International Finance Discussion Papers

Number 249

September 1983
Revised September 1984

MONETARY POLICY GAMES AND THE ROLE OF PRIVATE INFORMATION

by

Matthew B. Canzoneri

Monetary Policy Games and the Role of Private Information

by

Matthew B. Canzoneri*

I. Introduction

Consider the following scenario: Monetary aggregates surge ahead of expected or targeted rates; the Fed claims that it is accommodating a perceived increase in money demand in order to stabilize the price level; the private sector (or the Administration, or Congress) counters that the Fed is running an inflationary policy to expand employment; a period of credibility building ensues, which may focus upon the personalities of policy makers, targeting procedures, or even proposals for legislative reform. A key element in this scenario is that the private sector cannot verify the Fed's claim; the Fed's forecast of money demand is private information.

A plausible explanation for such a scenario is that the Fed is involved in a non-cooperative game with certain agents in the private sector, or with some other branch of government. Why else for example would the Fed seek to constrain its own behavior, with say a targeting procedure, unless it thought that this would bring a resolution to a credibility problem resulting from a non-cooperative game.[1/] And why does the Fed have a "credibility" problem, as opposed to an "information" problem; that is, why are its announcements not accepted at face value?

A game suggested by Kydland and Prescott (1977) has been receiving a lot of attention recently.[2/] In this game, wage setters have to specify the nominal wage in a labor contract before the Fed sets the money supply. The Fed wants a higher level of employment than the wage setters, but it is also inflation conscious. The wage setters know that the Fed will be tempted to inflate away some of their real wage, to achieve a higher level of employment, so they purposely set the wage high. The Fed weighs its employment and inflation objectives, and only finds it optimal to inflate the real wage down to the level wanted by the wage setters. The non-cooperative solution has an inefficient inflation bias, simply because the Fed has no credible way of precommitting itself to a non-inflationary policy. Adding a stabilization role for monetary policy, and private information, the Kydland-Prescott game can be used to model the scenario outlined above.

There are a variety of ways in which the non-cooperative solution described above might be improved upon. Congress could legislate institutional reforms. It could legislate a monetary policy rule directly, or it could constrain the behavior of the players by changing the rules of the game; wage-price guidelines and legislated targeting procedures are examples of the later. Thompson (1981) and Rogoff (1983a) suggest changing the incentive structure for Fed policy makers, or simply choosing a policy maker with "perverse" preferences. And Barro and Gordon (1983a) suggest that the players themselves may have already found a resolution to the problem in a reputation building mechanism.

An efficient resolution of the credibility problem must leave the Fed with the latitude or discretion to perform its stabilization role. Rogoff (1983a) and Barro and Gordon (1983a) discuss the trade-off between flexibility, which is required for stabilization, and constraint, which is required for a resolution of the credibility problem, but neither of their papers incorporates private information. An efficient resolution is much more difficult to achieve if the Fed's forecast of money demand is indeed private information. No stable resolution of the credibility problem can rely on the Fed's own announcement of its forecast if the Fed has an incentive to misrepresent; the resolution must be, in Hurwicz's (1972) terminology, incentive compatible.[3/]

The present paper proceeds as follows:

Section II outlines the Kydland-Prescott game in some detail and describes the modifications necessary to give the Fed a stabilization role. The importance of private information is highlighted. Without private information, Taylor's (1983) observations seem justified; resolutions of the precommitment problem would be too easy to come by for this game to be taken seriously as a real world problem.

Section III reviews the work of Rogoff (1983a) and Barro and Gordon (1983a) in the light of private information. Rogoff's (1983a) theory of the perverse policy maker is shown to have incentive compatibility problems since the policy maker's actions do not reveal his true preferences (or whether he is responding to an incentive structure laid down by Congress). However, this fact may only serve to make the theory more compelling as an explanation of the often dramatic comings and goings of individual policy makers. The rest of section III (and section IV) focuses exclusively on repeated game solution concepts. In a repeated game, honest behavior in the present period can be made to yield rewards or expanded opportunities in the next. Thus, there is an added strategic dimension; a player can invest in his reputation. Using Friedman's (1971) notion of reversionary periods, Barro and Gordon (1983a) develop a model that resolves the credibility problem in the absence of private information. Their model falls apart if the Fed's forecast of money demand is private, because wage setters cannot tell when the Fed has been honest. However, as Green and Porter (1983) have noted in a different context, this may be a blessing in disguise. The Friedman solution, like Rogoff's (1983a), is too stable to explain the episodic periods of breakdown that are actually observed. A reformulation of the model accommodates private information and actually predicts periodic bouts of inflation (though it would be inappropriate to think of these episodes as credibility breakdowns). It also suggests that, despite these periodic difficulties, the Fed and the wage setters may have already found a rather satisfactory resolution to the credibility problem.

Section IV discusses the legislative approach in a more sophisticated way than was possible in section II; the proper legislation can also make the Fed's opportunities in future periods depend upon honest behavior in the present. There are monetary policy rules that produce a better game outcome (generally at some expense to the stabilization effort) and are incentive compatible, even if the Fed's forecast of money demand is private information. These rules are, however, less intuitive than the typical stabilization rule; they are more easily motivated as the Fed's optimal response to a Congressionally mandated targeting procedure.[4/] Various targeting procedures are considered, some leaving the Fed with more latitude than others. Which is best depends on the Fed's ability to predict money demand disturbances and on the importance society attaches to the benefits of stabilization relative to the costs of the inflation bias resulting from the non-cooperative behavior of the players.

Section V summarizes the basic results of the paper and suggests where further research is needed.

II. The Game and the Importance of Private Information

Consider a wage contracting model like those of Gray (1976) and Fischer (1977). A typical firm's notional supply and demand for labor are illustrated in Figure 1. The basic assumption of this model is that the nominal wage must be set in a labor contract prior to the setting of the money supply, and thus the realization of the price level. Later, during the contract period, firms learn the price from the market, and the contract provides that they are allowed to maximize profits; they hire along the demand curve in Figure 1.[5/] The output supply function can be expressed as

$$(1)\quad y_t = \bar{y} + \theta(p_t - w_t)$$

where y_t, p_t and w_t are the logs of output, the price level and the contract wage; $\bar{y}$ is the equilibrium rate of output, corresponding to the rate of employment $\bar{n}$.

Wage setters want to achieve the equilibrium pictured in Figure 1. Their utility can be represented by[6/]

$$(2)\quad UW_t = -(p_t - w_t)^2 = -(1/\theta^2)(y_t - \bar{y})^2$$

Of course, p_t is not known when the contract is made; the expected utility maximizing strategy for wage setters is thus

$$(3)\quad w_t = p_t^e$$

where p_t^e is the wage setters' prediction of the price level that will obtain in the contract period. The wage setters' play is, essentially, their prediction.

