

Working Paper 9601

DYNAMIC COMMITMENT AND IMPERFECT POLICY RULES

by Joseph G. Haubrich and Joseph A. Ritter

Joseph G. Haubrich is a consultant and economist at the Federal Reserve Bank of Cleveland. Joseph A. Ritter is a senior economist at the Federal Reserve Bank of St. Louis. The authors wish to thank Rob Dittmar for programming assistance and Avinash Dixit and seminar participants at North Carolina State University, the University of Iowa, the University of Pittsburgh, and the Federal Reserve Bank of Atlanta for helpful comments.

Working papers of the Federal Reserve Bank of Cleveland are preliminary materials circulated to stimulate discussion and critical comment. The views stated herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or of the Board of Governors of the Federal Reserve System.

February 1996

Abstract

Examining the dynamics of commitment highlights some neglected features of time inconsistency. We modify the rules-versus-discretion question in three ways: 1) A government that does not commit today retains the option to do so tomorrow; 2) the government's commitment capability is restricted to some class of simple rules; and 3) the government's ability to make irrevocable commitments is restricted.

Three results stand out. First, the option to wait makes discretion relatively more attractive. Second, the option to wait means that increased uncertainty makes discretion more attractive. Third, because the commitment decision takes place in "real time," policy choice displays hysteresis.

I. Introduction

In recent years, there has been a drift in economists' way of thinking about policy rules versus discretion. Beginning with Kydland and Prescott (1977), a theoretical presumption has developed in favor of rules, which allow outcomes otherwise precluded by strategic behavior. This contrasts with the early monetarists, who proposed simple rules because the monetary authority could not handle the complexities of the actual economy. The theoretical case demonstrates how the precommitment to fully state-contingent rules solves the time inconsistency problem and is superior to discretion. In this paper we modify the rules-versus-discretion question in three important ways: 1) A government that does not commit today retains the option to do so tomorrow; 2) the government's commitment capability is limited to a class of simple rules; and 3) the government's ability to make irrevocable commitments is restricted.

The first of these modifications raises the possibility that the government might delay committing to a rule, with the outcome of the decision depending on the current state of the world and, in our most general model, on history. In this sense it becomes important that the decision be made in "real time." A broader implication, also unrecognized in the previous literature, is that choosing discretion today has an option value, since the government may still choose rules in the future. Previous work considers a once-and-for-all choice between rules and discretion and does not allow a

future government to adopt rules. If the option to wait indeed has positive value--as such options often do--it adds to the desirability of discretion.¹

Complexity makes commitment more difficult. Conceiving, specifying, and committing to all possible contingencies becomes prohibitively costly, if possible at all. No one is surprised that the Federal Reserve Act of 1914 discusses neither Internet cash nor monetary policy in the event of an oil crisis. This leads to our second modification, that policymakers must choose not between discretion and optimal state-contingent rules, but between discretion and comparatively simple and imperfect rules (as recently emphasized by Flood and Isard [1989] and Lohmann [1992]). Thus, it is logically possible for policymakers to "regret" their commitment to a rule. Regret makes questions of delay interesting, because a government would never delay committing to a rule that was always better than discretion in every state.

Our third modification stems from the observation that governments cannot make irrevocable commitments. Nevertheless, they do have a wide array of commitment mechanisms, from campaign promises to constitutions. These mechanisms involve a range of commitment costs (from zero to high) and, more important, a range of renegeing costs.

The remainder of this paper develops these themes in two main variations. In section II, we explore what happens when governments can and cannot commit to

¹ Our exploration of regret and the associated option value of waiting distinguishes this paper from similar efforts, such as Cukierman and Meltzer (1986) and Flood and Isard (1989). Cukierman and Meltzer discuss flexibility, but do not consider imperfect fixed rules and *a fortiori* miss the associated option value. Admittedly, for some models (Barro and Gordon [1983] for example) the optimal rule is simple, but generally the optimal state-contingent rules are rather complex. The option value of waiting, of course, plays a key role in the analysis of irreversible investment (see Pindyck [1991] or McDonald and Siegel [1986]).

fully state-contingent rules, using the standard simple model of monetary policy traditional in the time-consistency literature. In section III we extend the model to many periods and drop the assumption that commitment is a once-and-for-all decision. We trace the consequences when only "simple" rules are feasible and when choosing discretion today does not rule out choosing commitment in the future. Commitment, however, once made, remains irrevocable. Some numerical examples explore the significance of the results.

Section IV removes the rigid and unrealistic assumption that irrevocable commitment is feasible. It provides a very general way of thinking about policy, allowing costly commitment with costly reversal. In it we illustrate how decisions to commit or renege depend on the commitment and renegeing costs and on uncertainty in the environment.

In section V, we conclude by emphasizing three general results. First, the option to wait, which we have restored to the policymaker's decision problem, makes discretion relatively more attractive. Second, the option to wait means that increased uncertainty makes discretion even more attractive. This is the "bad news principle" of irreversible investment applied in a policy context. Third, by allowing the commitment decision to take place in "real time," we find that the policy choice process displays hysteresis; which policy is in force at a given time depends on history, not just the prevailing state.

II. Optimal Rules, Simple Rules, and Discretion

Most of the debate about rules versus discretion has taken place in the arena of monetary economics. We continue this tradition, and in this section set out our model. Though slightly specialized to highlight the main points, it derives from a fairly general framework based on Flood and Isard (1989). We first use the model to explore the distinctions between monetary policy under optimal rules, simple rules, and discretion.

A. Basic Specification

The growth of base money, b_t , relative to a velocity shock, v_t (ignored hereafter), determines the inflation rate, π_t :

$$(1) \quad \pi_t = b_t + v_t.$$

Output depends on unexpected inflation, with the Federal Reserve focusing on the deviation of output from a natural level. Because of distortions (for instance, unemployment insurance or imperfectly clearing labor markets, depending on your preferred ideology), that natural level may not be socially optimal.

