10 Optimal Control Analysis

10.1 Introduction

Optimal control techniques have a number of potentially important uses in macroeconometrics. Solving optimal control problems for a particular model may yield insights about the model that one would not pick up from multiplier calculations. Depending on the objective function, the solutions of optimal control problems are sometimes extreme in that they result in the predicted values being considerably away from the historical values, and this sometimes conveys new information about the properties of the model. Optimal control techniques can also be used to evaluate past policies in the light of particular objective functions. The techniques may also be useful in the long run in helping to make actual policy decisions, depending on how good an approximation to the structure of the economy models eventually become.

10.2 A Method for Solving Optimal Control Problems

10.2.1 The Method

Optimal control problems have historically been formulated in continuous time and have been looked upon as problems in choosing *functions* of time to maximize an objective function. This is particularly true in the engineering literature. Fairly advanced mathematical techniques are required to solve these problems. For discrete time models, however, which include virtually all macroeconometric models, optimal control problems can also be looked upon as problems in choosing *variables* to maximize an objective function. The number of variables to be determined is equal to the number of control variables times the number of time periods chosen for the problem. From this perspective, optimal control problems are straightforward maximization problems, and one can attempt to solve them using algorithms like the DFP algorithm discussed in Section 2.4.

Let the model be represented by (6.1), which is repeated here:

(6.1) $$f_i(y_t, x_t, \alpha_i) = u_{it}, \qquad i = 1, \ldots, n.$$

The variables in the x_t vector include both exogenous and lagged endogenous variables. Among the exogenous variables are variables that are under the control of the government and variables that are not. It will be useful to redefine x_t to include only noncontrolled exogenous variables. Let z_t denote the vector of control variables, and let q_{t-1} denote the vector of all lagged endogenous variables in the model, even variables lagged more than one period. Rewrite (6.1) to include these changes:

$$(10.1) \qquad f_i(y_t, q_{t-1}, x_t, z_t, \alpha_i) = u_{it}, \qquad i = 1, \ldots, n.$$

In the following discussion the coefficients α_i are assumed to be known with certainty.

The first step in setting up an optimal control problem is to postulate an objective function. Assume that the period of interest is $t = 1, \ldots, T$. A general specification of the objective function is

$$(10.2) \qquad W = h(y_1, \ldots, y_T, x_1, \ldots, x_T, z_1, \ldots, z_T),$$

where W, a scalar, is the value of the objective function corresponding to values of y_t, x_t, and z_t $(t = 1, \ldots, T)$. In most applications the objective function is assumed to be additive across time, which means that (10.2) can be written

$$(10.3) \qquad W = \sum_{t=1}^{T} h_t(y_t, x_t, z_t),$$

where $h_t(y_t, x_t, z_t)$ is the value of the objective function for period t. The function h has a t subscript to note the fact that it may vary over time. This will be true, for example, if future periods are discounted.

The optimal control problem is to choose values of $z_1, \ldots, z_T$ so as to maximize the expected value of W in (10.2) subject to the model (10.1). Consider first the deterministic case where the error terms in (10.1) are all zero. Assume that z_t is of dimension k, so that there are kT control values to determine, and let z be the kT-component vector denoting these values: $z = (z_1, \ldots, z_T)$. For each value of z one can compute a value of W by first solving the model (10.1) for $y_1, \ldots, y_T$ and then using these values along with the values for $x_1, \ldots, x_T$ and z to compute W in (10.2). Stated this way, the optimal control problem is a problem in choosing variables (the elements of z) to maximize an *unconstrained* nonlinear function. By substitution, the constrained maximization problem is transformed into the problem of maximizing an unconstrained function of the control variables:

$$(10.4) \qquad W = \phi(z),$$

where ϕ stands for the mapping $z \rightarrow z, y_1, \ldots, y_T, x_1, \ldots, x_T \rightarrow W$. For nonlinear models it is generally not possible to express y_t explicitly in terms of z_t and x_t, which means that it is generally not possible to write W in (10.2) explicitly as a function of z and $x_1, \ldots, x_T$. Nevertheless, given values for $x_1, \ldots, x_T$, values of W can be obtained numerically for different values of z.

Given this setup, the problem can be turned over to a nonlinear maximization algorithm like DFP. For each iteration, the derivatives of ϕ with respect to the elements of z, which are needed by the algorithm, can be computed numerically. Each iteration will thus require kT function evaluations for the derivatives plus a few more for the line search. Each function evaluation requires one solution (dynamic simulation) of the model for T periods plus the computation of W in (10.2) after $y_1, \ldots, y_T$ are determined.

