REPUTATION IN A MODEL OF MONETARY POLICY WITH INCOMPLETE INFORMATION*

Robert J. BARRO

University of Rochester, Rochester, NY 14627, USA

Previous models of rules versus discretion are extended to include uncertainty about the policy-maker's 'type'. When people observe low inflation, they raise the possibility that the policymaker is committed to low inflation (type 1). This enhancement of reputation gives the uncommitted policymaker (type 2) an incentive to masquerade as the committed type. In the equilibrium the policymaker of type 1 delivers surprisingly low inflation – with corresponding costs to the economy – over an extended interval. The type 2 person mimics this outcome for awhile, but shifts eventually to high inflation. This high inflation is surprising initially, but subsequently becomes anticipated.

1. Introduction

In previous research David Gordon and I (1983b) built on the work of Kydland and Prescott (1977) and others to distinguish rules versus discretion in monetary policy. When the policymaker could make binding commitments (rules), it was feasible to achieve low average inflation. This result was superior to that attainable with no commitments (discretion), where inflation tended to be high, but not surprisingly high. Because of the incentive to create inflation surprises ex post (stemming from a desire to relieve some existing distortions in the economy), only high inflation turned out to be incentive-compatible for the policymaker who was not bound by a rule.

In some subsequent work [Barro and Gordon (1983a)], we considered reputational forces that might substitute for formal rules. When future inflationary expectations were tied to current actions, the discretionary policymaker was motivated to keep inflation down. Thereby we found that an equilibrium entailed lower inflation, which corresponded to a weighted average of the outcome under a rule and that under discretion. However, there were two shortcomings of this approach. First, the reputational possibility hinged on an infinite horizon (otherwise the guaranteed cheating in the last period leads to an unraveling of the solution). Second, the equilibrium was not unique.[1] From the positive standpoint of using the model to predict the policymaker's

behavior – and thereby to predict monetary growth and inflation – this last feature is at least unfortunate.

Kreps and Wilson (1982) and Milgrom and Roberts (1982) have dealt with analogous problems in the area of industrial organization by introducing uncertainty about one of the player's objectives. In the present context there could be uncertainty about either the policymaker's preferences or about his technology for making commitments. Then the extent of reputation or credibility would correspond to an outside observer's subjective probability that the policymaker is of one type or another. In general, current policy actions influence this probability and thereby affect the way that people learn about the policymaker's true type. In addition, the policymaker takes this learning process into account when deciding how to perform. Aside from giving content to notions of reputation and learning, this approach has the advantage of not depending on an infinite horizon and of sometimes delivering a unique perfect equilibrium. In the following I apply this analysis to the setting of monetary policy.[2]

2. Setup of the model

As in previous models [Barro and Gordon (1983a, b)], costs for period t depend on actual and unexpected inflation,

$$z_t = z\left(\pi_t, \pi_t - \pi_t^e\right),\tag{1}$$

where π_t is actual inflation for period t and π_t^e is the representative person's forecast of π_t as of the beginning of period t. I assume $\partial z/\partial|\pi_t| > 0$ and $\partial z/\partial(\pi_t - \pi_t^e) < 0$ in the relevant range. Thus, costs rise with the magnitude of inflation, $|\pi_t|$, but fall with surprise inflation, $\pi_t - \pi_t^e$. The benefit from surprise inflation can reflect some existing distortions, such as taxation of market income or unemployment insurance, that lead to inefficiently low levels of employment and output. Through the standard mechanism of the expectational Phillips Curve, surprise inflation (reflecting surprise monetary expansion) raises output and is thereby beneficial. Alternatively, the use of surprise inflation as a capital levy on nominally-denominated government obligations is desirable because it lessens the need for distorting income taxes or other types of non-lump-sum taxes [see Barro (1983)]. For a given value of surprise inflation, $\pi_t - \pi_t^e$, the minimum of costs in eq. (1) occurs at $\pi_t = 0$. (The model can readily be modified so that the minimum obtains at an arbitrary value $\bar{\pi}$.)

---

[1] This problem appears also in the related literature from game theory – see Friedman (1971) and Green and Porter (1984).

[2] Previous discussions of this approach in the context of monetary policy include Barro and Gordon (1983a, p. 119), Backus and Driffill (1984, 1985), Tabellini (1983) and Horn and Persson (1985, sec. 6).

The policymaker strives to minimize the expected present value of costs,

$$\mathrm{E}\left[z_t + z_{t+1}/(1+r) + \cdots + z_T/(1+r)^{T-t}\right],$$

where $r > 0$ is the exogenous and constant real discount rate and T is the terminal period, which is discussed below. Since everybody agrees on the merits of this objective, there is no principal-agent problem in the model.

In the case where the policymaker makes a serious advance commitment to inflation – so that $\pi_t^e = \pi_t$ follows at once from general knowledge of this commitment – the cost-minimizing value of the inflation rate is $\pi_t = 0$ for all periods. The corresponding cost for each period is normalized to zero – that is, $z(0,0) = 0$.

In a discretionary situation – where commitments are precluded – the policymaker takes current and future inflationary expectations as given. Then the cost-minimizing inflation rate for period t is some value $\hat{\pi}_t$, which generally depends on π_t^e – that is,

$$\hat{\pi}_t = \phi(\pi_t^e). \tag{2}$$

Since people understand the policymaker's behavior, it also follows here that $\pi_t^e = \hat{\pi}_t = \phi(\pi_t^e)$, so that surprise inflation, $\hat{\pi}_t - \pi_t^e$, is zero. For a given form of the cost function in eq. (1) (with appropriate curvature properties), the equilibrium determines some number $\hat{\pi}$, which is such that $\hat{\pi} = \phi(\hat{\pi})$. This result is the discretionary equilibrium as discussed in previous research. The cost for each period in this equilibrium is

$$z_t = \hat{z} = z(\hat{\pi}, 0) > 0. \tag{3}$$

Hence, the outcome is higher costs than under the zero-inflation rule. This result follows because the absence of commitments leads to inflation that is high but not surprisingly high.

