

RULES, DISCRETION AND REPUTATION IN A MODEL OF MONETARY POLICY*

Robert J. BARRO

*University of Chicago, Chicago, IL 60637, USA
NBER, Cambridge, MA 02138, USA*

David B. GORDON

University of Rochester, Rochester, NY 14627, USA

In a discretionary regime the monetary authority can print more money and create more inflation than people expect. But, although these inflation surprises can have some benefits, they cannot arise systematically in equilibrium when people understand the policymaker's incentives and form their expectations accordingly. Because the policymaker has the power to create inflation shocks *ex post*, the equilibrium growth rates of money and prices turn out to be higher than otherwise. Therefore, enforced commitments (rules) for monetary behavior can improve matters. Given the repeated interaction between the policymaker and the private agents, it is possible that reputational forces can substitute for formal rules. Here, we develop an example of a reputational equilibrium where the outcomes turn out to be weighted averages of those from discretion and those from the ideal rule. In particular, the rates of inflation and monetary growth look more like those under discretion when the discount rate is high.

1. Introduction

In a discretionary regime the monetary authority can print more money and create more inflation than people expect. The benefits from this surprise inflation may include expansions of economic activity and reductions in the real value of the government's nominal liabilities. However, because people understand the policymaker's incentives, these types of surprises — and their resulting benefits — cannot arise systematically in equilibrium. People adjust their inflationary expectations in order to eliminate a consistent pattern of surprises. In this case the potential for creating inflation shocks, *ex post*, means that, in equilibrium, the average rates of inflation and monetary growth — and the corresponding costs of inflation — will be higher than otherwise. Enforced commitments on monetary behavior, as embodied in monetary or price rules, eliminate the potential for *ex post* surprises. Therefore, the equilibrium rates of inflation and monetary growth can be

*We have benefited from discussion at the conference and from seminars at Chicago, Northwestern and Iowa. We are particularly grateful for comments from Gary Fethke, Roger Myerson, José Scheinkman, and John Taylor. Part of this research is supported by the National Science Foundation.

lowered by shifts from monetary institutions that allow discretion to ones that enforce rules.

When monetary rules are in place, the policymaker has the temptation each period to 'cheat' in order to secure the benefits from inflation shocks. (Because of existing distortions in the economy, these benefits can accrue generally to private agents, rather than merely to the policymaker.) However, this tendency to cheat threatens the viability of the rules equilibrium and tends to move the economy toward the inferior equilibrium under discretion. Because of the repeated interactions between the policymaker and the private agents, it is possible that reputational forces can support the rule. That is, the potential loss of reputation — or credibility — motivates the policymaker to abide by the rule. Then, the policymaker foregoes the short-term benefits from inflation shocks in order to secure the gain from low average inflation over the long term.

We extend the positive theory of monetary policy from our previous paper [Barro and Gordon (1983)] to allow for reputational forces. Some monetary rules, but generally not the ideal one, can be enforced by the policymaker's potential loss of reputation. We find that the resulting equilibrium looks like a weighted average of that under discretion and that under the ideal rule. Specifically, the outcomes are superior to those under discretion — where no commitments are pertinent — but inferior to those under the ideal rule (which cannot be enforced in our model by the potential loss of reputation). The results look more like discretion when the policymaker's discount rate is high, but more like the ideal rule when the discount rate is low. Otherwise, we generate predictions about the behavior of monetary growth and inflation that resemble those from our previous analysis of discretionary policy. Namely, any change that raises the benefits of inflation shocks — such as a supply shock or a war — leads to a higher growth rate of money and prices.

2. The policymaker's objective

As in our earlier analysis, we think of the monetary authority's objective as reflecting the preferences of the 'representative' private agent. Ultimately, we express this objective as a function of actual and expected rates of inflation. Specifically, benefits derive from positive inflation shocks (at least over some range), but costs attach to higher rates of inflation.

2.1. *The benefits from surprise inflation*

We assume that some benefits arise when the inflation rate for period t , π_t , exceeds the anticipated amount, π_t^e . One source of benefits — discussed in Barro and Gordon (1983) and in an example from Kydland and Prescott (1977, p. 477) — derives from the expectational Phillips curve. Here,

unanticipated monetary expansions, reflected in positive values for $\pi_t - \pi_t^e$, lead to increases in real economic activity. Equivalently, these nominal shocks lower the unemployment rate below the natural rate. By the natural rate, we mean here the value that would be ground out by the private sector in the absence of monetary disturbances. This natural rate can shift over time because of supply shocks, demographic changes, shifts in governmental tax and transfer programs, and so on. The natural rate also need not be optimal. In fact, the benefits from surprise inflation arise when the policymaker views the natural rate as excessive. This can occur, for example, if the distortions from income taxation, unemployment compensation, and the like make the average level of privately-chosen work and production too low. Because of the externalities from these distortions, the government (and the private agents) would value stimulative policy actions that lower the unemployment rate below its natural value.

Other sources of benefits from surprise inflation involve governmental revenues. Barro (1983) focuses on the proceeds from inflationary finance. The expectation of inflation (formed the previous period), π_t^e , determines people's holdings of real cash, M_{t-1}/P_{t-1} . Surprise inflation, $\pi_t - \pi_t^e$, depreciates the real value of these holdings, which allows the government to issue more new money in real terms, $(M_t - M_{t-1})/P_t$, as a replacement. The policymaker values this inflationary finance if alternative methods of raising revenue — such as an income tax — entail distortions. Hence, the benefit from surprise inflation depends again on some existing externality. Calvo (1978) discusses the necessity of existing distortions in this type of model.