There is one important cost-saving feature regarding the method that should be noted. Assume that there are two control variables and that the length of the period is 30. The number of unknowns is thus 60, and therefore 60 function evaluations will have to be done per iteration to get the numerical first derivatives. In perturbing the control values to get the derivatives, one should start from the end of the control period and work backward. When the control values for period 30 are perturbed, the solution of the model for periods 1 through 29 remains unchanged from the base solution, so these calculations can be skipped. The model only needs to be resolved for period 30. Similarly, when the control values for period 29 are perturbed, the model only needs to be resolved for periods 29 and 30, and so on. This cuts the cost of computing the derivatives roughly in half.

10.2.2 Stochastic Simulation Option

Consider now the stochastic case where the error terms in (10.1) are not zero. It is possible to convert this case into the deterministic case by simply setting the error terms to their expected values (usually zero). The problem can then be solved as above. In the nonlinear case this does not lead to the exact answer because the values of W that are computed numerically in the process of solving the problem are not the expected values. In order to compute the expected values correctly, stochastic simulation would have to be done. In this case each function evaluation (that is, each evaluation of the expected value of W for a given value of z) would consist of the following.

1. A set of values of the u_{it} error terms in (10.1) would be drawn from an estimated distribution.

2. Given the values of the error terms, the model would be solved for $y_1, \ldots, y_T$ and the values of W corresponding to this solution would be computed from (10.2). Let $\tilde{W}^j$ denote this value.
3. Steps 1 and 2 would be repeated J times, where J is the number of trials.
4. Given the J values of $\tilde{W}^j$ $(j = 1, \ldots, J)$, the expected value of W would be taken to be the mean of these values.

$$\tilde{\tilde{W}} = \frac{1}{J} \sum_{j=1}^{J} \tilde{W}^j. \tag{10.4}$$

This procedure increases the cost of solving the control problem by roughly a factor of J, since the maximization algorithm spends most of its time doing function evaluations. It is probably not worth the extra cost for most applications. It was seen in Chapter 7 that the bias in predicting the endogenous variables that results from using deterministic rather than stochastic simulation seems to be small for most models, and thus the bias in computing the expected value of W is also likely to be small. At any rate, the stochastic simulation option is always open if computer time is no constraint.

10.2.3 Comparison of the Method to Other Procedures

There are two main advantages of the method just described. One is that it can handle very general objective functions; the objective function need not be quadratic and need not even be additive across time. The second is that the method is extremely easy to use. Assuming that a program is available for solving the model, which is almost always the case, all that needs to be supplied is a subroutine that computes W in (10.2) for a given set of RHS values. In a program that is structured like the Fair-Parke program in Appendix C, which allows one to move automatically from estimation to solution, this is an important advantage. Given a subroutine that computes W, one can move automatically from estimation to solving control problems. There are thus virtually no extra setup costs involved in using the method.

The method described above is "open-loop." The alternative type of method is "closed-loop," where closed-loop feedback control equations are derived. A feedback control equation is one that relates the current value of a control variable to the lagged values of the endogenous variables. In the case of a linear model and a quadratic objective function, it is relatively easy to compute the feedback equations. (Chow 1975 is a good reference for this.) One of the advantages of obtaining feedback equations is that they can be used to compute the optimal control values for all future periods without having to solve any further problems. Given the realizations of the endogenous variables for a given period, the optimal control values for the next period can simply be computed from the feedback equations. For open-loop methods, on the other hand, a new optimization problem has to be solved after each period's realization. Consider, for example, the problem presented above, where optimal values for periods 1 through T were computed. If this solution were used in practice, the optimal values for period 1 would be used, but the values for periods 2 through T would not. The latter values are needed only to compute the period 1 values. After the realization in period 1, where in general the endogenous variable values will not equal the values that were expected at the time the control problem was solved, a new control problem would have to be solved to get the optimal values for period 2.

In the linear-quadratic case, open-loop methods with reoptimization after each realization and closed-loop methods lead to the same control values being used each period. This is the certainty equivalence theorem. In the general nonlinear case, analytic expressions for the feedback equations are not available, so there is no known closed-loop solution. An interesting question is whether the current open-loop method with stochastic simulation to eliminate the bias in computing the expected value of W and with reoptimization after each realization leads to the correct answer aside from errors introduced by the stochastic simulation procedure. The answer is no. Maximizing the expected value of W simultaneously with respect to $z_1, \ldots, z_T$ fails to account for the fact that the optimal strategy is sequential rather than simultaneous. (See Chow 1975, pp. 295–296, for a discussion of this.) This is a subtle point, and it is an open question whether it is important quantitatively.

