Working Paper 9217

COMMITMENT AS IRREVERSIBLE INVESTMENT

by Joseph G. Haubrich and Joseph A. Ritter

December 1992

Abstract

Considering time inconsistency as a problem of irreversible investment brings some neglected points to the fore. Making a policy choice in real time and under current conditions emphasizes the importance of the timing of commitment, the regret over past decisions, and the option value of not committing. This paper applies these concepts to monetary policy, banking regulation, and capital taxation.

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 in certain contexts. From this point of view, rules allow policymakers to achieve outcomes otherwise precluded by the strategic behavior of the public. This theoretical emphasis contrasts with earlier arguments for rules based on practical considerations. According to the early monetarists, simple monetary rules were necessary precisely because the central bank was incapable of handling the economy's complexities well enough to make discretionary policy desirable.[1]

One aspect of this shift that has been overlooked to a large extent is the change in the nature of the rules under consideration. The pragmatic case for rules almost by definition requires that they be simple and easily implemented. The theoretical case, based on the time inconsistency of discretionary policies, presumes fully state-contingent rules. Although sometimes simple (see Barro and Gordon [1983]), the optimal state-contingent rules are generally rather complex.

Complex rules make commitment more difficult. Governments can make it costly (to themselves) to change a rule, and this may overcome their incentive to retract it or to change course in midstream. It is manifestly true, however, that governments' commitment mechanisms (ranging from campaign promises to constitutions) cannot be contingent on all possible states of the world. Policymakers must choose not between discretion and optimal state-contingent rules, but between discretion and comparatively simple and imperfect rules (as recently emphasized by Lohmann [1992]). Thus, it is logically possible for policymakers to "regret" their commitment to a rule.

In this paper, we reframe the rules-versus-discretion question along these lines and explore the consequences of the change in perspective. We find that this modification makes the rules-versus-discretion decision similar to the choice to make an irreversible investment (see Pindyck [1991]). Pursuing this analogy further raises the possibility that the government might want to delay committing to a rule, with the outcome of the

1

decision dependent on the current state of the world. In this sense, it becomes important that the decision be made in "real time." A broader implication, not recognized in previous literature, is that choosing discretion today has an option value, since the government may still choose rules in the future. If this option indeed has positive value -- as such options often do -- it adds to the desirability of discretion.

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 (1988). Cukierman and Meltzer discuss flexibility, but do not consider imperfect fixed rules and a fortiori miss the associated option value.

The remainder of this paper develops these themes, with several variations. Section II presents an expository example couched in terms of the too-big-to-fail doctrine. Section III presents a richer analysis in a framework of monetary policy based on Flood and Isard. Section IV uses some detailed numerical examples to explore the significance of the results, and section V applies these ideas to capital taxation. Section VI concludes.

II. A Preliminary Example: Too Big To Fail

In this section, we present a concrete example from the banking industry showing the option value of waiting, the "bad news principle," and the necessity of regret. By bringing banking into the analysis, we are able to link the discussion to policy and (we hope) to shed some light on current disputes within the industry.

Our model has two periods, 1 and 2, with no discounting between them, and two equally likely states of the world, good (G) and bad (B). After waking up on a particular date and in one of these states, a bank regulator must decide what to do when banks fail. He may choose to be weak (W) and bail them out, or he can be tough (T) and commit to refusing them.

In the good state, we prefer to have a tough regulator who eliminates the costly wealth transfers from the rest of society to bank investors. In the bad state, we prefer the

2

weak regulator. Think of the bad state as one with systemic risk (perhaps a recession) where being tough leads to a financial panic. We can express this in payoffs, or utility levels, for the regulator: In the good state, TG>WG; in the bad state, TB<WB. Figure 1 illustrates this game's structure.

Since we have suppressed any explicit choice by the banks, the above example is formally identical to an irreversible investment problem. This highlights the connections between the two cases, while the rules-versus-discretion language foreshadows the more important, complicated models in sections III and IV.

Because the policy decision takes place in real time, we have two cases to consider. Either the economy starts out in the good state, or it starts out in the bad state. Suppose the economy starts out in the good state. If the regulator is weak, he gets a payoff of WG today and chooses to be tough or weak again in period 2. If the regulator is tough, he gets a payoff of TG today and remains tough in period 2. The payoffs then look as follows:

Case I: Start in Good State

$$\text{Strategy W: WG} + 1/2(\text{TG+WB})$$
$$\text{T: TG} + 1/2(\text{TG+TB}).$$

The regulator picks the strategy with the higher expected return,

$$\text{W-T} = \text{WG-TG} + 1/2(\text{WB-TB}).$$

In the standard time-consistency literature, failure to commit means that the suboptimal, or "weak," choice is made in each period. In this simple example, we do not model that choice, but the natural analogue is that W is chosen in each period. Making such a decision forever seems simple-minded in this simple model, but it is precisely the

3

choice forced on regulators in most of the current time-consistency literature. The basic game described here brings out rather sharply the problem of committing in real time.

If the regulator decides at the beginning to adopt either rules or discretion forever (TF or WF), the new payoff for choosing WF is

$$WF = WG + 1/2(WG+WB).$$

Note that the payoff to choosing W today and retaining the option to choose T tomorrow exceeds the payoff to choosing W in both periods. This is because W-WF=1/2(TG-WG), and since TG>WG by assumption, W-WF>0. Similar calculations hold in the bad state:

Case II: Start in Bad State

Strategy W: WB + 1/2(TG+WB)

T: TB + 1/2(TG+TB)

WF: WB + 1/2(WG+WB)

W-WF= 1/2(TG-WG)>0.