Policymakers wish to minimize a social-loss function that reflects both output deviations and inflation:

$$(2) \quad L_t = (b_t - E_{t-1}b_t - K + u_t)^2 + ab_t^2.$$

The term $b_t - E_{t-1}b_t$ measures the unexpected base growth (or unexpected inflation), K measures distortion (or the divergence between the natural level of output and the socially optimal level), and u_t , i.i.d with $E u_t = 0$, measures the productivity shock. The parameter a measures the relative weight given inflation--as opposed to output--deviations.

The first step in finding the optimal policy is to minimize the loss function, L_t , under both rules and discretion.

B. Discretion

From the first-order conditions for minimizing (2), we find

$$(3) \quad b_t = \left(\frac{1}{1+a}\right)(E_{t-1}b_t + K - u_t).$$

This implies

$$(4) \quad E_{t-1} b_t = E_{t-1} b_t = \frac{K}{a}$$

As in Barro and Gordon (1983), the distortion term K determines the inflationary bias of discretion. Actual base growth under discretion is

$$(5) \quad b_t^D = -\frac{u_t}{1+a} + \frac{K}{a}.$$

From this, we can calculate both the expected and realized social loss using equation (2).

$$(6) \quad \text{Realized Loss: } L_t^D = \frac{1+a}{a} \left[-K + \frac{a}{1+a} u_t \right]^2 \text{ and}$$

$$(7) \quad \text{Expected Loss: } E_{t-1} L_t^D = \frac{1+a}{a} K^2 + \frac{a}{1+a} u_t^2.$$

The first term of equation (7) is the loss from the inflation bias of discretion, while the second is the loss caused by output variance, some of which shows up in the inflation rate via monetary policy.

C. Rules

Suppose the money supply, b_t , cannot respond to u_t , thus restricting the policymaker to rules that are not state contingent. Then money only causes inflation; it cannot reduce output variance. The best such rule sets $b_t = 0$ in all periods. This is the optimal rule without state contingency. If it were feasible, a better rule (which we derive below) would let the base react to productivity shocks, but would avoid the inflationary bias of pure discretion.

For the simple rule setting $b_t = 0$ for all t , we can substitute into the loss function.

$$(8) \quad \text{Realized Loss: } L_t^R = (u_t - K)^2 \text{ and}$$

$$(9) \quad \text{Expected Loss: } E_{t-1} L_t^R = K^2 + \frac{u^2}{a}.$$

Equations (8) and (9) show that the rule has a lower inflation bias than does discretion, but a higher output variance.

Discretion is better than the simple rule when $L^D - L^R < 0$. Substitution from equations (6) and (8) shows that this is the case when

$$(10) \quad u^2 > \left(1 + \frac{1}{a}\right) K^2.$$

Notice that discretion is preferable in extreme times (that is, for large u_t s), when the costs of shocks are especially high.

As inflation costs (a) increase, discretion is preferred in more and more states. This may seem counterintuitive, but in fact it makes sense. Consider, for example, the case of $u_t = 0$. For the simple rule setting $b_t = 0$, the loss due to inflation is 0. For

discretion, the corresponding loss is $a\left(\frac{K}{a}\right)^2 = \frac{K^2}{a}$. As a increases, this cost decreases. Because discretion weighs the inflationary costs of intervention, higher inflation costs reduce the inflation bias of discretion. In the limit, with inflation infinitely costly, discretion involves zero inflation.

Similarly, as K increases, discretion is preferred in fewer states. As the distortion worsens, the inflation bias rises and it becomes worthwhile to sacrifice discretion in favor of a rule. The relative return to the rule increases because the higher distortion increases the inflation bias.

If the government can commit to a state-contingent rule, it can replicate discretion's offset to productivity shocks while simultaneously eliminating the inflationary bias. When feasible, this rule would let the monetary base react to productivity shocks but avoid the inflationary bias of pure discretion. In our simple model, it is possible to find this optimal rule. Its form illustrates several points about the relationships among optimal rules, simple rules, and discretion.

To find the optimal state-contingent rule, we minimize the expected loss function from equation (2):

$$\min EL = \sum_{i=1}^n \left\{ \left[b_i - \left(\sum_{j=1}^n g_j b_j \right) - K + u_i \right]^2 + ab_i^2 \right\} g_i ,$$

where g_i denotes the probability of state i , b_i denotes money growth in state i , and n denotes the number of states.² We end up with:

² Nothing essential depends on using a discrete probability distribution. A continuous distribution leads to identical results, but necessitates needlessly cumbersome notation.

$$(11) \quad b_i = -\frac{1}{1+a}u_t.$$

Substituting into the loss function, we have

$$(12) \quad \text{Realized Loss: } L_t^F = \frac{a}{1+a}u^2 - 2\frac{a}{1+a}Ku + K^2 \text{ and}$$

$$(13) \quad \text{Expected Loss: } E_{t-1}L_t^F = \frac{a}{1+a}u^2 + K^2.$$

To understand the implications of fully state-contingent rules, it is important to look at several relationships. Since the optimal rule ties the policymaker's hands, it allows actions that would otherwise fall victim to time inconsistency. This makes the optimal rule better than discretion. Second, by construction, the optimal rule dominates a restricted or simple rule in *expected value*. In the current example, the optimal rule turns out to be linear. Our simple rule is trivially linear: a constant $b_t = 0$.

