NBER WORKING PAPERS SERIES

TIME-CONSISTENT POLICY AND PERSISTENT CHANGES IN INFLATION

Laurence Ball

Working Paper No. 3529

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 1990

I am grateful for suggestions from Alan Blinder, Alex Cukierman, Robert Gibbons, N. Gregory Mankiw, David Romer, Lawrence Summers, and seminar participants at Penn, Princeton, and VPI. An NBER Olin Fellowship provided financial support. This paper is part of NBER's research program in Financial Markets and Monetary Economics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #3529
December 1990

TIME-CONSISTENT POLICY AND PERSISTENT CHANGES IN INFLATION

ABSTRACT

This paper presents a model of dynamically consistent monetary policy that explains changes in inflation over time. In the model -- as in the postwar United States -- adverse supply shocks trigger persistent increases in inflation, and disinflation occurs when a tough policymaker creates a recession. The paper also proposes an approach to selecting a unique, plausible equilibrium in infinite-horizon models of monetary policy.

Laurence Ball
NBER
1050 Massachusetts Avenue
Cambridge, MA 02138

I. INTRODUCTION

The literature on time-consistent policy is indispensable for understanding inflation. In economies like the U.S., the costs of inflation appear to exceed the small seignorage revenues, and the long run Phillips curve is vertical. Thus optimal trend inflation is close to zero. The idea that optimal policy is not dynamically consistent is the most promising explanation for periods of high trend inflation. In addition, given the short run gains from demand expansion, the related idea that the Fed cares about its reputation is essential for understanding periods of low inflation.

Nonetheless, the time-consistency literature is unpopular among many economists. There are two broad problems. First, the models of dynamic games in which ideas are formalized have unappealing features. Often the games have multiple equilibria, and choosing among the equilibria appears arbitrary. In addition, agents' behavior in the models, which are borrowed from the oligopoly literature, seems too sophisticated to fit the Fed and the U.S. public. Alan Greenspan does not use a mixed strategy to choose M1 growth, and citizens, in forming expectations, do not use trigger strategies that punish Greenspan for misbehavior.

The second weakness of current models is that they do not explain changes over time in trend inflation. The models provide plausible explanations for both high inflation (the time-consistency problem) and low inflation (reputation). But a theory of inflation should also explain why the economy moves between high and low inflation regimes, and here the models are less convincing. In some, inflation rises because a policymaker nears the end of his term and decides to leave with a boom. In others, inflation rises because of exogenous shocks (which is more realistic), but then falls costlessly after a "punishment period." As documented below, these scenarios do not fit the U.S. experience.

This paper presents a model of monetary policy that captures the central insights of previous models while trying to avoid their problems. There are two main features of my approach. First, my specification includes realistic sources of movements in inflation. Exogenous macro shocks, such as OPEC price increases, cause inflation to rise. Inflation falls when the Fed shifts to tougher monetary policy — when it decides to create a recession to disinflate. An important detail is that shocks whose direct effects are temporary lead (along with monetary accomodation) to inflation that can be eliminated only through a recession. That is, temporary macro shocks generate inertial inflation.

Second, as an equilibrium concept I choose Maskin and Tirole's (1988a) "Markov perfect equilibrium." A Markov perfect equilibrium is a perfect Nash equilibrium in which actions depend only on variables that directly affect payoffs. As explained below, the concept rules out equilibria involving sophisticated threats of punishment. While the model possesses many perfect Nash equilibria, a range of parameter values implies a unique Markov perfect equilibrium. Agents' behavior in this equilibrium is simple and realistic.

The rest of the paper contains seven sections. Section II provides further motivation: I describe the stylized facts that a model of inflation should capture and review previous models. Section III describes the new model, and Section IV proposes rules for agents' behavior. Section V derives conditions under which this behavior is a Markov perfect equilibrium, and Section VI considers uniqueness.

While the qualitative features of the model mimic actual economies, I use strong simplifying assumptions. Supply shocks take on only two values, and "strong" policymakers do not care at all about unemployment. Section VII discusses the robustness of the results to relaxing these assumptions. Section VIII concludes.

II. MOTIVATION

Part A of this section reviews stylized facts about inflation in recent U.S. history. Part B evaluates the success of previous models in capturing these facts.

A. The Stylized Facts

Figure 1 plots U.S. inflation since 1960. This paper is based on five conclusions about this experience. These conclusions can be debated, but they are conventional and should appear reasonable to most macroeconomists.

(1) The trend inflation rate sometimes exceeds the social optimum. Indeed, since inflation is costly and the long run Phillips curve is vertical, optimal trend inflation is close to zero. It is surely lower than average inflation in the 1970's.[1]

(2) Higher inflation would often raise welfare in the short run. In many circumstances, an inflationary demand expansion would create a boom and thus raise welfare temporarily. And clearly disinflation, as in 1979-82, is harmful in the short run. Apparently policymakers perceive future benefits from keeping inflation below the short run optimum.

(3) The sources of large rises in inflation are supply and demand shocks exogenous to the Fed. Figure 1 shows large rises in inflation in the late 60's

[1]Optimal inflation may be slightly positive because inflation yields seignorage, and because it makes relative wage adjustments easier if workers resist nominal wage cuts (Tobin, 1972).

Figure 1

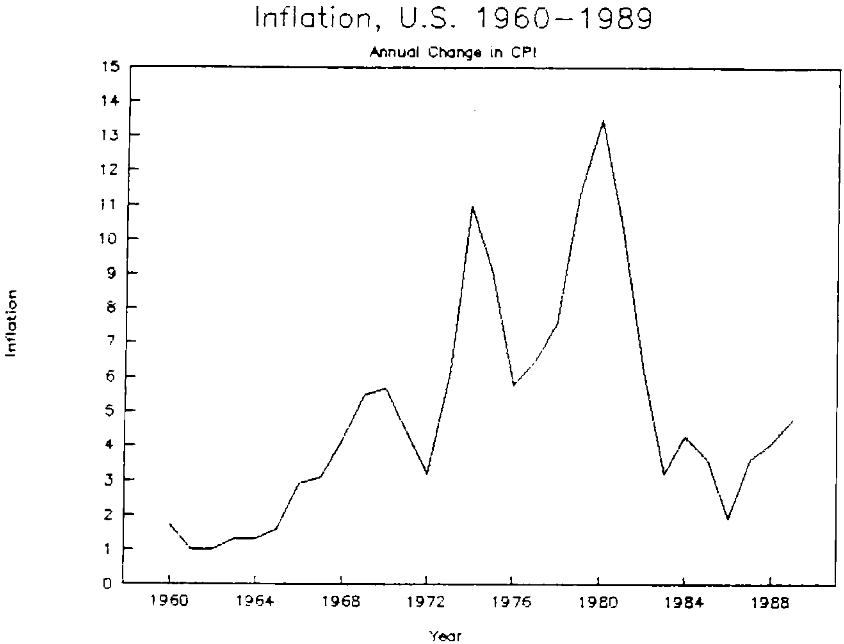
Inflation, U.S. 1960–1989
Annual Change in CPI
Inflation
Year
15
14
13
12
11
10
9
8
7
6
5
4
3
2
1
0
1960
1964
1968
1972
1976
1980
1984
1988

and at two points in the 70's. The consensus explanations are guns-and-butter fiscal policy in the 60's and food and energy shocks in the 70's (Blinder, 1982). While the shocks' direct effects are temporary, the increases in inflation persist. For example, the direct effect of a rise in oil prices is a one-time rise in the price level. But the OPEC shocks triggered inflation that lasted until the Fed created a recession (see point (4)).

