Michael J. O'Donnell
Computational Economics, HW4
February 24, 2010

Dynamic Optimization in MATLAB

The dynamic optimization problem in MATLAB involved the formulation and solution of a Quadratic Linear Problem (QLP) with dynamic programming. QLPs are deterministic control problems that can be formulated as continuous- or discrete-time models. The model explored in this paper, qlpabel.m, is a discrete-time model. The objective of this QLP is to find the optimal control path (u_k) that minimizes the criterion (J), a quadratic cost function, while subject to a linear system equation ($x_{k+1}(x_k, u_k)$) and an initial condition on the state parameter (x_0). The system equation in qlpabel.m represents a simple macroeconomic model, in which the state variables (x_k) are consumption (C) and investment (I), and the control variables (u_k) are government expenditures (G) and money supply (M):

$$x_{k+1} = \begin{bmatrix} 0.914 & -0.016 \\ 0.097 & 0.424 \end{bmatrix} x_k + \begin{bmatrix} 0.305 & 0.424 \\ -0.101 & 1.459 \end{bmatrix} u_k + \begin{bmatrix} -59.437 \\ -184.766 \end{bmatrix}$$

where,

$$x_k = \begin{bmatrix} C_k \\ I_k \end{bmatrix} \quad \text{and} \quad u_k = \begin{bmatrix} G_k \\ M_k \end{bmatrix}$$

In this sort of control problem, the policy maker (or modeler) specifies a target, or desired, path for the state and control variables, and the program computes the optimal path that minimizes J subject to the constraints (i.e., system equation, initial condition). In order to better understand this QLP model, I decided to explore the relationship between the target and actual paths of the state and control variables. As a first step, I updated the code to extract the target paths (i.e., target) and plot them against the optimal paths (i.e., path). Figures 1 and 2 show the results for default settings in qlpabel.m. In both figures, the lower trends (plotted in red) show the difference between the actual and target paths, while the upper trends plot these paths separately for each state or control variable. As shown in figure 1, C dips below the target path after the initial time period (t=0) and never returns to the target level. I increases above the target after the initial time period, but returns to the target in the terminal period (t=7). This trend results from the heavy weight placed on I in the terminal period (the 100 entry in matrix W_N):

$$W_N = \begin{bmatrix} 6.25 & 0 \\ 0 & 100 \end{bmatrix}$$

In figure 2, both G and M deviate from the target paths for all periods. G exceeds target spending, while M is less than the target supply. Deviation from the target paths is partially attributed to the small weights placed on the controls, as specified in matrix Λ_k:

$$\Lambda_k = \begin{bmatrix} 1 & 0 \\ 0 & 0.444 \end{bmatrix}$$

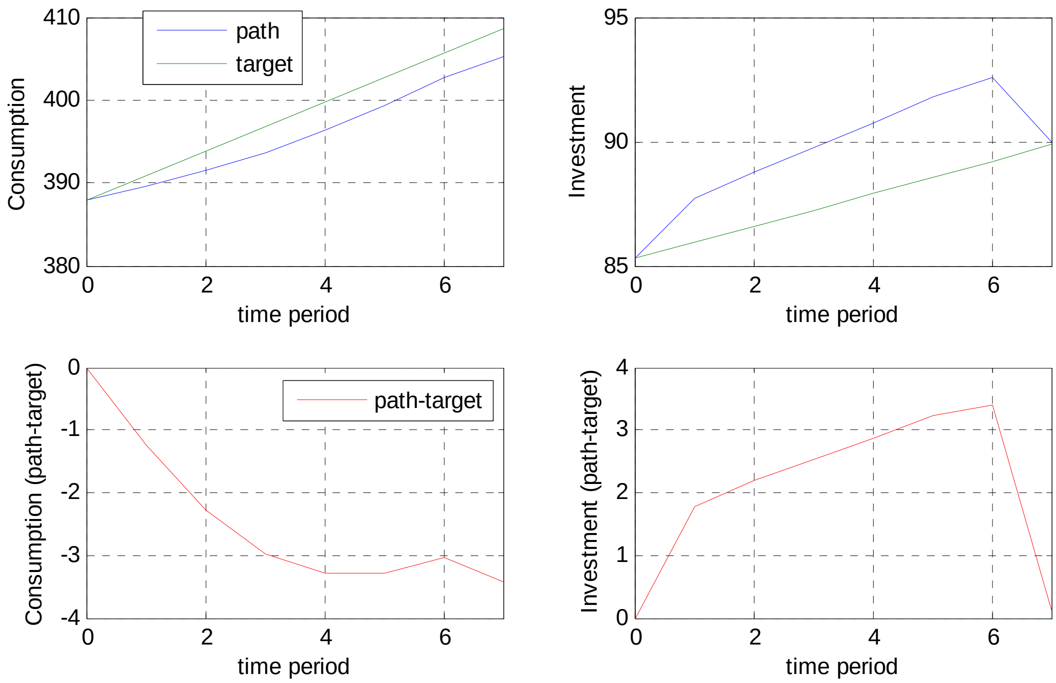

Figure 1. The actual and target paths of the state variables (*C, I*) are compared over time (for a default run of qlpabel.m).

