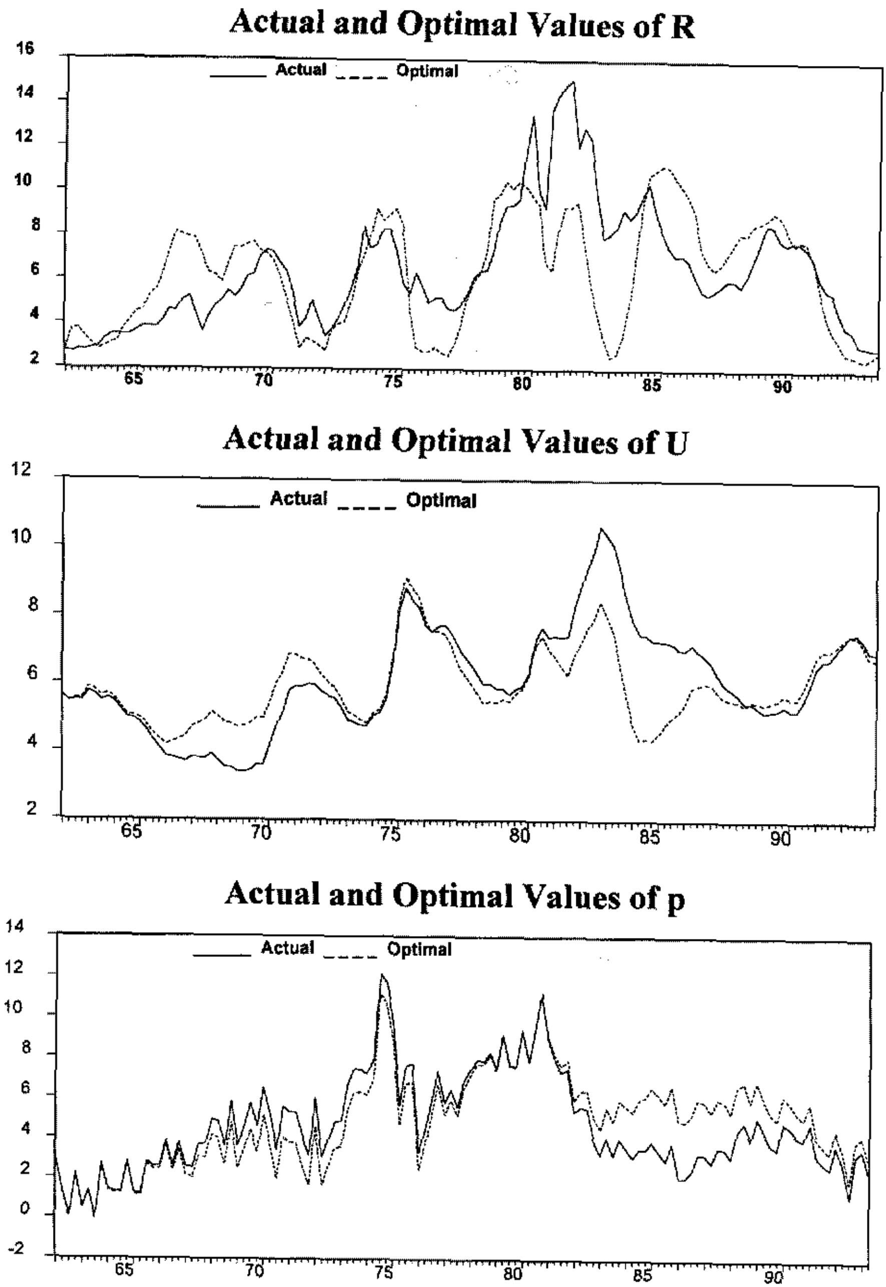

Fig. 1. Model S(R).