Suppose now that there are two types of policymakers. Type 1 is capable of commitments and binds himself to the outcome $\pi_t = 0$ for each period. Type 2 has no capacity to make commitments and simply strives at each date t to minimize the expected present value of costs, $\mathrm{E}(z_t + z_{t+1}/(1+r) + \cdots + z_T/(1+r)^{T-t})$, where costs for each period satisfy eq. (1).[3] Although the policymaker knows his own type, the private agents cannot discern this type directly. Instead they attempt to infer the true type from observed

---

[3] It is basically equivalent if the policymakers differ by their relative weights for the costs and benefits of inflation. But in the present formulation each policymaker has the same tastes, which just reflect the preferences of the 'representative person'. The suggestion of uncertainty about whether one player has or has not made a commitment appears in Milgrom and Roberts (1982, p. 303) and has also been used by Tabellini (1983).

performance, which means that they learn from experience. The policymaker of type 2, who is not bound by commitments, understands this learning process and may be motivated to exploit it when making choices of inflation rates. Throughout I assume that the parameters of the cost function in eq. (1), as well as the real interest rate r, are time-invariant and common knowledge.

The policymaker is installed at the beginning of period 0 and remains in power through period T. Thus, the game between the policymaker and the private agents has a known, finite horizon.[4] (The value of T is common knowledge.) In some circumstances the horizon can be identified with the term of office. However, that perspective may be too narrow, since various mechanisms can motivate lame-duck officeholders to behave. The horizon can also be interpreted in terms of the persistence of differences across policymakers with respect to capacity for commitments (or preferences about inflation). In other words, the distinction between type 1 and type 2 may not be permanent. Then a large value of T signifies that this designation by type holds up for a long time. The form of the subsequent results continues to apply when the horizon becomes arbitrarily long.

If people knew that the policymaker was of type 2, then the (perfect, Nash) equilibrium would be the discretionary one, $\pi_t = \hat{\pi}$, as described before.[5] Further, since the type 1 person always picks $\pi_t = 0$, the policymaker of type 2 can conceal his identity only by choosing zero inflation. The motivation for doing this is to hold down subsequent inflationary expectations, which helps to keep down future costs. However, since the game ends in period T, there is no gain from masquerading as a type 1 person after this period. Therefore, the policymaker of type 2 definitely picks $\pi_t = \hat{\pi}$ in period T.

Suppose that there is an interval $(\tau, T-1)$ for some $\tau \geq 0$, during which the policymaker of type 2 randomizes the choice of inflation between 0 and $\hat{\pi}_t = \phi(\pi_t^e)$. (Recall that the value $\hat{\pi}_t$ minimizes costs for period t when expectations are given.) The motivation for randomizing can arise only if the true type has yet to be revealed – specifically, if $\pi_t = 0$ applied in each prior period. Accordingly, let p_t be the probability that a type 2 policymaker sets $\pi_t = 0$, conditional on having chosen zero inflation in all previous periods. Then the conditional probability of picking $\pi_t = \hat{\pi}_t$ is $1 - p_t$. The period $(\tau, T-1)$ is the one where $0 < p_t < 1$ is supposed to apply. In this interval the probability of selecting $\pi_t = 0$ is neither zero nor one, but is in the interior of this range.

Let α_t be the representative person's subjective probability as of the start of period t that the policymaker is of type 1. The probability for period 0, α_0, is a

---

given value and is common knowledge. In the absence of other information, α_0 would be the fraction of the population of potential policymakers who are capable of commitments (and are therefore of type 1).

The two possible outcomes for inflation at date t are $\pi_t = 0$ or $\pi_t = \hat{\pi}_t = \phi(\pi_t^e)$. The former occurs if the policymaker is of type 1 (probability α_t) or if the policymaker is of type 2 (probability $1 - \alpha_t$) but masquerades as a type 1 person (conditional probability p_t). Hence, the value $\hat{\pi}_t$ emerges only if the policymaker is of type 2 (probability $1 - \alpha_t$) and does not masquerade (conditional probability $1 - p_t$). For later use, let $\tilde{p}_t$ be the representative person's perception of p_t. Then expected inflation for period t is

$$\pi_t^e = \hat{\pi}_t(1 - \alpha_t)(1 - \tilde{p}_t). \tag{4}$$

Eq. (4) indicates the best forecast of period t's inflation rate, π_t, given α_t and $\tilde{p}_t$. Note that, since there are a large number of private agents, no individual has the incentive to pretend that his expectations are different from this value of π_t^e. In particular, each person takes as given the policymaker's behavior and everyone else's method for formulating expectations. It follows that there is no reason for atomistic agents to behave strategically – each one just aims for the best forecast of inflation, given how everyone else is acting.[6] (Presumably, accurate forecasts of inflation aid in other individual decisions, which are unnecessary to detail for present purposes.)

As people observe 'good behavior' from the policymaker – that is, $\pi_t = 0$ – they revise upward the probability that the policymaker is of type 1. The adaptation formula follows from Bayes' Law as

$$\alpha_{t+1} \equiv \text{Prob}(\text{type } 1|\pi_t, \pi_{t-1},\ldots = 0)$$

$$= \frac{\text{Prob}(\text{type } 1|\pi_{t-1},\ldots = 0) \cdot \text{Prob}(\pi_t = 0|\text{type } 1)}{\text{Prob}(\pi_t = 0|\pi_{t-1},\ldots = 0)} \tag{5}$$

$$= \frac{\alpha_t \cdot 1}{\alpha_t + (1 - \alpha_t)\tilde{p}_t}.$$

Note that if $0 < \alpha_t < 1$ and $0 \le \tilde{p}_t < 1$, then $\alpha_{t+1} > \alpha_t$. In other words, the

---

[6]Backus and Driffill (1985) treat π_t^e as a dichotomous variable, which in the present context would take on only the values 0 or $\hat{\pi}_t$. In this case π_t^e cannot be interpreted as expected inflation. In another paper, Backus and Driffill (1984) consider a two-sided game in which the private sector and the policymaker each behave strategically. This setting is inappropriate if the private sector consists of a large number of independent agents, as I assume in the present paper. The assumption of a monolithic private sector might be applicable to a monopoly trade union that bargains for an economy-wide wage rate – which is the setting imagined by Horn and Persson (1985, sec. 6). However, if the private sector acts collectively, then it is unclear why there is an existing distortion (externality) that underlies the benefit from unexpected inflation in eq. (1).

observation of $\pi_t = 0$ raises the probability that the policymaker is of type 1. From the standpoint of a type 2 policymaker, this learning process means that masquerading as a type 1 person builds up one's image as someone who is committed to low inflation.