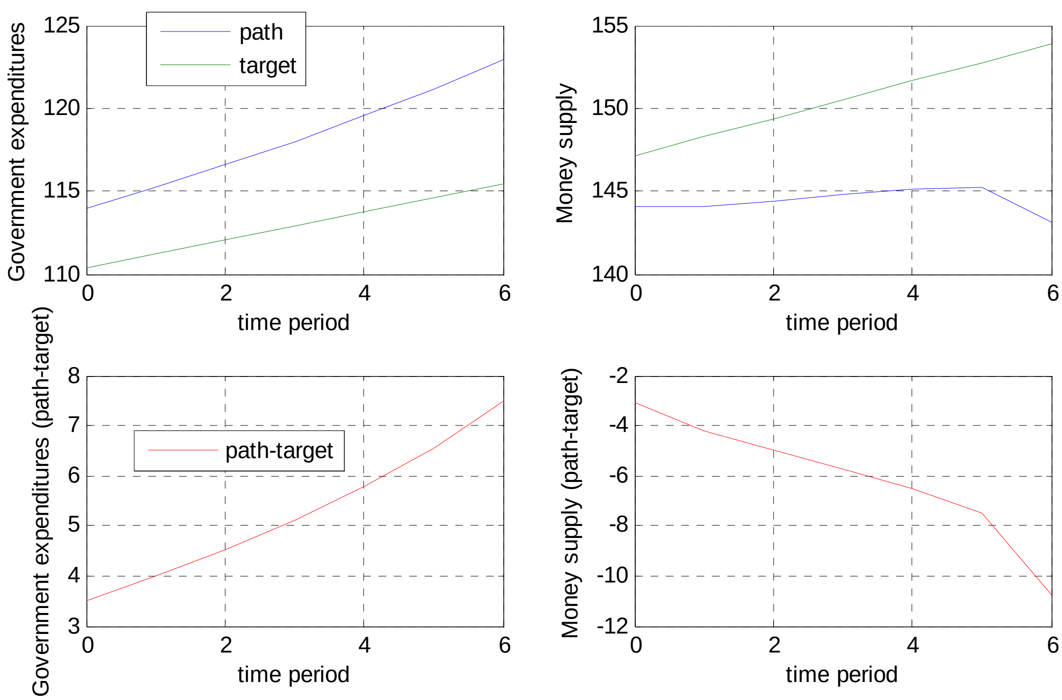

Figure 2. The actual and target paths of the control variables (*G, M*) are compared over time (for a default run of qlpabel.m). The control variables are determined only through the 6th time period since the controls at t=6 determine the state at t=7 (the terminal period).

Experimental Run

To further explore the model, I updated the code to give the user the ability to alter the desired, or target, path for the individual control variables (i.e., separate growth rates for *G* and *M*).[1] This change allows the user to specify a unique growth rate for the state variables (i.e., a single, unique growth rate for x_k). The code was updated such that the user can enter any growth rate ranging from -100% to 100% (-1 to 1) for each of the three growth rate parameters (i.e., `growth_x`, `growth_G`, `growth_M`).

This new code was first tested by setting each of the growth rate parameters to the default setting of 0.75% and reproducing figures 1 and 2. I then proceeded to experiment with the growth parameters, in combination with different weights on the state and control variables in W_N and Λ_k. As an example, consider an Administration that is seeking to rein in, and reduce, government spending due to backlash from the public. The Administration is seeking a quarterly reduction in *G* of 0.75%, while allowing *M* to expand at 0.75%. However, this new policy must permit the economy to expand at 0.75% (again, per quarter). With the focus on reducing *G*, the weight on this control variable is increased by two orders of magnitude in Λ_k:

$$\Lambda_k = \begin{bmatrix} 100 & 0 \\ 0 & 0.444 \end{bmatrix}$$

Therefore, the administration has little concern over *M*, or inflation, relative to the desired reductions in *G*. Results are presented in figures 3 and 4.