Chow (1975, chap. 9) has proposed an alternative method for solving optimal control problems in the nonlinear case. He suggests obtaining a linear approximation to the model and a quadratic approximation to the objective function and then solving the resulting linear-quadratic problem by standard methods. One then iterates on the approximations. This method also does not lead to the correct answer, although for a different reason than in the case of the open-loop method. The linearization of the model must be around the solution path of the deterministic control problem (since the future values of the error terms are not known), and therefore the linearization is not quite right. The computed optimal values are thus not truly optimal. The method has the advantage that feedback equations are obtained, although this is not as much of an advantage as it might at first appear. Even given the feedback equations, one may want to reoptimize after the realization for a given period because the linearization will change. One will not get the same optimal values for the period using the old feedback equations as one would get by reoptimizing based on an updated linearization.

From a computational point of view, Chow's method is somewhat messy because of the linear approximations. These approximations require considerable storage space for the matrices, and it is not as easy to adjust for changes in the model because for each adjustment the linearization must also be adjusted. In addition, if the model is large, a large matrix must be inverted in calculating the optimal values. An advantage of the method over the open-loop method is that the computational costs only increase linearly in T, the length of the control period, whereas they increase roughly as the square of T for the open-loop method. (The cost for the open-loop method increases as the square of T because an increase in T increases both the number of control values to determine and the cost of solving the model for a given function evaluation.) There are thus likely to be some applications for which Chow's method is better and some for which the open-loop method is better. Whether one will end up dominating for most applications remains to be seen.

The discussion so far has been based on the assumption that the coefficients are known with certainty. The question of how to handle coefficient uncertainty in the nonlinear case is difficult, and no exact solutions are available. This issue will be not be explored here; the interested reader is referred to Chow (1976), who presents an approximate solution.

The discussion so far has also been based on the assumption that the model is not a rational expectations model. The solution of optimal control problems for rational expectations models is discussed in Section 11.5.

10.2.4 Steps that a Policymaker Would Follow

For purposes of the discussion in the next section, it will be useful to review the steps that a policymaker would follow if he or she were setting policies by solving control problems. Assume that a policy decision is to be made at the beginning of period 1 and that at this time data for period 0 and all prior periods are available. Given the model (10.1) and, say, a horizon of length T, the steps that could be followed are:

1. Estimate the coefficients of the model over the sample period ending in 0.
2. Form expectations of the exogenous variables (other than the control variables) for periods 1 through T.
3. Form expectations of the values of the error terms for periods 1 through T.

4. Decide on the objective function (10.2) to be maximized.
5. Using some maximization algorithm (like DFP), maximize (10.2) with respect to $z_1, z_2, \ldots, z_T$. Let $z_1^*, z_2^*, \ldots, z_T^*$ denote the optimum values.
6. Use z_1^* as the vector of policy values for period 1.

After the values for period 1 have been realized, steps 1–6 can be repeated for period 2. As noted in Section 10.2.3, the optimal value of z_2 that is computed at this time is not in general equal to z_2^* in step 5. The actual values of the endogenous variables for period 1 are in general different from what they were predicted to be, and therefore the initial conditions for the problem beginning in period 2 are different from what the solution at the beginning of period 1 implied that they would be. Also, the coefficient estimates will have changed because of the reestimation through period 1. The actual values of the exogenous variables for period 1 will in general be different from what they were expected to be, and the expectations for periods 2 and beyond are likely to have changed.

If stochastic simulation is used, step 3 is replaced by a step in which the distribution of the error terms is chosen. This distribution is then used in step 5 in the manner discussed in Section 10.2.2.

If Chow's procedure is used to solve the control problem, step 5 is replaced with this procedure. It is still necessary in this case to form expectations of the error terms for periods 1 through T (step 3), because this is needed for the linearization. Also, as noted in Section 10.2.3, steps 1 through 6 would be performed again after the values for period 1 have taken place because these values affect the linearization and thus the feedback equations. The different coefficient estimates and exogenous variable values will also affect the linearization.

The reestimation of the model in step 1 means that the coefficient estimates are always based on the latest available data. This does not mean, however, that by doing this one has accounted for coefficient uncertainty in solving the optimal control problem. Nothing in this procedure informs the method in step 5 that the coefficient estimates are to be reestimated in the future, and so this information is not taken into account.

10.3 Use of Optimal Control Analysis to Measure the Performance of Policymakers

It is common practice in political discussions to hold policymakers accountable for the state of the economy that existed during their time in power.

Policymakers are generally blamed for high unemployment, low real growth, and high inflation rates during their time in power and praised for the opposite. Although at first glance this may seem to be a reasonable way of evaluating the economic performances of policymakers, there are at least two serious problems with it. The first is that this kind of evaluation does not take into account possible differences in the degree of difficulty of controlling the economy in different periods. The economy may be more difficult to control at one time than another either because of more unfavorable values of the uncontrolled exogenous variables or because of a more unfavorable initial state of the economy (or both). The second problem with the evaluation is that it ignores the effects of a policymaker's actions on the state of the economy beyond its time in power. If, for example, a policymaker strongly stimulates the economy in the year of an election, in, say, the belief that this might improve its chances of staying in power, most of the inflationary effects of this policy might not be felt until after the election. Any evaluation of performance that was concerned only with the time before the election would not, of course, pick up these effects.