3. The incentives of policymakers

Consider now the incentives of a policymaker of type 2. Let $V_t(\alpha_t)$ be the minimized expected present value of costs from date t onward, $E[z_t + z_{t+1}/(1+r) + \cdots + z_T/(1+r)^{T-t}]$, conditional on having chosen zero inflation in all previous periods. In this case the policymaker has a certain amount of current reputation, as summarized by the value of α_t. The overall cost, $V_t(\alpha_t)$, equals the expected cost for period t, $E(z_t)$, plus the expected present value of costs from date $t+1$ onward. With probability p_t the policymaker sets $\pi_t = 0$, which generates the amount of reputation α_{t+1} shown in (5). In this case the minimized expected present value of costs beginning at date $t+1$ is $V_{t+1}(\alpha_{t+1})$. With probability $1 - p_t$ the policymaker sets $\pi_t = \hat{\pi}_t$. Then the reputation is lost $(\alpha_{t+1} = \alpha_{t+2} = \cdots = 0)$ and the costs each period are the discretionary ones $\hat{z}$ from eq. (3). The present value (expressed in terms of period t units) of this flow from date $t+1$ through date T is $(\hat{z}/r) \cdot [1 - 1/(1+r)^{T-t}]$. Putting these results together implies

$$V_t(\alpha_t) = E(z_t) + \frac{1}{1+r} \cdot p_t V_{t+1}(\alpha_{t+1})$$

$$+ (1 - p_t)(\hat{z}/r)\left[1 - 1/(1+r)^{T-t}\right], \tag{6}$$

for $\tau \le t \le T - 1$ and where α_{t+1} satisfies eq. (5). The value for $E(z_t)$ is given by

$$E(z_t) = p_t \cdot z(0, -\pi_t^e) + (1 - p_t) \cdot z(\hat{\pi}_t, \hat{\pi}_t - \pi_t^e), \tag{7}$$

where $\pi_t^e = \hat{\pi}_t(1 - \alpha_t)(1 - \tilde{p}_t)$ from eq. (4). Note that $z(\hat{\pi}_t, \hat{\pi}_t - \pi_t^e) < z(0, -\pi_t^e)$, since $\hat{\pi}_t$ minimizes period t's costs for given inflationary expectations.

The policymaker of type 2 selects p_t at the start of period t in order to minimize the expected present value of costs. At the same time – that is, without seeing the value for p_t or the realization of π_t – the representative person perceives that the probability of a type 2 policymaker choosing zero inflation is $\tilde{p}_t$. However, there is nothing to constrain the policymaker at a point in time to set $p_t = \tilde{p}_t$. The policymaker does not commit himself in advance to pick zero inflation with some designated probability – in fact, there is no commitment even to pursuing a mixed strategy. Further, the private agents cannot verify ex post what probability the policymaker used. If people could observe this probability (and hence verify it), then they would know that the policymaker was of type 2. (The type 1 person does not pursue a mixed

strategy.)[7] Therefore, the analysis will go through only if the probability, p_t, is unobservable to the public, even *ex post*. The only thing that people observe – at the start of the next period – is the value of π_t from the previous period. At a point in time the uncommitted policymaker can choose any value of p_t that he wishes, which includes the pure strategies of picking $\pi_t = 0$ or $\pi_t = \hat{\pi}_t$ with probability 1. Thus, the outcome $p_t = \tilde{p}_t$ can emerge only if this choice for p_t is at least as good for the policymaker as any other value of p_t, *given people's perceptions* $\tilde{p}_t$, the value of α_t, and the structure of the optimization problem as laid out in eqs. (6) and (7).

Note that, for given $\tilde{p}_t$, $V_t(\alpha_t)$ is a linear function of p_t in eq. (6).[8] Thus, if $V_t(\alpha_t)$ were increasing in p_t, then the policymaker would set $p_t = 0$ ($\pi_t = \hat{\pi}_t$), whereas if $V_t(\alpha_t)$ were decreasing in p_t, then he would set $p_t = 1$ ($\pi_t = 0$). In order for the policymaker to be willing to randomize – that is, to set p_t in the interior where $0 < p_t < 1$ – it must be that $V_t(\alpha_t)$ is independent of p_t. This independence holds for a linear function if and only if

$$0 = \partial V_t(\alpha_t)/\partial p_t|_{\tilde{p}_t, \alpha_t} = z(0, -\pi_t^e) - z(\hat{\pi}_t, \hat{\pi}_t - \pi_t^e)$$

$$+ V_{t+1}(\alpha_{t+1})/(1 + r)$$

$$- (\hat{z}/r)\left[1 - 1/(1 + r)^{T-t}\right]. \tag{8}$$

Eq. (8) applies in the interval, $(\tau, T - 1)$, where randomization occurs ($0 < p_t < 1$).

The result in eq. (8) implies that the minimized expected present value of costs, V, must satisfy the condition

$$\frac{V_{t+1}(\alpha_{t+1})}{1 + r} = -z(0 - \pi_t^e) + z(\hat{\pi}_t, \hat{\pi}_t - \pi_t^e)$$

$$+ (\hat{z}/r)\left[1 - 1/(1 + r)^{T-t}\right], \tag{9}$$

for $\tau \le t \le T - 1$. Using this result to substitute out for V_t and V_{t+1} in eq. (6) leads to the basic condition that supports randomization of policy,[9]

$$z(\hat{\pi}_t, \hat{\pi}_t - \pi_t^e) + \frac{1}{(1 + r)}\hat{z} = z(0, -\pi_t^e) + \frac{1}{(1 + r)}z(\hat{\pi}_{t+1}, \hat{\pi}_{t+1} - \pi_{t+1}^e),$$

$$\tag{10}$$

---

[7] I am grateful to Michael Jones for this point.

[8] Given $\tilde{p}_t$, the term $V_{t+1}(\alpha_{t+1})$ is independent of p_t. Note that α_{t+1} in eq. (5) depends on $\tilde{p}_t$, rather than p_t, because p_t is unobservable even at the start of period $t + 1$.