Thus, the waiting option has value, something not accounted for by studies that ignore the real-time aspect of choosing rules over discretion. A short numerical example shows that this cost can be high enough to make discretion the optimal strategy. In the good state, the comparison is between

$$W-T=WG-TG + 1/2(WB-TB)$$

and

$$WF-T=3/2(WG-TG) + 1/2(WB-TB).$$

Setting TG=6.2, WG=4.2, TB=0, and WB=4.2 yields WF-T= -0.9, which means that adopting rules is the better choice. When option value is considered, however, we find that W-T=0.1, making discretion preferable.[2]

4

When future prospects change, this example illustrates the bad news principle, which we suspect lies behind the rhetorical tendency to accentuate the negative. To see this, we first distinguish payoffs with period subscripts. Now, in the good state,

$$W = WG_1 + 1/2(TG_2 + WB_2),$$

$$T = TG_1 + 1/2(TG_2 + TB_2),$$

and thus

$$W - T = [WG_1 - TG_1] + 1/2(WB_2 - TB_2).$$

Holding the period 1 payoffs (the bracketed term) fixed, because we are already in that period, consider changing the payoffs in period 2. An increase in TG_2 or WG_2 has no effect. What matters is the "regret spread," or $WB_2 - TB_2$. If we end up in the bad state, we'd like to choose W and get payoff WB_2, but if we have already committed, we play T and get TB_2. If the good state ensues and we want to commit, we can. Shifts in WG_2 and TG_2 do not matter when deciding between strategies.

The key here is that we sometimes regret the commitment to be tough. We never regret an initial decision to be weak, because if it pays to be strong later on, we can make that choice. Increasing the payoff to toughness does not affect the relative payoffs -- and thus the choice -- today. This illustrates the bad news principle: Only news about the bad outcomes affects the value of the option to wait.

One could make a timeless comparison, contrasting expected values before even the initial state of the world is realized, but the results would be irrelevant. Any current choice must take place in real time, since the economy already exists.

III. Monetary Policy

Most of the debate about rules versus discretion has taken place in the arena of monetary economics. The insights from thinking about policy commitment as irreversible

5

investment apply here as well. The central bank must make decisions in real time, and a failure to commit today does not preclude commitment in the future. This section models commitment as irreversible investment in a monetary policy context. Though slightly specialized to highlight the main points, the model derives from a fairly general and plausible framework based on Flood and Isard (1988).

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:

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

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:

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

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 measures the production shock. The term 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.

6

B. Discretion

From the first-order conditions for equation (1), we find

$$(3) \qquad b_t = (1/1+a)(E_{t-1}b_t + K - u_t).$$

This implies

$$(4) \qquad E_{t-1}\pi_t = E_{t-1}b_t = K/a.$$

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

$$(5) \qquad b_t^D = -(u_t/1+a) + K/a.$$

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

$$(6) \qquad \text{Realized Loss:} \quad L_t^D = (1+a/a)[-K + (a/1+a)u_t]^2.$$

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

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 (assuming $Eu_t = 0$), some of which shows up in the inflation rate via monetary policy.

7

C. Rules

If the money supply, b_t, cannot respond to u_t, money only causes inflation; it cannot reduce output variance. In that case, the best rule sets $b_t = 0$ in all periods. This is the optimal rule without state contingency. If it were feasible, a better rule 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.

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

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

Equations (8) and (9) have a lower inflation bias than discretion, but a higher output variance.

When is $L^D - L^R < 0$? Straightforward substitution from equations (6) and (8) shows that this is the case when

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

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

8

Similarly, as K increases, discretion is preferable in more states. As distortion worsens, there is a larger deviation of output from its natural level and thus more of an advantage to reducing shocks.

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):

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

where g_i denotes the probability of state i and b_i denotes money growth in state i.[3] The first-order conditions $\frac{\partial EL}{\partial b_i} = 0$ imply

$$\left(b_i - \sum_{j=1}^{n} g_j b_j - K + u_i \right)(1 - g_i) + a b_i + \sum_{k \neq i}^{n} \left(b_k - \sum_{j=1}^{n} g_j b_j - K + u_k \right)(-g_k) = 0,$$

or

$$\left(b_i - \sum_{j=1}^{n} g_j b_j - K + u_i \right) + a b_i + \sum_{k=1}^{n} \left[\left(b_k - \sum_{j=1}^{n} g_j b_j - K + u_k \right)(-g_k) \right] = 0.$$

Taking expected values, we get

$$\sum_{i=1}^{n} \left(b_i - \sum_{j=1}^{n} g_j b_j - K + u_i \right) g_i + a \sum_{i=1}^{n} b_i g_i$$

$$+ \sum_{i=1}^{n} g_i \sum_{k=1}^{n} \left(b_k - \sum_{j=1}^{n} g_j b_j - K + u_k \right)(-g_k) = 0,$$

9

or

$$a \sum_{i=1}^{n} b_i g_i = 0 \Rightarrow E b_i = 0.$$

This means that the optimal rule has no inflation bias. Setting Eb_i equal to 0 in the first-order conditions gives

$$(b_i - K + u_i) + ab_i + \sum_{k=1}^{n} [b_k - K + u_k](-g_k) = 0,$$

which implies

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

Substituting into the loss function, we have

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

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

To understand the implications of restricting rules to a subset of those that are fully state contingent, it is important to look at several relationships. First, 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 by assumption the class of linear rules that do not allow b_t to respond to u_t.

Second, the optimal rule can replicate the behavior of the policymaker under discretion, so it must be at least as good as discretion *in every state*. Since the optimal

10

rule ties the policymaker's hands, it allows actions that would otherwise fall victim to time inconsistency. This again makes the optimal rule better by construction.

