Carnegie-Rochester Conference Series on Public Policy 26 (1987) 141-182
North-Holland

REPUTATIONAL CONSTRAINTS ON MONETARY POLICY

KENNETH ROGOFF*
University of Wisconsin
and
National Bureau of Economic Research

I. INTRODUCTION

Research on strategic issues in macroeconomic policy design is proceeding at a rather rapid pace, and there is some risk involved in any attempt to survey it.[1] To a very limited extent, one can anticipate new developments in strategic macro policy by using the applied game theory literature as a leading indicator. The limitations stem from the fact that most of the applied game theory literature has evolved around applications to industrial organization, and the more literal translations of these results into macroeconomics have not been very successful. The problem is partly that the models can be quite sensitive to the specification of the institutional environment, and more importantly that abstractions which are plausible in the study of duopolies are not necessarily plausible in the study of macroeconomic policy.

It will be convenient to treat monetary policy as our generic example of macroeconomic policy, in part because much of the extant literature concentrates on monetary policy. But the issues raised here are clearly germaine to, say, taxation and government spending.

Early analyses of the "time consistency" problem of monetary policy demonstrated the possibility that the government might be able to increase its own welfare, and in some instances social welfare, if only it could tie

*I am grateful to David Backus, Maurice Obstfeld, Guido Tabellini, Marco Terrones, and Paul Wood for comments on an earlier draft.

[1]Fischer (1986) and Cukierman (1985) have recently provided excellent surveys of the monetary-policy credibility literature. The emphasis here is quite different, and therefore this study should be viewed as a complement rather than a substitute for these earlier surveys.

0 167 – 2231/87/$03.50 © 1987 Elsevier Science Publishers B.V. (North-Holland)

its hands and precommit to a (perhaps state-contingent) path for the money supply.[2] This can be the case even if there are no exogenous disturbances, and even if the government is trying to maximize the welfare of the representative individual. That is, the optimal money supply rule is not always subgame perfect.[3] A main theme of the recent literature is that by focusing on "one-shot" games, the early analyses <u>may</u> have overstated the government's credibility problems. Because monetary policy involves repeated interactions between the government and the public, reputational considerations can mitigate, or even eliminate, the time-consistency problem.

Whereas the current generation of reputational models of monetary policy have very appealing features, they also have one fundamental limitation. Typically, the models either yield a multiplicity of equilibria, or else yield an equilibrium which is extremely sensitive to the assumed information structure. This defect, which is inherited from antecedent game theory models, is well-known to careful readers of the policy credibility literature. But because many articles focus, perhaps excessively, on the most efficient attainable equilibria, casual readers may not fully appreciate how important the uniqueness question may be. It is true that the new reputation models suggest ways in which the government can be induced to behave "cooperatively," even when there is no legal mechanism for enforcing its good behavior. There is as yet no compelling argument, however, as to why, out of the continuum of reputational equilibria, the economy will coordinate on a "good" equilibrium and not a "bad" equilibrium. There is a real sense in which these repeated game models replace a cooperation problem with a coordination problem. Resolution of this question is central to understanding the implications of time consistency for government policy.

An extreme reputation view is that time consistency is not a serious issue in policy provided that the government places significant weight on the future. Hence, because it is virtually impossible to foresee every type of problem which will confront society (that is, there is qualitative

---

[2]See Kydland and Prescott (1977) and Calvo (1978). Phelps (1967) and Phelps and Pollak (1968) anticipated some of the basic themes underlying the modern time-consistency literature.

[3]For a discussion of the relationship between time consistency and subgame perfection, see Fershtman (1986). Subgame perfect equilibrium (or more generally, sequential equilibrium) will be the equilibrium construct used here.

uncertainty), it is unwise and unnecessary to try to legally bind the government. A less sanguine view is that all of the continuum of equilibria are equiprobable (since we have as yet no theory for choosing among them). Of course, if governments have high discount rates, the issue is moot; all the reputational equilibria will be qualitatively similar to the equilibrium of the one-shot game. Some who hold this view have suggested that time consistency problems imply a need to constrain monetary policy constitutionally. [An intermediate position is presented in Rogoff (1985b). I argue that the social institutions which evolve in response to time consistency problems represent a compromise between the benefits of complete flexibility and the need for precommitment.]

In Section II, I begin by reviewing the model of central bank reputation first proposed by Barro and Gordon (1983a). Their reputation mechanism is a variant of the infinite-horizon trigger-strategy equilibrium proposed by J. Friedman (1971). In a discrete-time version of the model, the equilibrium inflation rate can take on a range of values, with the range depending on the central bank's discount rate and on the length of the "punishment" period. The multiplicity of supergame equilibria can be drastically curtailed by imposing that the public's expectations of future inflation be continuous in current inflation. Although this assumption has some appeal, it is not clear how to rigorously justify it. Also, there is a danger of throwing out the baby (any good reputational equilibria) with the bathwater. To further illustrate the multiple equilibrium problem, I extend the analysis to admit "severe" punishment strategies analogous to those considered by Abreu (1982). By allowing for this class of equilibria, I show that it is possible, for a given discount rate and punishment interval, to sustain lower inflation rates than would be possible under the expectations rules considered by Barro and Gordon. It would seem important to recognize the existence of such equilibria in evaluating any casual arguments concerning how an equilibrium is chosen. Severe punishment strategy equilibria are also relevant when the central bank has private information, as in Section IV.

In Section III, I examine the case where the policymaker has a finite horizon. Unless the equilibrium of the one-shot game is unique, then it is still possible to have trigger-strategy equilibria analogous to those considered in Section II. Moreover, even if the range of one-shot game equilibrium inflation rates is very narrow, the range of repeated game equilibria can be very broad. The range again depends on the policymaker's time horizon and discount rate. It may even be possible to have an equilibrium where inflation is zero during the initial periods. It is interesting that these results do not appear sensitive to the assumption that there is literally a finite horizon. Qualitatively similar equilibria are shown to obtain (even when the one-shot game equilibrium is unique) where the policymaker has an infinite horizon but heavily discounts periods which come after the end of his term in office.

An alternative finite-horizon formulation has been developed by Tabellini (1983), Backus and Driffill (1985), Barro (1986), and by Horn and Persson (1985). In these analyses, which draw heavily on the work of Kreps and Wilson (1982), the public is not certain what type of policymaker they are facing. For this reason, policymakers who may be tempted by the transitory gains from unanticipated inflation have an incentive to pose as hard-money types. Once the public is certain that the policymaker is not a hard-money type, inflationary expectations will rise. This gives "soft-money" types an incentive not to reveal themselves too early. A possible advantage of this formulation is that for some variants, there is a unique equilibrium. There are drawbacks, however. The approach requires one to specify priors for the public, and it is not clear where these come from. Ideally, one would like to endogenize the evolution of priors across regimes. Also, superficially minor changes in the public's beliefs can significantly affect the nature of the equilibrium. For example, it matters whether the hard-money type is someone who places a higher weight on inflation than average, or whether he is a "robot" who is programmed never to inflate. In the free-will case, a type who is very tempted to inflate may never find it worthwhile to pose as a hard-money type. As Vickers (1985) has shown, the hard-money type may be able to take actions (deflate temporarily) which separate himself from high-inflation types.

Another possible problem with the adaptations of Kreps' and Wilson's model is that certain of the results might be sensitive to the assumption that there is a fixed finite horizon. I argue that it is difficult to find an example in which the finite-horizon assumption can be taken literally. Two other features of the models have attracted criticism, but these criticisms can be addressed. First, the existing models allow for only two types of policymakers. Second, the equilibria involve randomizing strategies. In a self-contained appendix, I illustrate one way to extend these models to allow for a continuum of types of policymakers, instead of just two. The model of the Appendix has an equilibrium in pure strategies with pooling. That model also illustrates why it is important how the hard-money type(s) are specified.

In Section IV, I present two views of how private information may impinge on the analysis. Canzoneri (1985) argues that it is impossible for the central bank to precisely control the price level. The Friedman-type reputational equilibria of Section II cannot be sustained if the public can never directly observe how much of any given price-level change was intentional, and how much was due to an incorrect forecast of money demand. Following Green and Porter (1984), and Barro and Gordon (1983a), Canzoneri demonstrates that it is still possible to have reputational equilibria, though to sustain them the economy must suffer periodic reversions to a high inflation equilibrium. There remain a multiplicity of equilibria. One reason Canzoneri's analysis is interesting is that it illustrates how the problem of coordinating on the best equilibrium seems to become more acute when there is private information. Canzoneri suggests that if private information is indeed the explanation behind the economy's periodic bouts with inflation, then attempts to achieve monetary policy credibility through legislation will have inherent limitations. Cukierman and Meltzer (1986) share the view that the central bank's private information is important. They argue, however, that the level of private information is endogenous. In their model, the central bank has incentives to adopt imperfect monetary control procedures so that it can mask its intentions. Obviously, their theory has somewhat different implications for institutional reform. Cukierman and Meltzer show that their model has a unique linear equilibrium but do not provide a complete resolution of the multiple-equilibrium problem.

In Section V, I summarize several issues which arise when there is more than one government controller. In the international context, this is relevant because of sovereign governments. In the domestic context, it may be important when there are two or more quasi-independent government agencies and when there are two or more political parties. The introduction of multiple controllers suggests a range of interesting applications and adds a new dimension of strategic complexity. In the Conclusions, I ask whether the models of reputation developed to date are compelling. Can we rely on reputational considerations to accomplish what we once thought could only be accomplished through institutional reform?

II. CREDIBLE MONETARY POLICIES IN THE INFINITE-HORIZON CASE

In the first part of this section, I review the "trigger-strategy" model of monetary policy reputation due to Barro and Gordon (1983a). As Barro and Gordon stressed, there are a multiplicity of equilibria of the type they consider. There are also other classes of equilibria, as I demonstrate by extending their analysis to allow for an analogue of the "severe" punishment strategies identified by Abreu (1982) (in a different context). I then speculate on how it may ultimately be possible to modify these models to produce more definite results.