A tangential but important point is that while the optimal rule dominates both discretion and the simple rule on average, it does not do so in all states. In some states the other policies do better, as a comparison of equations (6), (8), and (12) shows. For example, in states where $0 < u_t < 2K$, the simple rule does better because the optimal rule's response to u_t is not worth the (small) amount of inflation that ensues. A rule that attempts to exploit this inefficient response, however, changes expectations in a way that on average hurts more than it helps. To illustrate, suppose we attempt to revise the optimal rule by setting $b_t = 0$ whenever $0 < u_t < 2K$. This lowers expected inflation but increases the loss in states where state contingency is useful. The gain in states where $0 < u_t < 2K$ is offset by the loss in other states, even though policy is

unchanged in those other states. The response of individual behavior (in this case expectations) distinguishes an equilibrium problem from a simple control problem.

III. Waiting to Commit

The approach we take only begins to differ from the standard approach in a more dynamic setting. First we turn to a simple numerical example that helps to clarify the notions of regret, delay, and option value.

A. A Two-Period Example

The simplification begins with the productivity shock, assuming u_t is i.i.d. and equals $-x$, 0 , or $+x$ with probabilities g_1 , g_2 , and g_3 .

Now, suppose Alan Greenspan wakes up and finds that today, $u_1 = 0$. If he says, "I commit," then the two-period social loss function is

$$(14) \quad V_R(0) = K^2 + g_1(x^2 + 2Kx + K^2) + g_3(x^2 - 2Kx + K^2) + g_2K^2.$$

The first term, K^2 , measures the loss today, while the following three terms measure the next period's expected loss.

If the Chairman chooses discretion today, the loss function becomes more complicated because he may commit tomorrow, depending on the state:

$$(15) \quad V_D(0) = \frac{1+a}{a}K^2 + g_1 \left(\frac{a}{1+a}x^2 + 2Kx + \frac{1+a}{a}K^2 \right) \\ + g_3 \left(\frac{a}{1+a}x^2 - 2Kx + \frac{1+a}{a}K^2 \right) + g_2K^2.$$

Again, the first term is the loss today. The terms containing g_1 and g_3 are the expected loss when the government chooses discretion next period, as it does when $u = \pm x$. The last term, $g_2 K^2$, is the expected loss when the government commits to rules *tomorrow*, as it does when $u = 0$.

Choosing rules or discretion comes down to comparing equation (14) with (15) and then choosing the strategy with the smaller expected loss.

$$(16) \quad V_R(0) - V_D(0) = \frac{-1}{a} K^2 + (g_1 + g_3) \left[\frac{1}{1+a} x^2 - \frac{1}{a} K^2 \right].$$

Removing the option value, that is, forcing the government to make an irrevocable choice between rules and discretion, leads to a different expression because in the value of discretion forever, $V_{DF}(0)$, the $g_2 K^2$ term in (15) is replaced with $g_2 \frac{1+a}{a} K^2$.

This makes the difference between rules and discretion forever:

$$(17) \quad \begin{aligned} V_R(0) - V_{DF}(0) &= \frac{-1}{a} K^2 + (g_1 + g_3) \left[\frac{1}{1+a} x^2 - \frac{1}{a} K^2 \right] - g_2 \frac{1}{a} K^2 \\ &= V_R(0) - V_D(0) - g_2 \frac{1}{a} K^2. \end{aligned}$$

For a range of cases, $V_R(0) - V_D(0) > 0$ and $V_R(0) - V_{DF}(0) < 0$, so that correctly valuing the option leads to choosing discretion, while ignoring it leads to choosing rules. Taking account of the real-time aspect of decision making and properly valuing the waiting option can reverse the policy decision. As an example, let

$$a = 1, \quad x^2 = 6K^2 \text{ and } g_1 = g_3 = 1/4.$$

Then $V_R(0) - V_D(0) = -K^2 + \frac{1}{2}(3K^2 - K^2) = 0$, and

$$V_R(0) - V_{DF}(0) = -K^2 + \frac{1}{2}(3K^2 - K^2) - \frac{1}{4}K^2 = -\frac{1}{4}K^2 < 0.$$

This example shows both that the option can reverse the normal presumption of the superiority of rules and that the option value may be sizable. Using equation (14), the loss from adopting rules is $5K^2$, making the option difference $1/4K^2$, or 5 percent of the total value.

This example also illustrates the bad-news principle. It shows why policymakers may be correct in focusing on the problems of committing at an inappropriate time. In (16), a mean-preserving spread in the distribution, say a decrease in g_2 and corresponding increase in g_1 and g_3 , increases the relative attractiveness of discretion. Furthermore, it does so in a particular way. What counts in (16) is the effect in states 1 and 3, where we choose discretion instead of rules. The benefit arising in state 2 does not appear. What matters is the loss in states with big shocks (states 1 and 3) that are bad for rules.

B. *Many Periods*

Adequately capturing irreversibility requires a number of adjustments to the model. First, it clearly needs more than one period. Second, to better focus on the problems of regret, it is also helpful to revise the within-period time structure. In what follows, we let the government observe the shock before the public does and before it chooses to commit. The new time line, which leaves equations (1)-(12) intact, is as follows:

Government sees $u_t \rightarrow$ Government decides whether to commit, announces \rightarrow Economy revises expectations $E_{t-1}b_{it} \rightarrow$ Government chooses $b_t \rightarrow$; Economy sees u_t ; production.

The contrived aspect here concerns observing the shock. After seeing today's shock, the government chooses rules or discretion, but the public does not see u_t until much later. Some variant of this assumption appears in much of the literature. In Cukierman and Meltzer (1986), for instance, the government has information on a state variable that the public observes one period later. In Canzoneri (1985), the government observes (perhaps noisily) a random disturbance that the public cannot.

In general, this new timing sequence will change the public's behavior. Seeing what action the government takes provides information about the unseen shock to the economy. In our specific model, however, the quadratic loss function and the symmetry of the shocks mean that the public cannot extract useful information from

the government's decision to commit or not. People can infer the size, but not the sign, of the shock, so that $E(u_t | \text{government choice}) = 0$ and $E(b_t | \text{government choice}) = 0$.³

Once the government chooses a simple rule, it must stick with that decision forever, in effect setting $b_t = 0$ permanently. By contrast, choosing discretion today does not prevent choosing rules tomorrow.