Substituting (3) into (1) results in a familiar prediction error model of supply. Note also that (1) and (2) can be expressed in terms of inflation and expected inflation:

(4) $y_t = \bar{y} + \theta(p_t - p_t^e) = \bar{y} + \theta(\pi_t - \pi_t^e)$

(5) $UW_t = -(\pi_t - \pi_t^e)^2 = -(1/\theta^2)(y_t - \bar{y})^2$

where $\pi_t = p_t - p_{t-1}$ is the actual rate of inflation, and $\pi_t^e = p_t^e - p_{t-1}$ is expected inflation.

Equation (2) (or (5)) gives the wage setters' utility. The Fed maximizes social utility, which is given by

(6) $US_t = -(y_t - k\bar{y})^2 - s(\pi_t - \pi^*)^2$

where $k > 1$, and π^* is the optimal rate of inflation. $\bar{y}$, the rate of output sought by wage setters, is too small from a social point of view.[7/] This assumption can be motivated in several ways. Canzoneri (1980) assumes that the labor market is dominated by large unions. The labor supply curve in Figure 1 includes only union members, and wage setters' behavior systematically excludes other workers. By contrast, the social utility function includes all workers. Barro and Gordon (1983a, b) have a different view. They assert that tax policy and unemployment compensation distort the labor supply curve in Figure 1 and lead to a socially inefficient equilibrium. (Under this interpretation, equation (2) is a guide to wage setters' behavior, but not a measure of their welfare.) Either view can be used to rationalize the conflict between (2) and (6) over optimal levels of the real wage, employment and output.[8/] The second term in the social utility function penalizes deviations from the optimal rate of inflation. In a less than fully indexed economy, inflation redistributes income from one group to another and distorts incentives; in addition, there is seigniorage. Many factors go into the calculation of π^*; here, π^* is rather cavalierly taken to be a regime independent constant.

Some will prefer not to identify the Fed's utility function (6) with social utility. In this view, (6) merely reflects the biases of Fed policy makers. It is sometimes asserted for example that macro policy makers have tended to underestimate how high the "natural rate" of unemployment has actually been, and that while the "evils" of inflation are poorly documented, there is clearly a political mandate to sacrifice some employment and output to keep inflation in check. If this view of (6) and its conflict with (2) is adopted, the normative content of what follows must be appropriately reinterpreted.

To explain the Fed's interaction with the wage setters, the rest of the economic environment must be modeled. This is done in the most rudimentary way; that is, control theoretic aspects are trivialized so as to focus upon game theoretic aspects. The model is closed with a simple quantity equation,

$$(7)\ m_t - p_t = \bar{y}$$

where m_t is the log of the money supply. First differencing, (7) becomes

(8) $g_t - \pi_t = 0$

where $g_t = m_t - m_{t-1}$ is the growth rate of the money supply; g_t is taken to be the Fed's instrument.

The basic structure of the precommitment problem should now be apparent. The wage setters must act first; this is an assumption of the contracting model. They must specify w_t in labor contracts, not knowing what g_t will be, but fully aware of the Fed's motivations as embodied in the social utility function. The wage setters know that the Fed will be tempted to inflate away their real wage and achieve its higher output target, $k\bar{y}$. The wage setters can make this strategy expensive for the Fed, in terms of its inflation goal, by setting w_t high in the first place; in fact, if they set w_t high enough, the Fed will find it optimal to give them their preferred real wage and output, $\bar{y}$. This is the problem analyzed by Canzoneri (1980) and Barro and Gordon (1983a), and by Kydland and Prescott (1977) before them. As they all point out, the non-cooperative solution to this game is inefficient. If the Fed could somehow be credibly committed to setting g_t equal to π^*, then wage setters would not feel compelled to set a high rate of wage inflation. Wage setters would still get their desired rate of output, $\bar{y}$, but society would be spared the unnecessary inflation.

The message so far is that the inefficiencies brought about by the conflict over output goals and the timing of players' actions could be eliminated by tying the hands of the Fed. This would be the final message if there were no benefit to society from the Fed's retaining a degree of flexibility. There is just such a benefit to society if the Fed can play a stabilization role.

Suppose a stochastic disturbance is added to money demand in (7). If this disturbance follows a random walk, then (8) becomes

(9) $g_t - \pi_t = \delta_t$

where δ_t is the white noise innovation. Suppose further that wage setters do not see δ_t at the time they have to specify w_t, but that the Fed does have δ_t (or some forecast of it) when it sets g_t. More specifically, let δ_t be decomposed into e_t and ε_t,

(10) $\delta_t = e_t + \varepsilon_t$

where e_t is the Fed's forecast of δ_t, and ε_t is its forecast error. If the Fed is left with sufficient flexibility, it can accommodate e_t before it passes on to prices and output; this benefits both wage setters and society as a whole. The problem now is one of trading off the flexibility needed for stabilization with the constraint required for eliminating the inflation bias; this problem is touched upon in Barro and Gordon (1983a), and it is the main subject of Rogoff (1983a).

Two quite different information structures will be considered below, and it is convenient to be more explicit about them at this point. With "symmetric information," both wage setters and the Fed are assumed to observe e_t at the same time, though it is too late for the wage setters to incorporate this new information into the contract wage, w_t. With "private information," either the Fed has superior information that cannot be revealed or the policy making process is such that the private sector cannot reconstruct the Fed's forecast of money demand. At the end of the contract period, wage setters can calculate δ_t from (9), for they will then observe both g_t and π_t. However, they cannot decompose δ_t into e_t and ε_t. The Fed's forecast, e_t, is private information.

Before going on, it is also convenient to express the players' utility functions in terms of their actions, and to give a more formal description of the game. Using (9), the wage setters' utility function (5) can be expressed as

$$(11)\ UW_t = -(g_t - g_t^e - \delta_t)^2$$

where g_t^e is the wage setters' prediction of the Fed's action at the time they must set w_t. In view of (5) and the fact that $g_t^e = \pi_t^e$, it should be clear that g_t^e can be taken to be the wage setters' action. Using (9) and (1), the social utility function (6) can be expressed as

$$(12)\ UF_t = -(g_t - g_t^e - \delta_t - y^*)^2 - f(g_t - \delta_t - \pi^*)^2$$

where $UF_t = US_t/\theta^2$, $f = s/\theta^2$ and $y^* = (k-1)\bar{y}/\theta$; θy^* is the difference between the players' output goals. The game proceeds as follows: (1) wage setters form g_t^e on the basis of their knowledge of the Fed's utility function and their expectation that δ_t (and e_t and ε_t) is zero. (2) The Fed sets g_t to maximize its expected value of UF_t in (12), knowing g_t^e and its own forecast of δ_t. (3) Ex-post, wage setters observe δ_t, and if information is symmetric, they can decompose it into the Fed's forecast, e_t, and the forecast error, ε_t. If the Fed's forecast is private information, they cannot make this decomposition.