The framework for analyzing monetary policy credibility I will employ is a slight variant of a popular example due to Kydland and Prescott (1977). One justification for using this extremely simple model is that it forms the basis for virtually all the literature surveyed below. Obviously one would want to use a more fully articulated model for purposes of applied policy analysis. But the Kydland and Prescott example is very convenient for illustrating strategic factors, which may easily become obscured in a more complex model.[4]

Monetary policy can have real effects in our model because private agents form expectations of period t inflation, π_t, based on t - 1 information.[5] It is important to emphasize that the atomistic agents are "expectations takers." The aggregate inflation rate, π, is exogenous to the individual; he can only affect his own price prediction error, $\pi_t - (\pi^{ei})_t$. I stress this point because in the analysis below, it is easy to become confused into thinking that individuals are setting their expectations strategically. What is true is that there are equilibria in which the collective actions of private agents have a strategic effect on the government's choice of monetary policy. But these equilibria do not require any explicit cooperation within the private sector. Any individual who "defects" and tries setting his expectations differently will only be punishing himself.

The fact that the individual cannot affect the aggregate inflation rate or the aggregate prediction error does not necessarily imply that

---

[4] It is possible to restate the analysis below in terms of an overlapping generations version of the model presented in Section I of Fischer's (1986) survey. For efforts along these lines, see Atkeson (1986) or Kehoe (1985).

[5] In the underlying structural model, money can have real effects either because of confusion between local and aggregate disturbances, or because there are imperfectly-indexed wage contracts. In the former case, there must be a temporal lag in the diffusion of aggregate information.

these factors do not enter his utility function.[6] Consider the case, for example, where there is an externality arising from income taxation. When other citizens are "tricked" into working too much, or into holding too high a level of real money balances, the individual gains because government revenues rise. However, it never pays for the individual to try intentionally to guess wrong himself. Thus we will assume that an individual attempts to minimize[7]

$$J_t = [\pi_t - (\pi^{ei})_t]^2. \tag{1}$$

In most of the monetary-policy credibility literature, it is assumed that unanticipated inflation increases output (via a contracts' or an islands' model). Many parallel issues arise when unanticipated inflation matters because the government is trying to raise revenue by creating money or by inflicting capital losses on holders of nonindexed nominal bonds. I will assume that the loss function of the monetary authorities is given by

$$\Omega_t = \sum_{s=0}^{T} L_s \beta^{(s-t)}, \tag{2a}$$

$$L_s = f[\pi_s - (\pi^e)_s - k] + g(\pi_s), \tag{2b}$$

where $k > 0$, $f'(\hat{\ })$, $g'(\cdot) \gtrless 0$ for $(\cdot) \gtrless 0$, and $f'(\cdot)$, $g'(\cdot) \to 0$ as $(\cdot) \to 0$. b is the monetary authorities' subjective discount rate, and T is their time horizon. $\pi - \pi^e$ is the average level of private-sector price prediction errors. For now, we will assume that T is infinite and that $f''(\cdot)$, $g''(\cdot) > 0$; both assumptions will be relaxed in Section III. The basic structure underlying equations (1) and (2) has been extensively examined [see Barro and Gordon (1983a,b), Canzoneri (1985), Rogoff (1985b),

---

[6] It is difficult to argue that anticipated inflation has an effect of the same order of magnitude as unanticipated inflation. (There are the "shoe-leather" costs resulting from lower holdings of real money balances. Also, there may be some activities, such as income tax accounting, which are costly to properly index.) Our analysis does not require that the welfare effect of anticipated inflation be large.

[7] The specification (1), though used throughout the literature, is not entirely satisfactory for stochastic versions of the model, because if the level of inflation and the aggregate prediction error do affect the individual's utility function (because of externalities), he must take into account the covariance of these factors with his own price-prediction error.

or Tabellini (1983)]. Barro and Gordon discuss how, in the presence of externalities such as income taxation (our example above), the government's objective function can be interpreted as the social-welfare function, even though $k > 0$. If $k = 0$, then there is no externality, and it will turn out that the optimal monetary rule is subgame perfect.

Before considering repeated-game reputational equilibria, it is useful to first examine equilibria of the "one-shot" game (there is only one period). Because the private sector forms expectations about period t inflation based on t - 1 information, the central bank treats $(\pi^e)_t$ as given when setting π_t. Minimizing Ω over π, we obtain the first-order condition $-f'(\pi - \pi^e - k) = g'(\pi)$. Since the public forms expectations rationally, we require that $\pi = \pi^e$. Hence a necessary condition for a subgame perfect equilibrium in the one-shot game is

$$-f'(-k) = g'(\pi^*), \tag{3}$$

where π^* is the one-shot game equilibrium inflation rate. Given our assumptions that f'', $g'' > 0$, it is straightforward to show that the second-order conditions hold and that π^* is unique. The logic underlying the equilibrium characterized by (3) is well-known. The central bank always has the ability to inflict price-prediction errors on the private sector. But when $\pi^e = \pi^* > 0$, it will never choose to do so.[8] As inflation rises so, too, does the marginal cost of inflating. The time-consistent equilibrium level of inflation, π^*, is sufficiently high so that the marginal gain from surprise inflation equals the marginal cost.

In this nonstochastic model, the fact that the central bank can exercise discretion brings no benefits and only leads to a high rate of inflation. Recall that private agents do care about the aggregate inflation rate. However, because an individual's actions have only an infinitesimal effect on the aggregate price level, each agent acts as if he were concerned only with his own price-prediction error. Whether or not the economy will coordinate on a more favorable equilibrium (without imposing legal restraints) is the main focus of our investigation. The

---

[8] In a stochastic version of the model, the central bank may choose to cause price-prediction errors, though private-sector expectations will be correct on average. [See Barro and Gordon (1983a) or Rogoff (1985b).]

equilibrium characterized by (3) is of interest for a number of reasons. First, it is the unique subgame perfect equilibrium when the monetary authorities maximize over any finite horizon (if, as in our example, the equilibrium of the one-shot game is unique). Moreover, π^* remains an equilibrium when their horizon is infinite. Second, in the infinite-horizon case, the one-shot game equilibrium can serve as a credible threat to induce more "cooperative" behavior from the monetary authorities. We will now illustrate this point. This class of reputational equilibria was demonstrated by Barro and Gordon (1983a).

Consider a level of inflation, $\tilde{\pi}$, such that $0 \leq \tilde{\pi} < \pi^*$, and suppose that the public forms expectations according to

$$(\pi^e)_t = \begin{cases} \tilde{\pi} \text{ if } \pi_{t-1} = (\pi^e)_{t-1} \\ \\ \pi^* \quad \text{otherwise.} \end{cases} \tag{4}$$

Thus if $(\pi^e)_{t-1} = \tilde{\pi} < \pi^*$, the public will continue to expect low inflation as long as the central bank "cooperates" and does not try to fool them.[9] If the central bank ever inflates beyond $\tilde{\pi}$, the economy will be subjected to a "punishment" interval, which we have arbitrarily set at one period. (When $\pi_{t-1} > \tilde{\pi} = (\pi^e)_{t-1}$, π^e reverts to π^*. If the central bank then sets $\pi_t = \pi^*$, π^e reverts back to $\tilde{\pi}$.) It is very important to note that the public's expectations are rational in the subgame which would occur if the central bank were to "cheat." The central bank has absolutely no incentive to surprise private agents during a punishment period. For by setting $\pi = \pi^*$ during a punishment period, it minimizes both this period's loss function and next period's inflationary expectations.

We shall now confirm that there are indeed equilibria where the public forms its expectations according to (4), and where $\tilde{\pi} < \pi^*$. To determine whether $\tilde{\pi}$ is a trigger-strategy equilibrium level of inflation under (4), it is necessary to consider whether the central bank will have any incentive to defect and set $\pi \neq \tilde{\pi}$. This question turns on the magnitude of the maximum current-period gain from defecting, $B(\tilde{\pi})$, in comparison with the expected future cost to defecting, $C(\tilde{\pi})$. These magnitudes are given by

---

[9]Because f'', $g'' > 0$, the same results would obtain if we replaced (4) with the weaker condition: $(\pi^e)_t = \tilde{\pi}$ if $\pi_{t-1} \leq (\pi^e)_{t-1}$, and $(\pi^e)_t = \pi^*$ otherwise. In the private-information example studied in Section IV, it is necessary to make this modification.

$$B(\tilde{\pi}) = f(-k) + g(\tilde{\pi}) - f[\pi^D(\tilde{\pi}) - \tilde{\pi} - k] - g[\pi^D(\tilde{\pi})] > 0, \tag{5}$$

where $\pi^D(\tilde{\pi}) = \text{argmin}[f(\pi - \tilde{\pi} - k) + g(\pi)]$, and

$$C(\tilde{\pi}) = \beta[g(\pi^*) - g(\tilde{\pi})] > 0. \tag{6}$$

For a given level of $\tilde{\pi}$ to be an equilibrium it is necessary that $B(\tilde{\pi}) \leq C(\tilde{\pi})$; otherwise the central bank will always choose to defect. Though $\tilde{\pi} = 0$ may not be an equilibrium[10], it is possible to prove that there always exists some $\tilde{\pi}$ such that $0 \leq \tilde{\pi} < \pi^*$ and $B(\tilde{\pi}) \leq C(\tilde{\pi})$. [Proof: $\tilde{\pi} < \pi^D(\tilde{\pi}) < \pi^*$ by f'', $g'' > 0$. Let $\tilde{\pi} = \pi^* - \varepsilon$. Since $-f'(\pi^D - \tilde{\pi} - k) = g'(\pi^D)$, and since $\pi^D - \tilde{\pi} < \varepsilon$, then $B(\tilde{\pi})$ must become second order as ε becomes small (by an envelope theorem argument). Since $C(\tilde{\pi})$ remains first order for small ε then, by the continuity of f and g, there must exist some $\varepsilon > 0$ such that $B(\pi^* - \varepsilon) \leq C(\pi^* - \varepsilon)$.]