The revenue incentive for surprise inflation relates to governmental liabilities that are fixed in nominal terms, rather than to money, *per se*. Thus, the same argument applies to nominally-denominated, interest-bearing public debt. Suppose that people held last period the real amount of government bonds, B_{t-1}/P_{t-1} . These bonds carry the nominal yield, R_{t-1} , which is satisfactory given people's inflationary expectations over the pertinent horizon, π_t^e . Surprise inflation, $\pi_t - \pi_t^e$, depreciates part of the real value of these bonds, which lowers the government's future real expenditures for interest and repayment of principal. In effect, surprise inflation is again a source of revenue to the government. Quantitatively, this channel from public debt is likely to be more significant than the usually discussed mechanism, which involves revenue from printing high-powered money. For example, the outstanding public debt for the U.S. in 1981 is around \$1 trillion.¹ Therefore, a surprise inflation of 1%, lowers the real value of this debt by about \$10 billion. Hence, this channel produces an effective lump amount of revenue of about \$10 billion for each extra 1% of surprise inflation. By contrast, the entire annual flow of revenue through the Federal

¹For this purpose we should actually look at the privately-held component of the funded national debt, which is about \$700 billion in 1981.

Reserve from the creation of high-powered money is about the same magnitude (\$8 billion in 1981, \$13 billion in 1980).

The attractions of generating revenue from surprise inflation are clear if we view the depreciation of real cash or real bonds as an unexpected capital levy. As with a tax on existing capital, surprise inflation provides for a method of raising funds that is essentially non-distorting, *ex post*. Once people have built up the capital or held the real cash or real bonds, the government can extract revenue without disincentive effects. Of course, the distortions arise — for capital, money or bonds — when people anticipate, *ex ante*, the possibility of these capital levies, *ex post*. That's why these forms of raising revenue will not end up being so desirable in a full equilibrium where people form expectations rationally. But, for the moment, we are just listing the benefits that attach, *ex post*, to surprise inflation.

2.2. *The costs of inflation*

The second major element in our model is the cost of inflation. Costs are assumed to rise, and at an increasing rate, with the realized inflation rate, π_t . Although people generally regard inflation as very costly, economists have not presented very convincing arguments to explain these costs. Further, the present type of cost refers to the actual amount of inflation for the period, rather than to the variance of inflation, which could more easily be seen as costly. Direct costs of changing prices fit reasonably well into the model, although the quantitative role of these costs is doubtful. In any event the analysis has some interesting conclusions for the case where the actual amount of inflation for each period is not perceived as costly. Then, the model predicts a lot of inflation!

3. *The setup of our example*

We focus our discussion on the simplest possible example, which illustrates the main points about discretion, rules and reputation. Along the way, we indicate how the results generalize beyond this example.

The policymaker's objective involves a cost for each period, z_t , which is given by

$$z_t = (a/2)(\pi_t)^2 - b_t(\pi_t - \pi_t^e) \quad \text{where } a, b_t > 0. \quad (1)$$

The first term, $(a/2)(\pi_t)^2$, is the cost of inflation. Notice that our use of a quadratic form means that these costs rise at an increasing rate with the rate of inflation, π_t . The second term, $b_t(\pi_t - \pi_t^e)$, is the benefit from inflation

shocks. Here, we use a linear form for convenience.² Given that the benefit parameter, b_t , is positive, an increase in unexpected inflation, $\pi_t - \pi_t^e$, reduces costs. We can think of these benefits as reflecting reductions in unemployment or increases in governmental revenue.

We allow the benefit parameter, b_t , to move around over time. For example, a supply shock — which raises the natural rate of unemployment — may increase the value of reducing unemployment through aggressive monetary policy. Alternatively, a sharp rise in government spending increases the incentives to raise revenue via inflationary finance. In our example, b_t is distributed randomly with a fixed mean, $\bar{b}$, and variance, σ_b^2 .³ (Hence, we neglect serial correlation in the natural unemployment rate, government expenditures, etc.)

The policymaker's objective at date t entails minimization of the expected present value of costs,

$$Z_t = E[z_t + (1/(1+r_t)) \cdot z_{t+1} + (1/(1+r_t)(1+r_{t+1}))z_{t+2} + \dots], \quad (2)$$

where r_t is the discount rate that applies between periods t and $t+1$. We assume that r_t is generated from a stationary probability distribution. (Therefore, we again neglect any serial dependence.) Also, the discount rate is generated independently of the benefit parameter, b_t . For the first period ahead, the distribution of r_t implies a distribution for the discount factor, $q_t = 1/(1+r_t)$. We denote the mean and variance for q_t by $\bar{q}$ and σ_q^2 , respectively.

The policymaker controls a monetary instrument, which enables him to select the rate of inflation, π_t , in each period. The main points of our analysis do not change materially if we introduce random discrepancies between inflation and changes in the monetary instrument. For example, we could have shifts in velocity or control errors for the money supply. Also, the policymaker has no incentive to randomize choices of inflation in the model.

We begin with a symmetric case where no one knows the benefit parameter, b_t , or the discount factor for the next period, q_t , when they act for period t . Hence, the policymaker chooses the inflation rate, π_t , without observing either b_t or q_t . Similarly, people form their expectations, π_t^e , of the policymaker's choice without knowing these parameters. Later on we modify this informational structure.

²Our previous paper [Barro and Gordon (1983)] uses a term of the form, $[\phi_t - \lambda(\pi_t - \pi^*)]^2$, where $\phi_t > 0$ depends on the natural unemployment rate for the period. Then, the policymaker values inflation shocks — that is, $\pi_t - \pi_t^e$ — only over some range. But, the general nature of the results does not change if we substitute this more complicated form. Also, we could modify the cost of inflation to depend on $(\pi_t - \bar{\pi})^2$, where $\bar{\pi}$ is the optimal inflation tax on cash balances.

³In some models, such as Lucas (1973) and Barro (1976), the coefficient b_t depends on the forecast variance of inflation. Most of our results would not be affected if we allowed for this type of dependence. However, this element matters when we compare across regimes that have different forecast variances for inflation.

4. Discretionary policy

Our previous paper [Barro and Gordon (1983)] discusses discretionary policy in the present context as a non-cooperative game between the policymaker and the private agents. In particular, the policymaker treats the current inflationary expectation, π_t^e , and all future expectations, π_{t+i}^e for $i > 0$, as given when choosing the current inflation rate, π_t . Therefore, π_t is chosen to minimize the expected cost for the current period, Ez_t , while treating π_t^e and all future costs as fixed. Since future costs and expectations are independent of the policymaker's current actions, the discount factor does not enter into the results. The solution from minimizing Ez_t , where z_t is given in eq. (1), is

$$\hat{\pi}_t = \bar{b}/a \quad (\text{discretion}). \quad (3)$$

We use carets to denote the solution under discretion. (With other cost functions, π_t would depend also on π_t^e .)