In this setting, irreversibility introduces an option value whose worth is non-negative.⁴ With a simple, non-state-contingent rule, regret exists. For example, the government might regret committing to zero inflation and wish for discretion. This point does not depend merely on the rule's extreme simplicity. The analysis holds even with a more sophisticated state-contingent rule, as long as there are some states in which discretion is preferred. As mentioned before, in some states the government would even regret committing to the optimal state-contingent rule.

With many periods, policy choice comes down to comparing possible courses of action. This is most naturally done using dynamic programming (see Ross [1983]). For any policy (that is, for any set of b_t choices by the government, denoted π), we have a value function

³ This symmetry breaks down if we compare the optimal state-contingent rule with discretion. This happens because the states where discretion is preferred is not symmetrical around zero: the government's decision with regard to commitment would give information to the public, who would then update their expectations. The increased complexity adds to the signal extraction problem, without any corresponding gain in economic content. This is one reason why we do not pursue the comparison in this paper. The other, more important reason is that we consider simple, non-optimal rules more realistic.

⁴ Our definition emphasizes the option as an option to commit. An alternative, complementary, approach emphasizes the option as an option to undo a commitment (See Bernanke, 1983). The difference in perspective explains why the bad-news principle and the option value look different, even though both are aspects of the same phenomenon. Is it best to compare the optimal policy with discretion forever (as we do) or with rules? The answer depends on what is most convenient for the problem at hand.

$$V_p(u_t) = E \sum_{s=1}^{\infty} \beta^s L(b_{t+s}, u_{t+s}).$$

Here the factor β discounts the future. To rule out reputational equilibria, we restrict ourselves to nonrandomized policies and to those that depend only on today's shock and whether or not the government has committed in the past. The government begins this period by observing u_t . If it chooses to commit to zero inflation (the optimal simple rule), the loss is

$$(18) \quad \text{a) } V_R(u_t) = (-K + u_t)^2 + \beta E V_R(u_{t+1}), \text{ from which we arrive at}$$

$$\text{b) } V_R(u_t) = (u_t - K)^2 + \frac{\beta}{1-\beta} (K^2 + \frac{a}{1+a} u_t^2),$$

where $V_R(u)$ denotes the value function for rules. The first term measures today's loss, and the second gives the expected value of the problem tomorrow. Choosing discretion forever yields a loss of

$$(19) \quad \text{a) } V_{DF}(u_t) = \left(\frac{1+a}{a} \right) \left(-K + \frac{a}{1+a} u_t \right)^2 + \beta E V_{DF}(u_{t+1})$$

$$\text{b) } V_{DF}(u_t) = \left(\frac{1+a}{a} \right) \left(-K + \frac{a}{1+a} u_t \right)^2 + \frac{\beta}{1-\beta} \left(\frac{1+a}{a} K^2 + \frac{a}{1+a} u_t^2 \right)$$

The standard time-consistency literature, making the choice before any shocks are observed, asks whether rules are better than discretion by comparing the expected values of (18) and (19). The general case is more complicated because opting for discretion today leaves the door open for choosing rules tomorrow. The loss to choosing discretion today is

$$(20) \quad V_D(u_t) = \left(\frac{1+a}{a} \right) \left(-K + \frac{a}{1+a} u_t \right)^2 + \beta E V_D(u_{t+1}).$$

Two different representations of $EV_D(u_t)$ turn out to be useful. Without any simplifying, we can express this term as

$$(21) \quad EV_D(u_t) = \sum_{s=1}^{\infty} \left\{ \sum_{j \in UR_{t+s}} g_j (u_j - K)^2 + \sum_{j \in UR_{t+s}} g_j \left[\frac{1+a}{a} \left(-K + \frac{1}{1+a} u_j \right)^2 \right] \right\},$$

where g_j is the probability of state j and UR_{t+s} contains the period $t+s$ states in which the government Uses a Rule. Here, UR_{t+s} depends on history; that is, commitment to a rule implies commitment in all future states.

Simplifying this expression takes a little work. First, note that the set of states in which the government chooses to Commit to Rules, CR , does not vary with time. (This differs from UR_t in equation [21], where prior commitment does change the action. UR_t answers the question, "At time t , in which state does the government use rules?" CR_t answers the question, "At time t , given that it can still choose, in which states does the government commit to rules?") The time invariance of CR follows from the simple form of equation (20). Then, recursively using equation 18(b) yields

$$(22) \quad EV_D(u_t) = \sum_{j \in CR} g_j \left[(u_j - K)^2 + \frac{1}{1-a} (K^2 + \frac{2}{a} u_j) \right] + \sum_{j \in CR} g_j \left[\frac{1+a}{a} \left(-K + \frac{a}{1+a} u_j \right)^2 \right] \\ + \sum_{j \in CR} g_j \left\{ \sum_{j \in CR} g_j \left[(u_j - K)^2 + \frac{1}{1-a} (K^2 + \frac{2}{a} u_j) \right] \right\} + \sum_{j \in CR} g_j \left\{ \left[\sum_{j \in CR} g_i \frac{1+a}{a} \left(-K + \frac{a}{1+a} u_j \right)^2 \right] + \dots \right\}$$

The first term is the expected loss if we enter a state in which we choose rules and adhere to them forever. The second term represents the loss today from using discretion today only. The third term gives the loss from choosing rules in the period after discretion. The fourth term gives the loss from choosing discretion again, with this pattern repeating recursively.