[9] The result for period $T - 1$ uses the condition $p_T = 0$, which implies $V_T(\alpha_T) = E(z_T) = z(\hat{\pi}_T, \hat{\pi}_T - \pi_T^e)$.

for $\tau \le t \le T - 1$. The left side of eq. (10) is the present value of costs over period t and $t + 1$ if the policymaker reveals himself today by setting $\pi_t = \hat{\pi}_t = \phi(\pi_t^e)$. Then the cost next period is the discretionary one $\hat{z}$. The right side of eq. (10) is the present value of costs for the two periods if the policymaker masquerades today ($\pi_t = 0$) but reveals himself tomorrow [$\pi_{t+1} = \hat{\pi}_{t+1} = \phi(\pi_{t+1}^e)$]. [In either case the costs are $\hat{z}$ from date $t + 2$ onwards, so these terms do not appear in eq. (10).] In order to be willing to randomize over π_t – without having made a commitment to do so – the policymaker must be indifferent between 'cheating' today or cheating tomorrow. Hence, the equality of costs in eq. (10).

Rearranging terms in eq. (10) yields

$$z(0, -\pi_t^e) - z(\hat{\pi}_t, \hat{\pi}_t - \pi_t^e) = \frac{1}{(1 + r)}\left[\hat{z} - z(\hat{\pi}_{t+1}, \hat{\pi}_{t+1} - \pi_{t+1}^e)\right], \quad (11)$$

for $\tau \le t \le T - 1$. The left side of eq. (11) is the 'temptation to cheat' today [refer to Barro and Gordon (1983a, p. 107)] by setting $\pi_t = \hat{\pi}_t$ rather than $\pi_t = 0$. (The difference in costs in positive since the value $\hat{\pi}_t$ minimizes today's cost.) The right side is the 'enforcement power', which motivates the policy-maker to maintain a low-inflation reputation today by choosing $\pi_t = 0$. The gain from deferring high inflation is the difference between the discretionary cost $\hat{z}$ (which arises in period $t + 1$ if $\pi_t = \hat{\pi}_t$) and the lower cost from setting $\pi_{t+1} = \hat{\pi}_{t+1} = \phi(\pi_{t+1}^e)$ (which can arise only if $\pi_t = 0$).

4. Properties of the equilibrium

Eq. (11) prescribes a time path for expected inflation that must hold in order for the policymaker of type 2 to be willing to randomize at each date t. Thus far, I have explored in detail the implications of this condition only for the case of the simple cost function that I used previously [Barro and Gordon (1983a)]. Namely the cost function from eq. (1) is now specialized to

$$z_t = z(\pi_t, \pi_t - \pi_t^e) = \frac{a}{2}(\pi_t)^2 - b(\pi_t - \pi_t^e), \quad (12)$$

where $a, b > 0$ and common knowledge. As I discuss later, it appears that the main results are not very sensitive to modifications of this cost function.

Using the specification of costs from eq. (12), the inflation rate that minimizes costs for given expectations is

$$\hat{\pi}_t = \hat{\pi} = b/a, \quad (13)$$

which is independent of π_t^e. [This result follows because z_t in eq. (12) is linear

in π_t^e.] The terms that appear in eq. (11) are given by

$$z(0, -\pi_t^e) - z(\hat{\pi}_t, \hat{\pi}_t - \pi_t^e) = \tfrac{1}{2}b^2/a,$$

$$\hat{z} = \tfrac{1}{2}b^2/a,$$

$$z(\hat{\pi}_{t+1}, \hat{\pi}_{t+1} - \pi_{t+1}^e) = -\tfrac{1}{2}b^2/a + b\pi_{t+1}^e.$$

Then the condition for the path of expected inflation in eq. (11) becomes

$$\pi_t^e = \left(\frac{1-r}{2}\right)b/a, \tag{14}$$

for $\tau + 1 \le t \le T$, where recall that $b/a = \hat{\pi}$ in this model. For the cost function given in eq. (12), randomization requires that expected inflation be constant.

Eq. (4) implies that expected inflation must also satisfy

$$\pi_t^e = (b/z)(1 - \alpha_t)(1 - p_t). \tag{15}$$

Here I have substituted the actual probability p_t for the perceived probability $\tilde{p}_t$. That is, since the policymaker of type 2 is guaranteed to be indifferent over the choice of probabilities, I look at the behavior that is consistent with people's perceptions. The question of why an agent would be motivated to pursue just the right mixed strategy seems to arise whenever behavior is not committed and the equilibrium entails randomization. However, the resulting equilibrium will satisfy the condition that the policymaker not prefer any other course of action.

Eqs. (14) and (15) dictate a relationship over time between the probability α_t that the policymaker is of type 1 and the probability p_t that the type 2 person masquerades as a type 1, namely

$$(1 - \alpha_t)(1 - p_t) = \frac{(1-r)}{2}, \tag{16}$$

for $\tau + 1 \le t \le T$. It is apparent from eq. (16) that a rising path of reputation α_t [from the updating formula in eq. (5)] must be accompanied by a declining path of p_t. As the type 2 person builds reputation via good past performance (α_t rises as long as zero inflation is observed), it must be that the probability of continued good performance by a type 2 person, p_t, diminishes. Note also that the left side of eq. (16) is positive, which requires $r < 1$ on the right side. If $r \ge 1$, the discount rate is so high that the threat of lost reputation is insufficient ever to motivate the type 2 person to select zero inflation. Hence, $\pi_t = \hat{\pi} = b/a$ ($p_t = 0$) emerges from the start if $r \ge 1$. Some further discussion of this result appears below.

The determination of the equilibrium path of α_t and p_t follows by combining eq. (16) with Bayes' rule for updating α_t. The latter condition, stated in eq. (5), is

$$\alpha_{t+1} = \frac{\alpha_t}{\alpha_t + (1 - \alpha_t) p_t}, \tag{17}$$

for $\tau \le t \le T - 1$. With the addition of the boundary condition, $p_T = 0$, eqs. (16) and (17) determine the path of α_t and p_t for $t = \tau + 1, \ldots, T$,[10] up to the determination of the starting date for randomization τ.

The computations involve a first-order linear difference equation in α_t, which has the solution

$$\alpha_t = \left(\frac{1+r}{2}\right)^{T+1-t}, \tag{18}$$

$$p_t = \left[\left(\frac{1+r}{2}\right) - \left(\frac{1+r}{2}\right)^{T+1-t}\right] \bigg/ \left[1 - \left(\frac{1+r}{2}\right)^{T+1-t}\right], \tag{19}$$

for $\tau + 1 \le t \le T$. Note that α_t rises over time [since $(1 + r)/2 < 1$], while p_t falls. As $t \to T$, p_t approaches 0, while α_t tends toward the value $(1 + r)/2$.