Third, while the optimal rule dominates the simple rule on average, this is not true in every state, as can be seen by comparing equations (12) and (8). In states where $0 < u_t < 2K$, the response to u_t is not worth the (small) amount of inflation that ensues. However, a rule that attempts to exploit this inefficient response changes expectations in a way that hurts more than it helps on average. Suppose, for example, that we attempt to revise the optimal rule by setting $b_t = 0$ whenever $0 < u_t < 2K$. This lowers expected inflation and increases the loss in states where state contingency is useful. The gain in $0 < u_t < 2K$ states is offset by the loss in other states, even though policy is unchanged in the latter. The response of individual behavior, in this case expectations, distinguishes an equilibrium problem from a simple control problem.

D. Many Periods

Adequately capturing irreversibility requires a number of adjustments to the model. First, it clearly needs several periods. 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:

Gov't sees u_t → Gov't decides whether to commit, announces → Economy revises expectations $E_{t-1}b_{it}$ → Gov't chooses b_t → Production; economy sees u_t.

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. 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

11

the economy. In our specific model, however, symmetry of the shocks means 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

$$E(u_t \mid \text{gov't choice}) = 0 \text{ and } E(b_t \mid \text{gov't choice}) = 0.$$

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.

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.

We wish to illustrate two points regarding this framework. First, if the government can use the optimal state-contingent rule, there is no value to waiting, and the irreversibility makes no difference. Second, with a simple, non-state-contingent rule, irreversibility introduces an option value whose worth is non-negative.

The first point is readily apparent. With the optimal state-contingent rule, no regret occurs. In fact, the loss from discretion *exceeds* the loss from choosing the rule, so it is never worthwhile to wait. The key here is the state contingency of the rule, which allows enough flexibility to offset future shocks but eliminates the inflationary bias inherent in pure discretion. Committing to a long-term rule still allows flexibility in day-to-day decisionmaking. If such a contingent rule is politically and administratively feasible, no conflict arises between maintaining stable prices and responding to economic shocks.

With a simple, non-state-contingent rule, regret can exist. For example, the government might regret committing to zero inflation and wish for discretion. This point

12

does not depend merely on the rule's simplicity. The analysis holds even with a more sophisticated, less than fully state-contingent rule, as long as there are some states in which discretion is preferred.

The problem 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 new value function

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

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

(14) a) $V_R(u_t) = (-K + u_t)^2 + \beta E\, V_R(u_t + 1)$

  b) $V_R(u_t) = (u_t - K)^2 + \dfrac{\beta}{1-\beta}(K^2 + \sigma_u^2),$

where $V_R(u)$ denotes the value function for rules.

The first term measures today's loss, and the second gives the expected loss tomorrow. Choosing discretion forever yields a loss of

(15) a) $V_{DF}(u_t) = \left(\dfrac{1+a}{a}\right)\left(-K + \dfrac{a}{1+a}u_t\right)^2 + \beta E V_{DF}(u_{t+1})$

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

The general case is more complicated because opting for discretion today leaves the door open for choosing rules tomorrow. The loss to choosing discretion is

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

13

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

$$(17) \quad EV_D(u_t) = \sum_{s=1}^{\infty} \beta^s \left[\sum_{j \in UR_{t+s}} P_j(u_j - K)^2 + \sum_{j \in UR_{t+s}} P_j \left[\frac{1+a}{a} \left(-K + \frac{1}{1+a} u_j \right)^2 \right] \right],$$

where P_j is the probability of state j and UR_{t+s} is 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, CR, does not vary with time. (This differs from UR_t in equation [17], 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 (16). Then, recursively using equation 14(b) yields

$$(18) \quad EV_D(u_t) = \sum_{j \in CR} P_j \left[(u_j - K)^2 + \frac{\beta}{1-\beta}(K^2 + \sigma_u^2) \right] + \sum_{j \notin CR} P_j \left[\frac{1+a}{a} \left(-K + \frac{a}{1+a} u_j \right)^2 \right]$$
$$+ \sum_{j \notin CR} P_j \beta \left\{ \sum_{j \in CR} P_j \left[(u_j - K)^2 + \frac{\beta}{1-\beta}(K^2 + \sigma_u^2) \right] \right\} + \sum_{j \notin CR} P_j \left\{ \beta \left[\sum_{j \notin CR} P_j \frac{1+a}{a} \left(-K + \frac{a}{1+a} u_j \right)^2 \right] + ... \right\}.$$

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 the period after discretion. The fourth term gives the loss from choosing discretion again, with this pattern repeating recursively.

Equation (18) simplifies to

14

$$(19) \quad EV_D(u_t) = \frac{1}{1-\beta\left(\sum_{j\in CR} P_j\right)} \left\{ \sum_{j\in CR} P_j\left[(u_j - K)^2 + \frac{\beta}{1-\beta}(K^2 + \sigma_u^2)\right] \right.$$

$$\left. + \sum_{j\notin CR} P_j\left[\frac{1+a}{a}(-K + \frac{a}{1+a}u_j)^2\right]\right\}.$$

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 makes discretion today preferable, noting the importance of regret in that decision. A big increase in the attractiveness of discretion means that the government chooses 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 in every state the loss from rules is less than the loss from discretion, then it makes no sense to delay commitment or to choose discretion -- hence we rule out the optimal state-contingent policy. 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 $\left(K^2 > \sigma_u^2 < \frac{1+a}{a}K^2 + \frac{a}{1+a}\sigma_u^2\right)$, or

$$(20) \quad K^2 > \frac{a}{1+a}\sigma_u^2.$$

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

15

(21) Rules: $V_R(u_o) = (u_o - K)^2 + \dfrac{\beta}{1-\beta}\left(K^2 + \sigma_u^2\right)$

and

(22) Discretion: $V_D(u_o) = \dfrac{1+a}{a}\left(-K + \dfrac{1}{1+a}\,u_o\right)^2 + \beta EV(u_1),$

where $EV_R(u_1)$ is equation (17) or, equivalently, (19). 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 [15]).