Equation (22) simplifies to

$$(23) \quad EV_D(u_t) = \frac{1}{1 - \left(\sum_{j \in CR} g_j \right)} \left\{ \sum_{j \in CR} g_j \left[(u_j - K)^2 + \frac{1}{1-a} \left(K^2 + \frac{a}{1+a} u_j^2 \right) \right] \right. \\ \left. + \sum_{j \in CR} g_j \left[\frac{1+a}{a} \left(-K + \frac{a}{1+a} u_j \right)^2 \right] \right\}.$$

Finding the value function puts us in a position to examine the central issues of regret, option value, and delay. Of course, different parameters can make rules or discretion the better choice, but of interest here is what is unique to our model. To this end, we focus on parameter values for which an irrevocable choice between rules and discretion would favor rules. We then show that the possibility of future commitment can make discretion today preferable, noting the importance of regret in that decision. This increase in the attractiveness of discretion induces the government to choose discretion in more states, a policy shift perhaps best interpreted as a delay in commitment.

To rule out the trivial cases, we need some regret, so that simple rules do not dominate discretion in every state of the world. If the loss from rules is less than the loss from discretion in every state, it makes no sense to delay commitment or to choose discretion. To have any regret, it must be that for some (but not all) shocks u ,

$$(u - K)^2 > \frac{1+a}{a} \left(-K + \frac{1}{1+a} u \right)^2.$$

We also want rules to do better in expected value

terms than discretion forever, or else discretion forever is the obvious trivial choice.

This requires $\left(K^2 + \frac{a}{1+a} u^2 < \frac{1+a}{a} K^2 + \frac{a}{1+a} u^2 \right)$, or

$$(24) \quad K^2 > \frac{a}{1+a} \frac{2}{u}.$$

The problem for the government at $t = 0$ is to decide between

$$(25) \text{ Rules: } V_R(u_o) = (u_o - K)^2 + \frac{1}{1-a} \left(K^2 + \frac{2}{u} \right)$$

and

$$(26) \text{ Discretion: } V_D(u_o) = \frac{1+a}{a} \left(-K + \frac{1}{1+a} u_o \right)^2 + EV_D(u_1),$$

where $EV_D(u_1)$ is given by equation (21) or, equivalently, (23). It is also important to know how $V_D(u_o)$ and $V_R(u_o)$ compare with discretion forever, $V_{DF}(u_o)$ (given in equation [19]).

Since $V_R(u) - V_D(u) = [V_R(u) - V_{DF}(u)] - [V_D(u) - V_{DF}(u)]$, we can see that $V_D(u) - V_{DF}(u)$ gives the option value of discretion -- the value of the ability to abandon discretion and commit to rules. Since this quantity is ordinarily positive, we can have $V_R(u) - V_D(u) > 0$, even when $V_R(u) - V_{DF}(u) < 0$ (the standard criterion).

Note that since discounted future losses are lower for discretion,

$EV_D(u_t) < EV_R(u)$, (because at worst, the discretion regime could commit next period and attain equality) the government sometimes chooses discretion in states where the one-period return favors rules. This conceivably could create a paradox whereby we delay choosing rules forever, even though we prefer pure rules to pure discretion. Actually, we can use equation (23) to demonstrate that this never occurs in the case of irrevocable investment. Suppose the government never commits, so that

$CR = \emptyset$. Then (23) reduces to $\frac{1}{1-a} \left(\frac{1+a}{a} K^2 + \frac{1}{1+a} u^2 \right)$ which, from (19b), we know is $EV_{DF}(u_t)$. We assume, however, that discretion forever is worse than rules.

Along with eliminating such an "infinity paradox," the above calculation has another implication. The government commits with a fixed positive probability in each period, so with probability 1, the government eventually commits (by the Borel-Cantelli lemma).⁵

C. Numerical Example

To add a small degree of realism, the next example employs the infinite-horizon model, using parameter values we believe to be at least of the right order of magnitude. While it cannot be called a test, nor even a calibration exercise, we try to use plausible values for the effect of unanticipated money and the distribution of output shocks. In this scenario, the government chooses discretion in about half the states.

First differences of log GDP look somewhat like a standard normal. We therefore assume that u is drawn from a discrete distribution that approximates a normal. (For details, see section IV.) We choose a K value of 1.0, indicating that long-run output differs from the socially optimal rate by 1.0 percentage point. Following Barro (1987, p. 469), we make the assumption implicit in equation (2) that a 1 percent rise in money above expectations increases output by 1 percentage point.

⁵ For a very different view of commitment problems using similar stochastic commitment techniques, see Roberds (1987).

Two more parameter choices will fully specify the problem. Give inflation twice the weight of output in the social-loss function, choosing an a of two. Next, set β , the discount factor, to 0.95. We think of the policymaker as choosing between rules and discretion once a year.

Figure 1 shows the results of this example using these parameters.⁶ The top panel plots the difference between $V_R(u)$ and $V_D(u)$, or between the value of committing to rules and adopting discretion in a given state. Since we use a loss function, a positive value means discretion is better, and a negative value means rules are better.

Notice that for any u shock between -1.02 and +1.02, the social loss from discretion exceeds that from rules. Consequently, the monetary authority should commit to rules. For larger shocks, the monetary authority should choose discretion. For 30.

1 percent of the time, discretion is preferable to rules.

The bottom panel shows the importance of considering option value. If we compare using rules forever with using discretion forever, we would choose rules in every state. The possibility of future commitment and its associated option value changes discretion from a dominated policy to one preferred in a majority of states.

Another perspective is the "delay probability," or the expected time until a commitment is made. For example, if we interpret each decision as a yearly meeting date, the probability that the policymaker will go five years without committing to rules is $(1-0.699)^5 = 0.0025$. The independent nature of the shocks in this example means that even though commitment is chosen in fewer than half the states, the probability of ending up in those states at least once increases rapidly. In other words, we hit the "absorbing barrier" quickly.