Given rational expectations, people predict inflation by solving out the policymaker's optimization problem and forecasting the solution for $\hat{\pi}_t$ as well as possible. In the present case they can calculate exactly the choice of inflation from eq. (3) — hence, the expectations are

$$\pi_t^e = \hat{\pi}_t = \bar{b}/a. \quad (4)$$

Since inflation shocks are zero in equilibrium — that is, $\hat{\pi}_t - \pi_t^e = 0$ — the cost from eq. (1) ends up depending only on $\hat{\pi}_t$. In particular, the cost is

$$\hat{z}_t = (1/2)(\bar{b})^2/a \quad (\text{discretion}). \quad (5)$$

5. Policy under a rule

Suppose now that the policymaker can commit himself in advance to a rule for determining inflation. This rule can relate π_t to variables that the policymaker knows at date t . In the present case no one knows the parameters, b_t and q_t , at date t . But, everyone knows all previous values of these parameters. Therefore, the policymaker can condition the inflation rate, π_t , only on variables that are known also to the private agents. (The policymaker could randomize his choices, but he turns out not to have this incentive.) Therefore, the policymaker effectively chooses π_t and π_t^e together, subject to the condition that $\pi_t^e = \pi_t$. Then, the term that involves the inflation shock, $\pi_t - \pi_t^e$, drops out of the cost function in eq. (1). Given the way that we modeled the costs of inflation — namely, as $(a/2)(\pi_t)^2$ — it follows immediately that the best rule prescribes zero inflation at all dates,

$$\pi_t^* = 0 \quad (\text{rule}). \quad (6)$$

We use an asterisk to denote the results from a rule. Eq. (6) amounts to a constant-growth-rate-rule, where the rate of growth happens to be zero.

Finally, we can calculate the costs under a rule from eq. (1) as

$$z_t^* = 0 \quad (\text{rule}). \quad (7)$$

The general point is that the costs under the rule, z_t^* , are lower than those under discretion, $\hat{z}_t$, from eq. (5). The lower cost reflects the value of being able to make commitments — that is, contractual agreements between the policymaker and the private agents. Without these commitments, inflation ends up being excessive — specifically, $\hat{\pi}_t > 0$ — but, no benefits from higher inflation result.

6. Cheating and temptation

As noted by others [e.g., Taylor (1975), B. Friedman (1979)], the policymaker is tempted to renege on commitments. In particular, if people expect zero inflation — as occurs under the rule — then the policymaker would like to implement a positive inflation rate in order to secure some benefits from an inflation shock. Further, this desire does not stem from a peculiarity in the policymaker's tastes. Rather, it reflects the distortions that make inflation shocks desirable in the first place.

How much can the policymaker gain in period t by cheating? Assume that people have the inflationary expectation, $\pi_0^e = 0$, which they formed at the start of period t . If the policymaker treats this expectation as a given, the choice of π_t that minimizes z_t is the one that we found before under discretion⁴ — namely,

$$\tilde{\pi}_t = \bar{b}/a \quad (\text{cheating}). \quad (8)$$

We use tildes to denote values associated with cheating. The expected cost follows from eq. (1) as

$$E\tilde{z}_t = -(1/2)\bar{b}^2/a \quad (\text{cheating}). \quad (9)$$

The general point is that this expected cost is lower than that, $z_t^* = 0$, from following the rule. We refer to the difference between these expected costs as the temptation to renege on the rule — or simply as the temptation. In the present case we have

$$\text{temptation} = E(z_t^* - \tilde{z}_t) = (1/2)(\bar{b})^2/a > 0. \quad (10)$$

⁴With a different cost function, the result for $\tilde{\pi}_t$ generally differs from that under discretion, $\hat{\pi}_t$.

At the present stage we have three types of outcomes. Ranging from low costs to high, these are

- (1) cheating (with people expecting the rule), $E\tilde{z}_t = -(1/2)(\bar{b})^2/a$,
- (2) rule, $z_t^* = 0$,
- (3) discretion, $\hat{z}_t = (1/2)(\bar{b})^2/a$.

Discretion is worse than the rule because first, no inflation shocks arise in either case, but second, the commitment under the rule avoids excessive inflation. However, the rule is only a second-best solution. Cheating — *when people anticipate the rule* — delivers better results. That's because the inflation shock eliminates part of the existing distortion in the economy (which is worth the extra inflation). But, the cheating outcome is feasible only when people can be systematically deceived into maintaining low inflationary expectations. In our subsequent analysis this cannot happen in equilibrium. However, the incentive to cheat determines which rules are sustainable without legal or institutional mechanisms to enforce them. There is a tendency for the pursuit of the first best — that is, the cheating outcome — to generate results that are poorer than the second best (rules) and closer to the third best (discretion).

7. Enforcement of rules

Generally, a credible rule comes with some enforcement power that at least balances the temptation to cheat. We consider here only the enforcement that arises from the potential loss of reputation or credibility. This mechanism can apply here because of the repeated interaction between the policymaker and the private agents.⁵ Specifically, if the policymaker engineers today a higher rate of inflation than people expect, then everyone raises their expectations of future inflation in some manner. Hence, in a general way, the cost of cheating today involves the increase in inflationary expectations for the future.

Consider a rule that specifies the inflation rate, π_t^* , for period t . The rule might prescribe $\pi_t^* = 0$, as before, or it might dictate some non-zero rate of inflation. Generally, the rule can specify some dependence of π_t^* on the realizations of all variables through date $t-1$ — that is, the values for date t are still not observed when π_t^* is set.