The starting date for randomization, τ, depends mainly on the length of the horizon, T, and the initial probability α_0 of being type 1. Consider the value of t – not necessarily an integer – that would equate the result for α_t in eq. (18) to α_0. The result, denoted by t^*, is

$$t^* = T + 1 - \log(\alpha_0) \bigg/ \log\left(\frac{1+r}{2}\right). \tag{20}$$

The solution for τ is the largest integer contained in t^* – denoted int(t^*) – subject to the condition that $0 \le \text{int}(t^*) < T$, which corresponds to

$$\left(\frac{1+r}{2}\right)^T \le \alpha_0 < \left(\frac{1+r}{2}\right). \tag{21}$$

Hence,

$$\tau = \text{int}\left[T + 1 - \log(\alpha_0) \bigg/ \log\left(\frac{1+r}{2}\right)\right], \tag{22}$$

subject to the inequalities in (21).[11]

---

[10] This calculation does not yet use eq. (17) at $t = \tau$.

[11] This result corresponds to that in Backus and Driffill (1985, p. 536). Their other findings differ somewhat from mine because they view the private sector as randomizing its choice of expected inflation (see footnote 6).

Finally, the probability p_τ of masquerading in period τ is determined so that the updating formula in eq. (17) holds at date $t = \tau$, given that $\alpha_{\tau+1}$ satisfies eq. (18) and $\alpha_\tau = \alpha_0$. The resulting probability p_τ is somewhat above that indicated in eq. (19) (plugging in $t = \tau$), unless the value t^* in eq. (20) happens to be exactly an integer.

Assuming that $\tau > 0$, the policymaker of type 2 sets $\pi_t = 0$ with probability 1 at dates prior to τ (beginning with the starting date 0). In this situation the horizon is long enough so that the policymaker prefers $p_t = 1$ to randomization. Since $p_t = 1$ during this interval, expected inflation π_t^e is zero. Therefore, the costs are those associated with a full commitment to zero inflation, namely $z(0, 0) = 0$. Since $p_t = 1$, it also follows [from eq. (17)] that there is no updating of beliefs about the policymaker during this period. That is, $\alpha_t = \alpha_0$ for $0 \leq t \leq \tau$.

If the prior probability α_0 is very low and/or the horizon T is relatively short, then $\alpha_0 < ((1 + r)/2)^T$ would hold. In this case there is no interval during which $p_t = 1$ obtains. The starting point for randomization is then $\tau = 0$, but the formulas for α_t and p_t in eqs. (18) and (19) continue to apply for $0 < t \leq T$.

In this situation the initial probability of masquerading p_0 must be such as to build up the right amount of reputation for period 1 – namely $\alpha_1 = ((1 + r)/2)^T$ from eq. (18) – contingent on $\pi_0 = 0$ being realized. The smaller the value of α_0, the lower p_0 must be in order for the correct value of α_1 to be generated from the updating formula in eq. (17). The solution for this initial probability is

$$p_0 = \left(\frac{\alpha_0}{1 - \alpha_0} \right) \left[1 - \left(\frac{1 + r}{2} \right)^T \right] \Big/ \left(\frac{1 + r}{2} \right)^T. \tag{23}$$

In one sense the solutions for (α_t, p_t) in eqs. (18) and (19) apply no matter how small the initial probability α_0 that the policymaker is of type 1. However, as $\alpha_0 \to 0$, $p_0 \to 0$ in eq. (23). Therefore, it becomes increasingly likely that $\pi_0 = \hat{\pi} = b/a$, which means that the discretionary outcome would obtain in all succeeding periods. Further, as $\alpha_0 \to 0$ and $p_0 \to 0$, it also follows that $\pi_0^e = \hat{\pi}(1 - \alpha_0)(1 - p_0)$ tends to $\hat{\pi} = b/a$. Accordingly, for very small values of α_0, it becomes likely that a close approximation to the discretionary result will emerge from the outset.

A small value of α_0 does have to be weighed against the length of the horizon T, since the relevant condition is $\alpha_0 < ((1 + r)/2)^T$. For any finite α_0, an infinite horizon T rules out this inequality. Therefore, with an infinite horizon, the condition $\alpha_0 > 0$ implies that there will be a starting interval (of infinite length) during which $p_t = 1$ applies.

As the horizon T approaches infinity, the situation approaches that of a fully committed rule where $\pi_t = \pi_t^e = 0$ obtains for all t. In this case the present

value of costs approaches zero. The policymaker of type 2 retains the option of cheating today (or in any period) by setting $\pi_t = \hat{\pi}$, thereby attaining the lower one-period cost, $z(\hat{\pi}, \hat{\pi}) = -\frac{1}{2}b^2/a$. But the revelation of his identity implies that the subsequent outcomes are the discretionary ones, $\pi_t = \pi_t^e = \hat{\pi}$, with $z_t = \hat{z} = \frac{1}{2}b^2/a$. With an infinite horizon, the net effect of cheating in the current period on the present value of costs is therefore

$$-\tfrac{1}{2}b^2/a + \left(\tfrac{1}{2}b^2/a\right)\left[\frac{1}{1+r} + \frac{1}{(1+r)^2} + \cdots\right] = \left(\tfrac{1}{2}b^2/a\right)(1-r)/r.$$

The previously mentioned condition $r < 1$ ensures that cheating delivers a higher present value of costs than that from setting $\pi_t = 0$ for all finite t. In other words, if $T \to \infty$ and $r < 1$, the potential loss of reputation is sufficient to enforce the low-inflation outcome for a type 2 policymaker.[12]

The final possibility is a very high starting probability of being type 1 – specifically, $\alpha_0 > (1 + r)/2$ in expression (21). In this case (if $r < 1$), the policymaker sets $p_t = 1$ for the entire interval, $0 \le t \le T - 1$, before switching to $p_T = 0$ in the last period. Thus, there is no period of randomization – or of accumulation of reputation in the sense of a rising value of α_t – when the policymaker starts with a level of reputation α_0 that exceeds the critical level, $(1 + r)/2$.