(4) The sources of large disinflations are policy shifts: decisions by the Fed to pay the price of a recession. This is most clear for the Volcker disinflation of 1979-82, but a strong case can also be made for the milder disinflations of 1975, 1971, and so on (Romer and Romer, 1989). As discussed below, the timing of disinflation is the outcome of a complicated political process, and is hard to predict.

(5) In the absence of exogenous shocks or a decision to disinflate, inflation remains roughly constant. This is a corollary of (3) and (4). After the Volcker disinflation, for example, inflation remained low through the 80's because there was no large inflationary shock. (The stability of high inflation is less clearcut, because the Fed usually decides to disinflate within a few years.)

B. Previous Work

The large literature on dynamically consistent policy provides crucial insights but also has many unappealing features. Here I review the most relevant models, which fall into three categories.[2]

One-Period Models: These are the early models of Kydland-Prescott (1979) and Barro-Gordon (1983a). They show that fact (1) can arise from a time-consistency

[2]This discussion draws on surveys of the time-consistency literature by Fischer (1986) and Rogoff (1987).

problem. One period models cannot, however, explain the Fed's willingness to forgo short run gains (fact (2)), or the movements of inflation over time (facts (3) and (4)).

Repeated Game Models: Following James Friedman's (1971) oligopoly model, Barro and Gordon (1983b) present an infinite-period version of their model. The Fed keeps inflation low because high inflation would raise expected inflation in the future. This is a plausible explanation for (2). Following Green and Porter's (1983) oligopoly model, Canzoneri (1985) adds money demand shocks that cause inflation to rise occasionally despite the Fed's efforts to keep it low. This is a step towards capturing (3).

There are, however, four related problems with these models. First, there is a huge number of equilibria, and choosing among them appears arbitrary. Second, the equilibria that the authors emphasize have an unappealing bootstrap character. The models contain no link between periods, and so there is no fundamental reason for current policy to affect the future — it does so only because it is expected to. (The Markov criterion below rules out this kind of equilibrium.) Third, while exogenous shocks produce high inflation, inflation does not persist until the Fed creates a recession. Instead, expected inflation falls automatically after a "punishment period" of fixed length, and actual inflation falls costlessly at the same time.

Finally, Canzoneri's result that shocks raise expected inflation depends on his assumption that the Fed has private information about the shocks. (With private information, the public cannot tell whether a rise in inflation is intentional, and thus punishes the Fed for honest mistakes.) In reality, it is doubtful that the Fed has much private information about OPEC or fiscal policy. In any case, this issue is surely not crucial. The inflationary effects of an

OPEC shock would not decrease if the Fed could prove that it was surprised.[3]

Reputation Models: Backus and Driffill (1985) and Barro (1986) apply the Kreps-Wilson (1982) model of reputation. A "weak" policymaker forgos the short run gains from inflation (fact (2)) because masquerading as a "strong" policymaker keeps expected inflation low. Here, the channel through which current inflation affects future expectations — the Fed's reputation — is appealing. In addition, the models yield unique equilibria.

There are again four related problems. First, uniqueness is obtained through a finite horizon: a policymaker's actions do not affect the economy after the fixed end of his term. In reality, Volcker's policies permanently reduced expected inflation, making Greenspan's job easier. Second, the source of rises in inflation is not macroeconomic shocks (fact (3)), but deliberate decisions by weak policymakers to create a boom. Third, high inflation becomes more likely as a policymaker nears the end of his term: with a shrinking horizon, he is more tempted by short run gains. In Figure 1, there is no tendency for inflation to rise near the ends of terms. Finally, weak policymakers choose inflation through mixed strategies — by flipping coins.[4,5]

---

[3]The Fed might have private information about money demand because it observes bank deposits (Rogoff, 1987). In practice, however, shifts in money demand are not a major source of rises in inflation.

[4]Rogoff(1987) presents a reputation model in which policymakers choose pure strategies. In his model, however, policymakers' loss functions include an ad hoc cost of breaking commitments. It is better to assume that policymakers care about unemployment and inflation and derive the effects of broken commitments on these variables.

[5]Cukierman and Meltzer present another dynamic model of monetary policy in their 1986 paper on ambiguity. A weakness of this model as a general theory of inflation is that it attributes persistent rises in inflation to changes in policymakers' tastes. That is, as in reputation models, inflation rises because policymakers deliberately create a boom.

III. THE MODEL

This section presents a model that tries to capture the facts of inflation while avoiding the problems of previous models. Part A presents my assumptions and Part B describes the equilibrium concept.

A. Assumptions

The model is the usual Kydland-Prescott framework with two modifications. First, as in Canzoneri, there are exogenous shocks to the economy. Second, as in Backus-Driffill and Barro, there are two policymakers with different preferences, and the public does not observe who is in power. I depart from previous work by assuming that shocks are public information, and that the Fed chooses whether to accomodate them. I also assume that the policymaker in power changes stochastically.

The two policymakers are the strong, S, and the weak, W. W's loss function is

$$L = (U-U^{*})^{2} + a\pi^{2}, \quad a>0, \tag{1}$$

where U and U^{*} are actual and socially optimal unemployment, π is inflation, and a is a taste parameter. W minimizes the present value of (1) with discount factor $\beta<1$. As in Barro and Backus-Driffill, S sets π to zero every period. The simplest interpretation is that he minimizes (1) with $a\to\infty$ — that he cares about inflation but not about unemployment. Since this assumption is extreme, I consider more moderate tastes in Section VII.

The policymaker in power changes according to a Poisson process. If S is in power in period t, then with probability w he is replaced by W at t+1. W is replaced by S with probability s. One should think of s and w as fairly small, so there is strong serial correlation in policy. A change in policymaker should not be interpreted literally as the appointment of a new Fed chairman. Instead, it should be interpreted as a shift in policy regime, resulting for example from an FOMC meeting, that can occur without a change in personnel. These shifts result from a complicated political process involving the Fed and the pressure it receives from politicians and Wall Street. This process produces policy shifts at irregular intervals.

As an example of a policy shift, consider the late 1970's. During 1976–1978, policy was weak: the Fed tolerated high inflation. Given the political pressure for disinflation, it was clear that a shift would occur eventually, but the timing was unclear. The probability of disinflation during a given period (quarter?) was small. It turned out that strong policy arrived in October 1979.