Comparing $V_R(u)$ with $V_D(u)$ and $V_{DF}(u)$ shows the option value of discretion. Consider moving from (21) to a version of (19) in which the choice to commit is based solely on what's best this period (equation [6] versus [8]). This may not be the optimal policy, but it is certainly feasible. Such a policy differs from (21), but at each point of difference (state-time pair), the policy leading to (17) results in the smaller loss. This happens because a discretion term enters (17) only if it is lower than the corresponding rules term. A similar comparison with (15b) emphasizes the same point from a different perspective. Compared with discretion forever, (19) and (22) allow rules when rules are better. This is why we can have $EV_D(u_t) < \dfrac{\beta}{1-\beta}\left(K^2 + \sigma_u^2\right) = EV_R(u_t)$, even if

$EV_{DF} > EV_R$; that is, when $\dfrac{\beta}{1-\beta}\left(\dfrac{1+a}{a}K^2 + \dfrac{a}{1+a}\sigma_u^2\right) > \dfrac{\beta}{1-\beta}K^2 + \sigma_u^2.$

Note that since $EV_D(u_t) < \dfrac{1}{1-\beta}EV_R(u)$, the government may sometimes choose discretion even in states where the one-period return favors rules. This conceivably could create a paradox; that is, even though we prefer pure rules to pure discretion, we delay choosing rules forever. Actually, this never occurs, as demonstrated by equation (19). Suppose the government never commits, so that $CR = \varnothing$. Then (19) reduces to $\dfrac{1}{1-\beta}\left(\dfrac{1+a}{a}K^2 + \dfrac{1}{1+a}\sigma_u^2\right) = V_{DF}(u_t)$, the discretion-forever case. We assume, however, that discretion forever is worse than rules.

16

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).[4]

IV. Numerical Examples

To further illustrate our points, this section presents two numerical examples. The two-period case clarifies the notions of regret, delay, and option value, while the infinite-horizon example explores the quantitative importance of our results. While it cannot be called a test, nor even a calibration exercise, our approach strives to use plausible values for the effect of unanticipated money and the distribution of unemployment shocks. In this scenario, the government chooses discretion in about half the states.

A. Two-Period Example

This example also simplifies the productivity shock, assuming u_t is i.i.d. and equals $\{-x, 0, +x\}$. The probability takes a correspondingly simple form:

$$\text{Prob}(u_t = +x) = g_1$$

$$\text{Prob}(u_t = 0) = g_2$$

$$\text{Prob}(u_t = -x) = g_3.$$

Figure 2 illustrates the sequential probability structure.

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

$$(24) \quad V_R(o) = K^2 + g_1(x^2 - 2Kx + K^2) + g_3(x^2 + 2Kx + K^2) + g_2 K^2.$$

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

17

If the Chairman and the government choose discretion today, the loss is

$$(25) \quad V_D(o) = \frac{1+a}{a}K^2 + g_1 \left(\frac{a}{1+a}x^2 - 2Kx + \frac{1+a}{a}K^2 \right)$$

$$+ g_3 \left(\frac{a}{1+a}x^2 + 2Kx + \frac{1+a}{a}K^2 \right) + g_2 K^2.$$

Again, the first term is the loss today. The last term, $g_2 K^2$, is what happens if the government commits to rules tomorrow, when $u = 0$. If the government chooses discretion, that term would be $g_2 K^2(1+a/a)$. The difference represents the option value of waiting.

Deciding whether to choose rules or discretion comes down to comparing equation (24) with (25) and then choosing the strategy with the smaller expected loss.

$$(26) \quad V_R(o) - V_D(o) = \frac{-1}{a}K^2 + (g_1 + g_3) \left[\frac{1}{1+a}x^2 - \frac{1}{a}K^2 \right].$$

Removing the option value, that is, forcing the government to make an irrevocable choice between rules and discretion, leads to a different expression:

$$(27) \quad V_R(o) - V_{DF}(o) = \frac{-1}{a}K^2 + (g_1 + g_3) \left[\frac{1}{1+a}x^2 - \frac{1}{a}K^2 \right] - g_2 \frac{1}{a}K^2.$$

Conceivably, $V_R(o) - V_D(o) > 0$ and $V_R(o) - V_{DF}(o) < 0$, meaning that correctly valuing the option translates into choosing discretion, while ignoring it means choosing rules. That is, taking account of the real time aspect of decisionmaking and properly valuing the waiting option can reverse the policy decision. As an example, let

$a = 1$,

$x^2 = 6K^2$, and

$g_1 = g_3 = 1/4$.

18

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

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

This example shows that the option to wait can reverse the normal presumption of the superiority of rules and that the option value may be sizable. Using equation (24), the loss from adopting rules is $5K^2$, making the option difference $1/4K^2$, or 5 percent of the total value.

B. Infinite-Horizon Example

To add a small degree of realism, the next example employs the infinite-horizon model, using parameter values we believe to be plausible.

Monthly unemployment rates from January 1948 to August 1992 range from 2.5 percent to 10.8 percent, with a mean of 5.7 and a median of 5.6. Split into thirds, the mean of each third is 7.5, 5.6, and 3.9 percent. To approximate the distribution, we take the long-run average rate of unemployment to be 5.5 percent and assume a uniform distribution of shocks every tenth of a point between -2.0 and 2.0 percent. This gives a variance of 1.4 for unemployment shocks, somewhat below the actual value of 2.7. We choose a K value of 1.1, indicating that the long-run rate of unemployment differs from the socially optimal rate by 1.1 percentage points.