We postulate the following form of expectations mechanism, which we eventually show to be rational:

- (1) $\pi_t^e = \pi_t^*$ if $\pi_{t-1} = \pi_{t-1}^e$, and
 - (2) $\pi_t^e = \hat{\pi}_t$ if $\pi_{t-1} \neq \pi_{t-1}^e$.
- (11)

⁵This type of repeated game is discussed in J. Friedman (1971).

In other words if the previous inflation rate, π_{t-1} , accords with expectations, π_{t-1}^e , then people trust the government to perform in line with its announced rule for period t — that is, $\pi_t^e = \pi_t^*$. But, if the actual value departs from expectations last period, $\pi_{t-1} \neq \pi_{t-1}^e$, then people do not expect the government to follow its rule this period — hence, $\pi_t^e \neq \pi_t^*$. Rather, private agents anticipate that the policymaker will optimize subject to given expectations, which defines a discretionary situation. Hence, expectations are $\pi_t^e = \hat{\pi}_t$, where $\hat{\pi}_t$ is again the discretionary outcome.

If the government follows its rule in every period, then it also validates expectations each period. Then, the first part of eq. (11) says that the government maintains its reputation (or credibility) in each period. On the other hand, if the government cheats during period t , then the second part of eq. (11) says that the next period's expectations are the ones associated with discretion, $\hat{\pi}_{t+1}$. Then, if in period $t+1$ the government chooses the discretionary inflation rate, $\hat{\pi}_{t+1}$ (which is optimal given that expectations are $\hat{\pi}_{t+1}$), the actual and expected inflation rates coincide, although at the discretionary levels. Accordingly, the first part of eq. (11) says that people anticipate the rules outcome, π_{t+2}^* , for the following period. In other words the 'punishment' from violating the rule during period t is that the discretionary (non-cooperative) solution obtains during period $t+1$. But, credibility is restored as of period $t+2$ — that is, things carry on as of date $t+2$ as though no previous violation had occurred. Therefore, the mechanism in eq. (11) specifies only one period's worth of punishment for each 'crime'.⁶ Other equilibria exist that have punishment intervals of different length, as we discuss later on.

Consider our previous rule where $\pi_t^* = 0$. Suppose that the policymaker has credibility in period t , so that $\pi_t^e = 0$. If the policymaker cheats during period t , then his best choice of inflation is $\hat{\pi}_t = \bar{b}/a$ from eq. (8). [Note that eq. (11) says that the size and length of the punishment do not depend on the size of the crime.] Then, the policymaker gains the temptation, $E(z_t^* - \hat{z}_t) = (1/2)(\bar{b})^2/a$, from eq. (10).

The cost of this violation is that discretion, rather than the rule, applies for period $t+1$. Hence, the policymaker realizes next period the cost, $\hat{z}_{t+1} = (1/2)(\bar{b})^2/a$, from eq. (5), rather than that, $z_{t+1}^* = 0$, from eq. (7). Since costs for period $t+1$ are discounted by the factor $q_t = 1/(1+r_t)$ in eq. (2), the expected present value of the loss is

$$\text{enforcement} = E[q_t(\hat{z}_{t+1} - z_{t+1}^*)] = \bar{q} \cdot (1/2)(\bar{b})^2/a. \quad (12)$$

We use the term, *enforcement*, to refer to the expected present value of the loss from transgressions.

⁶Green and Porter (1981) use an analogous model for oligopoly pricing. There, the observation of a low price triggers $(T-1)$ periods of punishment, during which firms behave in a Cournot manner.

The policymaker abides by the rule during period t — that is, sets $\pi_t = \pi_t^*$ — if the enforcement is at least as great as the temptation. Otherwise, he opts for the cheating solution, $\pi_t = \tilde{\pi}_t$ (and suffers the consequences next period). But, when forming expectations for period t , π_t^e , people know whether the policymaker will find it worthwhile to cheat. Hence, if the cheating solution is preferable to the rule, then the expectation, $\pi_t^e = \pi_t^* = 0$, is irrational. Therefore, people would not stick with the expectation mechanism from eq. (11). The rules that can apply in equilibrium are those that have enough enforcement to motivate the policymaker to abide by them, given the expectations mechanism in eq. (11). Then, the equilibrium satisfies two properties. First, the expectations are rational. In particular, each individual's projection, π_t^e , is the best possible forecast of the policymaker's actual choice, π_t , given the way the policymaker behaves and given the way others form their expectations. Second, the policymaker's choice π_t , maximizes his objective, given the way people form their expectations.⁷

In equilibrium rules satisfy the enforceability restriction,

$$\text{temptation} = E(z_t^* - \tilde{z}_t) \leq \text{enforcement} = E[q_t(\hat{z}_{t+1} - z_{t+1}^*)]. \quad (13)$$

This condition says that the costs incurred today by following the rule, rather than cheating, are not greater than the expected value of having the cooperative (rules) outcome next period, rather than discretion. Consider now whether the proposed rule, $\pi_t^* = 0$, satisfies the enforceability restriction. From eq. (10), the temptation is $(1/2)(\bar{b})^2/a$, while from eq. (12), the enforcement is $\bar{q} \cdot (1/2)(\bar{b})^2/a$.⁸ Since $\bar{q} < 1$, the temptation is strictly greater than the enforcement. Hence, the ideal rule, $\pi_t^* = 0$, is not enforceable, at least given the expectations mechanism from eq. (11). Therefore, zero inflation is not an equilibrium in our model. [With a different form of cost function, rather than eq. (1), the ideal rule may or may not be enforceable.]

8. The best enforceable rule

We look here for the best enforceable rule — that is, the one that minimizes expected costs, subject to the constraint that the enforcement be at least as great as the temptation. In the present setting, where the parameters,

⁷The expectations mechanism from eq. (11) cannot be rational if the game has a known, finite endpoint. Then, no punishment arises for crimes in the last period. Working backward, the solution unravels period by period. Our framework assumes no known termination date for the game, although the game may end probabilistically. Then, a higher probability of termination shows up as a higher discount rate — that is, as a lower mean discount factor, $\bar{q}$. For some related game-theory literature, see Selten (1978), Kreps and Wilson (1980), and Milgrom and Roberts (1980).