The policymakers face the usual short run Phillips curve with the addition of a supply shock:

$$U = U^N - (\pi-\pi^e) + \eta \ , \qquad U^N = U^* + 1 \ , \tag{2}$$

where U^N is the natural rate of unemployment, π^e is expected inflation, and η is the shock. (The assumptions that $U^N-U^*=1$ and that the coefficient on $\pi-\pi^e$ is one are normalizations on the units of U and π.) A positive η is an adverse supply shock: it implies higher unemployment for a given inflation rate.

The shock η is observable. It is serially uncorrelated. For simplicity, η takes on only two values: it is zero with probability 1–q and $\bar{\eta}>0$ with probability q, where q is fairly small. Thus a shock is a discrete, occasional event, such as a major OPEC decision. While convenient, these assumptions can be relaxed; Section VII considers distributions for η that are continuous and symmetric around zero.

It is not essential that the disturbance is a supply shock. The crucial assumption is that the shock raises the cost of maintaining low inflation. Instead of supply shocks, one could add an ISLM demand side to the model and assume that policymakers dislike high real interest rates. In this case, an expansionary fiscal shock would raise the cost of keeping inflation low, because non-accomodative policy would produce high interest rates.

The public and the two policymakers play a simple game of asymmetric information. The public does not observe which policymaker is in power (the FOMC meetings that produce regime shifts are secret). Before each period, the public sets expected inflation given its imperfect information about who is in charge. Then the current value of η is determined, and the change in policymakers (if any) occurs. Finally, the current policymaker chooses inflation. Since the policymaker moves after η is known, he can choose whether to accomodate a shock.

B. The Equilibrium Concept

Infinite-horizon models of monetary policy have many perfect Nash equilibria. Backus-Driffill and Barro obtain uniqueness by assuming finite horizons, but this is unrealistic: actual economies never reach a final period in which inflation does not affect future expectations. The multiplicity of Nash equilibria is a natural feature of the repeated interaction of the Fed and the public. We should accept this feature and look for a reasonable way to choose among equilibria.

I focus on "Markov perfect equilibria" (Maskin and Tirole, 1988a). A Markov perfect equilibrium is a perfect Nash equilibrium in which actions depend only on variables that directly affect current payoffs. This usually rules out equilibria in which, as in Friedman and Barro-Gordon (1983b), collusion is supported by threats of punishment: whether a player cheated in the past does not affect current payoffs. The Markov criterion also rules out influences from extraneous "sunspots."[6]

Maskin and Tirole assume perfect information. I extend their concept slightly by assuming that, with imperfect information, actions depend on agents' best <u>estimates</u> of payoff-relevant variables. When the public sets π^e at the start of a period, the only payoff-relevant variable is the most recent identity of the policymaker. (The most recent η is irrelevant because η is serially uncorrelated.) Thus π^e equals $\pi^e(p)$, where p is the probability that W was in power in the previous period. When W chooses actual inflation, he conditions it on π^e — or equivalently on p — and on the current η. Thus inflation under W is given by

$$\pi = \pi(p,\eta) \ . \tag{3}$$

Inflation under S is always zero. Given the distribution of η and the process for policy switches, $\pi(p,\eta)$ and rational expectations determine $\pi^e(p)$. Thus an MPE is defined fully by the form of (3).

In the equilibrium below, p depends on the history of inflation and the supply shock, which helps the public infer who is in power. Thus, as in previous repeated game models, past inflation influences current inflation. The Markov criterion requires, however, that past conditions matter only through their effects on p — through the information they provide. In previous models, current inflation can be an arbitrary function of past inflation.

The main appeal of Markov perfect equilibria is simplicity: "actions depend on as little as possible while still being consistent with rationality" (Maskin-Tirole, p. 553). Maskin and Tirole argue that non-Markov punishment strategies are too sophisticated to describe the actual behavior of oligopolists. If so,

---

[6]Other papers that focus on Markov perfect equilibria include Maskin and Tirole (1988b) and Benabou (1989).

these strategies are certainly too sophisticated to describe how the public forms inflation expectations. Non-Markov equilibria are particularly unrealistic in this case because the "public" consists of millions of people. While two oligopolists might find their way to a collusive equilibrium, it is implausible that U.S. citizens coordinate on expectations that give the Fed the right incentives.

The Markov criterion is similar to the "minimum state variable" criterion that McCallum (1983) proposes for another set of macro models. McCallum eliminates equilibria in which a variable affects behavior "solely because it is (arbitrarily) expected to do so." He argues that it is often natural to focus on behavior determined by fundamentals.

IV. A PROPOSED EQUILIBRIUM

Sections IV-VI derive conditions under which a simple rule for W's behavior, $\pi(p,\eta)$, is a unique Markov perfect equilibrium. This section describes the equilibrium rule and the implied behavior of the economy.[7]

A. A Useful Lemma

The discussion is simplified by noting that any equilibrium must satisfy the following:

Lemma: For all (p,η), $\pi(p,\eta)$ equals either zero or π^d, where

$$\pi^d = \frac{\pi^e+\eta+1}{a+1} . \qquad (4)$$

π^d is the "discretionary" or "one-shot" inflation rate: it minimizes W's one-period loss for given π^e. The lemma holds because current inflation affects

[7]I assume that $\pi(\cdot)$ is non-stochastic. Mixed strategy equilibria can be ruled out by an extension of the uniqueness argument in Appendix B.

future losses only through its effect on the next period's p. Since S always chooses $\pi=0$, any $\pi\neq0$ raises p to one — it reveals that W is in power. If W chooses any $\pi\neq0$, he chooses the one that minimizes his current loss because all have the same effect on the future.

Recall that η equals either zero or $\bar{\eta}$. $\pi(p,\eta)$ is therefore summarized by a function $\pi(p)$ for each of these η's. The lemma implies that these functions have a simple form: π equals zero for some values of p and π^d otherwise. Thus an equilibrium is summarized by the sets of p's for which W chooses $\pi=0$ when $\eta=0$ and when $\eta=\bar{\eta}$. Denote these sets by X_0 and $X_{\bar{\eta}}$.

B. The Proposed Equilibrium

The proposed equilibrium is

$$X_0 = [0,\ 1)\ , \qquad X_{\bar{\eta}} = \emptyset\ . \tag{5}$$

When $\eta=0$, W sets $\pi=0$ unless $p=1$ — unless the public knows for sure that W was in power in the previous period. When $\eta=\bar{\eta}$, W never sets $\pi=0$; he sets $\pi=\pi^d$ for all p.