Following Barro (1987, p. 469), we assume that a 1 percent rise in money above expectations lowers unemployment by 0.6 percentage point. This makes the social loss function $L_t = [\alpha(b_t - E_{t-1}b_t) - K + u_t]^2 + ab^2_t$. Consequently, discretionary policy becomes $b_t = \frac{\alpha}{a}K - \frac{\alpha}{\alpha^2 + a}u_t$.

Adapting the work of section III to these differences allows us to make the computations.

Two more parameter choices will fully specify the problem. Give inflation and unemployment equal values in the social loss function and set a equal to one. Next, set β,

the discount factor, to 0.995. Compounded monthly, this yields an annual discount rate of 6 percent. Alternatively, if the decision period is seen as corresponding to the eight yearly meetings of the Federal Open Market Committee (FOMC), we get an annual discount rate of 4 percent.

Figure 3 shows the results of this example using these parameters. 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 a u shock between -0.9 and +0.9 (that is, for unemployment rates between 4.6 and 6.4 percent), 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. In 22 of the 41 possible states, 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 poor policy to one preferred in a majority of states.

Unemployment undoubtedly has more serial correlation than the i.i.d. structure of this model and example. Still, it is interesting to consider the "delay probability," or the expected time until a commitment is made. For example, if we interpret each decision time as an FOMC meeting date, the probability that the Fed will go a year without committing to rules is $(22/41)^8 = 0.007$. The independent nature of the shocks 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, once we hit the "absorbing barrier," we stay there.

V. Capital Taxation and Other Extensions

In some cases, there is no option value to waiting. The standard capital taxation model is one example of this. The basic difficulty is that once capital is in place, the government has an incentive to tax it at 100 percent, so no one invests. Issues of irreversibility may arise in more complex variants of the model, however. Assume that the government is constrained by an upper limit on its capital taxation rate. Perhaps the limit is fixed, or perhaps it is bounded by the labor tax rate. In any event, with discretion, capital is always taxed at that upper bound. Separate from the bound, the government may want to commit to a lower capital tax rate, perhaps reducing it directly, perhaps offering tax credits or accelerated depreciation. Commitment in this case offers benefits -- lower taxes and higher investment -- but it now also has a cost -- higher labor taxes in a national emergency.

In principle, the notion of commitment as irreversible investment can be applied to other areas, such as tariff agreements, deficit reduction, or tort reform. In this sense, our work complements recent studies focusing on the political economy problems behind the 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.

VI. Conclusion

The decision regarding rules versus discretion takes place in real time, not at some mythical starting date. That means opting for discretion today leaves open the possibility of adopting rules later on, making discretion look like the better choice. The option nature of this course of action also explains the tendency to accentuate the negative in arguing against rules. There is an economic rationale behind the rhetoric, and a concern with the rhetoric provides a springboard for the economics. But while the option-value results may explain the delay and refusal to adopt simple monetary targets or tax reforms

21

during recessions or wars, they do not justify permanently abandoning such rules. Eventually, when the time is right, the government should commit.

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

FOOTNOTES

1. This argument for simple rules represents just one part of the monetarist critique of activist policies.

2. In the bad state, both today's payoff and the waiting option work in the direction of allowing discretion today. Though for some values being tough may still be preferable, when $p=1/2$, the option value cannot reverse the optimal policy choice. Reversal in the good state requires two relations to hold: $WF-T<0$ and $W-T>0$. For $WF<T$ (making an irreversible choice), being tough is better than being weak. This corresponds to the standard rules-versus-discretion dilemma of committing to a rule or forever facing discretion. Here, adopting the rule is better. For $W>T$ (correctly considering the option value), discretion today looks better. Both conditions together imply

   $$1/2(WB-TB) > TG-WG > 1/3(WB-TB).$$

   For the bad state, the first condition for reversal, $WF-T<0$, now implies that

   $$TG-WG > 3(WB-TB),$$

   which violates the conditions for reversal in the good state.

3. Nothing essential depends on using a discrete probability distribution. A continuous distribution would lead to identical results, but would necessitate needlessly cumbersome notation.

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

REFERENCES

Alesina, Alberto, and Allan Drazen. "Why Are Stabilizations Delayed?"
American Economic Review, vol. 81, no. 5 (December 1991), pp. 1170-88.

Auernheimer, L. "The Honest Government's Guide to the Revenue from the Creation of
Money," *Journal of Political Economy*, vol. 82 (1974), pp. 598-606.

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.

Barro, Robert J. *Macroeconomics*, 2d ed. Toronto: John Wiley & Sons, 1987.

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.

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.

Fischer, Stanley. "Dynamic Inconsistency, Cooperation, and the Benevolent Dissembling
Government," *Journal of Economic Dynamics and Control*, vol. 2 (1980), pp. 93-107.

Flood, Robert, and Peter Isard. "Monetary Policy Strategies," National Bureau of
Economic Research, Working Paper No. 2770, 1988.

Friedman, Milton. "The Role of Monetary Policy," *American Economic Review*, vol. 58
(1968), pp. 1-17.

Kydland, Finn, and Edward Prescott. "Rules Rather Than Discretion: The Inconsistency
of Optimal Plans," *Journal of Political Economy*, vol. 85 (1977), pp. 473-91.

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

Modigliani, Franco. "The Monetarist Controversy, or, Should We Forsake Stabilization
Policies?" *American Economic Review*, vol. 67 (1977), pp. 1-19.

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

Prescott, Edward C. "Should Control Theory Be Used for Economic Stabilization?" in Karl Brunner and Allan H. Meltzer, eds., *Optimal Policies, Control Theory, and Technology Exports*. New York: North-Holland, 1977, pp. 13-38.