⁶ The most straightforward way to produce figure 1 is to use equation (23). We instead used a more general approach described in the next section.

There are really two vantage points on these numbers. One stresses the large number of states where the government prefers discretion. The other stresses the short horizon until commitment. A model with serial correlation would tend to reconcile the two vantages, because it would keep the economy in states with discretion for a longer time. Without having incorporated this into the formal model, though we do not hold obdurately to this point.

This model, too, illustrates the bad-news principle. A mean-preserving spread makes discretion preferable in more states. Increasing the variance of the distribution by 10 percent reduces the commitment region to the range of -0.957 to $+0.957$, so that rules are adopted only 64.6 percent of the time. The probability of delaying for five years rises to 0.0055.

IV. Entering and Exiting Commitment

The obvious impossibility of inescapable commitment (recently emphasized by McCallum, 1995) calls for a sophisticated approach to modeling commitment, not an abandonment of the insights generated by the time-inconsistency literature. The model so far has allowed only simple, inescapable commitment. We now generalize this, allowing the policymaker to enter and exit commitment (or, more generally, any policy regime) at a cost.

Mechanisms forcing governments to commit irrevocably are almost impossible to imagine. It is not difficult, however, to think of mechanisms that make it costly for a government to alter its policy. A constitutional amendment, for example, is difficult to put into place and difficult to repeal. Ordinary legislation has lower costs at both

ends. Governments can, in effect, tie their hands loosely or tightly, but can always escape if they have the will to bear the corresponding levels of pain.

It is important to understand that entry and exit do not destroy the possibility of commitment. Once the policymaker commits in time period t , the rule is in effect for that time period at least. Another decision is made at $t + 1$. Likewise for discretion. In this discrete time framework, we don't allow shifts in midstream, between FOMC meetings, or when Congress is out of session. Thus, commitment can tie the hands and reduce the possible choices of the policymaker long enough to influence the public's expectations.⁷

We maintain the traditional semantics of commitment and discretion, but we wish to highlight a bias in tone that creeps into the discussion when commitment is not irrevocable. This innovation forces us to words like “renege” and “weasel,” although they have clear negative connotations that we consider unfortunate. We interpret the results of this section as a model of optimal behavior and tolerate the terminology only to fit our paper into the literature on rules and discretion. The terminology does have one advantage, though: The emotional evocation reminds us of strategic aspects of the problem that might otherwise get overlooked in the formalism.

Thinking about the problem as entering or exiting commitment deepens the analogy to irreversible investment. Our extended model now resembles an extension of the irreversible-investment model, namely, Dixit's [1989] model of firm entry and exit. For many of these questions, the continuous-time approach set forth in Dixit and

⁷ This does not exclude the possibility that the commitment cost may be some sort of bond posted for credibility. But the foregoing analysis does say that if another commitment technology exists that does not require such a costly bond, it benefits the economy.

Pindyck (1994) generally proves more convenient. Rigorously formulating questions of time inconsistency, however, brings up serious difficulties in the theory of stochastic differential games. This is particularly true of the monetary-policy question, because the unanticipated-money model does not easily generalize to continuous time. Fortunately, the discrete-time approach, though less elegant, suffices for many important problems. In this we follow Lambson [1992], who used it to model entry-exit decisions.

A. Model Solution

We modify the model of section III by adding costs for entering and exiting commitment. A policymaker committing to rules in period t pays a cost C . Once committed, a policymaker may on renege -- or “weasel out” of -- rules and return to discretion by paying cost W . The problem becomes finding the boundaries where the policymaker switches between discretion and rules. The model produces *four* boundaries: an upper and a lower boundary for moving from rules to discretion, and an upper and a lower boundary for moving from discretion to rules. With i.i.d shocks, the zero mean of the u shocks and the quadratic loss function conspire to produce a rules region centered on zero. As before, small shocks imply that the government chooses (or stays with) rules, while large shocks imply it chooses (or stays with) discretion.

The optimal policy switches between the two quadratic loss functions with cost C of committing to rules, that is, of moving from discretion to rules, and cost W (for weaseling) of moving from rules to discretion.⁸

To solve the infinite-horizon model with switching costs, we use a discrete state-space approach. The shock u_t is a Markov chain with n states i . The probability of transition to state j from state i is

$$g_{ij} = \frac{f_j}{\sum_{k=1}^n f_k}$$

where

f_i is the normal density function with mean 0 and variance σ^2 . This produces a Markov chain that is similar to white noise with normal innovations. We set the range of possible states to include 6σ on each side of 0.

The policymaker is faced with a problem that has two state variables, u_t and current value--rules of discretion--of a policy variable. Thus the value function for this problem is an $n \times 2$ matrix where the columns correspond to rules and discretion.

Denote the columns by V^R and V^D . To solve the model, we choose an initial value function and iterate on the following mappings:

$$V^D(i) = \min_{C, NC} \{L^D(i) + E[V^D(\cdot) | i], L^R(i) + C + E[V^R(\cdot) | i]\}$$

⁸ Allowing weaseling adds a component similar to the “escape clause” models of Flood and Isard (1989) and Lohmann (1992), who consider a cost to renege. In one sense, we generalize those models by allowing a positive cost of recommitment and allowing delay in recommitment. In another sense, those models are more general in that they use more general state-contingent rules. Such rules can be embedded in our dynamic framework. We consider simple rules in order to focus on the dynamics.

$$V^R(i) = \min_{W, NW} \{L^R(i) + E[V^R(\cdot) | i], L^R(i) + W + E[V^D(\cdot) | i]\}.$$

Since the distribution of u is discrete, the following rules determine the regime switching points. The ~~upper~~ ~~lower~~ commitment boundary, \bar{c} , is the largest

$$V^R(\bar{c}) + C \leq V^D(\bar{c}).$$

The ~~lower~~ ~~upper~~ commitment boundary, \underline{c} , is the smallest

$$V^R(\underline{c}) + C \leq V^D(\underline{c}).$$

such that the upper weasel boundary is the smallest

$$V^D(\bar{w}) + W \leq V^R(\bar{w}).$$

such that the lower weasel boundary is the largest

$$V^D(\underline{w}) + W \leq V^R(\underline{w}).$$

To solve the model in the case of irrevocable commitment at zero cost, we set C to zero and W to an extremely large number.