To understand (5), consider the evolution of the economy. Start with a period in which p equals zero (as it sometimes will). If there is no supply shock, then either policymaker chooses zero inflation, and the public learns nothing about who is in charge. With no new information, the next period's p, $\hat{p}$, is given by

$$\hat{p} = p(1-s) + (1-p)w \tag{6}$$

(when $p=0$, $\hat{p}=w$). Intuitively, the public updates p to account for possible switches in policymakers, which occur with probabilities s and w. In the periods that follow, p continues to evolve according to (6) as long as $\eta=0$. p rises monotonically and approaches $\bar{p}=w/(s+w)$, the unconditional probability that W is in charge. Since p remains below one, neither policymaker inflates.

At some point a supply shock arrives ($\eta=\bar{\eta}$). W inflates and S does not, so the policymaker's identity is revealed (again, the public observes η ex post). If the policymaker is S, then the next period's p is zero and the above scenario begins again. If it is W, then the next period's p is one. In this case, W inflates again if he is still in charge, and this implies p=1 in the period after that. W continues to inflate, and p remains at one, until W is replaced by S. S sets $\pi=0$, and p drops to zero in the following period.

This equilibrium captures facts (3)-(5) in Section II: inflation rises when there is an adverse shock (and policy is accomodative), falls when policy turns tough, and remains constant otherwise. A one-time shock triggers a persistent rise in inflation because it provides information — W reveals his identity by accomodating it. This information is relevant to future behavior because of serial correlation in policymakers.

W's behavior fits "weak" policymakers in actual economies. When p<1, expected inflation is low. As long as there is no shock, W produces low inflation to maintain the status quo. When a shock arrives, W inflates because he is not willing to accept the high unemployment implied by nonaccomodation. And once inflation is high and expected to remain high, W is not willing to pay the cost of disinflation. Finally, note that the behavior of expectations is simple and roughly adaptive: when actual inflation rises or falls, expected inflation follows in the next period.

V. WHEN IS IT AN EQUILIBRIUM?

This section determines when (5) is a perfect Nash equilibrium and thus, since behavior depends on fundamentals, a Markov perfect equilibrium. (5) is a perfect Nash equilibrium if W cannot gain by deviating from it in any period. The analysis proceeds in several steps, which are sketched here with details in the Appendix. Part A determines the present value of W's loss when he obeys (5); Part B describes possible deviations; and Part C determines when all deviations are losers.

A. The Present Value of W's Loss

The first step is to derive W's one-period loss for all states of the economy. The loss depends on p, which determines π^e; on η; and on π, which is either zero or π^d. Let $L^0(p,\eta)$ and $L^+(p,\eta)$ denote the loss when π equals zero and π^d respectively. The Appendix derives these loss functions in terms of underlying parameters. The loss is increasing in η and p. For given η and p, L^0 exceeds L^+: choosing zero inflation raises the current loss.

The next step is to derive the expected present value of W's loss in the proposed equilibrium. I compute the present value at the start of a period, before η and the current policymaker are determined. The present value depends on p, the public's estimate of who was recently in charge, and on who was actually in charge. Let $V^W(p)$ and $V^S(p)$ denote the present value of W's loss when the most recent policymaker was W and S respectively. As described in the Appendix, these value functions are derived through dynamic programming, given the one-period losses in various states and the probabilities of moving to new states through shocks and policy switches.

B. Possible Deviations

I now describe possible deviations from (5) and their effects on W's loss. W always sets either $\pi=0$ or $\pi=\pi^d$ (this lemma holds for deviations as well as in equilibrium). Thus deviating means choosing π^d when (5) dictates zero, or vice-versa. I consider deviations in three cases that exhaust the possibilities: p<1 and $\eta=0$; p<1 and $\eta=\bar{\eta}$; and p=1.

$\underline{p<1,\ \eta=0}$: In this case, (5) dictates $\pi=0$. If W instead sets $\pi=\pi^d$ — if he creates surprise inflation — the effect on his loss is

$$\Delta_1 = [L^+(p,0) - L^0(p,0)] + \beta[V^w(1) - V^w(\hat{p})] , \quad (7)$$

where the difference in the first brackets in negative and the difference in the second is positive. $L^+(p,0)-L^0(p,0)$ is the gain in the current period from surprise inflation, which reduces unemployment. $V^w(1)-V^w(\hat{p})$ is the present value of W's loss from revealing his identity, so that p moves to one rather than $\hat{p}$. This loss occurs in the next period, and thus is multiplied by the discount factor β. Raising p to one is harmful because it moves the economy to persistently high inflation.

$\underline{p<1,\ \eta=\bar{\eta}}$: In this case, (5) dictates $\pi=\pi^d$. The effect on W's loss of choosing $\pi=0$ is

$$\Delta_2 = [L^0(p,\bar{\eta}) - L^+(p,\bar{\eta})] + \beta[V^w(0) - V^w(1)] . \quad (8)$$

There are again two effects given by the terms in the two brackets. The first effect, which is positive, is the short run cost of failing to accomodate the supply shock. The second, which is negative, is the gain from continuing to act strong and thus remaining at low inflation.[8]

$\underline{p=1}$: Here (5) implies $\pi=\pi^d$. If W chooses $\pi=0$ — if he disinflates — the effect is

$$\Delta_3 = [L^0(1,\eta) - L^+(1,\eta)] + \beta[V^w(0) - V^w(1)] . \quad (9)$$

The first effect is the current cost of disinflation, which raises unemployment. The second is the gain from acting strong and thus moving to low inflation.

---

[8]If W choose $\pi=0$, p falls to zero: this out-of-equilibrium move convinces the public that S is in power.

C. When Does W Lose from All Deviations?

(5) is a Markov perfect equilibrium if W loses from all deviations — if $\Delta_1, \Delta_2 > 0$ $\forall p<1$ and $\Delta_3 > 0$ for both values of η. To see when these conditions hold, I begin with a simple limiting case of the model: s, w, q $\to 0$. In this case, the arrival rates of supply shocks and policy shifts approach zero. As discussed above, one should think of these parameters as fairly small — shocks and policy shifts are occasional events. The behavior of the economy is continuous in s, w, and q, and so the limiting case is similar to cases in which the parameters are small but positive.

For the limiting case, the Appendix establishes that

$$\Delta_1 > 0 \ \forall p \ \ iff \ \ \delta > \frac{a}{1+a} \ ; \qquad (10a)$$

$$\Delta_2 > 0 \ \forall p \ \ iff \ \ \delta < \left(\frac{a}{1+a}\right)(\overline{\eta}+1)^2 \ ; \qquad (10b)$$

$$\Delta_3 > 0 \ \forall \eta \ \ iff \ \ \delta < \frac{1+a}{a} \ , \qquad (10c)$$

where $\delta = \beta/(1-\beta)$. These conditions have simple interpretations. Δ_1 is positive -- W loses from creating surprise inflation — if the cost of moving to persistently high inflation exceeds the short run gain from a boom. Not surprisingly, this holds if the discount factor β is large enough.[9] Δ_2 and Δ_3 are positive if the short run cost of a recession from nonaccomodation or disinflation exceeds the gain from keeping inflation low. These conditions hold if β is low enough.