Denote $\hat{\pi}$ as the lowest (positive) inflation rate which can be a trigger-strategy equilibrium level of inflation under (4). It is trivial to show that $\hat{\pi}$ is nonincreasing in β, the central bank's discount rate. It is also simple to show that if $\hat{\pi} > 0$, then it would be possible to have lower inflation if the expectations mechanism of the public embodied a punishment interval longer than just one period. If the discount rate β is small, however, even an infinite punishment interval may not be enough to sustain zero inflation.[11]

When $\hat{\pi} > 0$, there is another mechanism for sustaining a lower inflation rate, one which does not involve extending the punishment period. The alternative mechanism involves having a more severe punishment, instead of a more prolonged punishment. The more severe

---

[10] If $f(\cdot) = (\pi - \pi^e - k)^2$, $g(\cdot) = \pi^2$, and $\beta = 1$, then $\tilde{\pi} = 0$ is a trigger-strategy equilibrium under (4).

[11] It is easy to prove that zero inflation is always attainable if the time interval between periods is small enough, provided that the length of the punishment interval is unrestricted. In (1) and (2), we have arbitrarily set the time interval at one. As the interval approaches zero, the transitory output gains from defecting become very short-lived, whereas the punishment can be held constant. [Grossman and Van Huyck (1986) make this point in the context of an optimal seigniorage model. Their result requires the restriction that there be some maximum rate at which the central bank can print money. However, the constraint is never binding along the equilibrium path, unless one introduces private information as in Section IV.] Hence, if one chooses to rationalize (1) and (2) via an islands model, then a zero inflation rate is always attainable unless there is an explicit time lag in the diffusion of aggregate information.

punishment consists of reverting to an inflation rate higher than π^* whenever the central bank defects. In some applications, this alternative mechanism may be important because it is not intuitively appealing to have a long or infinite punishment interval. Also, severe punishments can play a role in the optimal equilibrium of the model with private information, which will be studied in Section IV. I do not, however, regard severe punishment equilibria as being particularly plausible in the present context. My primary motivation for introducing this alternative class of equilibria here is to underscore the severity of the multiple equilibrium problem. Thus it is sufficient merely to illustrate the equilibria, and we will not concern ourselves with deriving the optimal severe punishment equilibrium.

To make the mechanism underlying severe punishment equilibria more transparent, it is helpful first to demonstrate why it is possible for inflationary expectations to rise temporarily above π^*. Let $\delta > 0$, and consider the following path of expectations initiating in period t:

$$(\pi^e)_t = \pi^* + \delta,$$

$$(\pi^e)_{t+i} = \begin{cases} \hat{\pi} \text{ if } \pi_{t+i-1} = (\pi^e)_{t+i-1}, & i \geq 1 \\ \pi^* \text{ otherwise.} \end{cases} \tag{7}$$

Since $g(\pi^*) - g(\hat{\pi})$ is finite, it is clearly possible to choose a δ small enough so that $\pi^* + \delta$ is an equilibrium for period t, provided that the public's expectations are governed by (7). It is true that at $\pi^* + \delta$, the central bank would be willing to let output drop below the natural rate in order to achieve lower current-period inflation. But the central bank knows that it must be willing to suffer through exceptionally high inflation in period t if it wants inflation in t+1 to be $\hat{\pi}$, and not π^*. The fact that $\pi^* + \delta$ can be made a credible threat implies that it is possible to attain an inflation rate lower than $\hat{\pi}$ without extending the punishment interval. For example, consider an equilibrium analogous to (4):

$$(\pi^e)_t = \begin{cases} \tilde{\pi}' \text{ if } \pi_{t-1} = (\pi^e)_{t-1}, \\ \pi^* + \delta \text{ otherwise.}^{12} \end{cases} \tag{8}$$

It is clear that the lowest attainable inflation rate under (8) is lower than the lowest attainable inflation rate under (4), ($\hat{\pi}$), if $\hat{\pi} > 0$.

The trigger-strategy equilibria we have been analyzing do not require explicit cooperation between the private agents, or between the private sector and the central bank. If an atomistic private agent believes that other agents form inflationary expectations according to (4) [or (8)], then it is only rational for him to form expectations the same way (if the equilibrium is subgame perfect).[13] However, although these low-inflation equilibria do not require explicit cooperation across individual agents, there is a serious question of how agents coordinate on a particular equilibrium. First, what is to be the length of the punishment interval? Given the length of the punishment interval, is there any reason to suppose that the public will expect the lowest equilibrium inflation rate corresponding to this punishment interval? Even if we assume that the public can coordinate on the punishment interval and can agree to expect the lowest credible level of inflation, there is still a degree of indeterminacy. Will $\pi^e = \hat{\pi}$, the lowest attainable inflation rate under (4), or will the public expect the lower inflation rate attainable under a severe punishment strategy such as (8)?

It might be argued that by making pronouncements about its monetary policy, the government can focus the private sector's attention on a particular equilibrium. But I, for one, am extremely uncomfortable with this reasoning. The government has obvious incentives to make false announcements, and the public is not likely to pay attention to statements that are not backed by concrete measures. A somewhat more serious alternative is to explore whether the government can achieve some degree of coordination by placing external restraints on itself. (Such as making a

---

[12]Again, it is possible to define defections in terms of inequalities instead of equalities, but this does not make any qualitative difference in the case of symmetric information.

[13]Some have criticized Barro and Gordon's trigger-strategy equilibrium as being subject to a "free-rider" problem. This criticism is not well-founded. As the discussion in the text makes clear, the equilibrium is indeed Nash. There might be a free-rider problem if it were necessary to raise funds to improve coordination of private-sector expectations.

commitment to a fixed exchange-rate system which, if violated, would lead to the breakdown of a tariff agreement.) This resolution, of course, really amounts to changing the structure of the game so that there are less equilibria.

Perhaps the most implausible feature of the equilibria considered in this section is that (except for the one-shot game equilibrium), they require that the public's expectations about future inflation be discontinuous functions of current inflation. If the government defects by a small amount, expected inflation rises by just as much as if the government were to inflation massively. It would seem worthwhile to explore assumptions which imply that continuous changes in the environment lead to continuous changes in the public's beliefs about future inflation. Whereas it may still be possible to have reputational equilibria with continuous reaction functions, the severity of the multiple-equilibrium problem might be significantly diminished. [For a discussion along these lines within the context of I-O models, see Gul, Sonnenschein and Wilson (1986). Stanford (1986) provides some suggestive results.] Yet another approach to placing restrictions on the possible equilibria is to recognize that agents cannot make an unlimited number of calculations; see, for example, Rubinstein (1986).

The analysis above is readily generalized to the stochastic case, if information is symmetric. (In Section VI, we shall consider the case of asymmetric information.) Barro and Gordon (1983a) have illustrated some of the possibilities which can arise. The optimal trigger-strategy equilibria will involve having the public make (unforecastable) price-prediction errors. When a disturbance causes the benefits to unanticipated inflation to be unusually high, the monetary authorities engineer a surprise inflation. Expectations of inflation are still correct, on average, because the monetary authorities spring surprise deflations when the benefits are low. It should be noted that because the models studied in this section have multiple equilibria, they can generate variable or stochastic inflation even in a completely unchanging environment. It is possible to have trigger-strategy equilibria which bounce back and forth between different points either with certainty, or with reference to an extrinsic random variable (sunspots).

III. REPUTATION IN A FINITE-HORIZON CONTEXT

The trigger-strategy equilibria discussed in the preceding section break down if the monetary authorities maximize over a finite horizon. (Because only the one-shot game equilibrium can obtain in the last period, the strategies "unravel" backwards.) It has been suggested that this may be an important problem, because policymakers have finite terms in office. Before questioning the merits of this view, we will discuss mechanisms for modelling reputation in the finite-horizon case.

First of all, if there are multiple equilibria in the one-shot game, then there can exist trigger-strategy equilibria in the finite-horizon case. Benoit and Krishna (1985) and Friedman (1985) have demonstrated this general principle. In fact, even if only a narrow range of high-inflation rates can obtain in the one-shot game, it is still possible to sustain inflation rates very close to zero early in the policymaker's term (if his term is long enough and his discount factor low enough). Let $\pi_1 > 0$ be the lowest equilibrium-inflation rate in the one-shot game, and let π_2 be the highest. During the policymaker's final period in office, period T, the equilibria are the same as in the one-shot game. Hence $\pi_1 \leq \pi_T \leq \pi_2$. However, in period T - 1, it may be perfectly rational for the public to believe that $(\pi^e)_{T-1} = \pi_1 - \varepsilon$, $\varepsilon > 0$. For small enough ε, these expectations can be supported by the belief that if the government does not defect in T - 1 ($\pi_{T-1} = \pi_1 - \varepsilon$), then $(\pi^e)_T = \pi_1$. If it defects ($\pi_{T-1} \neq \pi_1 - \varepsilon$), then $(\pi^e)_T = \pi_2$. Hence defection is punished by going to the "bad" Nash equilibrium in the final period. Let $(\pi')_s$ denote the lowest subgame perfect equilibrium π which can be attained in period s. It is straightforward to show that if $(\pi')_s > 0$, then $(\pi')_s - (\pi')_{s-1} > 0$. (I am implicitly assuming that the discount rate β is constant.) The more periods that remain, the longer the punishment interval can be. Also, as we move back from date T, the maximum one-period punishment, $\pi_2 - (\pi')_s$, rises. To be more concrete, we present an example:

Suppose one replaces the central bank's inflation loss function, $g(\pi)$ [see equation (2b)], with the following loss function

$$h(\pi) = \begin{cases} g(\pi) & \text{for } \pi \leq \pi^* + k, \\ g(\pi^* + k) & \text{for } \pi^* + k < \pi \leq z, \\ g(\pi^* + k) + g(\pi - z) & \text{for } z < \pi, \end{cases} \tag{9}$$

where z is a sufficiently large constant such that $-f'[-(k+z)] \geq g'(\pi^* + k)$. If we replace $g(\cdot)$ with $h(\cdot)$ in equation (2b) then, as one can easily confirm, there are two equilibria in the one-shot game, π^* and $\pi^* + z$. Now let us assume that the central bank maximizes over a two-period horizon, and that it does not discount second-period welfare ($\beta = 1$). Consider what happens if the public forms expectations of π as follows:

$$
\begin{aligned}
(\pi^e)_{T-1} &= 0 \\
(\pi^e)_T &= \begin{cases} \pi^* & \text{if } \pi_{t-1} \leq 0, \\ \pi^* + z & \text{otherwise.} \end{cases}
\end{aligned}
\tag{10}
$$

When confronted with the inflation expectations mechanism (10), the central bank will ratify the public's beliefs and set $\pi_{T-1} = 0$. If it sets $\pi_{T-1} > 0$, then it will bear a cost in period T of $h(\pi^* + z) - h(\pi^*)$ which, by construction of h, is equal to $g(\pi^* + k)$. This cost outweighs any possible gain. The gain to inflating in period $T - 1$ is strictly less than $f(-k) - f(0)$, which is strictly less than $g(\pi^* + k) - g(\pi^*)$. Hence $\pi^e = 0$ is a subgame perfect equilibrium for period $T - 1$. As in the infinite-horizon case, there are a multiplicity of equilibria. For example, any inflation rate less than π^* but greater than zero can be equilibrium for period $T - 1$.

