Carnegie-Rochester Conference Series on Public Policy 39 (1993) 53-84
North-Holland

Designing institutions for monetary stability*

Torsten Persson[†]

Institute for International Economic Studies, S-10691 Stockholm, Sweden

and

Guido Tabellini

IGIER, Abbazia Mirasole, 20090 Opera, Italy

and

University of Brescia

Abstract

Lack of credibility or political feasibility are typically cited as serious obstacles to achieving monetary stability. We ask what kind of institutional reforms may help resolve such incentive problems in monetary policy, by help of the new theory of information, contracts and regulatory design. Credibility problems can be resolved by a remarkably simple performance contract that imposes a linear penalty for inflation on the central bank. Results extend when contracts cannot be conditioned on the state of the economy, in which case central bank announcements become important. When only incomplete and simple contracts can be embodied in the central banking institution, a tradeoff between incentives and information emerges; however, optimum institution design may circumvent the tradeoff between credibility and flexibility that has been stressed in the literature.

*We are grateful to Charles Freedman, Bob King, Allan Meltzer, Charles Plosser, Ricardo Ravelli, Lars E.O. Svensson, participants in the conference, and in seminars at IIES and the Stockholm School of Economics for helpful comments, to the Swedish Social Science Research Council for financial support, and to Kerstin Blomquist for secretarial assistance.

[†] *Correspondence to:* Torsten Persson, Institute for International Economic Studies, Stockholm University, S-10691 Stockholm, Sweden

0167-2231/93/$06.00

Introduction

Insufficient price stability is one of the most common problems that plague economic policy. Even though the extent of the problem has varied across countries and time, the ingredients are qualitatively similar: monetary policy is too expansionary, both on average and at the wrong points in time. Often most observers agree on how policy ought to be reoriented but deem such reorientation difficult. Lack of credibility or of political feasibility are mentioned as major obstacles.

In this paper we accept as a starting point that such *incentive problems* may indeed be major obstacles to price stability. We then ask what institutions may help resolve these problems. In doing so, we try not to compromise other tasks in monetary policy. In particular, we want to allow for a potentially valuable stabilization of output or employment.

If the incentive problems manifest themselves in too much and too variable inflation, common sense suggests that the most direct resolution would be to embed stronger incentives for stable prices in the institutions governing monetary policy. This is a common theme in our paper. Recent policy reform, as well as historical experience, suggest two particularly interesting, but different, routes to price stability along those general lines. One route would be a *legislative* approach, namely to create, by law, a very independent central bank, to give it, by law, control of monetary policy instruments, and to mandate it, by law, to direct its policies towards achieving price stability. What motivates interest in this approach is, of course, the successful policies of the German central bank. The proposed statutes of the European Central Bank (1990) are largely modeled after the law governing the Bundesbank.

Another route would be a *targeting* approach, namely to let the political principals of the central bank—the executive or the legislature—impose an explicit inflation target of price for monetary policy on top of the general mandate for price stability, and make the central bank leadership explicitly accountable to its principals for its success in meeting this target. Recently, New Zealand (1989) as well as Canada (1991) have started following, more or less, such a route. The only historical example we are aware of is Sweden in the 1930s, when the central bank successfully targeted the price level throughout the decade.

The traditional macroeconomic approach—in its Keynesian, Monetarist or Neoclassical variety—is useless in addressing questions of institutional reform, because it treats monetary policy as exogenous. The recent literature on macroeconomic policy games is more useful for our purposes, because it generally treats monetary policy as endogenous and emphasizes short-run and political incentives, which are embedded in the policy process.[1] While

[1] The pioneering papers by Kydland and Prescott (1977), Calvo (1978), and Barro and

this literature has been quite successful in coming up with positive models of policy, we believe it has been less successful in coming to grips with normative models of institution design. In the paper we try to shed new light on these design questions by borrowing heavily from the modern literature on contracts, principal-agent relationships, and regulatory design. This branch of microeconomics offers some powerful tools, which are designed to analyze precisely how to resolve incentive problems. We know of only one attempt to apply these tools to monetary policy problems. This is the recent paper by Walsh (1992), to which we owe a great deal of inspiration.