The non-cooperative solution can now be found by working backwards through the sequence of decisions that the players have to make. First, maximize EUF_t taking g_t^e as a fixed parameter; the first order condition gives the Fed's decision rule as a function of the wage setters' action, g_t^e. Then take the expectation of the Fed's decision rule (conditional on information available to the wage setters when they set w_t) to find the wage setters' utility maximizing action, $g_t^{e,nc}$. The resulting <u>Non-cooperative Solution</u> is

$$(13)\quad g_t^{nc} = \pi^* + e_t + y^*/f \qquad g_t^{e,nc} = \pi^* + y^*/f$$

$$y_t = \overline{y} - \theta\varepsilon_t,\ \pi_t = \pi^* + y^*/f - \varepsilon_t$$

$$EUF_t^{nc} = -(1 + f)\sigma_\varepsilon^2 - [1 + (1/f)]y^{*2},\ EUW_t^{nc} = -\sigma_\varepsilon^2$$

This is a Nash solution; neither player can expect to do better, given the strategy of the other. And the solution holds under either information structure; that is, one obtains the same non-cooperative solution whether or not the Fed's forecast of money demand is private information.

Wage setters get exactly what they want in this solution. They set wage inflation high enough that the Fed is not tempted to try to achieve its higher output goal, and the Fed fully accommodates predictable money demand disturbances, stabilizing output to the fullest extent possible about the wage setters' preferred rate. There is, of course, an inflation bias equal to y^*/f. Society is left (on average) with a level of employment that is too low and a rate of inflation that is too high; as Barro and Gordon (1983b) and Rogoff (1983a) point out, this is a model of stag-flation.

Now, it may be thought that the non-cooperative solution just described is a rather silly, destructive situation for the Fed and the wage setters to allow themselves to fall into. The players have conflicting output goals, but the inflation bias benefits no one. An Ideal Solution would eliminate the inflation bias without changing the (average) rate of output;

$$(14)\quad g_t^c = \pi^* + e_t,\; g_t^{e,c} = \pi^*$$

$$y_t = \bar{y} - \theta\varepsilon_t,\; \pi_t = \pi^* - \varepsilon_t$$

$$EUF_t^c = -(1+f)\sigma_\varepsilon^2 - y^{*2} > EUF_t^{nc},\; EUW_t^c = EUW_t^{nc}$$

Here, the Fed fully accommodates the predictable part of the money demand, and it does not try to inflate away θy^*. The solution is efficient in the sense that no other set of strategies, g_t and g_t^e, can make one player better off without making the other worse off; thus the Ideal Solution can be thought of as a cooperative solution. (That is why the superscript "c" is used to identify it.)

As Kydland and Prescott (1977), Canzoneri (1980) and Barro and Gordon (1983a) have all pointed out, the problem with this solution is that the Fed's policy is time inconsistent. Once the wage setters have committed themselves to $g_t^{e,c}$, g_t^c will no longer be the Fed's utility maximizing action. Letting $g_t^e = \pi^*$ in (12), and maximizing EUF_t with respect to g_t, the best <u>Cheating Solution</u> for the Fed is

$$(15)\quad g_t^{ch} = \pi^* + e_t + y^*/(1 + f), \quad g_t^{e,ch} = \pi^*$$

$$EUF_t^{ch} = EUF_t^c + y^{*2}/(1 + f), \quad EUW_t^{ch} = EUW_t^c - y^{*2}/(1 + f)$$

Wage setters cannot set w_t on the expectation that inflation will equal π^* and then count on the Fed to follow through with g_t^c. The Fed will be tempted to inflate a little more than π^*, lowering the real wage and inflating away some of θy^*. Unless the Fed can be credibly precommitted to the ideal policy, wage setters will almost surely play non-cooperative.

Can the Fed be credibly committed to the ideal policy at the time that wage setters have to set w_t in their contract? Views seem to differ as to whether this is a real world problem. Taylor (1983), commenting on Barro and Gordon (1983a and b), points out that society has found sensible ways out of other obvious time inconsistency problems; he cites the example of patent law. Taylor seems skeptical that the Kydland-Prescott game really captures the credibility problem the Fed is generally thought to face.

My own view is that Taylor would probably be right were it not for private information. Absent private information, resolutions of the precommitment problem would be relatively easy to find. Suppose information is symmetric, and wage setters observe e_t at the same time the Fed does. The Fed might be able to establish credibility in the ideal policy by simply running it for a number of periods. If wage setters actually observe e_t each period, they can verify that the Fed is implementing the policy g_t^c, and they might well be willing to go along.[9/] Alternatively, Congress could legislate the ideal policy rule and police the Fed's adherence to it.[10/]

If the Fed's forecast of money demand is private information, a resolution of the precommitment problem is much more difficult to come by. If e_t is not observed by the wage setters or Congress, then direct verification of the Fed's adherence to the ideal policy rule is not possible. The Fed is often urged to announce a forecast, e_t^a, at the time it sets g_t. However, this would be of little use in the present context, for in the language of Hurwicz (1972), the mechanism would not be incentive compatible; the Fed would have an incentive to misrepresent its information.[11/] By letting its announcement, e_t^a, be equal to the true forecast, e_t, plus $y^*/(1 + f)$, the Fed can make the cheating policy (15) look like the ideal policy; that is, $g_t^{ch} = \pi^* + e_t^a$. Straightforward reputational or legislative approaches to the problem seem more likely to result in the scenario outlined in the introduction. The Fed sets a high g_t, claiming to be accommodating an unexpectedly high money demand; the private sector charges that it is instead running an inflationary policy to expand output, and a credibility breakdown ensues.

The results of this section may be summarized as follows. The

Kydland-Prescott game of precommitment has a simple resolution if there is no benefit to the Fed's retaining discretionary powers; the Fed's hands should be tied with a k-percent rule. If the Fed can play a useful stabilization role, then an efficient resolution of the precommitment problem requires that the Fed should retain some discretionary power. However, straightforward resolutions still exist if information is symmetric; the Fed might be able to build a reputation for itself by simply running the ideal policy for a few periods, or Congress could legislate the ideal policy rule and police the Fed's adherence to it. If the Fed's forecast of money demand is private information, then the Fed's adherence to the ideal policy rule cannot be verified directly, and efficient resolutions of the precommitment problem are much more difficult to find.