It is interesting to observe that an analogous equilibrium arises in the case where (a) the one-shot game equilibrium is unique, and (b) the policymaker does not literally have a finite horizon, but heavily discounts events which will occur after he leaves office. Suppose, for example, that we replace the policymaker's loss function, (2a), with the alternative function

$$
\Omega_t = \sum_{s=t}^{\infty} L_s \beta(s)^{(s-t)}, \tag{11}
$$

where L_s is again given by equation (2b), but now $\beta(s) = 1$ for $s \leq T$, and $\beta(s) = \varepsilon$ for $s > T$, with ε being very small. In his final period in office, T, the lowest attainable trigger-strategy equilibrium level of π, $\hat{\pi}_T$, will be very close to π^*, since future periods are discounted very heavily. (This assertion is easily confirmed.) Nevertheless, the small wedge between $\hat{\pi}_T$ and π^* is sufficient to support a level of inflation $\hat{\pi}_{T-1} < \hat{\pi}_T$ via the same general argument as above. If the policymaker's term is long enough, it will be possible to credibly sustain a very low inflation rate during his initial periods in office.

A rather different model of monetary policy credibility is based on the assumption that the public is unsure about the policymaker's preferences or about his cost of breaking commitments. A number of researchers have adopted this approach, applying the framework of Kreps and Wilson (1982). Barro's (1986) version of the model is roughly as follows:

The policymaker has a fixed term in office. His horizon is finite either because has has no reason to care what happens to social welfare after his departure, or because he believes that his actions do not affect the credibility of the monetary authorities in future periods. Upon entering office, the policymaker makes a commitment never to inflate. The public thinks there is at least a small chance that the policymaker is a "type 1," for whom it is prohibitively expensive to break his commitment. Otherwise the policymaker is a "type 2," who bears no cost in breaking commitments. (In the Appendix, I extend the model to allow for a continuum of types.) Barro shows that if terms of office are long enough, then there will be no inflation in the early periods of a term regardless of which type the policymaker actually is. At some point, depending on the public's initial priors on the policymaker's type, the type-2 policymaker begins to randomize his behavior, inflating with a time-dependent, endogenously-determined probability. At the point randomization begins, private-sector expectations of inflation rise by a discrete amount (because there is now a chance inflation will occur). As long as the public continues to observe zero inflation, expected inflation fluctuates around a constant level. If the public observes any inflation, then it knows the policymaker is a type 2, and expected inflation rises to its one-shot game level, π^*.

Two features of this scenario seem odd at first glance but appear more reasonable upon closer inspection. First, it is not very appealing to think of the central bank as flipping a coin to decide whether or not to inflate. This aspect of the model is, in part, attributable to the fact that there are a discrete number of types. In the model of the Appendix, which has a continuum of types, there does exist a pure-strategy equilibrium.

A second "odd" characteristic of the model is that the expected rate of inflation rises only at the beginning of the randomization interval, and then remains constant until the public actually observes inflation. This is the result of two offsetting effects. On the one hand, the public knows that a type-2 policymaker is more likely to inflate as the end of the term approaches, and will certainly inflate in the last period. Offsetting this effect is the fact that each time a period passes and no inflation occurs, the public raises its probability that it is indeed facing a type-1 (precommitted) policymaker. This result is not general. In the model of the Appendix, it is possible to have expected inflation rise over the period in which the policymaker might break his commitment.

Also, if one conditions expectations only on information available to the public in period zero, then the path of expected inflation does indeed rise over the randomization interval. During the randomization interval, the cumulative probability that the type-2 policymaker will have revealed himself by the end of any given period rises over time. Once he reveals himself, then inflation rises to π^*.

For any initial set of public beliefs, the model discussed above yields a unique equilibrium. However, the public's priors are a "free parameter" which have an important effect on the predictions of the model. Also, as Vickers (1986) has shown, it makes a considerable difference whether a type-1 agent is someone who legally binds himself to a fixed low inflation rate, or whether he is simply a hard-money type who places a greater than average weight on inflation. In the latter case, the hard-money type may be able to signal his type of deflating. If there are not too many periods left, a type-2 policymaker may prefer to unmask himself rather than suffer a large deflation. A type-1 agent, on the other hand, finds the sacrifice worthwhile if it proves his type to the public. Thus in the separating equilibrium, type-2 policymakers set inflation at π^*. Type-1 policymakers initially deflate. Then, having separated themselves, they inflate at a low level (which depends on the relative weight they place on inflation versus unemployment). In a richer institutional setting, hard-money types might have other ways to send separating signals. (They can send signals via the budget deficit, they can appoint conservatives to govern the central bank, etc.) The issue of whether or not all policymakers have at least some discretion also arises in the model of the Appendix.

Possibly the least robust results for this class of models are those pertaining to the policymaker's final periods in office, the "endplay" of the model. These results may be sensitive to the assumption that the policymaker has a known finite horizon, an assumption I will now argue is seldom plausible. First of all, let us consider the case where type-2 policymakers genuinely care only about the social-welfare function, (2), and where society's horizon is infinite. It is true that for any one play of the game, the public will be better off when the policymaker actually turns out to be a type 2, [(since the one-time inflation surprise raises social welfare; see Barro (1986).] However, the public is also better off the higher its initial subjective probability that the policymaker is actually a type 1. And it seems implausible for a type-2 policymaker to think that his actions this term will have no effect on the public's probability distribution over types in future periods and hence on future social welfare.

It is true that a policymaker's actions would not affect the public's beliefs about future policymakers if his successor is drawn at random from a large population, and if the public has very strong priors about the relative distribution of type-1 and type-2 agents in this population. But it seems (to me) more plausible to think that the policymakers are chosen via some nonrandom process. The observation that the latest Fed chairman was a type 2 ought to influence the public's priors as to the nature of his successor. Thus if the policymaker really cares about social welfare, he should take into account the effects of his actions on future periods.

It is entirely possible that the policymaker does not care at all about social welfare and only aims to maximize his seigniorage revenues while in office. There may be many countries where this scenario is plausible, but in these countries, policymakers usually do not have fixed one-time terms in office. Moreover, the length of their terms in office is probably not exogenous.

Another rationale for the policymaker's fixed finite perspective might be the electoral cycle. Backus and Driffill (1985) note (only in passing) that their model yields something akin to a political business cycle. On average (that is, averaging over both type-1 and type-2 policymakers), inflation tends to be higher towards the end of a term. There is no tendency for output to be high on average before elections; the booms which occur during the regimes of type-2 agents are cancelled out by the recessions which take place under type-1 agents. A problem with an electoral-cycle interpretation of the model is that it contradicts empirical evidence that pre-election distortions in policy are most severe when the incumbent is up for re-election, not during his final term in office. Rogoff and Sibert (1986) present an equilibrium-signaling model in which the electoral cycle in macroeconomic policy arises precisely because the incumbent party is striving to stay in power.

IV. PRIVATE INFORMATION

Thus far, we have assumed that the public can perfectly monitor the central bank's actions. Canzoneri (1985) and Cukierman and Meltzer (1986) have analyzed the implications of relaxing this assumption. In this case, achieving the coordination necessary to attain optimal reputational equilibria seems even more problematic. The studies considered in this section are also interesting because they have implications for attempts at institutional reform. We will first consider Canzoneri's model.

Canzoneri analyzes an infinite-horizon model similar to the model of Section II. He assumes that the central bank does not discount the future ($\beta = 1$), so that (for a long enough punishment interval) there always exists a trigger-strategy equilibrium in which expected inflation is zero. He then introduces money-demand shocks into the model. These shocks are observed only after the central bank has set the money supply. However, the central bank is able to condition its actions on a forecast of the money-demand disturbance. Its forecast is imperfect, so the central bank would be unable to completely damp out price fluctuations even if it were trying to minimize the price-prediction errors of private agents.

If the public is able to observe both the money-demand disturbance and the central bank's forecast of it, then no new conceptual issues arise. There are trigger-strategy equilibria analogous to those of Section II (for stochastic versions of the models). As long as the public can directly confirm that any unanticipated inflation is entirely attributable to an error the central bank made in forecasting, there is no need to "execute" any punishment. (Technically, of course, private agents do not act strategically, and there is no explicit cooperation among them.) Canzoneri argues, however, that it might be very difficult for the public to directly confirm the central bank's forecast and that this forecast should be treated as private information.[14] He then shows, by applying Green and Porter's (1984) extension of Friedman's trigger-strategy model, that it is

---

[14]Whether the central bank actually has any private macroeconomic information is debatable. If the central bank's forecast of money demand is based entirely on publicly available data, then the private sector should be able to construct the same forecast. It might be argued that the central bank has much faster access to data on bank deposits, and that this information is only released to the public with a long lag. Of course, if the central bank does not discount the future too heavily, then even the lagged release of data is still sufficient to have trigger-strategy equilibria similar to those analyzed in Section II.

still possible to have an equilibrium which improves on the outcome of the one-shot game.[15]

In the equilibrium Canzoneri analyses, the public sets expected inflation equal to zero, as long as the economy is not entering a reversionary (punishment) period. The public then observes actual inflation and employs a one-tailed test. If inflation is above a certain threshold value, then there will be a one-period reversion to the inflation-rate expectations of the one-shot game. If the threshold is set at just the right level, the central bank can be induced to target zero inflation.[16] (In setting the level of the money supply, the central bank must trade off increases in current employment with increases in the probability of entering a reversionary period.) Even though the central bank does not cheat (in equilibrium), large money-demand forecast errors still occur periodically, thereby throwing the economy into periods of high expected inflation.