[9]Surprisingly, this condition is most likely to hold if a is small. That is, W is less likely to inflate if he attaches a small cost to inflation. The explanation is that, as stressed by Fischer and Summers (1989), π^d is decreasing in a. When a falls, π^d rises so much that W's loss at $\pi = \pi^d$ is higher even though inflation is less costly.

The three conditions can hold simultaneously, and so (5) can be an equilibrium, because the two upper bounds on β exceed the lower bound. For moderate discount rates, W forgoes a boom to avoid inflation but accepts inflation to avoid a recession. This behavior is possible because the cost of a recession exceeds the gain from a boom. The source of this asymmetry is the convexity of the loss function, (1) — a rise in unemployment has a larger absolute effect than a fall.[10]

When the assumption of s, w, q $\rightarrow$ 0 is relaxed, the conditions for (5) to be an equilibrium become complicated (see Appendix). However, starting at s, w, q $\rightarrow$ 0, one can establish the qualitative effects of increasing these parameters. An increase in s reduces both the lower bound on δ in (10a) and the upper bound in (10b). Intuitively, a rise in s raises the loss from high inflation by making a costly disinflation more likely; thus W is more inclined to choose π=0. Starting at s, w, q $\rightarrow$ 0, an increase in w has no effect on (10). The effects of an increase in q are ambiguous.

VI. UNIQUENESS

Part A of this section considers uniqueness of the Markov perfect equilibrium. Part B discusses some of the perfect Nash equilibria that are ruled out by the Markov criterion.

A. Uniqueness of the MPE

Note first that many combinations of X_0 and X_π imply the same behavior of inflation as (5), the regime considered above. For example, X_0 can be changed

[10]A convex loss function is equivalent to the assumption that W prefers stable unemployment at U^N to symmetric fluctuations around U^N. An alternative source of the asymmetry between booms and recessions is an asymmetry in the Phillips curve: disinflation has a larger effect on unemployment than an equal rise in inflation. These points are discussed further in Ball (1990).

from $[0,1)$ to $[0,p^*)$, where $\bar{p}<p^*<1$. This change in the upper bound is irrelevant because, in equilibrium, p never lies between $\bar{p}$ and one (recall that p rises from zero towards $\bar{p}$, jumps to one, and then returns to zero). When (5) is an equilibrium, there are usually other equilibria that imply the same behavior. I ignore these regimes and ask whether there are equilibria with different behavior. That is, I consider uniqueness of the equilibrium path of inflation, not of X_0 and X_η.

The Appendix derives conditions for uniqueness. I consider all combinations of X_0 and X_η that imply different behavior from (5) and determine when W deviates from each. The results are simplest in the limiting case of s, w, q $\rightarrow$ 0. In this case, the conditions that make (5) an equilibrium, (10), also rule out most other candidates. Perhaps surprisingly, the only exceptions are regimes with $X_\eta=0$ and $X_0=[0,p^*)$, where $0<p^*<\bar{p}$. With $p^*<\bar{p}$, such a regime does imply different behavior from (5): when p rises to p^*, W inflates even if there is no supply shock. Equilibria of this form coexist with (5) for a range of parameter values. However, as detailed in the Appendix, they are ruled out by a moderate strengthening of (10).

B. Non-Markov Equilibria

The model possesses many perfect Nash equilibria that do not satisfy the Markov criterion. For example, whenever (5) is an equilibrium, there is another perfect equilibrium in which W chooses $\pi=0$ if $\eta=0$ and $\pi=\pi^d$ if $\eta=\bar{\eta}$. A supply shock raises inflation in the current period, but inflation returns to zero in the next period (unless there is another shock). Low inflation when $\eta=0$ is supported by the threat that π^e will rise considerably if W inflates in this case.[11]

It is perhaps surprising that the Markov criterion, which is designed to tie behavior to fundamentals, rejects this equilibrium in favor of (5). Here, inflation under W depends only on the current η, which is part of fundamentals. In contrast, the persistence of inflation implied by (5) seems to involve sunspots. A shock triggers inflation that lasts until S disinflates even though the shock affects fundamentals only for one period.

(5) satisfies the Markov criterion because, as explained above, a shock provides information relevant to the future: it forces W to reveal his identity. The equilibrium with temporary effects of shocks is <u>not</u> Markov because of the out-of-equilibrium threats that support it. Zero inflation whenever η=0 cannot be supported by an unconditional expectation of this outcome; without punishments for deviations, W would repeatedly create surprise inflation. To put it differently, in the non-Markov equilibrium a rise in inflation raises expected inflation only if there is no supply shock — if W is cheating. The Markov criterion requires that any rise in inflation be treated the same, since any rise reveals that W is in power.

These results are significant because of fact (3) above: one-time shocks to inflation fundamentals, such as OPEC price rises, generate persistent inflation. A common explanation is that the initial rise in inflation produces high expected inflation, which induces the Fed to keep inflation high. In my model, the Markov criterion selects this behavior over a regime in which OPEC's effects are temporary.

---

[11]If (as in Canzoneri) the Fed has private information about η, then it cannot be an equilibrium for inflation to return to zero immediately after a shock. However, inflation can fall after a brief punishment period.

VII. ROBUSTNESS

For simplicity, the basic model assumes that η takes on only two values, and that S cares only about inflation. This section discusses the effects of relaxing these assumptions.

A. The Distribution of the Supply Shock

I first consider a discrete but symmetric distribution for η, and then a continuous distribution. To keep the discussion manageable, I consider only the limiting case of s, w, q → 0.

A Symmetric Distribution: Suppose that there are beneficial as well as adverse supply shocks. Specifically, let η equal zero with probability 1-q, $\bar{\eta}$ with probability q/2, and $-\bar{\eta}$ with probability q/2. A Markov perfect equilibrium is now defined by X_0, X_η, and $X_{-\eta}$, the range of p for which W chooses zero inflation when $\eta=-\bar{\eta}$. With reasoning parallel to Sections V-VI, one can prove existence of a unique equilibrium under the same conditions as in the basic model. As before, $X_0=[0,1)$ and $X_\eta=\emptyset$. There are two cases for $X_{-\eta}$:

$$X_{-\bar{\eta}} = [0,1) \quad if \quad \delta < \frac{(a+1-a\bar{\eta})^2}{a(a+1)}; \qquad = [0,1] \quad otherwise, \tag{11}$$

where again $\delta=\beta/(1-\beta)$. If β is in the lower part of the range defined by (10), W behaves the same when $\eta=-\bar{\eta}$ as when $\eta=0$. As in the basic model, inflation rises when W accomodates an adverse shock and falls when S disinflates. On the other hand, if β is relatively large, then W chooses zero inflation when p=1 and $\eta=-\bar{\eta}$. That is, when high inflation is expected, W takes the gains from a good shock in lower inflation rather than lower unemployment. In this case, high inflation ends <u>either</u> when S arrives or when there is a good shock.[12]

Disinflations in recent U.S. history were generally triggered by policy shifts rather than exogenous shocks. One interpretation is that the first case in (11) fits the U.S. economy. An alternative is simply that large beneficial shocks are rare.