⁸The two terms are equal when $\bar{q} = 1$ only because of the specific cost function from eq. (1). Generally, equality would arise for a value of $\bar{q}$ that is either above or below one.

b_t and q_t , are unobservable at date t , the best rule has the simple form,

$$\pi_t^* = \pi. \quad (14)$$

That is, the rule specifies constant inflation (a 'constant-growth-rate rule'). But, we already know that the ideal rule, $\pi=0$, is not enforceable. Given this, the enforceability restriction turns out to hold with equality for the best enforceable rule.

Using the procedures described before, we can calculate the temptation and enforcement associated with the rule, $\pi_t^* = \pi$. (Note that $\pi_t^e = \pi$ also applies here.) The results are

$$\text{temptation} = E(z_t^* - \hat{z}_t) = (a/2)(\bar{b}/a - \pi)^2, \quad \text{and} \quad (15)$$

$$\text{enforcement} = \bar{q} \cdot E(\hat{z}_{t+1} - z_{t+1}^*) = \bar{q} \cdot (a/2)[(\bar{b}/a)^2 - \pi^2]. \quad (16)$$

We graph the temptation and enforcement versus the inflation rate, π , in fig. 1. (This figure was suggested to us by John Taylor.) At $\pi=0$, the temptation follows from eq. (10) as $(\bar{b})^2/2a$. Then, as π rises, the gain from cheating diminishes — hence, the temptation falls. Finally, when π equals the discretionary value, $\bar{b}/a$, the temptation equals zero. That's because the cheating solution and the rule prescribe the same inflation rate, $\bar{b}/a$, at this point (As π increases further, the temptation increases, because — for given expectations — the policymaker prefers the discretionary inflation rate, $\bar{b}/a$, to higher rates of inflation.)

The enforcement equals $\bar{q}(\bar{b})^2/2a$ when $\pi=0$, from eq. (12). Then, as π rises, the enforcement declines. That's because the cost from losing reputation becomes smaller when the rule prescribes a higher rate of inflation. As with the temptation, the enforcement equals zero when π equals the discretionary value, $\bar{b}/a$. Here, when the policymaker cheats, people expect the same rate of inflation — namely, the discretionary amount $\bar{b}/a$ — as when the policymaker abides by the rule. Consequently, there is no enforcement. (When π increases further, the enforcement becomes negative — that is, the policymaker prefers the punishment, where people anticipate the inflation rate $\bar{b}/a$, to the maintenance of the rule, where people expect an inflation rate that exceeds $\bar{b}/a$.)

Notice that fig. 1 determines a range of announced inflation rates that can be equilibria. Specifically, the enforcement is at least as large as the temptation for values of π in the interval, $(\bar{b}/a)(1 - \bar{q})/(1 + \bar{q}) \leq \pi \leq \bar{b}/a$. Among these, we focus on the value of π that delivers the best results in the sense of minimizing the expected costs from eq. (2). We can rationalize this focus on one of the possible equilibria by allowing the policymaker to choose which value of π to announce in some initial period. Then, as long as this value is




Fig. 1. Temptation and enforcement.

in the enforceable range, we assume that the private agents follow along. That is, they all use the announced value of π as the basis for assessing the policymaker's performance [in accordance with eq. (11)].⁹ Within this setup, the policymaker will, in fact, announce the value of π that leads to a minimum of expected costs.

The best of the enforceable rules occurs where the curves for temptation and enforcement intersect in the interior of fig. 1. (The curves also intersect at the discretionary value, $\pi = \bar{b}/a$, but expected costs are higher at this point.)

⁹But, recall that the equilibrium is itself non-cooperative. In particular, each agent calculates the best forecast, π_t^e , of the policymaker's actions, while taking as given the way the policymaker behaves and the way other agents form their expectations.

Hence, the announced inflation rate is

$$\pi^* = (\bar{b}/a)(1 - \bar{q})/(1 + \bar{q}) \quad (\text{best enforceable rule}), \quad (17)$$

for which the expected cost in each period is

$$E\pi_t^* = [(1/2)(\bar{b})^2/a] [(1 - \bar{q})/(1 + \bar{q})]^2. \quad (18)$$

Notice that, with $0 < \bar{q} < 1$, the inflation rate, π^* , is intermediate between the ideal rule, 0, and discretion, $\bar{b}/a$. In fact, the best enforceable rule is a weighted average of the ideal rule and discretion, with the weights depending on the mean discount factor, $\bar{q}$. A relatively small value of $\bar{q}$, which means a high rate of discount on future costs, implies a relatively high weight on discretion — that is, a high value of π^* . That's because a decrease in $\bar{q}$ weakens the enforcement [eq. (16)], which requires π^* to increase in order to maintain the equality between the enforcement and the temptation.

Generally, an increase in the mean discount factor, $\bar{q}$, reduces π^* , with π^* tending toward zero (the ideal rule) as $\bar{q}$ tends to one.¹⁰ On the other hand, π^* tends to $\bar{b}/a$ — the discretionary result — as $\bar{q}$ tends to zero. (A zero discount factor means zero enforcement, so that only discretion is credible.) Notice that any force that influences the mean discount factor, $\bar{q}$, has a corresponding effect on inflation. For example, during a war we anticipate that $\bar{q}$ will be low, which triggers high inflation (via high monetary growth).

The expected cost from eq. (18) is also intermediate between that from the ideal rule, which is zero, and that for discretion, which is $(1/2)(\bar{b})^2/a$. Remember that the ideal rule is itself a second-best solution, which is inferior to cheating when people anticipate the ideal rule. But, cheating cannot occur systematically when people understand the policymaker's incentives and form their expectations accordingly. Rather, the lure of the better outcome from cheating creates the temptation, which makes the ideal rule non-enforceable. Hence, the attraction of the first best makes the second best unattainable. We end up with a cost that exceeds the second best (the ideal rule), but which is still lower than the third best (discretion).