Fig. 1 shows the equilibrium path for the probabilities α_t and p_t. These values apply for a policymaker of type 2, assuming that $((1 + r)/2)^T < \alpha_0 < (1 + r)/2$. Prior to date τ, the figure shows $p_t = 1$ and $\alpha_t = \alpha_0$. During this period $\pi_t^e = \pi_t = 0$. Then at date τ there is a discrete decline in p_t, but α_t has not yet changed. From date $\tau + 1$ onward the probability p_t of masquerading as a type 1 person follows the declining path shown in eq. (19), while the probability α_t of being type 1 follows the rising path given by eq. (18). Along this path expected inflation π_t^e is the constant $\frac{1}{2}(1 - r)b/a$. Finally, at date T, the value of p_T is zero, while α_t equals $(1 + r)/2$. [Note that $\pi_T^e = (1 - \alpha_T)b/a = \frac{1}{2}(1 - r)b/a$, which is the same value as in the interval between $\tau + 1$ and $T - 1$.][13]

---

[12] This result did not obtain in Barro and Gordon (1983a) because the loss of reputation lasted for only one period in that model. Since cheating reveals one's identity in the present case, the punishment interval is effectively infinite when the horizon is infinite.

[13] It is possible to rule out other equilibria in this model by showing that they are inconsistent with rational expectations, given the incentives of a type 2 policymaker. For example, if $\alpha_t > 0$ and $r < 1$, then $p_t > 0$ must hold for all $t < T$. Hence there is no interval – analogous to that before date τ – where the high-inflation outcome $\hat{\pi}$ occurs deterministically. If $\tilde{p}_t = 0$ were conjectured, then the choice $\pi_t = 0$ implies $\alpha_{t+1} = 1$ and $\pi_{t+1}^e = 0$. Then the choice $\pi_{t+1} = \hat{\pi}$ delivers a low cost for period $t + 1$ (with the discretionary outcomes arising thereafter). This behavior turns out to generate a lower present value of costs (if $r < 1$) than those from setting $\pi_t = \hat{\pi}$, which means going along at date t with the perception $\tilde{p}_t = 0$. The result demonstrates that $\tilde{p}_t = 0$ leads to $p_t = 1$. It follows that $\tilde{p}_t$ is irrational for all $t < T$ if $\alpha_t > 0$ and $r < 1$. Similarly, it is possible to rule out $\tilde{p}_t < 1$ for $t < \tau$ or $\tilde{p}_t = 1$ for $t \ge \tau$.
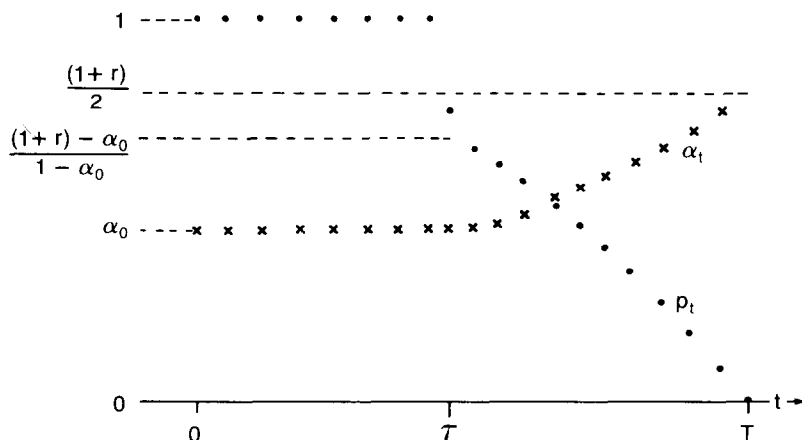
Fig. 1. Behavior of probabilities α_t and p_t.

Note: Before date τ, $p_t = 1$ and $\alpha_t = \alpha_0$. From date $\tau + 1$ until date T, α_t and p_t satisfy eqs. (18) and (19), respectively. For date τ, $\alpha_\tau = \alpha_0$ while p_τ is somewhat above the value implied by eq. (19).

The length of the interval, $T - \tau$, where randomization occurs follows from eq. (22) as

$$T - \tau = \text{int}\left[\log(\alpha_0)/\log\left(\frac{1+r}{2}\right)\right] - 1. \tag{24}$$

Since $0 < \alpha_0 < 1$ and $0 < r < 1$, this interval falls (or does not change) with a higher prior probability α_0 of being type 1. Accordingly, a higher value of α_0 implies a longer interval, $(0, \tau - 1)$, during which the zero-inflation outcome occurs with probability 1. It also follows from eq. (24) that a higher value of the discount rate r raises (or leaves unchanged) the randomization interval, $T - \tau$. In effect, a higher value of r lessens the value of reputation and thereby decreases the length of the period, $(0, \tau - 1)$, during which low inflation arises for sure. Table 1 shows the length of the interval, $T - \tau$ (in numbers of 'periods'), from eq. (24) for various values of the parameters α_0 and r. (The numbers apply subject to the condition that each be less than or equal to the value $T - 1$.)

Note that a change in α_0 alters the length of the randomization interval, $T - \tau$, but does not change the path of (α_t, p_t) during this interval. The main consequence of a higher value of α_0 is the increase in the length of the period, $(0, \tau - 1)$, during which actual and expected inflation equal zero.[14] Correspondingly, there is a decrease in the length of the period, $(\tau + 1, T)$, for which

---

[14] For a given τ, an increase in α_0 also raises the value of p_τ for the initial period of randomization. Hence π_τ^e declines when α_0 rises.

Table 1

Length of interval for randomization.[a]

r	α_0				
	0.05	0.10	0.25	0.50	0.75
0	3	2	1	0	0
0.05	3	2	1	0	0
0.10	4	2	1	0	0
0.25	5	3	1	0	0
0.50	9	7	3	1	0
0.75	21	16	9	4	1

[a] The table shows the value $T - \tau = \text{int}[\log(\alpha_0)/\log((1 + r)/2)] - 1$, where 'int' indicates the largest integer contained in the bracketed term. The numbers apply subject to the condition $\alpha_0 > ((1 + r)/2)^T$ from expression (21), which requires that each number in the table be less than or equal to $T - 1$.

expected inflation is the higher amount $\pi^e = \frac{1}{2}(1 - r)b/a$. These effects mean that the cost realized by either a type 1 or type 2 policymaker are lower the higher the value of α_0. Therefore, the greater the fraction of committed persons in the population of potential policymakers the better the outcomes for either type of policymaker.