<u>A Continuous Distribution</u>: Assume as above that η is zero with probability $1-q$, but let its distribution when non-zero be continuous. There is again a unique Markov perfect equilibrium for a range of parameter values. When $p<1$, W chooses positive inflation only if η exceeds a cutoff $\eta^{*}>0$. If $p=1$, high inflation is expected and W disinflates only if $\eta<-\eta^{**}$, where $\eta^{**}>0$. In other words, large shocks cause persistent changes in inflation but small shocks do not. The two cutoffs are defined by the condition that W is indifferent between postive and zero inflation; one can derive

$$\eta^{*} = \sqrt{\delta\left(\frac{a+1}{a}\right)} - 1\ ; \qquad \eta^{**} = \frac{a+1-\sqrt{\delta a(a+1)}}{a}\ . \tag{12}$$

(Conditions (10a) and (10c), which are required for this equilibrium, assure $\eta^{*},\eta^{**}>0$.) Note that the cutoffs are generally unequal. The case of $\eta^{**}>\eta^{*}$ is appealing because, as discussed above, decreases in inflation from good shocks

---

[12]When $p=1$ and a beneficial shock produces low inflation, p drops to $\hat{1}-(1-s)$ in the following period. (Since both S and W choose zero inflation, the public cannot tell whether W is still in power.) Since $p<1$, both policymakers continue to choose low inflation until there is an adverse shock; p evolves according to (6) and approaches $\bar{p}$ from above. When an adverse shock arrives, p jumps to zero or one.

appear less common that increases from bad shocks.[13]

B. S's Preferences

I follow Barro and Backus-Driffill in assuming that the strong policymaker cares only about inflation, and thus that his one-shot inflation rate is zero. In reality, even a "tough" policymaker like Paul Volcker cares somewhat about unemployment, and his one-shot rate is significantly positive. (Volcker endured the costs of disinflation for its long run benefits; he would have chosen higher inflation if he were minimizing his one-period loss.) It is realistic to assume that S minimizes the loss function (1) with the parameter a large but finite.

This version of the model is complicated, and so I leave it for future research. Possible results are suggested by the work of Vickers (1986) and Hoshi (1988). These authors study two-period models in which S's value of a is finite but larger than W's. For some parameter values, both policymakers choose zero inflation in the first period even though they have positive one-shot rates. Intuitively, S tries to differentiate himself from W by reducing inflation, but W follows. For other parameter values, S chooses zero inflation but W chooses his one-shot rate. In the modified version of my model, I conjecture that S chooses zero inflation and W imitates him in some states but not others. As in the basic model, W does not imitate S when zero inflation is costly because expected inflation is high or there is an adverse shock. If this conjecture is correct, then the behavior of inflation in the modified model is the same as in

[13]Equation (12) depends on the assumption that $q \to 0$. With a continuous distribution for η, it is natural to assume that $q=1$ — there is a (possibly small) shock every period. In this case, η^* is a function of the current p, which affects π^e. It appears, however, that my qualitative results carry over.

the basic model.[14]

VIII. CONCLUSION

This paper presents a model of inflation in the postwar U.S. There are two advances over previous work. First, the model explains movements between low- and high-inflation regimes. Inflation rises when there is an adverse supply shock and the Fed, unwilling to accept high unemployment, accomodates it. Once inflation rises, it stays high: a one-time shock generates inertial inflation. Disinflation occurs only when a tough policymaker arrives and creates a recession. These results fit experiences such as the inflation of the 1970's and the Volcker disinflation.

Second, the paper proposes an approach to selecting an equilibrium in infinite-horizon models of monetary policy. I focus on "Markov perfect equilibria" — equilibria in which actions depend only on variables that directly affect payoffs. For a range of parameter values, this concept yields uniqueness. In addition, the behavior of the Fed and the public is simple and realistic. The equilibrium is not supported by sophisticated punishment strategies.

---

[14]As in Hoshi and Vickers, there are likely to be multiple Markov perfect equilibria supported by various out-of-equilibrium beliefs. This case will require an additional refinement criterion, such as Cho-Sobel (1987), to rule out unreasonable beliefs.

APPENDIX

A. When Is (5) an Equilibrium?

Section V derives conditions under which the behavior in (5) is a Markov perfect equilibrium. Here I present details of this derivation.

One-Period Losses: Substituting (2) into (1) gives W's one-period loss in terms of π, π^e, and η:

$$L = (\pi - \pi^e - \eta - 1)^2 + a\pi^2 . \tag{A1}$$

In equilibrium, π^e is given by

$$\pi^e(p) = \hat{p}[(1-q)\pi(p,0) + q\pi(p,\overline{\eta})] , \tag{A2}$$

where $\hat{p}$ is the probability that W is in charge after the possible switch in policymakers (see (6)), and $\pi(\cdot)$ gives W's equilibrium choices of inflation for the two values of η. For given π^e, $\pi(\cdot)$ is determined by the definition of π^d, (4), and the rule for when W chooses π^d, (5). Substituting the expressions for $\pi(\cdot)$ into (A2) and solving for π^e yields

$$\begin{aligned} \pi^e(p) &= \frac{\hat{p}q(\overline{\eta}+1)}{a+1-\hat{p}q} \qquad for\ p<1 ; \\ &= \frac{\hat{p}(q\overline{\eta}+1)}{a+1-\hat{p}} \qquad for\ p=1 . \end{aligned} \tag{A3}$$

Substituting (A3) into (4) yields solutions for π^d for all (p,η):

$$
\begin{aligned}
\pi^d &= \frac{a+1+\hat{p}q\bar{\eta}}{(a+1)(a+1-\hat{p}q)}, && p<1,\ \eta=0\ ; \\
&= \frac{\bar{\eta}+1}{a+1-\hat{p}q}, && p<1,\ \eta=\bar{\eta}\ ; \\
&= \frac{a+1+\hat{p}q\bar{\eta}}{(a+1)(a+1-\hat{p})}, && p=1,\ \eta=0\ ; \\
&= \frac{\bar{\eta}(a+1+\hat{p}q-\hat{p})+a+1}{(a+1)(a+1-\hat{p})}, && p=1,\ \eta=\bar{\eta}\ .
\end{aligned}
\tag{A4}
$$

Finally, substituting (A3) and (A4) into (A1) defines $L^+(p,\eta)$ and $L^0(p,\eta)$, which give W's one-period losses. These expressions assume equilibrium expectations, but they cover cases in which W's choice between $\pi=0$ and $\pi=\pi^d$ deviates from the equilibrium.