As in the model of Section II, there is a multiplicity of other equilibria, and it is not clear how or why the public would coordinate on this particular one. It is not satisfactory to argue that this equilibrium is somehow "focal" because it is optimal. For one thing, the optimal equilibrium does not, in general, have such a simple structure. (Note that punishments actually occur in this model. So even if different punishment strategies yield the same level of expected inflation, they do not necessarily yield the same level of welfare.) Abreu, Pierce and Stacchetti (1985) have shown that optimal punishment strategies in the Green-Porter model typically involve an analogue of the severe punishment strategies discussed in Section II. They also show that the optimal strategies are not, in general, based on a simple one-tailed test (though the one-tailed test equilibrium is intuitively appealing). That the optimal trigger strategies can be so complicated, even when the underlying model has a relatively simple structure, is further reason to avoid loose arguments

---

[15]The Keynesian flavor of Canzoneri's analysis is not an essential ingredient. Barro and Gordon (1983a) discuss how to extend their model to the case where the monetary authority has imperfect control over the inflation rate, and where its control error is private information. They, too, consider the class of equilibria identified by Green and Porter (1984). In their paper, however, they do not present their formal results.

[16]Canzoneri presents only the first-order conditions necessary to sustain an equilibrium with zero-expected inflation. The second-order conditions obtain because of restrictions on the concavity of the distribution function of the monetary authorities' forecast error.

that the public will coordinate on the best equilibrium.

Canzoneri's model has both attractive features and, at least superficially, odd features. By introducing private information, Canzoneri is able to explain why there must be some inflationary bias (on average) even if the public can coordinate on the best attainable equilibrium. Also, the model illustrates how serially uncorrelated forecast errors can produce serially correlated inflation rates: reversionary periods follow a cooperative period in which inflation was high. On the negative side, the public's expectations mechanism does not seem particularly plausible. The public finds itself punishing the central bank periodically, even though it knows that the central bank would never cheat (in equilibrium). Whenever the central bank inadvertently allows inflation to slip above its threshold value, the public must punish it by discontinuously raising inflation-rate expectations. The punishment is necessary in order to induce the central bank to continue to target low inflation in the future. Note that the public never actually learns anything about the policymaker's type; it knows everything at the outset and knows that it would never pay for the central bank to defect from the equilibrium. Canzoneri's model should not be interpreted as one in which the central bank has private information about its preferences. His model has quite different properties than the models discussed in the latter part of Section III.

Canzoneri treats the information structure as exogenous. Cukierman and Meltzer (1986) suggest that the central bank may deliberately saddle itself with inefficient operating procedures in order to mask its intentions. Their basic argument may be illustrated in the following one-period model:

Let the central bank's objective function take the specific functional form:

$$L = -x(\pi - \pi^e) + \tfrac{1}{2}\pi^2. \tag{12}$$

(Cukierman and Meltzer do not interpret the central bank's loss function as a social-welfare function.) It is easily seen that in a one-shot game, the equilibrium inflation rate $\pi^* = x$. It is helpful to note that for the loss function (12), the central bank will set $\pi = x$ regardless of the value of π^e.

A key element of Cukierman and Meltzer's result is the assumption that the central bank's preferences are stochastic. Suppose, for example, that based on time t - 1 information, $x = 0$ in period t with probability ½, and that $x = 2$ in period t with probability ½. If the private sector is able to observe x before setting π^e, then $E_{t-1}(L_t) = ½(0) + ½(2) = 1$. [Thus $E_{t-1}(L)_t$ is the expected value of the central bank's period t objective function, prior to the realization of x.] If the private sector is unable to observe x before setting π^e, then $\pi^e = 1$, and $E_{t-1}(L_t) = ½(0) + ½(0) = 0$. Therefore if the central bank can precommit itself not to reveal x before the public sets π^e, it will choose to do so. Hiding its preferences does not allow the central bank to systematically fool the private sector, which can still guess inflation correctly, on average. Rather, the central bank gains because it is able to cause surprise inflation when the benefits to inflating are high and to save surprise deflations for periods when the benefits are low.

Cukierman and Meltzer's complete model involves an infinite horizon, with serially correlated preference shocks. The public never directly observes the central bank's preferences but is able to infer something about them from the path of the money supply. By intentionally adopting an imprecise monetary-control procedure, the central bank is able to obscure its preferences. It gains via the channel illustrated above, but it loses because inaccurate monetary control raises the variance of inflation. (The reader will have to look to Cukierman and Meltzer's article for further details.)

Cukierman and Meltzer focus on equilibria in which the public forms its beliefs using a linear feedback rule. They show that there is a unique equilibrium of this type; the equilibrium is analogous to a one-shot game equilibrium in the sense that all the dynamics come from the internal structure of the model. They do not prove that there are no other classes of equilibria, however. Given the fact that their model incorporates an infinite horizon, and given the pervasive information asymmetries in its structure, the issue of uniqueness would seem to be a serious one. However, a question for future research is whether (for some variant of the model) the linear equilibrium might be unique within a broader class of equilibria, say, for example, the class of equilibria in which the public's expectations are continuous in the observed variables. The results of Stanford (1986) suggest posing this question.

It seems possible that the central bank's decision to obfuscate its behavior might depend on whether or not the public can coordinate on a relatively cooperative equilibrium. As Canzoneri's model illustrates, private information impedes the attainment of the pareto-efficient equilibrium, since its enforcement requires occasional reversionary periods. So although obfuscation might be a good move for a central bank perpetually caught in the worst equilibrium (in which the inflationary bias is large), it might not be a good move in an economy which coordinates on a low-inflation equilibrium.

The models of Canzoneri and of Cukierman and Meltzer have somewhat different implications for monetary reform. If the information structure is exogenous, then one needs to be concerned about how to deal with the private-information problem in any legislative solution to the inflation problem. Cukierman and Meltzer's analysis, on the other hand, suggests that the extent to which private information is a problem may itself be a function of the legal and institutional structure.

V. COORDINATION WITH MULTIPLE CONTROLLERS

Here I briefly discuss a couple of questions which arise when the public faces more than one government agency or sovereign. Allowing for multiple controllers adds an interesting new dimension of strategic complexity to the analysis.

Alesina (1985) analyzes a model in which there are two political parties with different preferences over the relative importance of inflation and unemployment. There is exogenous uncertainty over the outcome of the election, making it difficult for voters to forecast the post-election inflation rate. The induced volatility in post-election output and inflation is such that both parties would be better off (on average) if they could agree on a consensus inflation policy. But if the winner of the election always behaves myopically, this will break any such agreement. Alesina demonstrates that in the infinite-horizon case, a more efficient outcome may be achieved via reputation. In fact, if the two parties have low enough discount rates, then it is possible to have an equilibrium in which post-election inflation-rate volatility is completely eliminated. What supports the "cooperative" equilibrium is that the winning party believes that if it defects, then the next time the opposition party gains power, it, too, will choose its own most preferred inflation rate.[17]

Rogoff (1985a) analyzes a two-country model in which each country's monetary authority faces a credibility problem vis-a-vis its own private sector. The credibility problem is similar to that of the model described in Section II above. However, there are also strategic interactions between the two sovereign monetary authorities. What creates an overlap in their objective functions is that when either country (unilaterally) increases its money supply (by more than the private sectors anticipated), it causes its real exchange rate (vis-a-vis the other country) to depreciate. (This is a robust result which obtains across a broad class of open-economy macroeconomic models.) The central banks tend to regard this as an undesirable consequence of conducting an unanticipated inflation. First, real-exchange-rate depreciation affects output adversely if the foreign good enters as an intermediate good in the production function, or if wages are indexed to a basket which includes the foreign good. (The main results obtain with either an islands model or a contracting model.) Second, a depreciation in the real exchange rate raises the rate of CPI inflation.

In the one-shot game, both governments can actually make themselves worse off by coordinating their monetary policies (via a legally-fixed exchange rate or via a monetary union). By coordinating their monetary policies, they remove a check on themselves. The private sectors recognize this when forming their inflationary expectations, and the time-consistent rate of inflation actually rises. The same result holds in a repeated game if the public is unable to coordinate on the optimal trigger-strategy equilibrium, or if the monetary authorities have high discount rates.[18] Thus institutional reforms aimed at promoting government to government cooperation must be designed with private-sector responses in mind. Kehoe (1985) extends these results to show that government to government tax-policy cooperation can also be counterproductive. Again, the main theme is

---

[17]Roberds (1985) gives a political interpretation to his model of "stochastic replanning," in which the preferences of the government randomly evolve over time. Alesina and Tabellini (1985) provide a framework for analyzing the case where competing political interests control different aspects of macroeconomic policy.

[18]Oudiz and Sachs (1985) consider repeated game-solution concepts in a two-country framework. In their example, and in the stochastic version of Rogoff (1985a), cooperation between governments can be beneficial.

that competition between governments may be beneficial by mitigating their credibility problems vis-a-vis the private sector.