The other feature of our results is the dependence of the inflation rate, π^* , on the ratio of cost parameters, $\bar{b}/a$. This ratio pertains to the benefit from inflation shocks, which depends on $\bar{b}$, relative to the costs of inflation, which depends on the parameter a . An increase in the ratio, $\bar{b}/a$, raises the temptation, relative to the enforcement, which requires π^* to increase. In particular, if inflation is not very costly, so that the parameter a is small, then we end up with a lot of inflation. Also, anything that raises the mean benefit attached to an inflation shock, $\bar{b}$, leads to higher inflation (but, not to

¹⁰This last condition depends on the specifics of our example. However, the direction of effect for $\bar{q}$ on π^* applies generally.

more benefits from inflation shocks). In our previous paper [Barro and Gordon (1983)], which focused on the results under discretion, we discussed some changes in the economy that can affect the benefits from inflation shocks. For example, the parameter $\bar{b}$ tends to be high in the following cases:

- when the natural unemployment rate is high,
- during a recession,
- during a war or other period where government expenditures rise sharply,
- when the deadweight losses from conventional taxes are high, and
- when the outstanding real stock of nominally-denominated public debt is large.

In each case we predict that the high value of $\bar{b}$ triggers a high value of π^*

— that is, a high rate of monetary expansion by the policymaker. This view accounts for¹¹

- a rise in the mean inflation rate along with a rise in the natural unemployment rate (as in the U.S. over the last 10–15 years),
- countercyclical response of monetary policy,
- high rates of monetary expansion during wartime,
- high rates of monetary growth in some less developed countries, and
- an inflationary effect from the outstanding real stock of public debt.

9. Contingent rules

We get some new results when we modify the informational structure in ways that motivate the policymaker to employ a contingent rule. Then, the inflation rate varies each period in accordance with the state of the economy.

Suppose that the policymaker knows the values of the benefit parameter, b_t , and the discount factor, q_t , when choosing the inflation rate, π_t . If people also condition their expectations, π_t^e , on b_t and q_t , then the results change little from those already presented. So, we focus on the case where π_t^e is still generated without knowledge of the contemporaneous variables, b_t and q_t .

One possibility is that the policymaker receives information more quickly than private agents. However, our setup does not require this informational asymmetry. For example, when setting demands for real money balances or holdings of real government bonds, people have to forecast rates of inflation. Once people hold the government's nominal liabilities, their real wealth changes when subsequent inflation shocks occur. Therefore, although the government and private agents may have the same information at any point

¹¹Some of these results can also be explained by changes in the optimal tax rate on cash balances, which applies to the systematic part of inflation. For example, this effect is probably important for monetary growth during wartime and in less developed countries.

in time, the agents' decisions (on how much real money and bonds to hold) depend on expectations of inflation that were formed earlier. Therefore, we can think of π_t^e as not being conditioned on the realizations, b_t and q_t . However, these realizations can influence the actual inflation rate, π_t .

The situation is less clear for the example of the Phillips curve. In models where only unperceived nominal disturbances matter for real variables — as in Lucas (1972, 1973) and Barro (1976) — the pertinent value for π_t^e is the one based on contemporaneously available information. However, some models with long-term nominal contracting [Gray (1976), Fischer (1977), Taylor (1980)] suggest that inflationary expectations formed at earlier dates will matter for today's choices of employment, production, etc. Then, the situation resembles that from above where people choose their holdings of money and bonds based on forecasts of inflation. However, the rationality of the Gray-Fischer-Taylor contracts has been questioned [Barro (1977)].

9.1. Discretion

We find the results under discretion in the same way as before. Specifically, we get

$$\hat{\pi}_t = b_t/a. \quad (19)$$

Now, the policymaker reacts to the actual value of the benefit parameter, b_t , rather than to its mean, $\bar{b}$. However, people's expectations — not conditioned on b_t — are $\pi_t^e = \bar{b}/a$. Therefore, although $\pi_t^e = E\hat{\pi}_t$, the realizations for b_t generate departures of inflation from its expectation. Therefore, the inflation shocks — and the corresponding benefits from them — are sometimes positive and sometimes negative.

The costs under discretion are now

$$\hat{z}_t = (1/2)(b_t)^2/a - (b_t/a)(b_t - \bar{b}). \quad (20)$$

The results correspond to those from before [eq. (5)] if $b_t = \bar{b}$. Looking one period ahead, we can calculate

$$E\hat{z}_{t+1} = (1/2a)[(\bar{b})^2 - \sigma_b^2]. \quad (21)$$

The new term is the variance of the benefit parameter, σ_b^2 .

9.2. The ideal contingent rule

When b_t is observed, the ideal rule no longer prescribes zero (or constant) inflation at all times. Rather, the policymaker conditions the inflation rate on

the realizations for the benefit parameter. The present example is simple enough to write out the ideal contingent rule in closed form. Specifically — abstracting from enforcement problems — the best rule turns out (after a large amount of algebra) to be

$$\pi_t^* = (1/a)(b_t - \bar{b}). \quad (22)$$

As before, the prior mean of inflation, $\pi_t^e = E\pi_t^*$, is zero. But, realized inflation exceeds its expectation — and benefit from inflation shocks arise — when b_t exceeds its mean, $\bar{b}$. Conversely, inflation is below its expectation — so that costs from unexpectedly low inflation occur — when b_t is below its mean.

Note that inflationary expectations are always zero, but the policymaker creates surprisingly low inflation (i.e., deflation) when the benefits parameter takes on relatively low values. These realizations may show up as a recession or as other costs from a negative inflation shock. Yet, *ex post*, it would clearly be preferable to have zero, rather than negative, inflation. Then, we avoid the negative inflation shock and also have less costs due to inflation [which are $(a/2)(\pi_t)^2$]. So, the negative inflation shocks may appear pointless. Yet, the ideal rule says that the policymaker should ‘bite the bullet’ — that is, cause a recession through contractionary monetary policy — under some circumstances. That’s because the surprisingly low rule of inflation when the benefit parameter, b_t , is low is the counterpart of the surprisingly high rate of inflation when the benefit parameter is high. Choosing zero, rather than negative, inflation for the low states means that the prior expectation of inflation is higher than otherwise. Then, the policymaker achieves lower benefits in the states where b_t is relatively high. In fact, it is worthwhile to incur some costs in the low states — namely, bite the bullet through unexpectedly low inflation — in order to ‘buy’ the unexpectedly high inflation and the corresponding benefits in the high states. In effect the policymaker invests in credibility when it is relatively cheap to do so — namely, when b_t is low — in order to cash in on this investment when it is most important — that is, when b_t is high.