In the limiting case of s, w, q $\rightarrow$ 0, the one-period losses simplify to

$$
\begin{aligned}
L^0(p,\eta) &= (\eta+1)^2, && p<1\ ; \\
L^0(1,\eta) &= \frac{(1+a+a\eta)^2}{a^2}\ ; \\
L^+(p,\eta) &= \frac{a(\eta+1)^2}{a+1}, && p<1\ ; \\
L^+(1,\eta) &= \frac{(1+a+a\eta)^2}{a(a+1)},
\end{aligned}
\tag{A5}
$$

The Present Value of the Loss: $V^S(p)$ and $V^W(p)$ give the expected present value of W's loss at the start of a period given p and the initial policymaker. These functions are defined implicitly by

$$
\begin{aligned}
V^{s}(p) &= R^{s}(p) + \beta[(1-w)(1-q)V^{s}(\hat{p}) + (1-w)qV^{s}(0) \\
&\quad + w(1-q)V^{w}(\hat{p}) + wqV^{w}(1)] \quad \text{for } p<1 ; \\
V^{w}(p) &= R^{w}(p) + \beta[(1-s)(1-q)V^{w}(\hat{p}) + (1-s)qV^{w}(1) \\
&\quad + s(1-q)V^{s}(\hat{p}) + sqV^{s}(0)] \quad \text{for } p<1 ; \qquad \text{(A6)} \\
V^{w}(1) &= R^{w}(1) + \beta[(1-s)V^{w}(1) + sV^{s}(0)] ,
\end{aligned}
$$

where $R^{s}(\cdot)$ and $R^{W}(\cdot)$ are the expected losses in the current period, derived below. ($V^{s}(1)$ is not defined, because $p=1$ implies that W is initially in charge.) The present value of the loss equals the current loss plus the discounted present value in the next period. The present value in the next period is an average over the four possible combinations of supply shocks and policymakers. In the first line, for example, with probability $(1-w)(1-q)$ there is no shock and S remains in power, so p rises to $\hat{p}$ in the next period; with probability $(1-w)q$, S remains but there is a shock, so S is revealed and p drops to zero; and so on.

The current losses are

$$
\begin{aligned}
R^{s}(p) &= (1-q)L^{0}(p,0) + q(1-w)L^{0}(p,\bar{\eta}) + qwL^{+}(p,\bar{\eta}) \\
&\qquad \text{for } p<1 ; \\
R^{w}(p) &= (1-q)L^{0}(p,0) + q(1-s)L^{+}(p,\bar{\eta}) + qsL^{0}(p,\bar{\eta}) \\
&\qquad \text{for } p<1 ; \\
R^{w}(1) &= (1-q)(1-s)L^{+}(1,0) + (1-q)sL^{0}(1,0) + q(1-s)L^{+}(1,\bar{\eta}) \\
&\quad + qsL^{0}(1,\bar{\eta}) .
\end{aligned}
\qquad \text{(A7)}
$$

Again, these are averages over the possible shocks and policymakers.

In the limiting case of s, w, $q \to 0$, the present value of the loss simplifies to

$$V^s(p) = V^w(p) = \frac{1}{1-\beta} L^0(p,0) \; ; \qquad V^w(1) = \frac{1}{1-\beta} L^+(1,0) \; , \tag{A8}$$

where I use the fact that $\hat{p}=p$ when s, w $\rightarrow$ 0. Intuitively, in the limiting case shocks and policy shifts never occur, and so either low inflation (p<1) or high inflation (p=1) continues forever. $L^0(p,0)$ and $L^+(1,0)$ are the constant one-period losses in the two regimes.

Ruling Out Deviations: In the limiting case, substituting (A5) and (A8) into (7)-(9) yields (10), the conditions under which W does not deviate from the equilibrium. Note that Δ_1 and Δ_2 are the same for all p<1. Δ_3 is decreasing in η, so I set $\eta=0$ in (9) to determine when $\Delta_3>0$ for all η.

When s, w, q > 0, the conditions for (5) to be an equilibrium are defined by combining (A6), (A7), and the general expressions for $L^+(\cdot)$ and $L^0(\cdot)$ with (7)-(9). Differentiating (7)-(9) yields the results about the effects of s, w, and q reported in the text. One could further interpret the general conditions by numerically calculating ranges of parameter values for which they hold.

B. Uniqueness

Here I derive conditions under which the Markov perfect equilibrium, (5), is unique. More precisely, as described in the text, the conditions assure that any equilibrium X_0 and X_η imply the same inflation path as (5). I focus on the limiting case of s, w, q $\rightarrow$ 0; again, by continuity the results hold for a range of positive values of these parameters. For the limiting case, there are two major steps in the argument. First, I show that conditions (10), which make (5) an equilibrium, also assure that any equilibrium is of the form $X_0=[0,p^*)$, $X_\eta=\emptyset$, where $0<p^*\leq 1$. Second, I derive a moderate strengthening of (10) under which

$p^* \geq \bar{p}$. As described in the text, any regime with these properties implies the same behavior as (5).

To establish $X_0=[0,p^*)$, $X_\eta=\emptyset$, I show that any equilibrium has three properties. The approach is to show that, given (10), W deviates from a regime that lacks any of the properties. For convenience, let $Z^+(r,\eta)$ and $Z^0(r,\eta)$ denote W's one-period losses when $\pi=\pi^d$ and $\pi=0$, given that the public expects $\pi=\pi^d$ with probability r. (In the limiting case, r depends on p and on X_0 in the regime under consideration.) The three properties are the following:

(1) $\underline{X_0 \not\ni 1}$. That is, in any equilibrium W chooses $\pi=\pi^d$ if $\eta=0$ and the public knows that he was previously in power. This property holds trivially for all parameter values. If the public expected zero inflation when p=1, W would create surprise inflation. A boom would occur, and there would be no cost: p would remain at one and π^e at zero. (This argument uses the assumption of $q\to0$, which implies that π^e depends only on X_0.)

(2) $\underline{X_0 \ni 0}$. That is, W does not inflate when p=0 and $\eta=0$. If W were expected to inflate in this situation, deviating would raise his current loss by $Z^0(0,0)-Z^+(0,0)$. (Since p=0, r=0 even though W is expected to inflate.) Returning to expected behavior in the next period, W would gain $\beta[Z^+(1,0)-Z^+(0,0)]$ because his identity was not previously revealed. (Here I use property (1): W inflates if p=1.) Equations (4) and (A1) imply $Z^+(0,0)=a/(a+1)$, $Z^0(0,0)=1$, and $Z^+(1,0)=(a+1)/a$. Along with condition (10a), these results imply that W gains overall from deviating.