CONCLUSIONS

It would seem reasonable to suppose that reputational considerations temper the government's incentives to conduct surprise inflations. However, while considerable progress has been made in introducing reputation into models of monetary policy, there are still important unanswered questions. The most disturbing feature of the models proposed to date is that either the equilibrium is very sensitive to changes in the informational structure, or/and there are a multiplicity of equilibria. There would appear to be substantial coordination problems involved in achieving the most favorable reputational equilibria. Crawford (1985) has observed that in some situations, strategic uncertainty -- uncertainty about which equilibrium strategy other agents are adopting -- may be just as important as uncertainty about exogenous factors. [See also Axelrod (1984).] The fact that the current generation of repeated game models does not place sharp restrictions on the data makes it difficult to apply these models with confidence.

It may be possible to construct an argument that certain equilibria are "focal." Perhaps the government can aid in the coordination problem by, for example, announcing monetary targets. But this line of reasoning is tenuous and, in some sense, runs counter to the whole thrust of the rational expectations revolution. The public is much more likely to be influenced by governmental announcements which are backed by concrete measures. Another, perhaps more promising, approach to eliminating certain equilibria in these models would be to search for refinements in the equilibrium concept. If, for example, one could provide a strong argument for assuming that the public's expectations about future inflation are continuous in current inflation (an assumption which seems quite plausible), this might rule out many of the reputational equilibria discussed in the text. Other approaches include allowing for more heterogeneity in the private sector and introducing more institutional detail. Until reputational models of monetary policy can be refined to yield sharper predictions, it would seem premature to focus attention only on the most favorable equilibria. It is certainly too soon to conclude that reputational constraints substantially vitiate the case for imposing legal constraints on monetary policy, as some have inferred.

Although there are reasons why one cannot yet be satisfied with extant models of strategic monetary policy, they represent a clear improvement over early rational expectations models, in which the government's behavior was treated as exogenous. Whereas it may be constructive to ignore strategic factors in studying certain macroeconomic phenomena, they are central to the analysis of the government's role in the economy. A major appeal of strategic macro models is that they allow one to formally model political and institutional relationships which previously could only be discussed informally.

In my effort to highlight certain general modelling issues, I have not focused on institutional details. But it is ultimately important to take these details into account in constructing an applied strategic monetary-policy model.

APPENDIX

In Section III, we considered some attempts to adapt Kreps and Wilson's (1982) sequential-equilibrium model of reputation to analyze monetary policy credibility. One mildly unattractive feature of these models is that their equilibria involve randomizing strategies. It seems implausible to think that the central bank decides when to start inflating based on the outcome of a sequence of coin flips. Another drawback to the models is that they allow only for two types of policymakers. In this (essentially self-contained) Appendix, I present an alternative formulation in which there are a continuum of types, and for which there does exist a pure-strategy (sequential) equilibrium.[19] Otherwise, the model generally yields qualitatively equivalent results to the models surveyed in the text, though there are also other differences. For example, the results concerning the path of conditional expected inflation differ somewhat from those of Barro (1986). Also, the analysis illustrates how there}can be multiple sequential equilibria in this type of model, unless one places restrictions on beliefs concerning off-the-equilibrium-path behavior.

In the models of Tabellini (1983) and Barro (1986), policymakers differ according to how much it costs them to break a commitment never to inflate. These studies do not go into detail about just what form such a commitment might take. One natural possibility would be for the legislature to pass a law dictating the rate of growth of a monetary aggregate. Such a commitment should not necessarily be treated as absolute; there are any number of reasons why the central bank might inflate in spite of such a law. First of all, the central bank (or some special interest group) might be able to challenge the law's constitutionality. Or, if the penalties are not sufficiently severe, the central bank may simply be willing to pay the price for violating the law. Even if it proves impossible to revoke or ignore the law, the central bank could still try to circumvent it via regulatory changes which influence the transactions' demand for money. Such regulatory changes might involve deleterious microeconomic side-effects, but the central bank may be willing to tolerate such inefficiencies in order to reap the

[19]The general approach builds on that of Milgrom and Roberts (1982). Wood (1986) has applied the Milgrom-Roberts model to the problem of speculative attacks in the foreign-exchange market.

benefits of unanticipated inflation. Finally, depending on how the law is structured, there is always the possibility that the legislature will decide to back off and repeal the law if the central bank actually inflates. In each of the scenarios described above, one could argue that the central bank has private information about the disutility it will receive if it breaks the law.

In Barro's model, the cost to the central bank of breaking its zero-inflation commitment takes on one of two extreme values: zero or prohibitive. The central bank knows its cost type, whereas the public knows only the distribution of types. Here I modify Barro's model to allow the cost of reneging to take on a continuum of values.

The policymaker has a finite-horizon loss function given by

$$\Omega_0 = \sum_{0}^{T} L_s(\pi,\pi^e,c)\beta^s, \qquad 1/2 < \beta < 1, \tag{A1}$$

$$L_t[\pi_t,(\pi^e)_t,c] = -[\pi_t - (\pi^e)_t] + \tfrac{1}{2}(\pi_t)^2 + \tfrac{1}{2}Z(c, \pi_t, \pi_{t-1}, \pi_{t-2},\ldots), \tag{A2}$$

where $Z = c$ if $\pi_t \neq 0$ <u>and</u> $\pi_{t-i} = 0$ for all $i > 0$; $Z_t = 0$ otherwise. In other words, the central bank bears a fixed one-time cost to reneging on its commitment never to inflate. (This cost might be associated with the cost of repealing the law. One way to justify the finite horizon would be to assume that the law has a known expiration date. Or, perhaps it is known that at some future date, a new transactions technology will come on line which will render meaningless the definition of money embodied in the law.)[20] As in Barro's and Tabellini's analyses, the public does not directly observe c. At time zero, the public knows only that $c \in [0,\mu]$, where $\mu > 1$; it has uniform priors over this interval. In subsequent periods, the public uses Bayes' rule to update its priors in a manner we shall specify shortly.

[20]The analysis would only have to be modified slightly if the fixed cost to inflating had to be paid every period that the central bank inflates. To allow for the case where the cost of inflating depends on the level of inflation would have to involve a more substantial modification.

Because the central bank bears only a one-time fixed cost in repealing its zero-inflation commitment, then it must be true that $(\pi^e)_t = \pi^* = 1$ if $\pi \neq 0$ for any $s < t$, where π^* is the (here unique) equilibrium level of inflation in the one-shot game. After the central bank inflates once, the Z function for future periods (Z = 0) becomes common knowledge, and there is a unique equilibrium because of the finite horizon. Thus if the central bank is going to inflate at all in the current period t, then it should choose π_t to minimize its current-period loss function; expected future inflation will equal π^* for any $\pi_t > 0$. Thus one can deduce

<u>Proposition 1</u>: $\pi_t = 0$ or $\pi_t = 1$ for all t.

Proposition 1 holds, of course, only because of the special form of the loss function (A2), in which $\pi - \pi^e$ enters linearly, and in which Z can only take on one of two values, c or 0.

<u>Definition</u>: $\sigma_t \equiv \{(\pi^e)_t | \pi_{t-1}, \pi_{t-2}, \ldots = 0\}$
is the public's expectation of inflation in period t, given that the central bank has not broken its commitment prior to t.

Clearly $\sigma_t \leq 1$, since by Proposition 1, π will equal zero or one in any period t. Temporarily treating the path of σ as exogenous, we can calculate the loss to the central bank if it follows a strategy of setting $\pi = 0$ in periods 0 through t-1, and $\pi = 1$ in periods t through T:

$$\Gamma(t,c) = \sum_{0}^{t-1} \beta^s \sigma_s + \beta^t[(c-1)/2 + \sigma_t] + \tfrac{1}{2}\sum_{t+1}^{T} \beta^s. \qquad \text{(A3)}$$

Note that $\Gamma(t,c)$ is continuous and increasing in c. From (A3), we can immediately deduce two important facts: First,

<u>Proposition 2</u>: If $\pi_t = 0$ for all $t < T$, then $\pi_T = 1$ if $c < 1$, and $\pi_T = 0$ if $c \geq 1$.

Proposition 2 follows from (A3); $\Gamma(T,c) \gtrless \sum_{0}^{T} \beta^s \sigma_s$ as $c \gtrless 1$. In other words, if the cost to the central bank of breaking its commitment is less than one, then it will certainly be inflating by the final period T. Conversely, we can similarly deduce that no type $c > 1$ will inflate for the first time in the final period T. However, without further restrictions on

$\{\sigma_s\}$, we cannot yet rule out the possibility that types $c > 1$ will begin inflating before period T. This point will turn out to be crucial in our later discussion of the uniqueness of equilibrium.

We can also deduce from (A3) that a high-cost type would never begin inflating in an earlier period than a low-cost type would. Holding $\{\sigma_s\}$ fixed, the higher the cost to the central bank of breaking its commitment, the more incentive it has to wait to incur this cost.

Proposition 3: $[r(t_1,c_2) - r(t_2,c_2)] > [r(t_1,c_1) - r(t_2,c_1)]$, for $c_2 > c_1$ and $t_2 > t_1$.

Proposition 3 follows immediately from the fact that

$$[r(t_1,c_1) - r(t_2,c_1)] - [r(t_1,c_2) - r(t_2,c_2)] = (1 - \beta^m)(c_1 - c_2) < 0, \tag{A4}$$

for $\beta < 1$, where $m \equiv t_2 - t_1$.

With the above results, we are now prepared to discuss the evolution of σ_t. In a sequential equilibrium, the public's beliefs must evolve according to Bayes' rule, so that

$$\sigma_t = (\hat{c}_t - \hat{c}_{t-1})/(\mu - \hat{c}_{t-1}), \tag{A5}$$

where $\hat{c}_t \equiv \sup\{\hat{c} \in [0,\mu] | c < \hat{c}$ implies $r(t + i,c) > r(t,c)$ for all i such that $0 < i \leq T - t\}$.

Thus all types $c < \hat{c}_t$ will begin inflating in period t if they have not already begun inflating in an earlier period. The denominator of (A5) represents the range of cost types which the public believes would not have inflated prior to period t. The numerator represents the range of cost types which the public believes will inflate for the first time in period t (in which case they will set $\pi = 1$). (A5) gives the expected inflation rate, conditional on past inflation being uniformly zero, because the public has uniform priors over c. Since Proposition 3 implies that $\hat{c}_t$ must be nondecreasing in t, (A5) is the only possible form for rational expectations.