Throughout the paper we ask the following normative question of institution design. What are the features of an optimal contract from the point of view of society's *ex ante* preferences? Thus, the framer of the monetary institutions is society at large. However, we also consider the case in which the principal is a politically motivated body, possibly with a private agenda different from social welfare. The normative question in this case amounts to how much power over the central bank contract should be given to this politically motivated body. Thus, our normative approach to institution design is not necessarily inconsistent with a positive analysis of evolving central bank legislation.

Section 2 of the paper sets the stage by introducing our notation and our analytical framework. We illustrate the credibility problem that forms the basis for much of the recent literature in the context of our model. The model is general enough to encompass many important contributions as special cases.

Section 3 shows how the credibility problem may be altogether resolved by a remarkably simple performance contract that imposes a linear penalty for inflation on the central bank. We show that the result is robust to imperfect monetary control and imperfect monitoring by the private sector. We also argue that this kind of contract has some resemblance to real-world institutions.

Section 4 shows how these results can be extended to the case when some of the information underlying the optimal performance contract is nonverifiable and therefore noncontractible. The analysis suggests an important role for policy announcements by the central bank.

Section 5 studies the choice between the legislative and targeting routes to price stability as described above. Here, only incomplete and simple performance contracts can be written, and the central bank's political principal may have a private agenda, perhaps tied to partisan or electoral incentives.

Gordon (1983) stress the credibility problem in monetary policy, while Alesina (1987) and Rogoff and Sibert (1988) stress political incentive problems. Textbook treatments and surveys of the literature can be found in Persson and Tabellini (1990) and Cukierman (1992).

A trade-off between incentives and information emerges.

Section 6 also looks at simple incomplete contracts, but now in a situation where the central bank controls the inflation outcome only indirectly and imperfectly and where its control may covary with the extent of the incentive problem.

Section 7 concludes by summarizing some of the general insights from the analysis and by suggesting open issues for future work.

A general framework

The model

Suppose society evaluates the results of monetary policy according to a welfare function of the following general form:

$$E_{\theta,\epsilon} W(p, p - p^e; \theta, \epsilon). \qquad (1)$$

In (1) p and p^e denote actual and expected inflation, respectively, $W(\cdot)$ is a concave and continuously differentiable function and E denotes the unconditional expectations operator. The random variables θ and ϵ denote shocks to the economy that change the welfare effects of a given inflation rate or unexpected inflation rate. Throughout the paper this welfare function, or a slightly different version of it, represents the normative criterion according to which alternative equilibrium policies are compared.

We start with the assumption that the policymaker controls p directly. Later on we also consider the possibility of control errors. As in the existing literature, the government has an incentive to create inflation surprises. Let $x = p - p^e$ and let subscripts denote partial derivatives. Then we assume:

A 1 *Inflation surprises are welfare-improving: $W_x \geq 0$ for all $x \leq 0$, with strict inequality for at least some θ, ϵ if $x = 0$.*

The expectations of inflation here are meant to reflect the behavior of wage-setters, as in the Phillips-curve model of Kydland and Prescott (1977) and Barro and Gordon (1983). However, the model can also readily be interpreted with reference to public-finance models like those of Calvo (1978) or Barro (1983), in which the incentive to surprise is associated with the government desire to collect seigniorage revenue without distorting money demand or savings decisions.

Higher realizations of the random variables θ and ϵ are associated with stronger incentives to engage in inflation surprises.

A 2 *Shocks change the desirability of inflation surprises: $W_{x\theta}, W_{x\epsilon} \geq 0$.*

Assumption A2 has no particular importance and really amounts to a definition of the two random variables. The distinction between θ and ϵ is captured by the following assumptions:

A 3 *The policymaker has an information advantage with respect to ϵ but not with respect to θ: The policymaker sets inflation having observed the realization of both θ and ϵ. Thus, inflation is a function $P(\theta, \epsilon)$. Wage-setters set nominal wages having observed the realization of θ, but not of ϵ. Thus, expected inflation is a function $P^e(\theta) = E_{\epsilon/\theta} P(\theta, \epsilon)$ where $E_{x/y}$ denotes the expectation over x conditional on y.*

Thus, ϵ-shocks give the policymaker an informational advantage over wage-setters and provide a role for stabilization policies. θ-shocks on the other hand are public information and ideally should only induce a policy change if they affect the marginal cost of inflation. An example of an ϵ-shock would be a sudden oil-price change while an example of a θ-shock would be a change in the nominal debt outstanding, or in the degree of wage indexation (that is, something more easily observable at the time of signing labor contracts).