(3) $\underline{X_\eta=\emptyset}$. That is, W inflates whenever there is a supply shock. If W were expected not to accomodate a shock, deviating would yield a current gain of $Z^0(r,\bar{\eta})-Z^+(r,\bar{\eta})$, which is bounded below by $Z^0(0,\bar{\eta})-Z^+(0,\bar{\eta})=(\bar{\eta}+1)^2/(a+1)$. The future cost of the deviation is bounded above by $[\beta/(1-\beta)][Z^+(1,0)-$

$Z^0(0,0)]=\beta/[(1-\beta)a]$, the cost of moving permanently from low to high inflation. These results and condition (10b) assure that W gains overall.

Given these three properties, any equilibrium implies the same behavior as a regime of the form $X_\eta=\emptyset$, $X_0=[0,p^*)$, where $0<p^*\leq 1$. An equilibrium is of this form as long as X_0 is convex. There can also be equilibria with $X_0=\Theta$, where Θ is not convex. In this case, the equilibrium implies the same behavior as $X_\eta=\emptyset$, $X_0=[0,p^*)$, where p^* is the smallest p not contained in Θ.

Given that an equilibrium is of the form $X_0=[0,p^*)$, $X_\eta=\emptyset$, the second major step in establishing uniqueness is to derive conditions that rule out $p^*<\bar{p}$. Again, any equilibrium with $p^*\geq\bar{p}$ implies the same behavior as (5). If $p^*<\bar{p}$, p sometimes rises to p^* and W inflates even without a supply shock. Conditions (10) are <u>not</u> sufficient to rule out equilibria of this type, and so they coexist with (5) for some parameter values. Intuitively, the expectation that W will inflate when p reaches p^* can be self-fulfilling. I now show, however, that a moderate strengthening of (10) rules out such behavior.

There are two possible deviations from a regime with $p^*<\bar{p}$. The first is for W not to inflate when p reaches p^*. The cost in the current period is $Z^0(p^*,0)-Z^+(p^*,0)$: there is a recession, but it is mild because inflation is expected only with probability p^*. There is a gain of $[\beta/(1-\beta)][Z^+(1,0)-Z^0(0,0)]$ from convincing the public that S is in power and thus (in the limiting case) maintaining zero inflation forever. For a given p^*, W gains overall if

$$\delta > \frac{a(a+1)}{(a+1-p^*)^2} . \qquad \text{(A9)}$$

The second possible deviation is to inflate the period __before__ p reaches p^*. W gains from a boom, and the cost is small: the economy moves to high inflation, but it would have done so anyway in the next period. Formally, the gain from the deviation is $Z^+(0,0)-Z^0(0,0)$, and the cost in the next period is $\beta[Z^+(1,0)-Z^+(p^*,0)]$. W gains overall if

$$\beta < \frac{a(a+1-p^*)^2}{(1+a)^2[(a+1-p^*)^2-a^2]} . \tag{A10}$$

A necessary and sufficient condition to rule out equilibria with $p^*<\bar{p}$ is that at least one of (A9) and (A10) holds for all $p^*<\bar{p}$. The right side of (A9) is increasing in p^*, and so a sufficient condition is simply that (A9) holds for $p^*=\bar{p}$. Suppose, for example, that $\bar{p}=1/2$ (in the limiting case this means that s and w approach zero at the same rate). In this case, the sufficient condition is

$$\delta > \frac{a(a+1)}{[a+(1/2)]^2} , \tag{A11}$$

which is a moderate strengthening of (10a). Intuitively, under this condition W keeps inflation low by accepting a recession when $p=p^*$ even though, by (10c), he is unwilling to accept the deeper recession caused by disinflation when p=1. An alternative sufficient condition is that, for some $\tilde{p}<\bar{p}$, (A9) holds for all $p^*\leq\tilde{p}$ and (A10) holds for all $p^*>\tilde{p}$ (note that the right side of (A10) is also increasing in p^*). In this case, for small p^*'s W accepts a mild recession, while for larger p^*'s he inflates the period before he is expected to.

REFERENCES

Backus, David and John Driffill, "Inflation and Reputation," American Economic Review 75 (1985): 530-538.

Ball, Laurence, "Why Does High Inflation Raise Inflation Uncertainty?," NBER Working Paper 3224 (January 1990).

Barro, Robert, "Reputation in a Model of Monetary Policy with Incomplete Information," Journal of Monetary Economics 17 (1986): 3-20.

________ and David Gordon, "A Positive Theory of Monetary Policy in a Natural Rate Model," Journal of Political Economy 91 (1983a): 589-610.

________ and ________, "Rules, Discretion, and Reputation in a Model of Monetary Policy," Journal of Monetary Economics 12 (1983b): 101-121.

Benabou, Roland, "Optimal Price Dynamics and Speculation with a Storable Good," Econometrica 57 (1989): 41-80.

Blinder, Alan, "The Anatomy of Double-Digit Inflation," in Robert Hall (ed.), Inflation: Causes and Effects, University of Chicago Press (1982).

Canzoneri, Matthew, "Monetary Policy Games and the Role of Private Information," American Economic Review 75 (1985): 1056-1070.

Cho, In-Koo and Joel Sobel, "Strategic Stability and Uniqueness in Signalling Games," Mimeo, University of Chicago, 1987.

Cukierman, Alex and Allan Meltzer, "A Theory of Ambiguity, Credibility, and Inflation Under Discretion and Asymmetric Information," Econometrica 54 (1986): 1099-1128.

Fischer, Stanley, "Time Consistent Monetary and Fiscal Policies: A Survey," Mimeo, MIT, 1986.

________ and Lawrence Summers, "Should Nations Learn to Live with Inflation?," American Economic Review 79 (May 1989): 382-387.

Friedman, James, "A Noncooperative Equilibrium for Supergames," Review of Economic Studies 38 (1971): 861-874.

Green, Edward and Robert Porter, "Noncooperative Collusion Under Imperfect Price Information," Econometrica 52 (1984): 87-100.

Hoshi, Takeo, Government Reputation and Monetary Policy, MIT Dissertation, 1988.

Kreps, David and Robert Wilson, "Reputation and Imperfect Information," Journal of Economic Theory 27 (1982): 253-279.

Kydland, Finn and Edward Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," Journal of Political Economy 85 (1977): 473–492.

Maskin, Eric and Jean Tirole, "Models of Dynamic Oligopoly I: Overview and Quantity Competition with Large Fixed Costs," Econometrica 56 (1988a): 549–569.

________ and ________, "Models of Dynamic Oligopoly II: Competition Through Prices," Econometrica 56 (1988b): 571–599.

McCallum, Bennett, "On Non-Uniqueness in Rational Expectations Models: An Attempt at Perspective," Journal of Monetary Economics 11 (1983): 139–168.

Rogoff, Kenneth, "Reputation, Coordination, and Monetary Policy," Mimeo, Stanford University, 1987.

Romer, Christina and David Romer, "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz," NBER Macroeconomics Annual, 1989.

Tobin, James, "Inflation and Unemployment," American Economic Review 62 (1972): 1–18.

Vickers, John, "Signalling in a Model of Monetary Policy with Incomplete Information," Oxford Economic Papers 38 (1986).