III. Adding Private Information to Rogoff (1983a) and Barro and Gordon (1983a)

Rogoff (1983a) and Barro and Gordon (1983a) discuss more sophisticated approaches to the precommitment problem. However, neither allows for private information. If the Fed's forecast of money demand is private information, both approaches yield more interesting interpretations of the credibility problem, but Rogoff's does not offer a resolution and Barro and Gordon's does not explain the scenario outlined in the introduction.

The Perverse Policy Maker:

It is sometimes asserted that monetary policy makers have a conservative or anti-inflation bias, or that they have an unnatural proclivity towards monetary targets. Rogoff (1983a) suggests that this might be one of society's ways of trying to improve upon the Non-Cooperative Solution. Society may have found that choosing a policy maker with preferences other than those expressed in (12) leads to a non-cooperative solution that is closer to the Ideal Solution.

Suppose the Fed is run by a perverse policy maker with the utility function

$$(16)\quad UP_t = -(g_t - g_t^e - \delta_t - y^*)^2 - \phi f(g_t - \delta_t - \pi^*)^2$$

UP_t puts more weight on inflation than UF_t if $\phi > 1$. The <u>Non-Cooperative Solution with a Perverse Policy Maker</u> is

$$(17)\quad g_t^p = \pi^* + e_t + y^*/\phi f \qquad g_t^{e,p} = \pi^* + y^*/\phi f$$

$$EUP_t^p = -(1 + \phi f)\sigma_\varepsilon^2 - [1 + (1/\phi f)]y^{*2}, \; EUF_t^p = EUF_t^c - y^{*2}/\phi^2$$

The more perverse is the policy maker (that is, the bigger is ϕ), the closer is solution (17) to the optimal solution (14). The policy maker's sensitivity to inflation allows the wage setters to get their way with less inflation.

Here again, adding private information both complicates and enriches the model. If e_t were directly observable at the end of the period, then wage setters could verify that the policy g_t^p had indeed been implemented, and that UP_t truly represented the policy maker's preferences, rather than UF_t. The policy maker's actions would reveal his preferences.

If e_t is private information, this is not possible. The policy maker could announce a forecast, e_t^a, when he set g_t, but once again the mechanism would not be incentive compatible.[12/] To see this, suppose the policy maker represents his preferences as UP_t (with $\phi > 1$), but actually his preferences are UF_t. Then, substituting $g_t^e = \pi^* + y^*/\phi f$ and $g_t = \pi^* + e_t^a + y^*/\phi f$ into (12) and maximizing EUF_t with respect to e_t^a, the policy maker's optimal misrepresentation, $e_t^a - e_t$, can be seen to be $[(\phi-1)/\phi(1 + f)]y^*$. It can be shown that a marginal increase in ϕ over unity will raise EUF_t above EUF_t^{nc} and lower EUW_t below EW_t^{nc}.

Thus a policy maker with true social preferences has an incentive to look conservative, and if the Fed's forecast of money demand is private information, wage setters have no way of telling for sure. The policy maker's actions and announcements will not reveal his actual preferences. The implications of this observation would appear to be two fold. The model may now provide a more plausible interpretation of history; without private information, Rogoff's resolution of the credibility problem is too stable to explain the comings and goings of individual policy makers, with all of the attendant hoopla, posturing and changes in public perception. However, the very same instability that may help explain history also detracts from the appeal of this approach as a final resolution of the credibility problem.

Rogoff (1983a) does not consider private information, but he does offer an alternative interpretation of his formal analysis that may be helpful in addressing the issue. He suggests that the added weight on inflation in (16) may be viewed as a commitment to targeting; in this case, the policy maker commits himself to targeting the inflation rate. Rogoff also suggests that the policy maker's commitment should be enforced through a system of rewards and punishments; Congress, for example, might be able to achieve an institutional reform by legislating appropriate bureaucratic incentives. If Fed forecasts are indeed private information, then incentive compatibility is one consideration that will make the design of such a system of rewards and punishments more difficult, but the general approach seems worth pursuing. The targeting discussion in section IV may be viewed in this light.

With this alternative interpretation, Rogoff (1983a) also has a plausible explanation for the existence of other forms of targeting. Suppose the policy maker's perversity is a penchant for targeting g_t on π^*, so that

$$(18)\quad UP_t = -(g_t - g_t^e - \delta_t - y^*)^2 - f(g_t - \delta_t - \pi^*)^2 - \tau(g_t - \pi^*)^2$$

In this case, the EUP_t maximizing policy is

$$(19)\quad g_t^p = \pi^* + [(1+f)/(1+f+\tau)]e_t + y^*/(f+\tau)$$

and τ can be chosen to maximize EUF_t.

The size of τ determines the tightness of the targeting procedure. Note that if $\tau > 0$, the policy maker will not try to fully accommodate money demand distrubances. The choice of τ must trade-off the benefits of reducing the inflation bias with the costs of constraining the stabilization effort. The choice of ϕ in the last example would involve a similar calculation if productivity disturbances were added to the model. Rogoff (1983a) considers the relative merits of money, interest rate, inflation rate and nominal income targeting, all in a much more interesting control setting than was practical here.

A Reputation Building Approach:

The credibility problem in the Kydland-Prescott game is one of precommitment; the Fed would like to be able to promise the wage setters that it will run the ideal policy, and then have no reason for reneging on its promise later. Barro and Gordon (1983a) show that Friedman's (1971) solution concept for repeated games offers a resolution to this problem.

Suppose the Fed and the wage setters have arrived at the following set up: Wage setters believe the Fed's promise to run the ideal policy in the current period if the Fed did not renege on its promise last period. (Assume for the moment that there is no private information.) If however the Fed reneged last period, then wage setters revert to their non-cooperative wage settings this period; in subsequent periods, they again believe the Fed's promise, until shown otherwise. Barro and Gordon (1983a) show that the wage setters are not irrational in going along with this set up, because the Fed will never have an incentive to renege.[13/]

To see this, note that the Fed's temptation to renege in period t is measured by

(20) $EUF_t^{ch} - EUF_t^c = y^{*2}/(1 + f)$

while its reward in period t + 1 for not reneging in period t is

(21) $EUF_{t+1}^c - EUF_{t+1}^{nc} = y^{*2}/f$

Since the reward for honesty outweighs the temptation to cheat, it pays the Fed to invest in its reputation; it has no incentive to renege.[14/]

There are two weaknesses in this resolution of the credibility problem that can be dealt with in ways suggested by Green and Porter (1983). (There are some others that are not so easily remedied; they will be discussed at the end of the section.) The first weakness is that the Barro - Gordon solution is, like Rogoff's (1983a), too stable. If taken literally, it implies that one should never expect to see reversions to inflationary periods. The second weakness is that if the Fed's forecast of money demand is private information, then the wage setters can never know for sure when the Fed has reneged. The wage setters do not observe the Fed's forecast at the end of the period, and as was shown in section II, it would do no good to have the Fed announce a forecast, for it would not be believed.