A measure of performance is proposed in this section that takes account of these problems. It is based on the solutions of optimal control problems. This performance measure requires that a welfare function be postulated and that the economy be represented by an econometric model. The welfare function must be additive across time. It will be convenient to take the objective function to be a loss function to be minimized rather than a welfare function to be maximized.

Let P denote either the entire period that policymaker p is in power or some subset of this period. The measure, denoted M, is as follows (low values of M are good):

(10.5) M = expected loss in P given p's actual behavior
− expected loss in P if p had behaved optimally
+ expected loss beyond P given p's actual behavior and given optimal behavior of future policymakers
− expected loss beyond P if p had behaved optimally and given optimal behavior of future policymakers
$= a - b + c - d.$

The term $a - b$ is the expected loss that could have been avoided during P if p had behaved optimally. The term $c - d$ is the potential expected loss to future policymakers from the fact that p did not behave optimally. If P is a subset of the entire period that p is in power, then "future policymakers" in the definition above may include p.

M takes account of the two problems mentioned earlier. If the economy is difficult to control for p, then b will be large, which will offset more than otherwise a large value of a. The term $c - d$ measures the effects of p's policies on the economy beyond period P, where these effects are measured under the assumption that future policymakers behave optimally.

The Computation of M

If a policymaker follows steps 1–6 in Section 10.2.4, he or she will be said to behave optimally. Remember, however, that the policy choice z_1^* in step 6 is not truly optimal because (1) the solution method is open-loop, (2) coefficient uncertainty has not been taken into account, and (3) deterministic rather than stochastic simulation has been used to compute the expected value of the objective function. As in (10.3), let $h_t(y_t, x_t, z_t)$ denote the objective function for period t, but now assume that it is a loss function rather than a welfare function. The loss function for the control problem is thus $\Sigma_{t=1}^{T} h_t(y_t, x_t, z_t)$.

In order to compute M, the period beyond P must be specified. Let 1 be the first period of P, and let P' be the length of P. The period beyond P will be assumed to run from $P' + 1$ to T'. The symbol T will continue to be used to denote the length of the horizon for the control problem. T is assumed to be larger than T'. It should be a number that is large enough so that further increases in T have a negligible effect on the optimal values for the first period of the horizon. Since only the values for the first period ever get used, the only criterion that needs to be used in deciding on the length of the horizon is the effect of this choice on the first-period values.

The procedure for computing M is as follows. (Steps 1–6 always refer to the steps in Section 10.2.4)

(i) Perform steps 1–6 for period 1. This requires choosing values for the expectations of the exogenous variables and error terms for periods 1 through T. These values should be estimates of what the policymaker actually knew at the beginning of period 1. The optimal values $z_1^*, z_2^*, \ldots, z_T^*$ minimize the expected value of $\Sigma_{t=1}^{T} h_t(y_t, x_t, z_t)$, where the expected value is computed by means of deterministic simulation. Let $\hat{h}_1^*$ denote the first term in the optimal sum, let x_1^e denote the values chosen for the expectations of the exogenous variables for period 1, and let u_1^e denote the values chosen for the expectations of the error terms for period 1. $\hat{h}_1^*$ is computed by solving the model for period 1 using z_1^*, x_1^e, u_1^e, and q_0 and then using these solution values (denoted $\hat{y}_1^*$) plus x_1^e and z_1^* to compute $\hat{h}_1^*$. The vector q_0 is the vector of initial conditions. $\hat{h}_1^*$ is h_1 evaluated at $\hat{y}_1^*$, x_1^e, and z_1^*. It is the part of b in (10.5) that corresponds to period 1.

(ii) Let z_1 denote the actual value of the control vector for period 1. Given z_1, x_1^e, u_1^e, and q_0, solve the model for period 1 and then use these solution values (denoted $\hat{y}_1$) plus x_1^e and z_1 to compute the value of the loss function for period 1 (denoted $\hat{h}_1$). $\hat{h}_1$ is h_1 evaluted at $\hat{y}_1$, x_1^e, and z_1. It is the part of a in (10.5) that corresponds to period 1.

(iii) Let u_1 denote the actual values of the error terms for period 1, and let x_1 denote the actual values of the exogenous variables for period 1. Given z_1^*, x_1, u_1, and q_0, solve the model for period 1. These solution values (denoted y_1^*) are estimates of what would have been observed in period 1 had the policy-maker behaved optimally. Let q_1^* denote the vector that includes y_1^*. (If there are lagged control variables in the model, then these variables should also be in q_{t-1} in Eq. 10.1. In this case z_1^* is in q_1^*.)