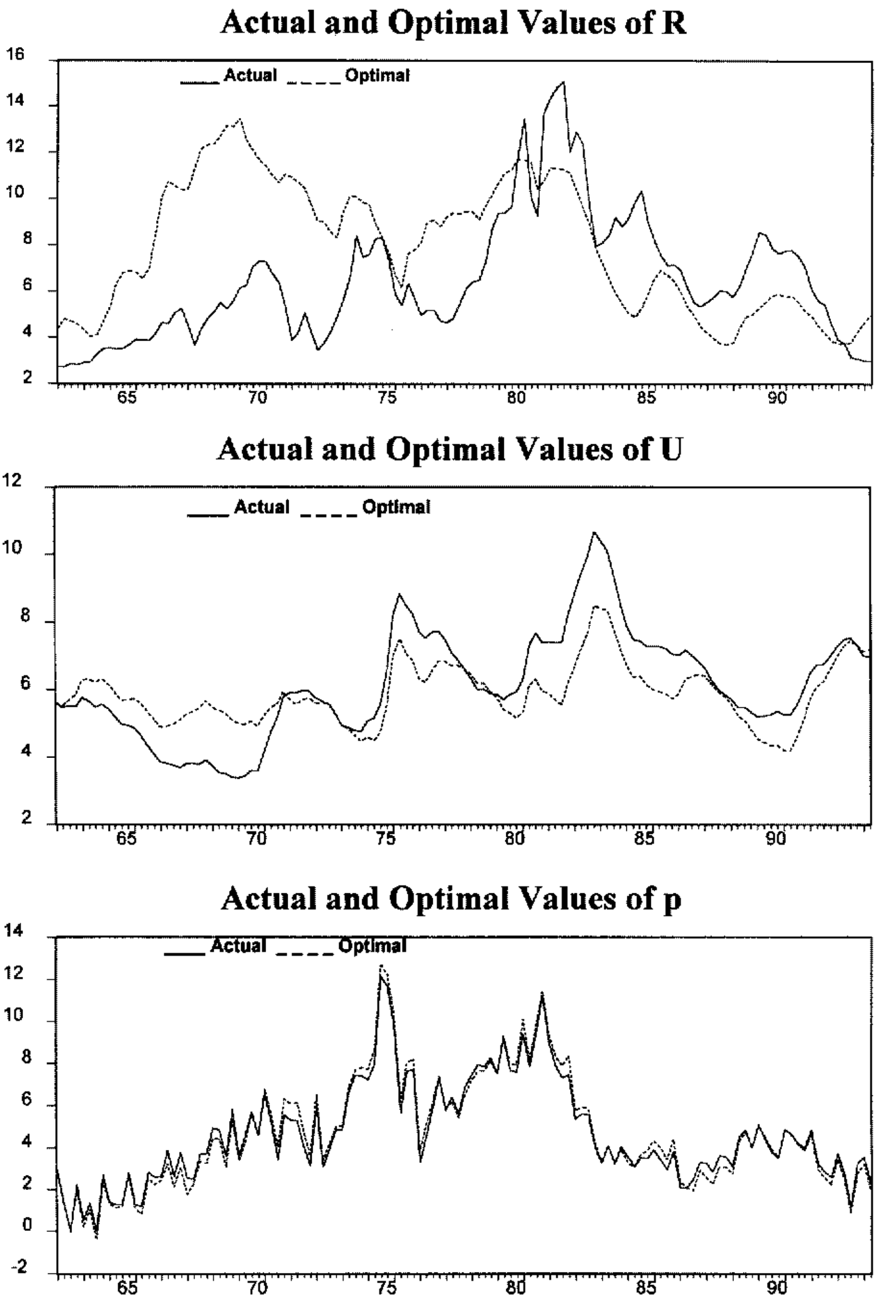

Fig. 2. Model L(R).

when the expected value of the true loss function is minimized for Models S(R) and L(R).

Another way to look at optimal control results is to examine feedback equations. For a linear model like S(R) and a quadratic objective function, the coefficients in the feedback equation can be obtained (using the solution values as data) by regressing the value of the policy instrument (R) on its lagged value and on y_{t-1}, p_{t-1}, p_{t-2}, p_{t-3}, U_{t-1}, U_{t-2}, and a constant. These are the variables in the optimal feedback equation. If bounds are imposed on the instrument, as is done here, the optimal feedback equation is no longer linear, and so the regression is only approximate.

The regression using the solution values for Model S(R) for the minimization of the expected value of the true loss function is

$$R_t = 5.20 + 0.077 y_{t-1} + 0.234 p_{t-1} + 0.098 p_{t-2} + 0.048 p_{t-3}$$
$$-1.605 U_{t-1} + 0.747 U_{t-2} + 0.651 R_{t-1}. \tag{18}$$

The fit of this regression is quite good ($R^2 = 0.991$), which suggests that the approximation is good. The coefficients on the y and p values show that the interest rate responds positively to output growth and inflation. The unemployment terms can be written $-0.858U_{t-1} - 0.747(U_{t-1} - U_{t-2})$, and so the unemployment coefficients show that the interest rate responds negatively to both the level and change of the unemployment rate. When Eq. (18) was used in place of solving the 126 deterministic control problems, the value of $Q_{p,U}$ was the same to two decimal places as the value in Table 1 (1.83), and so the equation is obviously a good approximation to the true feedback equation.