The model can be reformulated to deal with both weaknesses as follows: The wage setters do see the money demand disturbance, δ_t, at the end of the period; they just can't decompose it into the Fed's forecast, e_t, and the residual, ε_t. Thus while it is not feasible to have the wage setters revert to their inflationary wage settings when g_t is greater than $g_t^c = \pi^* + e_t$, it is feasible to have a reversionary period triggered if g_t is greater than $\pi^* + \delta_t + \bar{\varepsilon}$, where $\bar{\varepsilon}$ is some appropriately chosen constant. The probability of a reversion in period t + 1 is [15/]

$$(22)\quad P(g_t - g_t^c - \bar{\varepsilon}) = Pr[\varepsilon_t < g_t - g_t^c - \bar{\varepsilon}]$$

where $P(\cdot)$ is the c.d.f. of the Fed's forecast error, ε_t. If the Fed reneges and sets g_t greater than g_t^c, it increases the probability of a reversion in period t + 1. Similarly, decreasing the constant $\bar{\varepsilon}$ increases the probability of a reversion. Wage setters are not irrational in going along with this set up if $\bar{\varepsilon}$ is small enough that the Fed has no incentive to renege.[16/]

The problem now is to find an $\bar{\varepsilon}$ such that if the Fed considers raising g_t marginally above g_t^c, it will conclude that the expected gain in period t is offset by the expected loss in period t + 1. Suppose g_{t-1} was less than $\pi^* + \delta_{t-1} + \bar{\varepsilon}$ (so that the Fed is not currently in a reversionary period), and let $U(g_t, \bar{\varepsilon})$ be the Fed's expected utility over the next two periods;

$$(23)\quad U(g_t, \bar{\varepsilon}) = EUF_t + PEUF_{t+1}^{nc} + (1 - P)EUF_{t+1}^{c}$$

The Fed has no incentive to renege if $U_{g_t}(g_t^c, \bar{\varepsilon}) < 0$, where

$$(24)\quad U_{g_t}(g_t^c, \bar{\varepsilon}) = 2y^* - P'(-\bar{\varepsilon})[EUF_{t+1}^{c} - EUF_{t+1}^{nc}]$$

The first term is the expected gain in the current period from inflating away some of the output differential, θy^*;[17/] the second term is the expected loss next period from increasing the probability of a reversion. Note that the increase in the probability of a reversion (evaluated at $g_t = g_t^c$) is just the probability that ε_t equals $\bar{\varepsilon}$;

(25) $P'(0 - \bar{\varepsilon}) = p(-\bar{\varepsilon})$

where $p(\cdot)$ is the p.d.f. of the forecast error, ε_t. From (13) and (14), $EUF_{t+1}^c - EUF_{t+1}^{nc} = y^{*2}/f$; so

(26) $U_{g_t}(g_t^c, \bar{\varepsilon}) < 0$ iff $p(-\bar{\varepsilon}) > 2/(y^*/f)$

If y^*/f is greater than 2 (so that $0 < 2/(y^*/f) < 1$), then $\bar{\varepsilon}$ can be chosen to eliminate any incentive for the Fed to renege.

The trigger mechanism for a reversionary period, $g_t > \pi^* + \delta_t + \bar{\varepsilon}$, was designed to be applicable in the private information case, where wage setters cannot tell whether or not the Fed has reneged. As a result, it has a very interesting property. Even if the Fed is always running the ideal policy, there will still be periodic inflationary episodes associated with large negative prediction errors.[18/] Wage setters will see a large g_t and a small δ_t, and they will revert to inflationary wage settings. The frequency of these inflationary reversions depends inversely upon the size of y^*/f. If y^*/f

is large, then from (26), a big $\bar{\varepsilon}$ can be specified; consequently, the probability of a reversion is small, and so is the frequency of inflationary episodes.

So the modification for private information allows the model to explain periodic inflationary reversions.[19/] It also suggests that, despite the apparent instability, Taylor (1983) might be right; the Fed and the private sector may have already hit upon a rather good resolution of the problem. In fact, the worse is the problem to begin with, the better is the resolution; θy^* is the difference between the players' output goals, and if y^*/f (= $\theta y^*/(s/\theta)$) is big, a large $\bar{\varepsilon}$ can be specified, and the frequency of inflationary reversions is small. It should also be noted that in this resolution of the credibility problem, the Fed retains full flexibility for stabilization.

There still remain some weaknesses or unanswered questions: (1) The game cannot be finite, or as Barro and Gordon (1983a) have noted, the solution unravels backward. It is not clear how this fact jibes with the finite terms of office of Fed policy makers. (2) It is not clear that the inflationary reversions modeled here correspond to the policy scenario outlined in the introduction. It is tempting to interpret the triggering of reversions as credibility breakdowns, and the inflationary bouts as times of credibility building or learning. But this would not be appropriate. The reversionary periods modeled here are necessary evils that make the Fed's promise of a non-inflationary policy credible. Their existence are evidence that the credibility problem of precommitment has been resolved. The instability inherent in a straightforward interpretation of Rogoff's (1983a) model may more closely correspond to the policy scenario outlined in the introduction.

IV. The Legislative Approach with Private Information

The legislative approach might provide a resolution that better fits the game at hand. The social contract enacted can have a finite duration, matching the policy maker's term of office. And a dominant player with something to gain, namely the government, can take the active role in initiating the reform.

Congress could legislate incentive compatible policy rules directly, or it could legislate new rules of the game, rules that would constrain the strategies of the players; targeting procedures are an example of the latter. The solutions described in this section could be legislated either way.[20/] However, as was suggested in section II, it is not straightforward to design incentive compatible rules for monetary policy if the Fed's forecast of money demand is private information. It is easier to motivate the solutions described here in terms of targeting procedures.