B. Regime Switching

The actual numeric solutions are less interesting than the comparative statics.

Starting from a baseline of: $K = 1, a = 2,$

$\beta = 1, C = W = 1, \rho = 0.95,$ figures 2 to 5 depict

the solutions as we vary parameters one at a time. The state space has 401 nodes

evenly distributed from -6^2 to $+6^2$.

Figure 2 plots the commitment and weasel thresholds as the commitment cost changes, keeping the weasel cost fixed at 1. Notice that for any particular commitment

cost, the Fed adopts rules for “small” shocks on either side of zero, as it did in the discrete time model. For larger shocks the Fed adopts discretion. This is a natural consequence of the quadratic loss function.

The probability of being outside the area where rules are better for the current period does not change as C increases. Thus, as the cost of committing to rules increases, the range over which the policymaker is willing to commit shrinks. It will disappear altogether if C is high enough; the option to commit is worthless if its exercise price is too high. Thus positive commitment cost destroys the result that commitment will happen in finite time with probability one.

Another prominent feature is that the weasel boundary is further out than the commit boundary. Were there no cost of switching between regimes, the boundaries would be the same, at $L^D(u) = L^R(u)$, where the expected loss from continuing discretion just matches the expected loss from using rules. Adding a commitment cost drives a wedge between the two value functions, and requires that the policymaker gain *even more* from rules. This means moving the boundary further into the area where rules are better, i.e., closer to zero. Similarly, a cost to backing out of rules (weaseling your way out) means shifting the boundary even further into the area where discretion is preferred, i.e., away from zero.

Figure 2 shows that as the cost of commitment increases, the Fed is less likely to commit. As the cost increases, the relative benefits of rules over discretion must also increase, and so the commitment boundary shrinks towards zero. For high enough cost, commitment never occurs.

Figure 3 highlights a key point not easily noticed in section III's model; namely, the importance of history. Because the weasel and commit boundaries differ, in some states of the economy (levels of u) current policy depends on past policy. For anything above the upper commit line and below the upper weasel line, a policymaker committed to rules sticks with rules, and a policymaker using discretion sticks with discretion. Quite apparently, then, it is incorrect to judge policy simply on the current state of the economy, and particularly inappropriate to naively contrast current policy with past policies at a similar state of the economy or stage of the business cycle. In a word, our model predicts hysteresis in monetary policy.

Implicit in the hysteresis is something so obvious as to possibly escape notice: Over time, the policymaker switches from rules to discretion, and from discretion to rules. Regimes shift. Discretion, commitment, and weaseling out of commitment will all occur. As an example, figure 3 shows one path for independent shocks, the commitment and weasel boundaries, and shades the time spent committed to rules. The figure makes it easy to see both the historical dependence and the switches between rules and discretion. This shifting reemphasizes a point stressed by Flood and Garber (1984) in their work on the gold standard: To evaluate a policy rule, the entire dynamic policy sequence must be analyzed, including those periods where discretion reigns.

Not surprisingly, as W increases, the weasel boundaries move out (see figure 4). It is somewhat more surprising that the commitment boundaries are insensitive to W . This is because eventually crossing a weasel boundary is a low-probability event and therefore has little impact on the decision to commit. The commitment boundaries

are completely flat only because the state space is not fine enough; the commitment boundaries narrow slightly with a very fine state space.

Increasing the variance of the shocks σ^2 causes the policymaker to narrow the commitment ranges because it lowers the probability that later periods' shocks will fall in the range where the rules loss is less than the discretion loss, and thus increases the value of the waiting option. Two things can happen (depending on parameter values) when the variance gets large: 1) Rules may suddenly become inferior to discretion forever, meaning the commitment range cuts off and drops to zero or 2) the commitment range gradually disappears. We show the latter case in figure 5. The weasel boundaries tend to be relatively insensitive to changes in the variance, mostly because renegeing is a low-probability event.

At first, it seems that a ought to increase the likelihood of commitment, but, as mentioned above, increasing a decreases the inflation bias. This effect dominates the direct influence of increased desire to avoid inflation, so that commitment never occurs if a is large enough. For some values of K and σ^2 , discretion forever may be strictly preferred to rules.

C. Beyond Monetary Policy

We have noted, the points made here apply generally to questions of time inconsistency, not just the particular class of Barro-Gordon models. Bordo and Kydland (1992), for example, interpret the gold standard as a rule containing contingencies in case of wars and financial panics. Even with such contingencies, they recognize the possibility of regret, because a fully contingent rule would create “a lack of transparency and possible uncertainty among the public regarding the will to

obey the original plan” (p. 8). One advantage of a simple rule like Bordo and Kydland's interpretation of the gold standard is that the contingencies---wars and financial panics---are readily verified. This makes credible commitment easier. A complicated rule may lose some of the benefits of commitment because it is more costly to verify the government's compliance.

In our framework, the gold standard can have two slightly different interpretations. It may be seen as an imperfectly state-contingent rule that has been abandoned in favor of discretion since the advent of the Bretton Woods system World War II. Alternatively, because the gold standard did not bind government's hands in times of war, these could be seen as times when the government abandoned the rule in favor of discretion, returning to rules at a later time. Our own view is that the lack of wartime constraints points more to abandonment of a standard, and thus to the sort of entry and exit considerations we have analyzed in this paper.