(iv) Perform steps 1–6 for period 2 using q_1^* as the vector of initial conditions. This will in general require choosing new values for the expectations of the exogenous variables. Given z_2^*, x_2^e, u_2^e, and q_1^*, solve the model for $\hat{y}_2^*$ and then compute $\hat{h}_2^*$. $\hat{h}_2^*$ is the part of b in (10.5) that corresponds to period 2.

(v) Given z_2, x_2^e, u_2^e, and q_1, solve the model for period 2 and then compute $\hat{h}_2$. q_1 is the vector of actual values of the initial conditions. $\hat{h}_2$ is the part of a in (10.5) that corresponds to period 2.

(vi) Repeat steps (iii), (iv), and (v) for periods 3 thrugh P'.

(vii) a in (10.5) is equal to $\Sigma_{t=1}^{P'}\hat{h}_t$, and b is equal to $\Sigma_{t=1}^{P'}\hat{h}_t^*$.

(viii) Given the optimal values for period P' from step (vi), $z_{P'}^*$, and given $x_{P'}$, $u_{P'}$, and $q_{P'-1}^*$, solve the model for period P'. Denote the solution values $y_{P'}^*$, and let $q_{P'}^*$ denote the vector that includes $y_{P'}^*$.

(ix) Perform steps 1–6 for period $P'+1$ using $q_{P'}^*$ as the vector of initial conditions. Given $z_{P'+1}^*$, $x_{P'+1}^e$, $u_{P'+1}^e$, and $q_{P'}^*$, solve the model for $\hat{y}_{P'+1}^*$ and then compute $\hat{h}_{P'+1}^*$. $\hat{h}_{P'+1}^*$ is the part of d in (10.5) that corresponds to period $P'+1$. This step is the same as step (iv) except for a different period.

(x) Repeat step (viii) for period $P'+1$, and then repeat step (ix) for period $P'+2$. Keep repeating through period T'. d in (10.5) is equal to $\Sigma_{t=P'+1}^{T'}\hat{h}_t^*$.

(xi) Perform steps 1–6 for period $P'+1$ using the actual value of $q_{P'}$ as the vector of initial conditions. To distinguish these optimal values from the optimal values computed in step (ix), let $z_{P'+1}^{**}$ rather than $z_{P'+1}^*$ denote them. Given $z_{P'+1}^{**}$, $x_{P'+1}^e$, $u_{P'+1}^e$, and $q_{P'}$, solve the model for $\hat{y}_{P'+1}^{**}$ and then compute $\hat{h}_{P'+1}^{**}$. $\hat{h}_{P'+1}^{**}$ is the part of c in (10.5) that corresponds to period $P'+1$.

(xii) Given $z_{P'+1}^{**}$, $x_{P'+1}$, $u_{P'+1}$, and $q_{P'}$, solve the model for period $P'+1$. Denote the solution values $y_{P'+1}^{**}$, and let $q_{P'+1}^{**}$ denote the vector that includes $y_{P'+1}^{**}$.

(xiii) Perform steps 1 – 6 for period $P' + 2$ using $q^{**}_{P'+1}$ as the vector of initial conditions. Given $z^{**}_{P'+2}$, $x^{e}_{P'+2}$, $u^{e}_{P'+2}$, and $q^{**}_{P'+1}$, solve the model for $\hat{y}^{**}_{P'+2}$ and then compute $\hat{h}^{**}_{P'+2}$.

(xiv) Repeat step (xii) for period $P' + 2$, where $q^{**}_{P'+1}$ is used as the vector of initial conditions, and then repeat (xiii) for period $P' + 3$. Keep repeating through period T'. c in (10.5) is equal to $\Sigma^{T'}_{t=P'+1} \hat{h}^{**}_t$.

This completes the computational steps. M is equal to $a - b + c - d$, where a and b are defined in step (vii), c is defined in step (xiv), and d is defined in step (x). The only difference between the steps involved in computing c and those involved in computing d is that for d the series of control problems begins from the initial conditions that would have prevailed had optimal policies been followed during P, whereas for c the series of control problems begins from the initial conditions that actually prevailed.