The tightest form of money targeting is a k-percent rule requiring g_t to be equal to π^* each period. The non-cooperative, Strict Targeting Solution is[21/]

$$(27)\quad g_t^{st} = \pi^* \qquad g_t^{e,st} = \pi^*$$

$$EUF_t^{st} = -(1+f)\sigma_\delta^2 - y^{*2}$$

$$EUW_t^{st} = -\sigma_\delta^2 < EUW_t^{nc}$$

Comparing (27) and (13), it turns out that $EUF_t^{st} > EUF_t^{nc}$ if $y^{*2}/f > (1 + f)\sigma_e^2$. y^{*2}/f is equal to $EUF_t^c - EUF_t^{nc}$, the social disutility resulting from the non-cooperative behavior of the players; $(1 + f)\sigma_e^2$ is the social utility derived from having the Fed fully accommodate perceived money demand disturbances.

Legislating a k-percent rule is obviously incentive compatible, and it is better than doing nothing if the disutility of the inflation bias due to non-cooperative behavior outweighs the utility of stabilization. Of course, legislating a k-percent rule and doing nothing are the two polar extremes. There are better, less rigid, forms of targeting that limit the Fed's temptation to inflate, but leave it with some latitude for stabilization.

For example, Congress might just require the Fed to meet its target on average. In a two period social contract, beginning in period $t = 1$, average targeting requires that

$$(28)\quad g_1 + g_2 = 2\pi^*$$

This constrains the Fed's behavior intertemporally; in fact, once g_1 is set, g_2 is known with certainty, by everyone, including the wage setters writing contracts for period $t = 2$. The non-cooperative, Average Targeting Solution is

$$(29)\quad g_1^{at} = \pi^* + \rho e_1 + y^*/2f \qquad g_1^{e,at} = \pi^* + y^*/2f$$

$$g_2^{at} = \pi^* - \rho e_1 - y^*/2f \qquad g_2^{e,at} = g_2^{at}$$

$$EUF_t^{at} = -(1 + f)[\sigma_\delta^2 - (\rho/2)\sigma_e^2] - [1 + (1/4f)]y^{*2}$$

$$EUW_t^{at} = -\sigma_\varepsilon^2 - (1/2)[1 + (1 - \rho)^2]\sigma_e^2 < EUW_t^{nc}$$

where $\rho = (1 + f)/(1 + 2f) < 1$, and EUF_t^{at} and EUW_t^{at} are expected utilities averaged over the two periods.

Unlike the ideal rule $g_t^c = \pi^* + e_t$, the policy rules in (29) are incentive compatible. Congress could legislate them directly, instead of the targeting procedure; the Fed would announce e_1^a at the time it set g_1, and would have no incentive to make e_t^a greater than e_t.[22/]

Here, the Fed will do some stabilizing, but the stabilization will only be partial, and it will only occur in the first period. In addition, the new flexibility results in some inflation bias in the first period; the averaging requires the bias to be reversed in the second period, and this curbs the Fed's temptation to inflate. Compared with strict targeting, average targeting's advantage is in its flexibility for stabilization; compared with no targeting, average targeting's advantage is in its reduced inflation bias.

In fact, it can be shown that

(30) $EUF_t^{at} > EUF_t^{st} \qquad \text{if} \quad R > \phi$

$$EUF_t^{at} > EUF_t^{nc} \qquad \text{if} \quad R < (1 + \eta)\phi$$

where

$$R = \frac{(1+f)\sigma_e^2}{y^{*2}/f} \qquad \text{and} \qquad \begin{array}{l} \phi = (1 + 2f)/(1+f) \\ \eta = 2/(1 + 3f) \end{array}$$

R measures the utility of stabilization relative to the disutility resulting from non-cooperative behavior (that is, the inflation bias). Note that R depends upon the ability of the Fed to predict money demand disturbances (as measured by σ_e^2), as well as preferences. If $R < \phi$, the Fed's hands should be tied completely with the k-percent rule. If $\phi < R < (1 + \eta)\phi$, the Fed should only be constrained by average targeting. If $R > (1 + \eta)\phi$, then the Fed's stabilization efforts should not be constrained at all, and the full inflation bias should be tolerated.

Even here, the full range of posibilities has only begun to be considered. Average targeting is a rather heavy handed way of achieving flexibility; it is inefficient because it does not link flexibility directly to the size of the disturbances that the Fed ought to be allowed to accommodate. An ideal form of targeting would replace (28) with

$$(31) \quad g_1 + g_2 = 2\pi^* + e_1 + e_2$$

However, legislating (31) would not provide an incentive compatible resolution to the credibility problem if the Fed's forecast is private information; the Fed would have an incentive to announce e_1^a and e_2^a larger than e_1 and e_2.[23/] Legislating the ideal form of targeting would do no more good than legislating the ideal rule itself.

There is however an incentive compatible form of targeting that is more flexible than average targeting. Flexible targeting requires that

$$(32)\quad g_1 + g_2 = 2\pi^* + c\delta_1 \qquad c > 0$$

where δ_1 is of course observed by the wage setters at the end of the first period. With c set equal to one, the non-cooperative Flexible Targeting Solution (c = 1) is

$$(33)\quad g_1^{ft} = \pi^* + e_1 + y^*/2f \qquad g_1^{e,ft} = \pi^* + y^*/2f$$

$$g_2^{ft} = \pi^* + \varepsilon_1 - y^*/2f \qquad g_2^{e,ft} = g_2^{ft}$$

$$EUF_t^{ft} = -\frac{1}{2}(1+f)(\sigma_\delta^2 + \sigma_e^2) - \frac{1}{2}f\sigma_\varepsilon^2 - [1 + (1/4f)]y^{*2}$$

$$EUW_t^{ft} = -\frac{1}{2}(\sigma_\delta^2 + \sigma_\varepsilon^2) > EUW_t^{nc}$$

where EUF_t^{ft} and EUW_t^{ft} are expected utlities averaged over the two periods. The policy rules in (33) are incentive compatible and could be legislated directly; the Fed would have no incentive to misrepresent its forecast.[24/]

With c set equal to one, the Fed would find it optimal to stabilize fully in the first period. However, it turns out that this is again a polar case. Suppose c is chosen optimally; that is, let c^* be the value of c that maximizes EUF_t^{ft}. It can be shown that

$$(34) \quad c^* \to 1 \quad \text{as} \quad \sigma_\varepsilon^2/\sigma_e^2 \to 0 \quad \text{and} \quad c^* \to 0 \quad \text{as} \quad \sigma_\varepsilon^2/\sigma_e^2 \to \infty$$

and that if c is less than one, the Fed will not find it optimal to fully accommodate e_1.[25/] Full stabilization in the first period should be allowed only if the Fed is very good at forecasting money demand. In fact, with c set equal to one, flexible targeting beats average targeting only if $\sigma_\varepsilon^2/\sigma_e^2$ is less than $(1 + f)/(1 + 2f)$; with c set optimally, flexible targeting always beats average targeting.