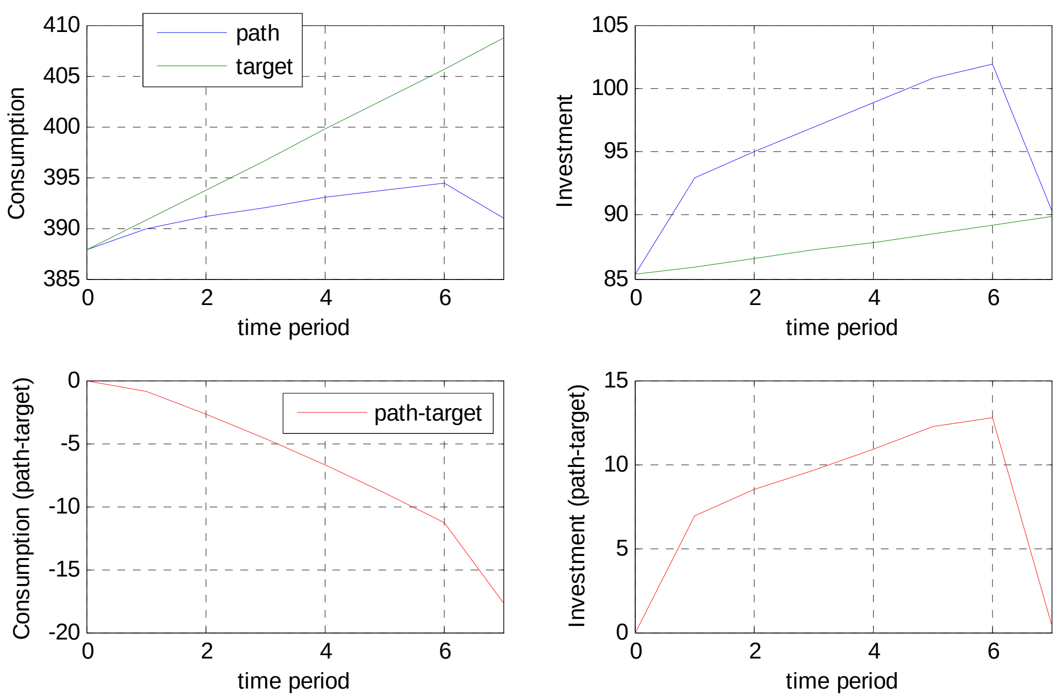

Figure 3. The actual and target paths of the state variables (*C, I*) are compared over time for the experimental run (i.e., fiscal policy aimed at reducing *G*).

[1] In the default code, the desired path is hard-coded as 3% growth per year (0.75% per quarter).

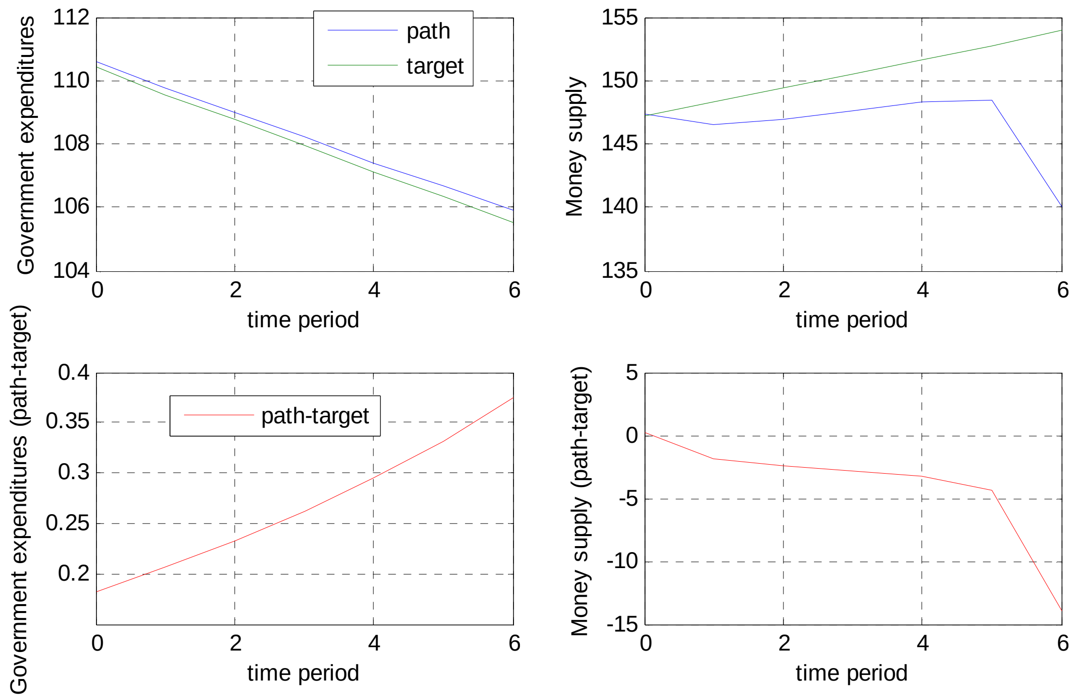

Figure 4. The actual and target paths of the control variables (G*, M*) are compared over time for the experimental run (i.e., fiscal policy aimed at reducing *G*). Again, the control variables are determined only through the 6th time period since the controls at t=6 determine the state at t=7 (the terminal period).

Turning first to figure 4, the control paths, we see that the desired fiscal policy has been accomplished, in that the government tightly follows the desired quarterly reduction in spending (*G*). As expected, *M* deviates rapidly from the target path after the initial period due to the small weight placed on this control variable. However, it is not clear why M falls relative to the target path. In figure 3, we see that this new policy has an impact on the desired trajectory of the economy. *C* remains essentially flat, while *I* increases dramatically between the initial and terminal periods. The government may have satisfied its goal of reining in spending, but the economy stagnated (as reflected by *C*).

Although this model was useful for exploring and understanding the QLP and deterministic control problems in economics, it is very simplistic in its treatment of the economy (with a linear difference system equation) and, like any deterministic problem, does nothing to address uncertainty. I look forward to later lessons in the course that specifically address the uncertainty in our understanding of the economy and how we think we can control it.