More generally, the tractability of the quadratic loss model makes it a natural approximation for many time inconsistency problems (along with many other economic problems as well) Thus, additional examples like restraining the lender of last resort from bailing out insolvent institutions (with regret in a true financial crisis), granting patents for the exclusive use of new technology (with regret in cases such as AZT), or allowing constitutions to bind future legislatures could illustrate of our main point.

For many applications, a continuous-time approach is more powerful, particularly when the dynamic game of policy choice takes a simple form. We explore these issues more deeply in a companion paper (Ritter and Haubrich, 1995).

V. Conclusion

Sometimes the right answer is inherent in the right question. The standard analysis of the choice between rules and discretion has not asked the right question. This failure may underlie the frustration felt on both sides of the of issue, by both the starry-eyed theorists and hard-nosed practitioners, who have mostly talked past each other. The decision regarding rules versus discretion occurs in real time, not at some mythical starting date. That means that, because opting for discretion today leaves open the possibility of adopting rules later on, it is often the better choice. Previous work, by ignoring this option, has ignored an important advantage of discretion.

Like other options, the option to wait increases in value as uncertainty increases--and so the value of discretion increases as well. Policy, then, has a “bad-news principle” because the ability to avoid regret leads us to wait: Only news about increased regret matters for the policy choice. But while the option-value results may explain delay and refusal to adopt simple monetary targets or tax reforms during recessions or wars, they do not generally justify permanently abandoning such rules. Eventually, when the time is right, the government should commit--at least for a while.

When commitment to rules is no longer an irrevocable choice made at the beginning of time, optimal policy looks more dynamic. Periods of rules alternate with periods of discretion, depending both on the state and the history of the economy . Policy at a given point in the business cycle may look quite different from policy at a similar point in an earlier cycle. Such seeming confusion nevertheless reflects a coherent, optimal choice.

In principle, the notion of commitment as irreversible investment can be applied to other areas like tariff agreements, deficit reduction, and tort reform. In this sense, our work complements recent studies on the political economy of resistance to reforms (Fernandez and Rodrick [1991]), as well as on the delay in their implementation (Alesina and Drazen [1991]). Our approach emphasizes delay and resistance as an optimal response to an uncertain future. It also suggests the possibility of hysteresis resulting from

Our findings are by no means the last word on the rules-versus-discretion debate. We hope that by clarifying some neglected issues--regret, future commitment, and the bad-news principle--we will contribute to clearer insight and a more focused dialogue.

REFERENCES

- Alesina, Alberto, and Allan Drazen. "Why Are Stabilizations Delayed?" *American Economic Review*, vol. 81, no. 5 (December 1991), pp. 1170-88.
- Barro, Robert J. *Macroeconomics*, 2d ed. Toronto: John Wiley & Sons, 1987.
- Barro, Robert J., and David Gordon. "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, vol. 81 (1983), pp. 589-610.
- Bernanke, Ben S., "Irreversibility, Uncertainty, and Cyclical Investment," *Quarterly Journal of Economics* (February 1983), pp. 85-106.
- Bordo, Michael D. and Finn E. Kydland, "The Gold Standard as a Rule," Federal Reserve Bank of Cleveland Working Paper No. 9205, March 1992.
- Canzoneri, Matthew B. "Monetary Policy Games and the Role of Private Information," *American Economic Review*, vol. 75, no. 5 (December 1985), pp. 1056-70.
- Cukierman, Alex, and Allan H. Meltzer. "A Positive Theory of Discretionary Policy, the Cost of Democratic Government, and the Benefits of a Constitution," *Economic Inquiry*, vol. 24 (1986), pp. 367-88.
- Dixit, Avinash K. "Entry and Exit Decisions under Uncertainty," *Journal of Political Economy*, vol. 97, number 3 (June 1989), pp. 620-38.
- Dixit, Avinash K., and Robert S. Pindyck, *Investment under Uncertainty*. Princeton, N.J.: Princeton University Press, 1994.
- Fernandez, Raquel, and Dani Rodrick. "Resistance to Reform: Status Quo Bias in the Presence of Individual-specific Uncertainty," *American Economic Review*, vol. 81, no. 5 (December 1991), pp. 1146-55.
- Flood, Robert, and Peter Isard. "Monetary Policy Strategies," *IMF Staff Papers* vol. 36, no. 3 (September 1989), pp. 612-32.
- _____, and Peter Garber, "Gold Monetization and Gold Discipline," *Journal of Political Economy* vol. 92.1 (February 1984), pp. 90-107.
- Kydland, Finn E., and Edward Prescott. "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, vol. 85 (1977), pp. 473-91.

Lambson, Val Eugene, "Competitive Profits in the Long Run," *Review of Economic Studies*, vol. 59 (1992), pp. 125-42.

Lohmann, Susanne. "Optimal Commitment in Monetary Policy: Credibility versus Flexibility," *American Economic Review*, vol. 82, no. 1 (March 1992), pp. 273-86.

McCallum, Bennet T. "Two Fallacies Concerning Central-Bank Independence," *American Economic Review*, vol. 85, no. 2 (May 1995), pp. 207-11.

McDonald, Robert and Daniel R. Siegel, "The Value of Waiting to Invest," *Quarterly Journal of Economics*, vol. 101 (November. 1986), pp. 707-27.

Pindyck, Robert S. "Irreversibility, Uncertainty, and Investment," *Journal of Economic Literature*, vol. 29 (1991), pp. 1110-48.

Ritter, Joseph A. and Joseph G. Haubrich, "Commitment as Investment under Uncertainty," manuscript, 1995.

Roberds, William. "Models of Policy under Stochastic Replanning," *International Economic Review*, vol. 28, no. 3 (October 1987), pp. 731-55.

Ross, Sheldon M. *Introduction to Stochastic Dynamic Programming*. New York: Academic Press, 1983.