This general framework captures several special cases studied in the literature. In a seminal paper, Rogoff (1985) studies a model which ultimately reduces to an objective function of the form:

$$W(p, p - p^e; \theta, \epsilon) \equiv -[p^2 + \lambda(p - p^e - \epsilon - \theta)^2]/2, \qquad (2)$$

where θ is a known parameter interpreted as the natural rate of employment, while ϵ is a random variable interpreted as a supply shock, the realization of which is known to the bank but not to wage-setters. Rogoff concludes that to overcome the credibility problem arising under discretion (see below), it may be optimal for society to delegate monetary policy to an agent (an independent and conservative central banker), with a value of λ lower than that of society. That delegation, however, entails costs, because a central banker with a low λ responds too little to ϵ-shocks.

Canzoneri (1985) writes down an objective function like (2), except that he interprets p as the rate of money growth (imperfectly controlled by the policymaker) and the random variable ϵ as a shock to velocity. The central bank is allowed to have a private forecast of the realization of ϵ, based on private information – or, equivalently, knows the realization of ϵ. The parameter θ is known to everybody and, as in Rogoff (1985), denotes the natural rate of employment. Canzoneri argues that the lack of monitoring makes credibility problems hard to handle and goes on to study a dynamic reputation model similar to that in Green and Porter (1984).

Cukierman and Meltzer (1986) write the objective function as:

$$W(p, p - p^e; \theta, \epsilon) \equiv (\theta + \epsilon)(p - p^e) - p^2/2, \qquad (3)$$

where again θ is a known parameter, while ϵ is a random variable whose realization is known to the policymaker but not to wage-setters. The price level p is also here imperfectly controlled by the policymaker. The expression $(\theta + \epsilon)$ represents the marginal benefit from unexpected inflation, where the random variable ϵ follows an autoregressive stochastic process. Hence, the expectation of inflation in period t depends on realized inflation in period $t-1$, with a coefficient that depends on how imperfect is the policymaker's control of inflation. Cukierman and Meltzer show that imperfect control of inflation may be more desirable for the policymaker than perfect control. The reason is that with imperfect control the link between current inflation and future expected inflation is looser, because the signal-extraction problem faced by wage-setters assigns lower informational content to the observed value of p. As a consequence, the policymaker can benefit to a larger extent from his private information about ϵ.

Commitment

We are now ready to characterize the *ex ante* optimal policy as the solution to the problem of maximizing (1) by choice of the function P, subject to the constraint that expectations are formed rationally. This problem corresponds to an environment in which the policymaker commits to a contingent policy before wages are set, and thus takes into account the effect of policy on expectations. This timing is illustrated in Figure 1.

The Lagrangean for this problem is:

$$E_{\theta,\epsilon}\{W(P(\theta,\epsilon), P(\theta,\epsilon) - P^e(\theta); \theta, \epsilon) + \mu(\theta)[P^e(\theta) - E_{\epsilon/\theta}P(\theta,\epsilon)]\} \qquad (4)$$

where $\mu(\theta)$ is the multiplier associated with the expectations constraint. The corresponding first-order conditions can be written as:

$$W_p + W_x = E_{\epsilon/\theta}W_x. \qquad (5)$$

The left-hand side of (5) is the sum of the marginal gain (or cost, if negative) of inflation (W_p) plus unexpected inflation (W_x) given p^e. The right-hand side is the (average) marginal cost of expected inflation. Thus, under commitment, the policymaker realizes that the benefit of a higher value of $P(\theta,\epsilon)$ brings with it the cost of a higher expected inflation on average. At the optimum the two must be equal. It is understood that the function W is evaluated at the point $p^e = P^e(\theta)$.