This finding does not mean that the results improve if the policymaker turns out, ex post, to be of type 1. Before date τ the outcomes are identical (at $\pi_t = \pi_t^e = 0$) for either type. From date τ onward the results continue to be the same for each as long as the policymaker of type 2 masquerades by choosing zero inflation. But at some point the type 2 person 'cheats' by setting $\pi_t = \hat{\pi} = b/a$, which exceeds the expectation $\pi_t^e = \frac{1}{2}(1 - r)b/a$. Thereafter the outcomes for this policymaker are the discretionary ones, where $\pi_t = \pi_t^e = \hat{\pi}$. To see the implications for overall costs, recall that a policymaker of type 2 employs a mixed strategy only because the present value of costs is invariant to the time of cheating, except that cheating surely occurs by date T. Therefore, pretend that cheating ($\pi_t = \hat{\pi}$ for the first time) occurs exactly at date T. In this case the outcomes for type 1 and type 2 coincide through date $T - 1$. Then at date T, where inflationary expectations are $\pi_T^e = \frac{1}{2}(1 - r)b/a$, the type 1 person sets $\pi_T = 0$, while type 2 sets $\pi_T = \hat{\pi} = b/a$. Thereby the type 2 person secures lower costs at date T. (The value $\hat{\pi}$ minimizes costs for given expectations, which is the case here.) It follows that the overall present value of costs is lower if the policymaker turns out to be of type 2. Because of the benefit from the one period of surprisingly high inflation, the outcomes are better if the randomly selected policymaker happens to be the kind that is incapable of commitments.

Ex ante, the expected costs weigh the type 1 realization by the probability α_0 and the type 2 by $1 - \alpha_0$. It turns out that these expected costs are lower the

higher the value of α_0.[15] Outcomes are better on average if the policymakers come from a pool that contains a higher fraction of those who are capable of commitments. This result is consistent with the previous conclusion that costs are smaller if the realization for the choice of policymaker is type 2.

I have assumed that the policymaker of type 1 always sets inflation at the committed value 0. Zero inflation is optimal with the assumed cost function if commitments are not only made but are also fully believed. In the present context credibility is tempered by the possibility that the policymaker is type 2. In this case the best value to commit to need no longer be zero inflation. It would be possible to determine the sequence of committed values, $\pi_t = \pi_t^*$, which minimizes the overall expected present value of costs (weighing type 1 outcomes by α_0 and type 2 by $1 - \alpha_0$). Conceivably these values would then be the ones announced each period by both types of policymakers, where type 1 makes a serious commitment and type 2 only masquerades with some probability. However, I have not yet made much progress in figuring out the properties of the resulting path of π_t^*.

5. Alternative cost functions

A surprising aspect of the result is that expected inflation remains constant during the period after date τ. That is, $\pi_t^e = \pi^e = ((1 - r)/2)(b/a)$ as long as people observe zero inflation. The choice of zero inflation does enhance credibility in the sense that the probability α_t of being type 1 rises over time. But the falling probability p_t of good behavior by a type 2 policymaker offsets this effect and keeps expected inflation constant. Thus, the results are discouraging from the perspective of relying on reputation to lower inflationary expectations.

In order to see whether these results depend on the simplified cost function in eq. (12), I modified the function to include a quadratic term in unexpected inflation. The new specification is

$$z_t = \frac{a}{2}(\pi_t)^2 - b(\pi_t - \pi_t^e) + \frac{c}{2}(\pi_t - \pi_t^e)^2. \tag{25}$$

The last term, with $c > 0$, could represent the distortion from incorrect information about the general price level. Now unexpected inflation conveys benefits on net (by relieving other distortions) only if the surprise is not too large – specifically, if $\pi_t - \pi_t^e < 2b/c$.

The new cost function in eq. (25) does imply that the discretionary inflation rate, $\hat{\pi}_t$, rises with expected inflation – specifically,

$$\hat{\pi}_t = (b + c\pi_t^e)/(a + c). \tag{26}$$

---

[15] This can be shown by direct calculation of the present value of costs for the two types, pretending now that the type 2 person cheats in period τ (that is, sets $\pi_\tau = \hat{\pi}$).

The condition for randomization of policy in eq. (11) becomes a first-order difference equation in expected inflation. The equation is quadratic in π_{t+1}^e and π_t^e. There is one value for expected inflation, call it π^e, such that expected inflation would be constant over time.[16] That is, $\pi_t^e = \pi^e$ implies $\pi_{t+1}^e = \pi_t^e = \pi^e$. Taking a linear expansion of the quadratic terms around the stationary value π^e, it is possible to solve the difference equation explicitly. The result is that π_t^e follows damped oscillations around the rest point π^e.[17] Further, the departure of inflationary expectations from π^e for the initial period of randomization τ – and hence for all subsequent dates t – derives from the discrete length of periods and the associated integer restriction on τ. Otherwise, $\pi_t^e = \pi^e$ would apply at each date.

The main conclusion is that the modification of the cost function produces no tendency for π_t^e to fall during the period of randomization. The condition that π_t^e is constant no longer holds precisely, but is a reasonable approximation to the solution. My conjecture is that this result would hold even with more complicated specifications of costs. It seems that expected inflation would approximate the (single?) value that corresponds to the rest point, $\pi_{t+1}^e = \pi_t^e$, in eq. (11).