It is beyond the scope of the present paper to try to identify the most efficient form of targeting. That would require a more interesting control setting (with dynamics, productivity shocks, etc.) and a multi-period targeting horizon (presumably corresponding to the policy maker's term of office). In addition, Townsend's (1983) simple example of multi-period contracting suggests that a complete specification of the most efficient social contract would be difficult, if not impossible.

The examples provided here do however provide two insights that should be helpful in designing efficient targeting procedures. (1) The procedure should be defined in terms of variables that are observed by all market participants; actual money demand was used in the examples above, but actual inflation would have served just as well. Credibility problems will remain if latitude for stabilization is incorporated in any other way. (2) The tightness of the targeting procedure should depend upon the ability of the Fed to stabilize (modeled above as the Fed's ability to predict money demand distburances that weren't reflected in existing wage settlements) and the importance of stabilization relative to the disutility resulting from the credibility problem (modeled above as an inflation bias).

V. Conclusion

Kydland and Prescott's (1977) game between wage setters and the Fed provides a plausible explanation of the scenario outlined in the introduction, especially if the Fed's forecast of money demand is private information. Without private information, Taylor's (1983) observations are persuasive; it is difficult to see why straightforward resolutions of the precommitment problem would not have already been found. But with private information, the Fed cannot demonstrate its forebearance by simply running a non-inflationary policy for a few periods, and it is not feasible to simply legislate the "optimal" feedback rule. There are incentive compatible policy rules that could be legislated, but they are more complicated than what is usualy envisioned; the rules suggested here can be equivalently encoded as targeting procedures. Private information also enriches Rogoff's (1983a) model of the perverse policy maker and Barro and Gordon's (1983a) reputation model; both are too stable, without private information, to explain periodic inflationary breakdowns.

There is a recurring tension in this paper that provides a new structure for the old controversy over rules versus discretion. Some rule or discipline must be placed on the Fed's behavior to resolve the credibility problem and achieve a better, non-inflationary, game outcome.

The trick is to do this in a way that leaves the Fed with the maximum amount of discretion for stabilization. The modified Barro-Gordon model resolves the credibility problem in a very efficient manner, as does Rogoff's perversely inflation conscious policy maker (in the absence of productivity distrubances and private information, anyway). In either case, the Fed is left with the latitude to fully accommodate percieved money demand disturbances. The targeting resolutions discussed in section V do not appear to be as successful; a decrease in the inflation bias is achieved at the expense of discretion, the optimal amount of discretion depending upon the ability of the Fed to stabilize and importance of stabilization relative to the elimination of the inflation bias.

More generally, private information should be viewed as a constraint, much like a technology constraint. The job of welfare economists is to develop utility maximizing mechanisms, given just such constraints. The trade-off between the discretion needed for stabilization and the rule required for a resolution of the credibility problem is not well understood and deserves more attention.

Footnotes

*/ I would like to thank J. Friedman, E. Green, D. Henderson, K. Rogoff, J. Taylor and an anonymous referee for their helpful comments. However, the views presented here are my own. They are not necessarily shared by the Federal Reserve Board or any other member of its staff.

1/ Targeting money or interest rates is not likely to be efficient in terms of meeting final targets; see, for example, Canzoneri (1977). This observation is the principal theme of Rogoff (1983a).

2/ See Barro and Gordon (1983a, b), Taylor (1983), Green (1983), Rogoff (1983a, b), Baccus and Driffill (1984), and more generally, Lee (1981). Phelps (1967) had many of the elements of this game.

3/ This paper borrows heavily from Hurwicz (1972), as does much of the recent contracting literature; see Townsend (1983), Hart (1983), and the Quarterly Journal of Economics, 98, (supplement, 1983).

4/ They could also be motivated along the lines of Peel's Bank Act of 1844; see Leijonhufvud's (1982) discussion of the Peel-Friedman system.

5/ This assumption has been criticized; see the exchange between Barro and Fischer in Journal of Monetary Economics 3, 1977, and also Waldo (1981).

6/ Aizenman and Frenkel (1983) provide a rationale for this utility function. The analysis is simplified by abstracting from productivity disturbances; see Gray (1976) and Rogoff (1983a).

7/ It is not necessary to specify conflicting utility functions in order to generate the time inconsistency on which this game rides; see Fischer (1980) and Canzoneri and Henderson (1985) for expositions of the time consistency problem that are mathematically unencumbered.

8/ While the two views are equivalent for present purposes, they do have different implications for the first-best solution to the policy problem posed here. Canzoneri's (1980) view suggests labor legislation, while Barro and Gordon's (1983a, b) points to fiscal policy reform.

9/ Some experimental studies suggest that players may be willing to go along in this way; see Roth and Schoumaker (1983).

10/ In a more realistic setting, the ideal policy rule could well be more complicated, but this should not in and of itself create a credibility problem. Credibility problems arise only when scope has been left for cheating.

11/ Wage setters can tell if the Fed is telling the truth "on average", for then the sum $\Sigma(\delta_t - e_t^a) = \Sigma\varepsilon_t$ should tend to zero. If the relationship is an enduring one, and if utility is not discounted over time, then this approach might work. See Radner (1981).

12/ Footnote 11 applies here as well.

13/ Actually, Barro and Gordon (1983a) postulate a different utility function for the Fed, one that is linear (rather than quadratic) in the output term, and consequently, for them the Ideal Solution is not an achievable outcome. However, they can improve upon the Non-Cooperative Solution. See also the next footnote.

14/ For simplicity, the Fed is assumed not to discount future utility. If future utility is heavily discounted, two or more reversionary periods might be required; if future utility is very heavily discounted, the approach may not work at all. See Friedman (1971).

15/ To see this, note that $g_t > \pi^* + \delta_t + \bar{\varepsilon} = \pi^* + e_t + \varepsilon_t + \bar{\varepsilon}$ $= g_t^c + \varepsilon_t + \bar{\varepsilon}$ is equivalent to $\varepsilon_t < g_t - g_t^c - \bar{\varepsilon}$.

16/ The reformulation of the model can also be explained in terms of incentive compatibility. Suppose the Fed announces e_t^a and sets $g_t = \pi^* + e_t^a$. Since $g_t^c = \pi^* + e_t$, $g_t = g_t^c$ if and only if $e_t^a = e_t$; no reneging, as defined above, is equivalent to incentive compatibility. Suppose the wage setters revert to inflationary wage settings if $e_t^a > \delta_t + \bar{\varepsilon}$. It is easy to show that $e_t^a > \delta_t + \bar{\varepsilon}$ if and only if $g_t > \pi^* + \delta_t + \bar{\varepsilon}$, and that (22) can be expressed as $P(e_t^a - e_t - \bar{\varepsilon}) = Pr[\varepsilon_t < e_t^a - e_t - \bar{\varepsilon}]$. If the Fed sets $e_t^a > e_t$, it increases the probability of a reversion in period t + 1. The arguement exactly parallels that in the text.