What about a feedback equation for Model L(R)? The model is nonlinear, and so even ignoring the penalty terms in the loss function, the feedback equation is nonlinear. Also, there are over a hundred predetermined variables in the model, and so the feedback equation is huge. It may be, however, that even though the true feedback equation is nonlinear and large, a few right-hand-side variables provide a good approximation. To check this, a number of regressions were run (using the solution values as data) with the policy instrument (R) on the left-hand side and various variables on the right-hand side. The aim was to find variables that contributed significantly to the fit. The following, quite simple, equation turned out to give a good fit:

$$R_t = 1.40 + 1.353 R_{t-1} - 0.387 R_{t-2} + 0.0544(400Y_{t-1})$$
$$-0.0549(400Y_{t-2}), \tag{19}$$

where Y is the log of real output. This regression has an R^2 of 0.980. This equation roughly says that $\Delta R_t = 1.4 + 0.35\Delta R_{t-1} + 0.055 y_{t-1}$, where y_{t-1}, the growth rate of real output in period $t-1$, equals $400(Y_{t-1} - Y_{t-2})$. This rule simply says that change in the policy instrument depends positively on its previous change and on the lagged real growth rate. Other variables, including lagged values of P and U, that were added to Eq. (19) tended not to be significant. It may be with more diligent searching that other variables could be found for inclusion in (19), but for present purposes we have stopped with (19).

When Eq. (19) was used in place of solving the 126 deterministic control problems for model L(R), the value of $Q_{p,U}$ was 2.03, which compares to the optimal value of 1.96 in Table 1. The loss is thus 3.6 percent higher using (19). This suggests that (19) is not a bad approximation, and an argument might be made for using it on grounds of simplicity. On the other hand, the extra computer cost in performing the complete optimization is trivial, and so the 3.6 percent increase in loss could not be justified on computational grounds. It is, however, interesting that the simple equation does as well as it does, and this might be an area for future work.

8. Stochastic simulation

The variability measures in Table 1 do not require stochastic simulation – they are based on the use of the historical shocks. Although for historical comparisons these are the most appropriate shocks to use, it is interesting to consider how alternative measures could be computed using stochastic simulation.

Consider the use of the above numerical procedure and Model L(R). There are 29 stochastic equations in L(R), and the length of the control period is 126 quarters. The results in a given row in Table 1 for L(R) are thus based on $126 \times 29 = 3654$ historical error terms. (After each of the 126 solutions of the deterministic control problem, 29 error terms are used.) A stochastic-simulation alternative to the use of the historical error terms is as follows. 1) Using the estimated distribution of the structural error terms, draw 29 error terms 126 times. 2) Go through the entire numerical procedure as above, using the drawn error terms instead of the historical error terms. This yields a value of Q_q for each variable q (as in Table 1). 3) Do steps 1) and 2) J times, where J may be around 100. This gives J values of Q_q, from which its mean, variance, and other statistics can be computed. The means of Q_q would be alternatives to the values in Table 1, and they could be compared across loss functions.

What would correspond to the 'Actual' row in Table 1 when stochastic simulation was used? The most obvious possibility for Model L(R) would be to use the estimated interest rate reaction function discussed above as the monetary policy rule. This is the rule the Fed is estimated to have followed historically. For Model V(R) it would be the estimated equation for R. Stochastic simulation is easy to do numerically in this case because the rule replaces the need to compute the optimal values of R. Each repetition involves solving the model with the rule for the particular set of error terms drawn.[10] The mean value of Q_q from this exercise could then be compared to the other mean values.

Since it takes about 12 hours of computer time to compute one value of Q_q for Model L(R), computing, say, 100 values is not yet practical. One could, however, experiment with smaller models at the present time, and computer chips are getting faster. In the future it will be interesting to see if the conclusions reached from Table 1 are sensitive to the use of stochastic simulation.