6. Predictions for inflation and other variables

Return now to the simplified model where costs are given by eq. (12). When viewed as a positive theory, the model has predictions for the behavior of inflation (and underlying monetary growth), which depend on whether the policymaker turns out to be of type 1 or type 2. With probability α_0 the policymaker is of type 1 and thereby makes a serious commitment to zero (or more generally low) inflation. There may be an interval of length τ over which expected inflation is also zero. [This applies if $r < 1$ and $\alpha_0 > ((1 + r)/2)^T$.] But [if $r < 1$ and $\alpha_0 < (1 + r)/2$] there is a subsequent interval of length $T - \tau$ where expectations are higher – namely $\pi^e = \frac{1}{2}(1 - r)b/a$ – but the policymaker continues to choose zero inflation. Correspondingly, the economy suffers from a string of surprisingly low inflation, which might show up as a recession or as an increase in distorting taxes. Thus, the committed policymaker continually 'bites the bullet' in the sense of tolerating the losses from surprisingly low inflation in order to maintain the low-inflation reputation. Nevertheless, this process does not succeed in reducing inflationary expecta-

---

[16] The second root of the quadratic equation turns out to be inadmissible because it is inconsistent with $\alpha_t \geq 0$ and $p_t \geq 0$. The admissible root is

$$\pi^e = \frac{b}{c}\left\{1 + \frac{c}{a}[a(1-r)+c]/[a+c(1+r)]\right\}^{1/2} - \frac{b}{c}.$$

[17] The solution is $\pi_t^e - \pi^e = (\pi_\tau^e - \pi^e)(-\gamma)^{t-\tau}$, where $\gamma = (c/a)(1+r)$ and τ is the first period for randomization. The oscillations of π_t^e around π^e are damped as long as $(c/a)(1+r) < 1$.

tions. Credibility does rise over time in the form of a growing belief α_t that the policymaker is of type 1. But the offsetting reduction over time in the probability p_t that a type 2 person would select low inflation keeps expected inflation constant.

If the policymaker turns out to be of type 2 (which occurs with probability $1 - \alpha_0$), then [if $r < 1$ and $\alpha_0 > ((1 + r)/2)^T$] there is again a period of length τ where $\pi_t = \pi_t^e = 0$ obtains. Subsequently, as long as $\pi_t = 0$ is chosen, the path mimics that of the type 1 person. But, with the rising probability $1 - p_t$, the policymaker opts for high inflation, $\hat{\pi}$, which generates some short-run benefits from a positive inflation shock (which reflects an underlying shock to money). From then on the outcomes are the discretionary ones where actual and expected inflation are both high ($\pi_t = \pi_t^e = \hat{\pi}$).

One characteristic of the equilibrium is a string of inflation rates that are below expectations. In the period after date τ, where expectations are $\pi^e = \frac{1}{2}(1 - r)b/a$, the type 1 person surely picks $\pi_t = 0 < \pi^e$, while the type 2 person has this realization with conditional probability p_t. The observation of a sequence of forecast errors of one sign may make it appear that expectations are irrational. In fact, these outcomes must be weighed against the probability $- (1 - \alpha_t)(1 - p_t) -$ that the policymaker is of type 2 and will engineer a large positive surprise for inflation. Taking this element into account, the expectations are rational in the equilibrium solution.

According to the model, a long history of data – with occasional changes in the identity of the policymaker (or of the policy 'regime') – would display a large number of relatively small negative inflation shocks that are offset by a small number of large positive surprises. On average, unexpected inflation is zero, but there are substantial runs of negative realizations. The result resembles the 'peso problem' for a 'fixed' exchange rate. In these cases the occasional discrete devaluations of a currency offset the strings of errors when the expected devaluation exceeds the realized devaluation of zero (which applies when the exchange rate remains fixed).[18] For a discussion of the devaluation model under rational expectations, see Blanco and Garber (1985).

7. Concluding observations

The introduction of uncertainty about the policymaker's type allows for meaningful notions of reputation and learning. Thereby the approach avoids some difficulties with multiple equilibria. Also, it is no longer necessary that policymakers plan over an infinite horizon.

The results are interesting since they show how surprisingly low or high inflation can emerge as part of the equilibrium. The extended interval where

---

[18] This case applies when the exchange rate either remains fixed or is devalued by a discrete amount. Appreciations of the currency are not considered in this model.

inflation is below expectations seems to correspond to notions of the policy-maker's 'bitting the bullet'. That is, the costs from surprisingly low inflation are accepted in order to enhance one's reputation for low inflation. (But the failure of expected inflation to fall here is a puzzling finding.) The uncommitted policymaker manages to create surprisingly high inflation for a period. This result corresponds to the idea that a surge in inflation can provide benefits in the 'short run'. But, as is also usually supposed, the long-term cost is that people raise their expectations of inflation. Thus – as in the case of the discretionary policymaker in previous models – the uncommitted policymaker ends up with an interval where inflation is high, but not surprisingly high.

One shortcoming of the approach is that it relies on differences in types of policymakers. In the present context these differences relate to capacities for making commitments. But divergences in preferences for 'inflation versus unemployment' would generate similar results. It would seem preferable to generate predictions for inflation that depended less on individual traits of policymakers and more on basic institutional features. But if all potential policymakers were the same, there would be nothing to learn from seeing their actions. Then the model reduces to ones studied previously, which had the shortcomings mentioned before.

References

Backus, D. and J. Driffill, 1984, Rational expectations and policy credibility following a change in regime, Unpublished paper, June (Queen's University, Kingston).

Backus, D. and J. Driffill, 1985, Inflation and reputation, American Economic Review, June.

Barro, R.J., 1983, Inflationary finance under discretion and rules, Canadian Journal of Economics, Jan.

Barro, R.J. and D.B. Gordon, 1983a, Rules, discretion and reputation in a model of monetary policy, Journal of Monetary Economics, July.

Barro, R.J. and D.B. Gordon, 1983b, A positive theory of monetary policy in a natural rate model, Journal of Political Economy, Aug.

Blanco, H. and P.M. Garber, 1985, Recurrent devaluation and speculative attacks on the Mexican peso, Unpublished paper, Jan. (University of Rochester, Rochester, NY).

Friedman, J.W., 1971, A non-cooperative equilibrium for supergames, Review of Economic Studies, Jan.

Green, E. and R. Porter, 1984, Noncooperative collusion under imperfect price information, Econometrica, Jan.

Horn, H. and T. Persson, 1985, Exchange rate policy, wage formation, and credibility, Unpublished paper, April (Institute for International Economic Studies, Stockholm).

Kreps, D.M. and R. Wilson, 1982, Reputation and imperfect information, Journal of Economic Theory, Aug.

Kydland, F.E. and E.C. Prescott, 1977, Rules rather than discretion: The inconsistency of optimal plans, Journal of Political Economy, June.

Milgrom, P. and J. Roberts, 1982, Predation, reputation, and entry deterrence, Journal of Economic Theory, Aug.

Tabellini, G., 1983, Accommodative monetary policy and central bank reputation, Unpublished paper, Oct. (University of California, Los Angeles, CA).