Ritter, Joseph A., and Joseph G. Haubrich. "Notes on Does Capital Taxation Discourage Investment? Insights from a General Model," unpublished manuscript, 1992.

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.
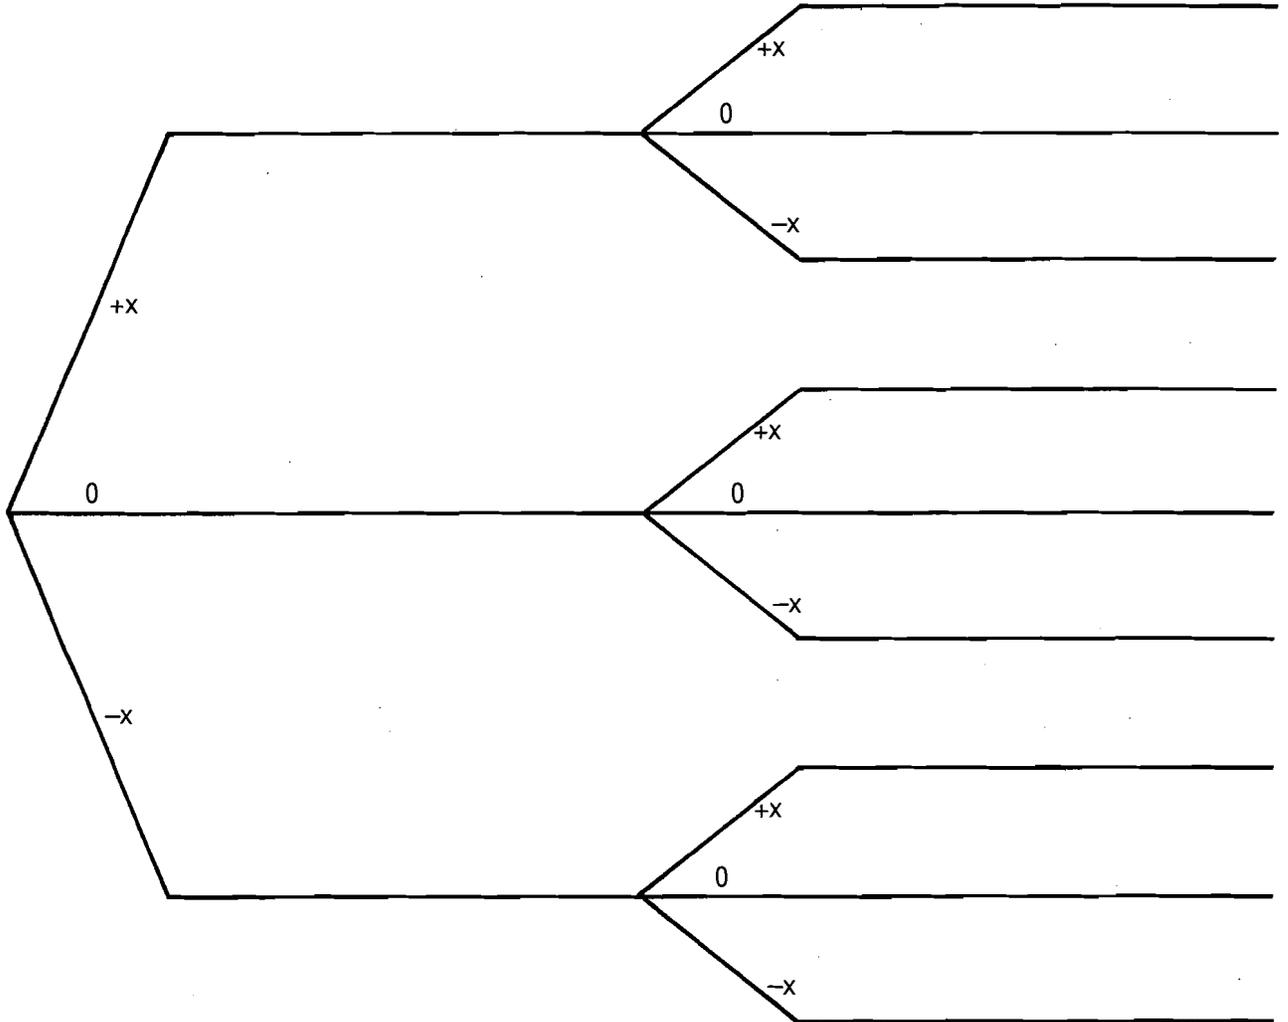
Schelling, Thomas C. "An Essay on Bargaining," *American Economic Review*, vol. 46 (1956), pp. 281-306. Reprinted in *The Strategy of Conflict*. New York: Oxford University Press, 1962.