It is natural to look for an equilibrium in which $\hat{c}_T = 1$. For if the time horizon were only one period ($T = 0$), then the unique sequential equilibrium would obviously be [by Proposition 2 and (A5)] $\sigma = 1/\mu$ and $\hat{c} = 1$. If the cost to the central bank of breaking its commitment is greater than one, then it will not inflate even in a one-shot game. In a multi-period game, the central bank must bear an additional cost if it inflates before period T, for then the central bank not only has to bear the one-time reneging cost, c, but it also must live with high expected inflation ($\pi^e = 1$) in all future periods.

We will proceed by showing that a necessary and sufficient condition for a sequential equilibrium with $\hat{c}_T = 1$ is that the path of $\hat{c}_t$ must be governed by the recursion relationship

$$\begin{aligned} r(t,\hat{c}_t) - r(t+1,\hat{c}_t) &> 0, \text{ if } \hat{c}_t = 0, \\ &= 0, \text{ if } \hat{c}_t > 0. \end{aligned} \tag{A6}$$

That the recursion (A6) is a sufficient condition for equilibrium is a straightforward consequence of Proposition 3 together with the continuity of r in c; the proof will be omitted. [The proof involves showing that when faced with expectations of inflation governed by (A5) and (A6), all types $c < \hat{c}_t$ would prefer to inflate first in period t over first inflating in some future period, and that no type $c > \hat{c}_t$ would prefer to begin inflating on or before date t.]

To demonstrate that (A6) is a necessary condition for a sequential equilibrium with $\hat{c}_T = 1$, we first prove that if there is a sequence of periods during which some types first inflate, then this sequence cannot be followed by a period(s) where no type would begin inflating. We then show that (A6) must hold with equality during any period prior to T in which some types would begin inflating.

Lemma 1: In any sequential equilibrium where $\mu > \hat{c}_t > \hat{c}_{t-1}$ for any $t < T$, then $\hat{c}_{t+1} > \hat{c}_t$.

Comment: The proof of Lemma 1 requires our assumption in (A1) that $\beta > 1/2$.

Proof: Suppose, in contradiction to the Lemma, that there is some time $t < T$ such that $\hat{c}_{t+1} = \hat{c}_t > \hat{c}_{t-1}$. Now this cannot be an equilibrium unless

$r(t,\hat{c}_t) \le r(t+1,\hat{c}_t)$ and, by (A5), $\sigma_{t+1} = 0$. But by (A3), if $\sigma_{t+1} = 0$, then

$$r(t,\hat{c}_t) - r(t+1,\hat{c}_t) =$$

$$(-1 + \hat{c}_t + \beta)/2 - \beta(-1 + \hat{c}_t)/2. \tag{A7}$$

But the RHS of (A7) must be positive for $\hat{c} \ge 0$ if $\beta > 1/2$.

Proposition 4: (A6) must hold with equality in any period t such that $\hat{c}_{t-1} < \hat{c}_t < \hat{c}_{t+1}$, for $t < T$.
Proof: In any such period t, there must exist some $\alpha > 0$, s.t. for $0 < \delta < \alpha$, any type $\hat{c}_t - \delta$ inflates first in period t, and any type $\hat{c} + \delta$ inflates first in period t + 1. The proposition then follows from the fact that $r(t,c)$ is continuous and increasing in c.

Lemma 1 and Proposition 4 together imply that (A6) must hold during any period $t < T$ such that $\hat{c}_t > 0$. That (A6) must hold with inequality over the initial periods where $\hat{c} = 0$ can be confirmed directly from (A3).

We are now ready to characterize the equilibrium governed by (A6). Combining (A3) and (A6), we obtain

$$[-1 + \hat{c}_t(1-\beta)]/2 = \beta(\sigma_{t+1} - 1). \tag{A8}$$

Substituting (A5) into (A8) yields[21]

$$\hat{c}_{t+1} = F(\hat{c}_t) =$$

$$[(\beta - 1)/2\beta](\hat{c}_t)^2 + \{[\mu(1-\beta) + 1]/2\beta\}\hat{c}_t + (\mu - \mu/2\beta). \tag{A9}$$

Via direct differentiation, we can confirm that paths which obey (A9) also satisfy Lemma 1:

---

[21] If we modify the analysis so that $Z(\pi) = c$ for $\pi \ne 0$, for any history of π, the recursion equation (A8) becomes

$$\hat{c}_{t+1} = \mu[1 - (1/2\beta)] + [(\mu + 1)/2\beta]\hat{c}_t - (\hat{c}_t)^2/2\beta. \tag{F1}$$

There exists a sequential equilibrium very similar to the one analyzed in the main text.

$$d\hat{c}_{t+1}/d\hat{c}_t = [\mu(1-\beta) + 1-2(1-\beta)\hat{c}_t]/2\beta > 0, \tag{A10}$$

since $\mu > 1$, and $\hat{c}_t \leq 1$ if $\hat{c}_T = 1$.

To prove existence and uniqueness of the equilibrium characterized by the fundamental recursion (A6) (together with the terminal condition $\hat{c}_T = 1$), we invert equation (A9) to solve for $\hat{c}_t$ in terms of $\hat{c}_{t+1}$:

$$\hat{c}_t = F^{-1}(\hat{c}_{t+1}) = [\mu(1-\beta) + 1]/(2 - 2\beta) \pm$$

$$\{[\mu(1-\beta) + 1]^2 - 4(1-\beta)[2\beta\hat{c}_{t+1} - \mu(2\beta-1)]\}^{\frac{1}{2}}/(2 - 2\beta). \tag{A11}$$

From (A11), it is readily confirmed that if $\mu > 1$ and $\beta > 1/2$, then for $0 \leq \hat{c}_{t+1} \leq 1$, F^{-1} has exactly one real root less than one. Moreover, $F^{-1}(\hat{c}_{t+1}) < \hat{c}_{t+1}$ [this also follows directly from (A10)]. By (A6), an equilibrium path is governed by (A9) only where $\hat{c} > 0$. Tracing the equilibrium backwards from time T, if there comes a point where $\hat{c}_{t+1}$ is such that $F^{-1}(\hat{c}_{t+1}) < 0$, then $\hat{c}_s = 0$ for all $s \leq t$. That this is indeed an equilibrium is readily confirmed. By (A9), or by (A11), $F^{-1} < 0$ if and only if

$$\hat{c}_{t+1} < \mu - \mu/2\beta. \tag{A12}$$

By setting $\hat{c}_t = 0$ and $\sigma_{t+1} = \hat{c}_{t+1}/\mu$ in (A8), we see that condition (A12) provides the maximum level of $\hat{c}_{t+1}$ such that a type zero would not choose to inflate for the first time at t. The equilibrium is unique because along the path obtained by using (A11) to trace back from $\hat{c}_T = 1$, there is only one $\hat{c} > 0$ such that (A12) holds. This rules any path where a type zero waits even longer to first inflate, so that $\hat{c}$ jumps directly from $\hat{c} = 0$ to some higher $\hat{c}$ on the path leading up to $\hat{c}_T = 1$.

To prove existence, it is sufficient to show that when we trace (A11) backwards from time T, there can only be a finite number of periods where $\hat{c} > 0$. From (A10), we have

$$d^2\hat{c}_{t+1}/d\hat{c}_t^{\,2} = -(1 - \beta)/\beta < 0. \tag{A13}$$

By (A10) and (A13), $F^{-1}{}''(\hat{c}_{t+1}) > 0$. This implies that $\hat{c}_{T-k} - \hat{c}_{T-k-1}$ is increasing in k; as we solve for the path of $\hat{c}$ working backwards from the terminal condition $\hat{c}_T = 1$, $\hat{c}$ must indeed become zero in finite time. Thus an equilibrium always exists. It also follows that if the time horizon T is long enough, then there will be no chance of inflation early in the policymaker's term.

Figure 1 presents a graph of equation (A9), and traces out the equilibrium path leading to the terminal condition, $\hat{c}_T = 1$. The condition $\beta > 1/2$ insures that the curve intersects the $\hat{c}_{t+1}$ axis at a positive value. The condition $1 > \mu - \mu/2\beta$ insures that the $\hat{c}_{t+1}$ intercept is less than one. [From (A12), this is the necessary and sufficient condition for $\hat{c}_{T-1} > 0$.]

An immediate implication of Lemma 1 together with equation (A8) is $d\sigma_t/dt > 0$ for $\sigma_t > 0$. Whether or not a given cost type chooses to initiate inflation in any given period depends on the opportunity cost. The higher σ_{t+1}, the lower the opportunity cost to inflating in the current period. It does not take as high a σ_{t+1} to tempt a low-cost type to begin inflation. Over time, σ must rise so}that higher and higher cost types are tempted to break away. This scenario is somewhat different than the one in Barro (1986), which was discussed in Section III. Observe that by (A8) and (A13), σ rises at a decreasing rate.