Figure 1

Timing and information

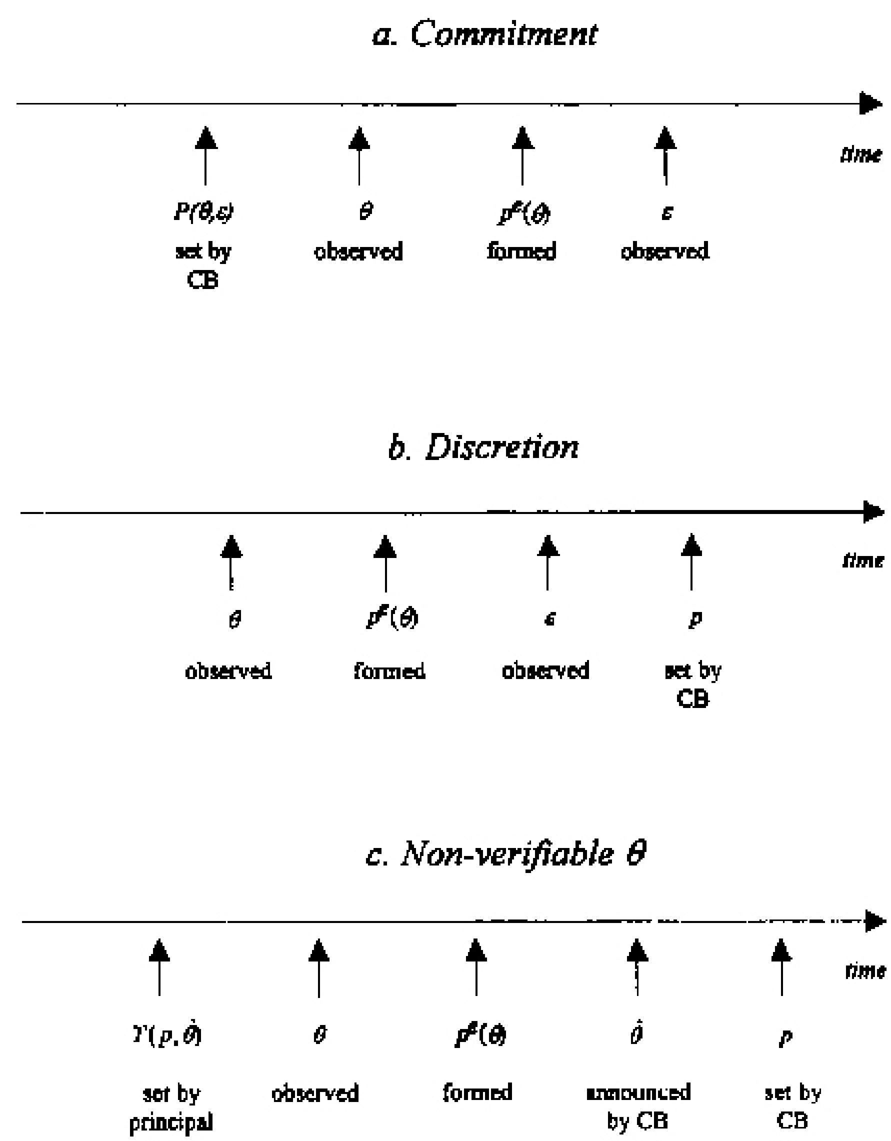

Equation (5) implicitly defines the optimal contingent policy: $p = P^*(\theta, \epsilon)$. This function has the property that on average it minimizes the marginal cost of inflation (by (5), $E_{\epsilon/\theta}W_p = 0$).[2,3]

Discretion

Consider a different timing, also illustrated in Figure 1, when policy is chosen under discretion, after expectations have been formed and wages set. The equilibrium policy must be *ex post* optimal. Hence, the policymaker maximizes (1) by choice of $P(\theta, \epsilon)$, but now takes p^e as given. The first-order conditions for an optimum implies:

$$W_p + W_x = 0, \tag{6}$$

where, as before, it is understood that the function W is evaluated at the equilibrium point, $p^e = P^e(\theta)$. Here, policy is chosen once expectations have been formed. This means that the effect of policy on expected inflation and hence on welfare—the right-hand side of (5)—is correctly neglected. As a result, the *ex post* optimal inflation rate equates the marginal benefit of unexpected inflation (W_x) to the marginal cost of inflation ($-W_p$).

Equation (6) implicitly defines the equilibrium inflation rate under discretion, $P^d(\theta, \epsilon)$. Contrasting (6) and (5), and noting that by A1 the right-hand side of (3) is strictly positive, we get immediately that $P^*(\theta, \epsilon) \leq P^d(\theta, \epsilon)$, for all θ and ϵ. This is the well-known result that the equilibrium inflation rate under discretion is too high compared to the *ex ante* optimal rate.[4]

Optimal central bank contracts

How can society design institutions that overcome the incentive problem discussed in the previous section and enforce the optimal policy? To answer this question we now study the design of optimal contracts between society in the form of a constitutional framer (the principal) and an independent central bank (the agent). This section draws quite a bit on an insightful recent paper by Walsh (1992) and, more generally, on the literature dealing with design of regulatory institutions—see Baron (1989) and Laffont and Tirole (1993).