It is clear that the work involved in computing M is substantial. Assume, for example, that one is interested in measuring the performance of a presidential administration in the United States during its four-year period in office. If the model is quarterly, then 16 control problems need to be solved to compute b in (10.5). If the period beyond P is taken to be, say, 24 quarters, then 24 control problems need to be solved to compute c and 24 need to be solved to compute d. Computing M thus involves solving $16 + 24 + 24 = 64$ control problems, each of length T, where T should probably be some number like 40 (a 10-year horizon). Each of the 16 problems and each pair of the 24 problems require choosing values of the expectations of the exogenous variables. Even though this is a substantial amount of work, it is not completely out of the question. It might not be unreasonable to use autoregressive equations to generate the expectations of at least some of the exogenous variables, which would mechanize this part of the problem. The cost then would merely be the computer time to solve the 64 control problems. Although it is not feasible to do this for the results in this book, it should be possible in the future with faster and cheaper computers.

Since the first step of steps 1 – 6 is to estimate the coefficients over the latest available data, computing M also requires that the model be estimated a number of times. The model was estimated a number of times for the results in Chapter 8, and this is not that expensive. Note with respect to steps (iv), (ix), (xi), and (xiii) that the estimation must be over the actual data, not the data that would have existed had the policymaker behaved optimally. This is one unavoidable difference between what a policymaker could do in practice and what can be done after the fact in measuring performance. Note also that reestimation can occur right before steps (iii), (viii), (x), and (xii) rather than right after them. In other words, the model can be reestimated before the "actual" values of the error terms are computed. If this were done, u_1 in step (iii) would be based on the model estimated through period 1 rather than through period 0.

The problem of data revisions that was discussed in Section 8.2 regarding the evaluation of ex ante forecasts is also a problem here. A policymaker must make decisions on the basis of preliminary data, not the latest revised data that are generally used in econometric work. One possible solution to this would be to construct separate data sets for each starting point (that is, for the solution of each control problem), where each data set contains the preliminary data that were used as initial conditions and estimated data for the future periods that are consistent with the preliminary data. This is, however, a very tedious task, and it is unlikely to be done very often in practice. Most often it will merely be assumed that the latest revised data are good approximations to the data that the policymakers actually used.

Comparison to Chow's Measure

Chow (1978) has proposed a measure of performance that is almost identical to M if the model is linear and P consists of only one period. If the length of P is greater than one period, the two measures differ more. For P length 2, Chow's measure in words is as follows.

(10.6) M' = expected loss in period 1 given p's actual behavior in period 1
− expected loss in period 1 if p had behaved optimally in period 1
+ expected loss in periods 2 and beyond given p's actual behavior in period 1 and given optimal behavior in periods 2 and beyond
− expected loss in periods 2 and beyond if p had behaved optimally in period 1 and given optimal behavior in periods 2 and beyond
+ expected loss in period 2 given p's actual behavior in periods 1 and 2
− expected loss in period 2 if p had behaved optimally in period 2 but not in period 1
+ expected loss in periods 3 and beyond given p's actual behavior in periods 1 and 2 and given optimal behavior in periods 3 and beyond

– expected loss in periods 3 and beyond if p had behaved optimally in period 2 but not in period 1 and given optimal behavior in periods 3 and beyond.

The first four terms in (10.6) are the same as those in (10.5) if P is of length 1, and therefore in this case M and M' are identical if the expected losses are computed in the same way. In fact, however, Chow bases his computations of expected loss on the closed-loop approach, whereas the computations for M are based on the open-loop approach with reoptimization. This means that the expected losses are computed slightly differently even in the linear-quadratic case. This difference is fairly subtle, and it is not likely to be of much practical importance.

For P length of 2 it is clear that M and M' differ more than merely in how the expected losses are computed. Although there is no right or wrong answer regarding which measure is better, the question that M$'$ answers does not seem to be as relevant for policy evaluation as the question that M answers. Consider a presidential administration and a 16-quarter period. M compares the administration's actual behavior over the 16 quarters to the behavior that it would have followed had it optimized over the 16 quarters. M' compares first the administration's actual behavior in quarter 1 to the behavior that it would have followed in quarter 1 had it optimized, then its actual behavior in quarter 2 to the behavior that it would have followed had it started optimizing in quarter 2, then its actual behavior in quarter 3 to the behavior that it would have followed had it started optimizing in quarter 3, and so on through quarter 16. M seems more relevant for policy evaluation since it simply compares how well an administration did to how well it could have done had it optimized from the beginning. The question that M' answers is more complicated and also seems to resemble less the kinds of questions that are asked in practice about an administration's performance.

10.4 Solution of an Optimal Control Problem for the US Model

This section contains an example of solving an optimal control problem for the US model. The example is not realistic in the sense that the postulated loss function is too simple to approximate well the preferences of policymakers. The example is primarily meant to illustrate the properties of the model regarding the trade-off between real output and inflation.