We have shown that there exists a unique sequential equilibrium with $\hat{c}_T = 1$. Can there exist other sequential equilibria with $\hat{c}_T \neq 1$? Candidate equilibria with $\hat{c}_T > 1$ are ruled out by Proposition 3, and candidate equilibria with $1 < \hat{c}_T < \mu$ are ruled out by applying Proposition 3 together with Lemma 1. Proposition 3 states that no type $c > 1$ will ever first inflate in the final period, and Lemma 1 proves that if $\hat{c}_T < \mu$, then some types must inflate for the first time in the final period. However, the results contained in this Appendix are insufficient to preclude sequential equilibrium paths in which $\hat{c}_T = \mu$. Suppose, for example, that the time horizon, T, is very large, and the public's beliefs are: $(\pi^e)_0 = 1$, and $\sigma_s = 1$ for all $1 \leq s < T$. If the policymaker does not inflate in period 0, he will, by (A2), suffer a loss of 1 in period zero. If $c > 1$, this loss is less than the loss he receives by inflating in period 0, $1/2 + c/2$. However, if the public maintains its beliefs, he will lose 1 again in period 1. This is greater than his loss would be in period 1 had he inflated in period zero (1/2). He will be worse off (by 1/2) in each ensuing period until he gives in. Hence, if the policymaker does not inflate in period 0, it may still be rational for the public to believe that he will inflate for certain in period 1. For any value of μ, there is a T large enough so that this type of equilibrium cannot be ruled out as sequential. (Conversely, for any T, there is a μ large enough so that the equilibrium is unique.) Note that we could definitely preclude such

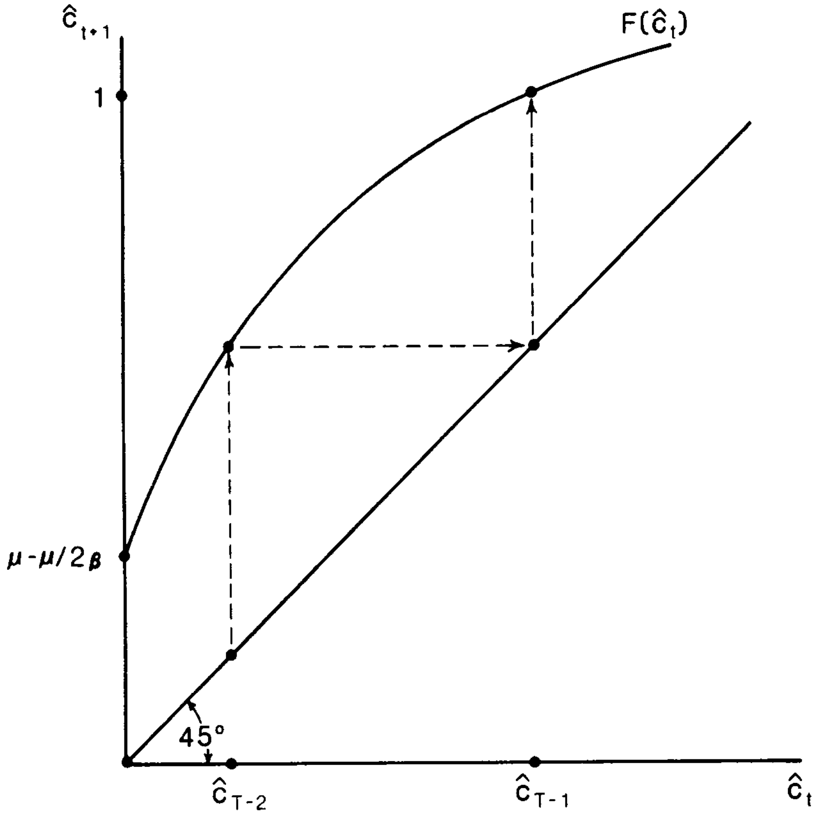

a $\hat{c}_{T-3}$, $\hat{c}_{T-4}$, ... $\hat{c}_0 = 0$.

Figure 1[a]

equilibria if the public's priors are that c is uniformly distributed on [0,1] with probability $1/(1+\mu)$, and $c = \infty$ with probability $\mu/(1+\mu)$. In this case, which would be more directly analogous to the model of Milgrom and Roberts (1982), the unique sequential equilibrium is the one illustrated in Figure 1.) If there is a second equilibrium in pure strategies, then it may not be possible to rule out mixed-strategy equilibria.

The basic problem is that the sequential-equilibrium concept does not place constraints on how the public must interpret events which occur with probability measure zero on the equilibrium path. Unless some types $c > 1$ are "robots" (as is the case where $c = \infty$ is a possibility), then their actions can be influenced by bizarre public beliefs. It is possible, however, that one may be able to rule out the "perverse" equilibrium described above by appealing to a refinement of sequential equilibrium.[22]

In conclusion, I should point out that although the finite-horizon model discussed in this Appendix improves on earlier models in a couple of dimensions, it still shares some of their major deficiencies (as discussed in Section III of the text). In particular, if different types of policymakers have different preferences instead of different costs in breaking commitments, then there might exist separating equilibria in which some types deflate. In a more general setting, there may well be other ways for types to separate themselves. Also, it is difficult to take the finite-horizon assumption literally, and some of the results may be sensitive to this assumption.

---

[22]See, for example, Cho and Kreps (1986).

REFERENCES

Abreu, D.

(1982) Repeated Games with Discounting: A General Theory and an Application to Oligopoly. Mimeographed, Princeton University.

_________, Pierce, D., and Stacchetti, E.

(1985) Optimal Cartel Equilibria with Imperfect Monitoring. Mimeographed, Harvard University.

Alesina, A.

(1985) Macroeconomic Policy in a Two-Party System as a Repeated Game. Mimeographed, Harvard University.

_________ and Tabellini, G.

(1985) Rules and Discretion with Non-coordinated Monetary and Fiscal Policies. Mimeographed, Harvard University.

Atkeson, A.

(1986) Reputation and Time Consistency in a Macro Model. Mimeographed, Stanford University.

Axelrod, R.

(1984) *The Evolution of Cooperation*. New York: Basic Books.

Backus, D. and Driffill, J.

(1985) Inflation and Reputation, *American Economic Review*, 75: 530-538.

Barro, R.J.

(1983) Inflationary Finance under Discretion and Rules, *Canadian Journal of Economics*, 16: 1-16.

____

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

Barro R.J. and Gordon, D.B.

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

_________ and _________

(1983b) A Positive Theory of Monetary Policy in a Natural Rate Model, *Journal of Political Economy*, 91: 589-610.

Benoit, J.P. and Krishna, V.

(1985) Finitely Repeated Games, *Econometrica*, 53: 905-922.

Calvo, G.

(1978) On the Time Consistency of Optimal Policy in a Monetary Economy, *Econometrica*, 46: 1411-1428.

Canzoneri, M.B.

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

Cho, I.K. and Kreps, D.

(1986) Signalling Games and Stable Equilibria, *Quarterly Journal of Economics*, forthcoming.

Crawford, V.P.

(1985) Dynamic Games and Dynamic Contract Theory, *Journal of Conflict Resolution*, 29: 195-224.

Cukierman, A.

(1985) Central Bank Behavior and Credibility - Some Recent Developments. Mimeographed, Tel-Aviv University.

_________ and Meltzer, A.H.

(1986) A Theory of Ambiguity, Credibility and Inflation Under Discretion and Asymmetric Information, *Econometrica*, forthcoming.

Fershtman, C.

(1986) Fixed Rules and Decision Rules: Time Consistency and Subgame Perfection. Mimeographed, Hebrew University.

Fischer, S.

(1986) Time Consistent Monetary and Fiscal Policies: A Survey. Mimeographed, Massachusetts Institute of Technology.

Friedman, J.W.

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

———

(1985) Cooperative Equilibria in Finite Horizon Noncooperative Supergames, *Journal of Economic Theory*, 35: 390-398.

Green, E.J. and Porter, R.H.

(1984) Noncooperative Collusion Under Imperfect Price Information, *Econometrica*, 52: 87-100.

Grossman, H.I. and Van Huyck, J.B.

(1986) Seigniorage, Inflation, and Reputation, *Journal of Monetary Economics*, forthcoming.

Gul, F., Sonnenschein, H., and Wilson, R.

(1986) Foundations of Dynamic Monopoly and the Coase Conjecture, *Journal of Economic Theory*, 38, forthcoming.

Horn, H. and Persson, T.

(1985) Exchange Rate Policy, Wage Formation and Credibility, *Institute for International Studies*, Seminar Paper No. 325, Stockholm.

Kehoe, P.

(1985) Sequential Coordination of Policies May Make Things Worse. Mimeographed, University of Minnesota.

Kreps, D.M. and Wilson, R.

(1982) Reputation and Imperfect Information, *Journal of Economic Theory*, 27: 253-279.

Kydland, F.E. and Prescott, E.C.

(1977) Rules Rather than Discretion: The Inconsistency of Optimal Plans, *Journal of Political Economy*, 85: 473-492.

Milgrom, P. and Roberts, J.
(1982) Predation, Reputation, and Entry Deterence, *Journal of Economic Theory*, 27: 280-312.

Oudiz, G. and Sachs, J.
(1985) International Policy Coordination in Dynamic Macroeconomic Models, W. Buiter and R.C. Marston (eds.), *International Economic Policy Coordination*. Cambridge: Cambridge University Press.

Phelps, E.S.
(1967) Phillips Curves, Expectations of Inflation and Optimal Employment Over Time, *Economica*, 34: 254-281.

_________ and Pollak, R.A.
(1968) Second-Best National Saving and Game-Equilibrium Growth, *Review of Economic Studies*, 2: 185-199.

Roberds, W.
(1985) Models of Policy under Stochastic Replanning, *Federal Reserve Bank of Minneapolis Research Department*, Working Paper 282.

Rogoff, K.
(1985a) Can International Monetary Policy Cooperation be Counterproductive? *Journal of International Economics*, 18: 199-217.

____
(1985b) The Optimal Degree of Commitment to an Intermediate Monetary Target, *Quarterly Journal of Economics*, 100: 1169-1189.

_________ and Sibert, A.
(1986) Elections and Macroeconomic Policy Cycles, *National Bureau of Economic Research*, Working Paper No. 1838.

Rubinstein, A.
(1986) Finite Automata Play the Repeated Prisoner's Dilemma, *Journal of Economic Theory*, 38, forthcoming.

Stanford, W.G.

(1986) On Continuous Reaction Function Equilibria in Duopoly Supergames with Mean Payoffs, *Journal of Economic Theory*, 38, forthcoming.

Tabellini, G.

(1983) Accommodative Monetary Policy and Central Bank Reputation. Mimeographed, UCLA.

Vickers, J.

(1986) Signalling in a Model of Monetary Policy with Incomplete Information. Mimeographed, Oxford University.

Woods, P.R.

(1986) Defending a Fixed Exchange Rate: The Central Bank's Reputation and the Deterrence of Speculative Attacks on its Foreign Reserves, Ph.D. dissertation, University of Wisconsin.