Finally, one could also draw coefficients as well as error terms for the stochastic simulations. For each repetition (i.e., each draw of the 126×29 error terms) one would draw a set of coefficients from an estimated distribution. The results with and without drawing coefficients could then be compared. This would be a way of examining how robust the conclusions are to alternative versions of the model, the alternative versions in this case being alternative sets of coefficients.

9. Conclusion

The conclusions from the results in Table 1 have been presented in Section 7, and they will not be repeated here. The numerical procedure presented in Section 3 is general enough to handle a variety of loss functions and models. The loss functions need not even be quadratic. Given that a number of the conclusions that hold for Model L(R) do not hold for the small models, it would be of interest in future work to try other models, particularly other large structural models.

The results in Fair (1994, pp. 320–328) suggest that monetary policy is becoming less effective for a given change in interest rates because of the growing size of the federal government debt. Most of the debt is financed by the household sector, and the larger is the debt, the larger is the change in interest revenue of the household sector for a given change in interest rates. As noted in the discussion of Model L(R) in Section 6, a fall in interest revenue from a fall in interest rates has, other things being equal, a negative effect on household expenditures, and this offsets part of the other (positive) effects on expenditures of a decrease in interest rates. The size of this offset is getting larger over time because of the growing size of the government debt, and so on net monetary policy is becoming less effective in the sense that there is a smaller change in output for a given change in the interest rate. The less effective monetary policy becomes in this sense, the less able will the Fed be to minimize any loss function that penalizes changes in the interest rate. The Fed may thus do less well in the future.

[10] One would not draw error terms for the rule, although the historical error terms should probably be added to the rule and taken as fixed. These historical error terms are the part of Fed policy that the estimated rule does not capture, which should probably be added in. The implicit assumption that would be made in doing this is that the Fed used a different constant term in the rule each period.

Another issue that can be analyzed using the numerical procedure in Section 3 is targeting levels versus targeting growth rates in the loss function. (As noted above, the loss function can be quite general.) One loss function might be squared deviations of real output from some measure of potential output, and another might be squared deviations of the price level from a target price level that grows at, say, one or two percent per year. In the process of writing this paper, some early experimentation was done using these types of loss functions, but for present purposes it was decided to stay with the loss functions in terms of growth rates. An added complication when using loss functions in levels is pushing the economy into situations that are substantially different from what existed historically. In a model like L(R), this can have large cumulative effects and change both the real stock and financial stock variables substantially. It is always dangerous to push a model too far from historical experience, and using loss functions in levels sometimes does this. At any rate, further experimentation with alternative loss functions would be of interest. At a minimum, it might be of interest to experiment with loss functions like (3) and (4) in which the target values for p and U were not necessarily constant throughout the whole period.

References

Chow, G.C., 1981, Econometric analysis by control methods (Wiley, New York, NY).

Clark, T.E., 1994, Nominal GDP targeting rules: Can they stabilize the economy?, Federal Reserve Bank of Kansas City Economic Review 3, 11–25.

Croushore, D. and T. Stark, 1994, Evaluating McCallum's rule for monetary policy, Federal Reserve Bank of Philadelphia mimeo.

Fair, R.C., 1974, On the solution of optimal control problems as maximization problems, Annals of Economic and Social Measurement 3, 135–154.

Fair, R.C., 1994, Testing macroeconometric models (Harvard University Press, Cambridge, MA).

Feldstein, M., and J.H. Stock, 1993, The use of a monetary aggregate to target nominal GDP, NBER working paper no. 4304.

Hall, R.E., and N.G. Mankiw, 1993, Nominal income targeting, NBER working paper no. 4439.

Judd, J.P. and B. Motley, 1993, Using a nominal GDP rule to guide discretionary monetary policy, Federal Reserve Bank of San Francisco Economic Review 3, 3–11.

McCallum, B.T., 1988, Robustness properties of a rule for monetary policy, Carnegie–Rochester Conference Series on Public Policy 29, 173–204.

Taylor, J.B., 1979, Estimation and control of a macroeconomic model with rational expectations, Econometrica 47, 1267–1286.

Taylor, J.B., 1985, What would nominal GDP targeting do to the business cycle?, Carnegie–Rochester Conference Series on Public Policy 22, 61–84.

Thornton, S.R., 1995, Adaptive monetary policy rules utilizing either MZM or M2 and the effects of positive bias in money control errors, presented at the AEA Meetings, Washington, DC, January 1995.

Tobin, J., 1980, Stabilization policy ten years after, Brookings Papers on Economic Activity 1, 19–72.