Figure 1: Too-Big-to-Fail Decision Tree
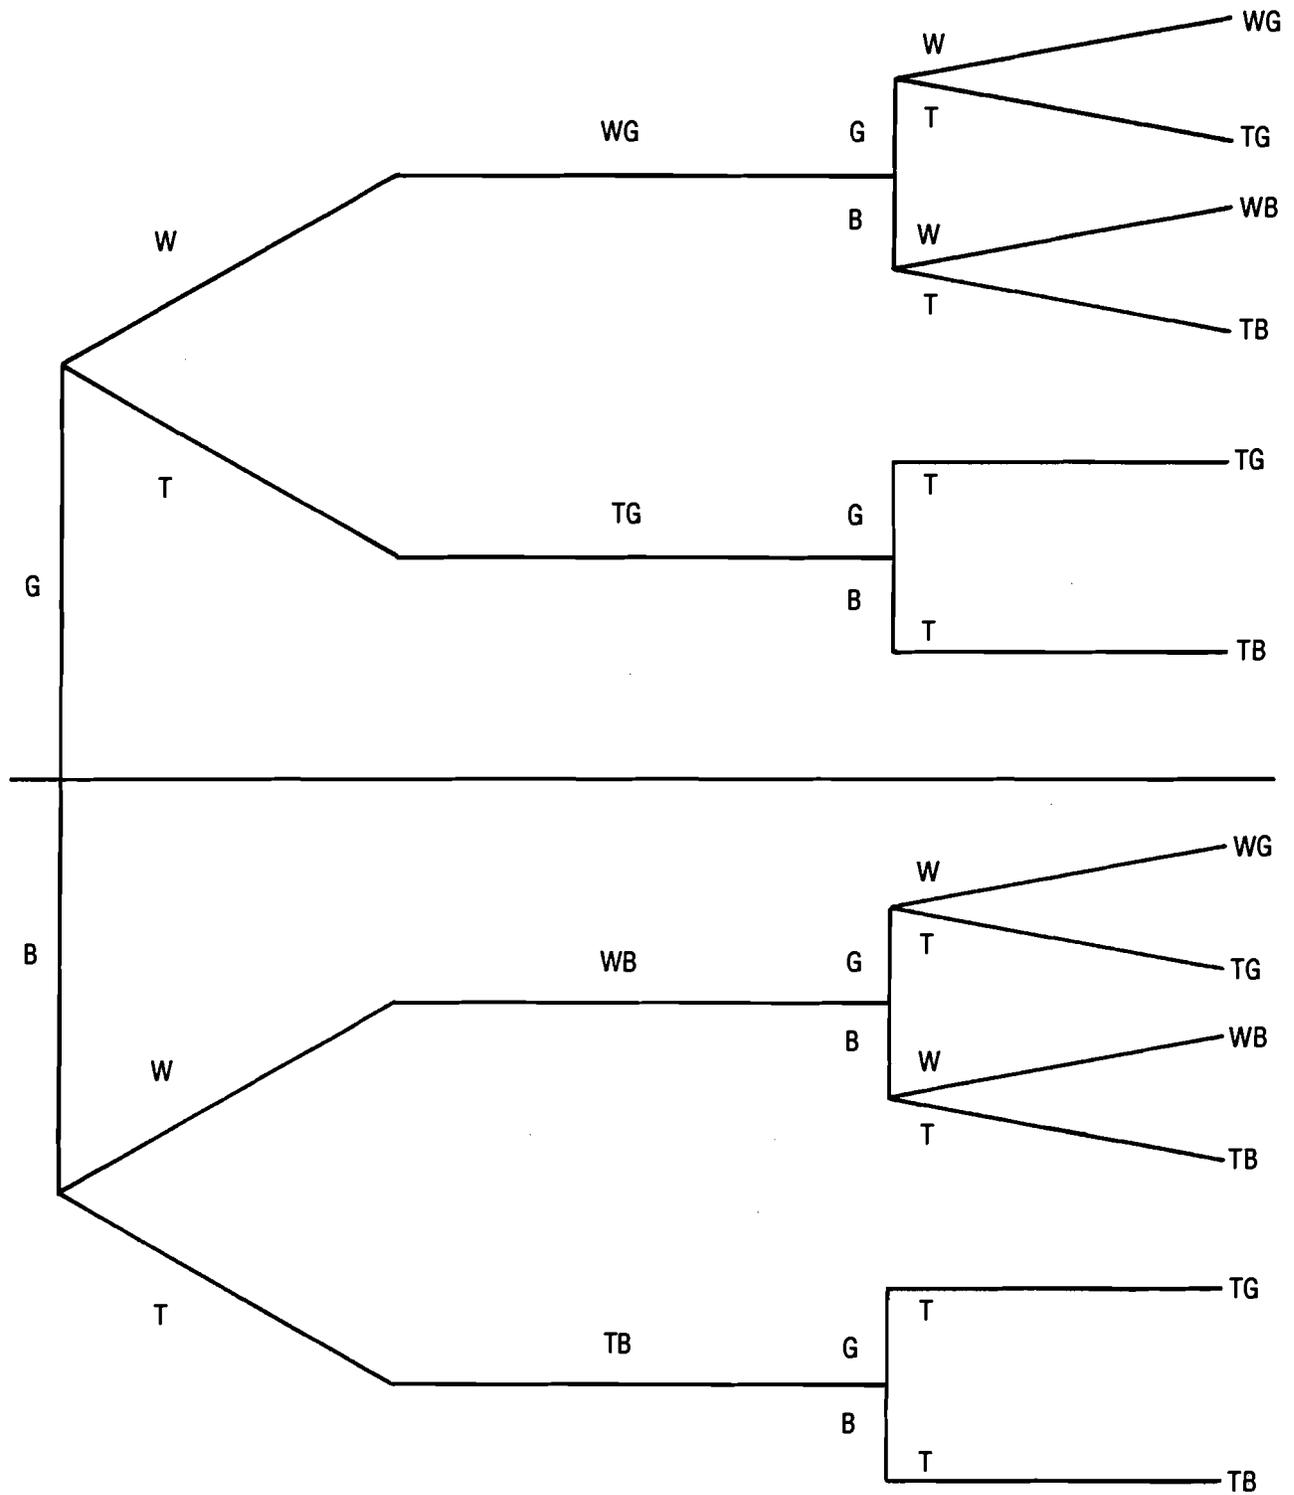
Figure 2: Two-Period Probability Structure