The costs associated with the ideal rule turn out to be

$$z_t^* = -(1/2a)[(b_t)^2 - (\bar{b})^2]. \quad (23)$$

Again, we get our previous results [eq. (7)] if $b_t = \bar{b}$. Looking ahead one period, the expectation of these costs is

$$Ez_{t+1}^* = -(1/2a)\sigma_b^2. \quad (24)$$

Because the policymaker can match the variations in b_t with appropriate

responses in π_t , the expected costs fall with an increase in the variance of the benefit parameter, σ_b^2 .¹²

As before, we can show that the ideal rule is not enforceable in our model.¹³ Therefore, we go on now to examine the best enforceable, contingent rule.

9.3. Enforceable contingent rules

We look at rules that express the inflation rate, π_t^* , as a stationary function of the state, which specifies the values of the two variables, b_t and q_t . Given that the ideal rule is unattainable, the best enforceable rule in our model turns out to equate the temptation to the enforcement for all realizations of b_t and q_t .¹⁴ The temptation cannot exceed the enforcement for any of these realizations in order for the rule to be credible. Further, if the enforcement exceeds the temptation in some state, then we can do better by changing the inflation rate for that state. That is, we bear more costs than necessary by having excessive enforcement.

The present example is sufficiently simple to work out the results in closed form. The solution for inflation turns out to be a linear function of b_t and of $\sqrt{q_t}$ — that is,¹⁵

$$\pi_t^* = c_1 + c_2 b_t + c_3 \sqrt{q_t}, \quad (25)$$

where the c 's are constants, which have to be determined. If we conjecture that the rule for inflation takes the form of eq. (25), then we can work out the temptation and enforcement as functions of the parameters, c_1 , c_2 , c_3 , and the realizations for b_t and q_t . Then, we determine the value of the c -coefficients in order to equate the temptation to the enforcement for all values of (b_t, q_t) . Since eq. (25) has the correct form, this operation turns out to be feasible. The results are

$$c_1 = 0, \quad c_2 = 1/a, \quad c_3 = -2(\bar{b}/a)\sqrt{\bar{q}}/(1 + \bar{q}), \quad (26)$$

where $\sqrt{\bar{q}}$ is the mean of $\sqrt{q_t}$. Hence, the best enforceable contingent rule

¹²However, we did not enter the variance of inflation directly into the cost function of eq. (1). If we had, this result could change.

¹³When considering the ideal rule, the temptation and enforcement turn out to be independent of the realization for b_t . Further, the temptation exceeds the enforcement for all discount factors, q_t , that are less than one.

¹⁴With other cost functions, the enforcement may exceed the temptation for some realizations. In particular, we then find that the inflation rate does not react to variations in q_t in some regions.

¹⁵The enforcement is linear in q_t . But, the temptation involves the square of the inflation rate. Therefore, if π is linear in $\sqrt{q_t}$, then the temptation also involves terms that are linear in q_t .

for inflation is ¹⁶

$$\pi_t^* = (b_t/a) - 2(\bar{b}/a)\sqrt{\bar{q}} \cdot \sqrt{q_t}/(1 + \bar{q}). \quad (27)$$

The enforceable rule can again be viewed as a weighted average of the ideal rule — eq. (22) — and discretion — eq. (19). In particular, the mean rate of inflation is positive, but lower than that associated with discretion, which is $\bar{b}/a$. The relative weights depend on the discount factor — both on the parameters of the probability distribution for q_t and on the realized value. Given the parameters of the distribution, a higher realization for q_t means a lower inflation rate, π_t^* .¹⁷

Note that the realization of the discount factor does not affect current benefits and costs from inflation, but does influence the amount of enforcement. Thus, the ideal rule does not depend on q_t in eq. (22). But, for low realizations of q_t , low inflation rates are not credible, because the temptation would exceed the enforcement. Therefore, the best enforceable rule does depend on q_t in eq. (27).

The inflation rate now moves around with fluctuations in the benefit parameter, b_t , or in the discount factor, q_t . In particular, relatively high realizations for b_t and relatively low ones for q_t lead to unexpectedly high inflation. Conversely, the policymaker 'bites the bullet' — that is, creates negative inflation shocks — when the benefit parameter is lower than normal or the discount factor is higher than normal. The reasoning here is similar to that from before. It is worthwhile to suffer negative inflation shocks in some cases — that is, for low values of b_t or high values of q_t — in order to sustain low prior expectations of inflation. Then, large gains are attained in the cases where the benefit parameter, b_t , is high or the discount factor, q_t , is low. These last cases are likely to be emergencies — such as wars or other times where economic activity or government revenues are valued especially highly. In effect, the policymaker bites the bullet during the non-emergencies in order to invest in credibility — an investment that yields returns during the emergencies.

10. The length of the punishment interval

So far, our results apply when the length of the punishment interval is fixed at one period. That is, the length of time for which the discretionary outcome obtains, conditional on cheating, equals the length of time over which the policymaker can enjoy the results of his cheating. (The last interval

¹⁶The solution reduces to the previous one in eq. (17) if there is no random variation in b_t and q_t . Then, $b_t = \bar{b}$, $q_t = \bar{q}$, and $\sqrt{\bar{q}} = \sqrt{\bar{q}}$.