10.4.1 The Loss Function and the Experiments

The period considered is 1973I–1977IV, and the loss function is

$$(10.7) \qquad L = \sum_{t=1973I}^{1977IV} \left\{ \left[\frac{GNPR_t - GNPR_t^*}{GNPR_t^*} \right]^2 + \lambda \left[\left(\frac{GNPD_t}{GNPD_{t-1}} \right)^4 - 1 \right]^2 \right\}.$$

This loss function is additive across periods and is quadratic. The first term is the square of the percentage deviation of real GNP from the high-activity-level $GNPR^*$, and the second term is the square of the percentage change in the GNP deflator at an annual rate. $GNPR^*$ is defined in Table A-4 in Appendix A. The parameter λ is the weight attached to inflation in the loss function.

One control variable was used: C_g, federal government purchases of goods. Monetary policy was assumed to be accommodating in the sense that the bill rate was taken to be exogenous and equal to its actual value each quarter. This means that the interest rate reaction function is not used. The Fed, for example, does not respond to any fiscal policy stimulus by raising short-term interest rates. Actual values for the exogenous variables and zero values for the error terms were used.

The objective is to choose C_g to minimize L subject to the US model. There are 20 values of C_g to determine, one per quarter. The problem was solved using the DFP algorithm. Actual values of C_g were used as starting values. Two problems were solved, one for $\lambda = 1$ and one for $\lambda = 2$. The results are presented in Table 10-1.

10.4.2 The Results

The first column in Table 10-1 presents the actual values of C_g, and the next two columns present the predicted values of the output gap and the rate of inflation that are based on the use of the actual C_g values. The predicted values are not equal to the actual values because zero error terms have been used. The 1974–1975 period was one of low output and high inflation, and the predicted values in the table are consistent with this.

The first set of optimal values is for $\lambda = 1$. The value of the loss function was lowered from .1470 to .1411. The output part of the loss was lowered from .0179 to .0069, and the inflation part was raised from .1291 to .1342. The optimal values of the output gap (the numbers in the a columns) are smaller than the base values for the 1974–1976 period, and the optimal inflation values are larger except for 1974I and 1976IV. The optimal values of C_g are

TABLE 10-1. Results of solving an optimal control problem for the US model

	Predicted values using actual C_g			Predicted values using optimal C_g					
				$\lambda = 1$			$\lambda = 2$		
	C_g	a	b	C_g	a	b	C_g	a	b
1973 I	12.45	-.81	7.15	11.36	-1.22	7.16	8.28	-2.34	7.20
II	11.65	-.29	8.70	9.80	-1.12	8.60	7.63	-2.14	8.26
III	11.45	-.78	8.60	13.18	-.44	8.17	12.16	-1.04	7.84
IV	12.12	-.78	6.72	11.73	-.90	7.01	9.69	-1.81	6.88
1974 I	11.67	-1.63	7.51	13.71	-.88	7.27	12.46	-1.65	7.10
II	12.10	-2.00	7.24	12.39	-1.63	7.81	9.46	-2.95	7.45
III	12.07	-3.22	9.17	15.84	-1.78	9.25	12.21	-3.41	8.70
IV	12.23	-4.78	9.05	19.39	-1.86	9.57	16.17	-3.46	9.19
1975 I	12.07	-5.34	7.24	18.85	-2.26	7.98	14.50	-4.17	7.54
II	12.03	-5.19	9.31	17.33	-2.49	10.71	12.59	-4.63	10.03
III	12.45	-4.32	9.05	16.46	-2.14	9.79	12.48	-4.06	9.26
IV	12.47	-3.45	7.98	16.68	-1.34	8.51	14.74	-2.56	8.13
1976 I	12.00	-3.30	6.14	14.93	-1.83	6.64	10.55	-3.53	6.38
II	11.95	-3.47	8.01	16.27	-1.72	8.19	12.57	-3.28	7.87
III	12.07	-3.18	6.47	15.55	-1.62	6.94	12.30	-3.04	6.59
IV	12.20	-2.83	6.63	11.84	-2.72	6.62	5.64	-4.95	6.23
1977 I	12.27	-2.41	10.14	12.14	-2.52	9.49	6.58	-4.78	8.90
II	12.92	-1.54	7.31	12.05	-1.76	7.19	9.64	-3.26	6.90
III	13.37	-.95	7.47	7.70	-3.22	6.71	-.71	-6.02	6.31
IV	13.25	-.85	9.27	12.99	-1.36	8.57	11.96	-2.60	8.27

	Value of L			
	$\lambda = 1$		$\lambda = 2$	
	Actual	Optimal	Actual	Optimal
Output loss	.0179	.0069	.0179	.0244
Inflation loss	.1291	.1342	.2582	.2448
Total loss	.1470	.1411	.2761	.2692

Notes: $a = 100 \cdot (GNPR - GNPR^*)/GNPR^*$.
$b = 100 \cdot [(GNPD/GNPD_{-1})^4 - 1]$.
• Cg is in units of billions of 1972 dollars.

larger for this period except for 1976IV. The overall results thus say that given the particular loss function and model, the optimal policy would have been for more stimulus in 1974–1976 than actually existed.