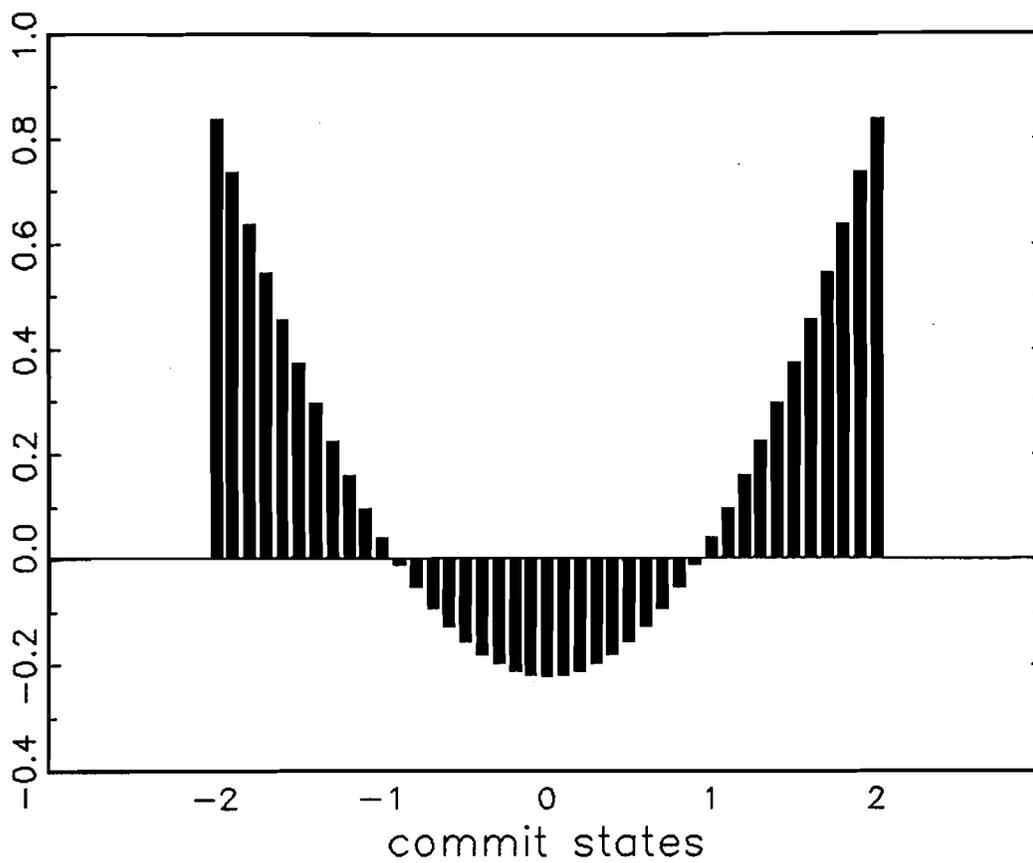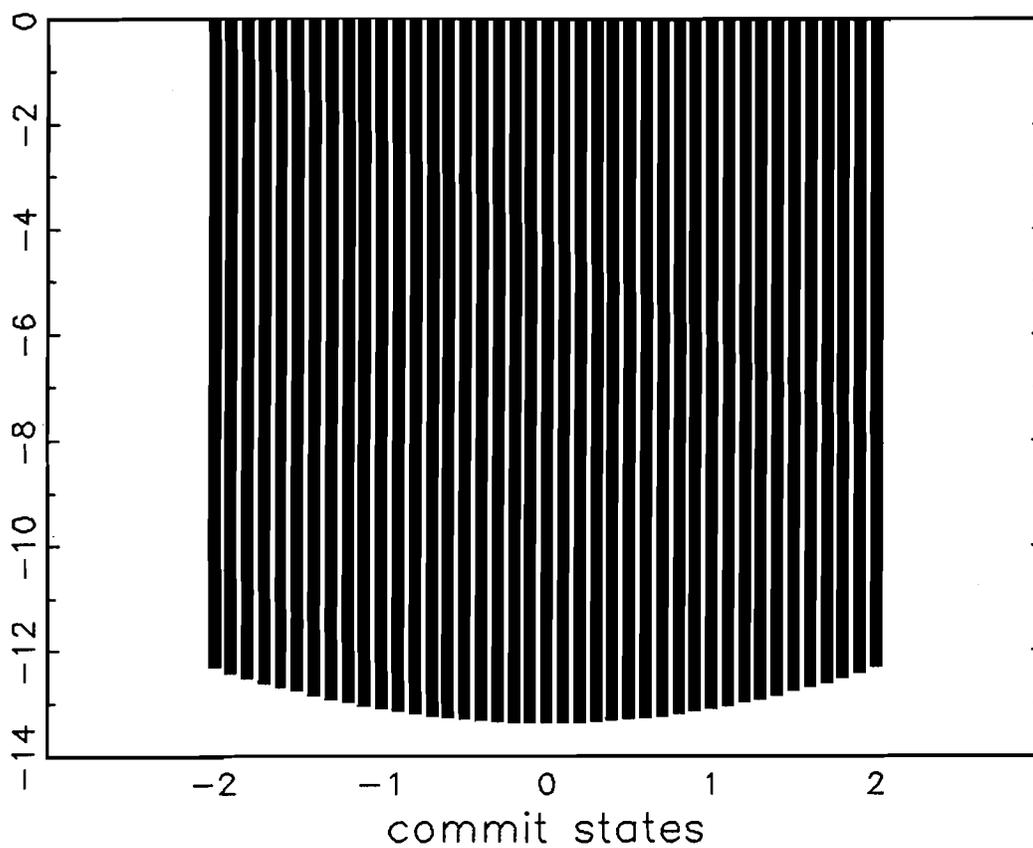Source: Authors.

VR−VD

commit states

VR−VDF

commit states

Source: Authors.