¹⁷Given the variance for $\sqrt{q_t}$ and the realized value of $\sqrt{q_t}$, a higher value of $\sqrt{\bar{q}}$ also lowers π_t^* . This follows by using the formula, $\bar{q} = \text{var}(\sqrt{q_t}) + (\sqrt{\bar{q}})^2$.

essentially defines the length of the period.) Given the length of the punishment interval, we obtained a unique reputational equilibrium by allowing the policymaker to announce the best one. But, if we look at different punishment intervals (which can be either greater or smaller than one period), then we find an array of reputational equilibria. At this point, we have no satisfactory way to resolve this problem of multiple equilibria. However, we have some observations.

First, we know that the length of the punishment interval cannot be zero. That is, the policymaker cannot instantly restore a lost reputation. If he could, there would be no enforcement, which means that the only equilibrium is the discretionary one.

We can calculate the effect of longer punishment intervals on expected costs. In the present model the punishments — that is, discretionary outcomes — never occur as part of a reputational equilibrium. Hence, we always do at least as well if we increase enforcement, which corresponds here to raising the length of the punishment interval. In particular, it always looks desirable in this model to have an infinite interval, which amounts to a form of 'capital punishment'.

We can modify the model so that punishments take place occasionally.¹⁸ For example, suppose that inflation depends partly on the policymaker's actions and partly on uncontrollable events. Further, assume that people cannot fully sort out these two influences on inflation, even *ex post*. Then, people adopt a form of control rule where the policymaker loses reputation if the observed inflation rate exceeds some critical value. But, because of the uncontrollable element, this loss of reputation — and hence, the punishment — actually occur from time to time. Then, in contemplating a more severe form of punishment, we have to weigh the losses when punishments occur against the benefits from greater enforcement. Thus, it is likely that the optimal punishment interval would be finite. (However, from a positive standpoint, it does not necessarily follow that the equilibrium with this punishment interval will be selected.)

Finally, another possibility is to introduce uncertainty about the policymaker's preferences. Then, people try to learn about these preferences by observing behavior. Further, the policymaker knows that people learn from his actions, and acts accordingly. Kreps and Wilson (1980), and Milgrom and Roberts (1980), who uses this general type of model, show that unique equilibria sometimes obtain.¹⁹ But, we have not yet pursued this route in our context, because it relies on differences in tastes among potential policymakers. Unfortunately, we have nothing interesting to say about the sources of these differences. But possibly, this idea would become meaningful

¹⁸Green and Porter (1981) have this feature in their model of oligopoly pricing.

¹⁹Also, the solution does not necessarily degenerate to the discretionary equilibrium when the game has a known, finite endpoint. See footnote 7, above.

if we identified policymakers with shifting interest groups, each of which were affected differently by variations in inflation.

11. Concluding observations

Our analysis provides an example of a reputational equilibrium for monetary policy. The results amount to a combination of the outcomes from discretion with those from the ideal rule. Previously, we analyzed discretion and rules as distinct possible equilibria. Now, the relative weights attached to the discretionary and rules solutions depend on the policymaker's discount rate and some other factors. From a predictive standpoint for monetary growth and inflation, the results modify and extend those that we discussed previously.

In some environments the rules take a contingent form, where inflation depends on the realization of the benefit parameter or the discount factor. Here, the policymaker sometimes engineers surprisingly low inflation, which is costly at a point in time. Thus, the monetary authority 'bites the bullet' and pursues a contractionary policy, given some states of the world. By acting this way, the policymaker sustains a reputation that permits surprisingly high inflation in other states of the world.

We have difficulties with multiplicity of equilibria, which show up also in the related game-theory literature. Here, the problem arises in determining how long a loss of reputation persists. In an extended version of the model, we can figure out the optimal length for this interval of punishment. But, from a positive standpoint, it is unclear which equilibrium will prevail.

References

- Barro, R.J., 1976, Rational expectations and the role of monetary policy, *Journal of Monetary Economics* 2, Jan., 1-32.
- Barro, R.J., 1977, Long-term contracts, sticky prices, and monetary policy, *Journal of Monetary Economics* 3, July, 305-316.
- Barro, R.J., 1983, Inflationary finance under discretion and rules, *Canadian Journal of Economics*, Jan.
- Barro, R.J. and D.B. Gordon, 1983, A positive theory of monetary policy in a natural-rate model, *Journal of Political Economy* 91, Aug.
- Calvo, G., 1978, On the time consistency of optimal policy in a monetary economy, *Econometrica* 46, Nov., 1411-1428.
- Fischer, S., 1977, Long-term contracts, rational expectations and the optimal money supply rule, *Journal of Political Economy* 85, Feb., 191-205.
- Friedman, B., 1979, Optimal expectations and the extreme information assumptions of 'rational expectations' macromodels, *Journal of Monetary Economics* 5, Jan., 23-42.
- Friedman, J.W., 1971, A non-cooperative equilibrium for supergames, *Review of Economic Studies* 38, Jan., 861-874.
- Gray, J.A., 1976, Wage indexation: A macroeconomic approach, *Journal of Monetary Economics* 2, April, 221-236.
- Green, E.J. and R.H. Porter, 1981, Noncooperative collusion under imperfect price information, unpublished, California Institute of Technology, Pasadena, CA, Jan.

- Kreps, D.M. and R. Wilson, 1980, On the chain-store paradox and predation: Reputation for toughness, unpublished, Stanford University, Stanford, CA, July.
- Kydland, F.E. and E.C. Prescott, 1977, Rules rather than discretion: The inconsistency of optimal plans, *Journal of Political Economy* 85, June, 473-491.
- Lucas, R.E., 1972, Expectations and the neutrality of money, *Journal of Economic Theory* 4, April, 103-124.
- Lucas, R.E., 1973, Some international evidence on output-inflation tradeoffs, *American Economic Review* 63, June, 326-334.
- Milgrom, P. and J. Roberts, 1980, Predation, reputation and entry deterrence, unpublished, Northwestern University, Evanston, IL.
- Selten, R., 1978, The chain-store paradox, *Theory and Decision* 9, 127-159.
- Taylor, J., 1975, Monetary policy during a transition to rational expectations, *Journal of Political Economy* 83, Oct., 1009-1022.
- Taylor, J., 1980, Aggregate dynamics and staggered contracts, *Journal of Political Economy* 88, Jan., 1-23.