17/ It is calculated by substituting $g_t^e = g_t^{e,c} = \pi^*$ into UF_t in (12), differentiating EUF_t w.r.t. g_t, and evaluating the result at $g_t = g_t^c$.

18/ $g_t^c = \pi^* + e_t > \pi^* + \delta_t + \bar{\varepsilon} = \pi^* + e_t + \varepsilon_t + \bar{\varepsilon}$ is equivalent to $\varepsilon_t < -\bar{\varepsilon}$.

19/ This is the basic insight of Green and Porter (1983).

20/ They could also be achieved by institutional reform; see footnote 4.

21/ All of the non-cooperative solutions in this section are calculated by working backwards, as in section II.

22/ This can be verified by substituting $g_1 = \pi^* + \rho e_1^a + y^*/2f$ and $g_2 = 2\pi^* - g_1$ into the Fed's two period optimization problem and maximizing w.r.t. e_t^a.

23/ To prove this, first calculate the non-cooperative solution with no private information; then replace e_1 and e_2 with e_1^a and e_2^a in the Fed's policy rules for g_1 and g_2, and calculate the Fed's utility maximizing values for e_1^a and e_2^a.

24/ This can be verified by substituting $g_1 = \pi^* + e_1^a + y^*/2f$ and $g_2 = 2\pi^* + \delta_1 - g_1$ into the Fed's two period optimization problem and maximizing w.r.t. e_1^a.

25/ The general Flexible Targeting Solution is

$$g_1^{ft} = \pi^* + [(1 + f + cf)/(1 + 2f)]e_1 + (1/2f)y^*$$

$$g_2^{ft} = 2\pi^* + c\delta_1 - g_1^{ft}$$

$$EUF_t^{ft} = -(1 - c)^2[f(1 + f)/2(1 + 2f)]\sigma_e^2 - c^2(f/2)\sigma_\varepsilon^2 - [(1 + f)/2](\sigma_\delta^2 + \sigma_\varepsilon^2) - [1 + (1/4f)]y^{*2}$$

REFERENCES

Aizenman, J. and J. Frenkel, "Wage Indexation and the Optimal Exchange Rate Regime," unpublished manuscript, Aug. 1983.

Baccus, D. and J. Driffill, "Rational Expectations and Policy Credibility Following a Change in Regime," unpublished manuscript, May 1984.

Barro, R. and D. Gordon, "Rules, Discretion and Reputation in a Model of Monetary Policy," Journal of Monetary Economics, July 1983a, p. 101-122.

Barro, R and D. Gordon, "A Positive Theory of Monetary Policy in a Natural-Rate Model," Journal of Political Economy 91, Aug. 1983b, p. 589-610.

Canzoneri, M. "The Intermediate Control Problem," Journal of Money Credit and Banking, May 1977, p. 368-371.

Canzoneri, M. "The Time Inconsistency Problem," unpublished manuscript, Aug. 1980.

Canzoneri, M. and D. Henderson, Strategic Aspects of Policy Making in Open Economies, monograph forthcoming in 1985.

Fischer, S. "Long-Term Contracts, Rational Expectation and the Optimal Money Supply Rule," Journal of Political Economics, Feb. 1977, p. 191-205.

Fischer, S. "Dynamics Inconsistency, Cooperation and the Benevolent Dissembling Government," Journal of Economic Dynamics and Control, Feb. 1980, p. 93-108.

Friedman, J. "A Noncooperative Equilibrium for Supergames," Review of Economic Studies, 38, Jan. 1971, p. 861-74.

Gray, J.A., "Wage Indexation: A Macroeconomic Approach," Journal of Monetary Economics 2, 1976, p. 221-235.

Green, E. and R. Porter, "Noncooperative Collusion Under Imperfect Price Information," Econometrica, Jan. 1984.

Green, S. "Unattainable Objectives, Incomplete Information and Regrettable Macroeconomic Policy," unpublished manuscript, April 1983.

Hart, O. "Optimal Labor Contracts Under Asymmetric Information: An Introduction," Review of Economic Studies, Jan. 1983, p. 3-36.

Hurwicz, L. "On Informationally Decentralized Systems," in Decision and Organization: A Volume in Honor of Jacob Marschak, edited by C.B. McGuire and Ray RAdner, Amsterdam: North-Holland, 1972.

Kydland, F. and E. Prescott, "Rules Rather than Discretion: the Inconsistency of Optimal Plans," Journal of Political Economy, 85, June 1977, p. 4730-91.

Lee, W., Optimal Policy Selection and Evaluation with Rational Expectations, Columbia University, Ph.D. dissertation, 1981.

Leijonhufvud, A. "Constitutional Constraints on the Monetary Powers of Goverment," presented at the Conference on Constitutional Economics: The Emerging Debate sponsored by The Heritage Foundation, Nov. 1982.

Phelps, E. "Phillips Curves, Expectations of Inflation and Optimal Employment Over Time," Economica, Aug. 1967, p. 254-281.

Radner, R., "Monitoring Cooperative Agreements in a Repeated Principal - Agent Relationship," Econometrica II, 1981, p. 1127 - 1148.

Rogoff, K. "The Optimal Degree of Commitment to an Intermediate Monetary Traget: Inflation Gains Versus Stabilization Costs," unpublished manuscript, June 1983a.

Rogoff, K. "Cooperative Monetary Policies with Flexible Exchange Rates," unpublished manuscript, Sept. 1983b.

Roth, A. and F. Schoumaker, "Expectations and Reputations in Bargaining: An Experimental Study," American Economic Review, June 1983, p. 362-372.

Taylor, J. "Comments on 'Rules, Discretion and Reputation in a Model of Monetary Policy'", Journal of Monetary Economics 12, July 1983, p. 123-125.

Thompson, E. "Who Should Control the Money Supply," American Economic Review, May 1981, p. 356-361.

Townsend, R. "Optimal Multiperiod Contracts and the Gain from Enduring Relationships under Private Information," Journal of Political Economy, vol. 90, no. 6, 1982.

Waldo, D., "Sticky Nominal Wages and the Optimal Employment Rule", Journal of Monetary Economics 7, 1981, p. 339-353.

FIGURE 1