The optimal C_g values in the last two or three quarters are not to be taken seriously because they are trading on the fact that there is no tomorrow after the end of the horizon. These values have very little effect on the optimal value for C_g for the first quarter, which is the only quarter that matters for carrying out actual policy.

The optimal C_g values show fairly large fluctuations from quarter to quarter, and this is one of the reasons the example is not realistic. In practice there are constraints on the degree to which fiscal policy variables can be changed. The way in which this would be handled in the present context would be to add a term like $\gamma(C_{gt} - C_{gt-1})^2$ to the loss function. This would penalize large quarter-to-quarter changes in C_g. If this were done, other fiscal policy variables might also be taken as control variables (with similar penalties in the loss function) to increase the ability to minimize the loss with respect to the basic target variables. With no penalties on the control variables in the loss function, little is gained by using more than one control variable. The fiscal policy variables work roughly the same way with respect to their effects on output and inflation, and thus the use of one to minimize a loss function in output and inflation does about as well as the use of many. In this sense the control variables are collinear if there are no penalties on them in the loss function.

The second set of optimal values is for $\lambda = 2$, which is a higher weight on inflation in the loss function. The value of the loss function was lowered from .2761 to .2692. The output part of the loss was raised from .0179 to .0244, and the inflation part was lowered from .2582 to .2448. On average the optimal values of the output gap for 1974 and 1975 are not much different from those for the base run. The second loss function is thus one for which the optimal policy is not for more stimulus than actually existed in these two years. Overall, the optimal policy is for less stimulus, since the output part of the loss increases from the base solution to the optimal solution. The comments made above about the fluctuations in C_g pertain to both sets of optimal values, as do the comments about the values at the end of the horizon.

It should be stressed again that this example is not realistic, not only because no penalty on C_g fluctuations was imposed, but also because of the use of the actual values of the exogenous variables. If one were trying to approximate what could have been done during this period, estimated values should be used. In addition, the model should be estimated only up to the beginning of the control period, and separate control problems should be solved at the beginning of each quarter. In other words, this example is not what would be done if one were trying to compute the measure of performance discussed in Section 10.3.

10.4.3 Computational Experience

The program that I wrote for the DFP algorithm, which is discussed in Section 2.5, was used to solve the optimal control problem. The accuracy of the answer depends on the tolerance criteria used for the Gauss-Seidel technique in solving the model. The criteria that are discussed in Section 7.5.1 were used. Given this, the DFP algorithm essentially converged after six iterations for the $\lambda = 1$ problem. The use of two-sided derivatives resulted in a value of the loss function of .141150 after six iterations. Further iterations did not lower this value. The use of one-sided derivatives resulted in a value of the loss function of .141175 after six iterations, and further iterations did not lower this value. The use of two-sided derivatives thus gave a slightly more accurate answer. Each iteration required about 50 function evaluations when two-sided derivatives were used, 40 for the derivatives and 10 for the line search. The number of function evaluations was 20 less per iteration when one-sided derivatives were used.

The procedure that was discussed at the end of Section 10.2.1 for saving computer time was not used for the present results, which means that each function evaluation required solving the model for 20 periods. Although the cost-saving procedure was not used, the problem was programmed in such a way that the starting values for the Gauss-Seidel algorithm were always the solution values from the previous function evaluation. These are generally very good starting values in the sense of being close to the final answer. (When, for example, the derivative with respect to the control value for quarter 10 is being computed, with the derivative with respect to the quarter 9 control value having been computed in the previous function evaluation, the number of passes through the model per quarter for the first 8 quarters is merely one, since the solution for the first 8 quarters is the same for both derivatives.) As a result, the Gauss-Seidel technique required on average fewer passes through the model to achieve convergence for a given quarter than are required for other problems. The average cost per solution per quarter for the control problem was about .1 seconds on the IBM 4341, which compares to about .2 seconds for other problems. The cost per function evaluation was thus about .1 seconds $\times$ 20 quarters = 2 seconds, and so the cost per iteration of the DFP algorithm when two-sided derivatives were used was about 2 seconds $\times$ 50 function evaluations = 1.67 minutes.

The BFGS algorithm was also used to solve the control problem, and the results were almost identical to those for the DFP algorithm. The BFGS algorithm also converged to the allowed accuracy after six iterations.

The computational experience for the $\lambda = 2$ problem was almost identical to that for the $\lambda = 1$ problem. The only notable difference is that seven rather than six iterations were needed for convergence.