Throughout the paper we assume that the principal commits to a contract before wages are set and is not allowed to renegotiate it later on. This is a reasonable assumption if we view the contract as a statute for the central bank. Clearly, it is possible to change the central bank law, but only according to a preset procedure which requires time. Any such change is therefore likely to be incorporated in expectations. Moreover, in this and the next section the principal is identified with society at large. Under these assumptions, the principal chooses the optimal contract to maximize the *ex ante* welfare function defined in (1), taking into account the effect of the contract on inflationary expectations.

We also assume that the agent operates under discretion, taking expectations as given. He maximizes the welfare function spelled out above plus any "transfer" conditional on the central bank performance, as specified by the central bank contract. Such a transfer can be interpreted literally, as a monetary compensation or penalty conditional on some publicly verifiable outcome. Alternatively, it can be interpreted more loosely as affecting the central bank prestige as an institution, or the reputation of the individuals with the highest responsibility in that institution. The recent reforms of the central bank law in New Zealand, for instance, specify that the governor can be dismissed if actual inflation falls outside a target range (0–2%). A governor with career concerns would internalize the loss of prestige resulting from such a dismissal.

Throughout the paper we do not attempt to formalize precisely these alternative interpretations, but simply assume that the transfer can be expressed in the same utility metric as social welfare. The central bank leadership thus perceives a clear tradeoff between social welfare and the "transfers" spelled out by the central-bank contract. Moreover, the central bank is assumed to be risk neutral, in that the transfer enters linearly in its utility function. Thus, the central bank's *ex post* objective is to maximize:

$$W(\cdot) + t, \tag{7}$$

for a given p^e, where $W(\cdot)$ is defined in (1) and where t is the transfer specified by the contract. This transfer is possibly dependent on inflation and on the random variables θ and ϵ.

Finally, we assume that the principal is completely unaffected by the transfer *per se* and that the central bank faces no binding participation constraint. Both seem reasonable assumptions. Even if the transfers have monetary implications for society (for instance, a greater or smaller budget for

[2] If in addition to A2 we assume $W_{p\theta} = W_{p\epsilon} = 0$ (i.e., the shocks *only* effect the incentive to surprise), we get a further prediction. Namely, $P^*_\theta = 0$, while $P^*_\epsilon \geq 0$. That is, the optimal policy responds only to ϵ-shocks, over which the policymaker has an informational advantage.

[3] In the example given by (2), the *ex ante* optimal policy reduces to: $P^*(\theta, \epsilon) = \frac{\lambda}{1+\lambda}\epsilon$. See also Persson and Tabellini (1990), chap. 2.

[4] If A2 is modified, as indicated in footnote 2, it is easy to show that $P_\theta, P_\epsilon > 0$. That is, under discretion the policymaker responds to *all kinds* of shocks, irrespective of whether or not he has an information advantage over private agents. In the example given by (2), we get: $P^d(\theta, \epsilon) = \lambda\theta + \frac{\lambda}{1+\lambda}\epsilon$, see Persson and Tabellini (1990). This difference between the equilibrium policy under commitment and discretion could form the basis of empirical tests aimed at assessing whether observed policies suffer from a credibility problem. If lack of credibility is not a binding constraint, policy should react only to variables which are *not* in the information sets of wage-setters.

the central bank), they are likely to be trivial compared to the macroeconomic consequences of monetary policy. Moreover, the ego-rents from being in charge of monetary policy are likely to be sufficiently large to compensate for any financial loss implied by the central bank contract.

Relaxing these assumptions would imply that the optimal monetary policy can no longer be implemented. Intuitively, society would face a tradeoff between enacting the optimal policy and giving up (costly) rents to the central bank. Hence, it would accept deviation from the optimal policy to reduce those rents. This kind of tradeoff is well-known in the literature of regulation (see Laffont and Tirole (1993), and Baron (1989)). Adding these considerations to the optimal contracts discussed below would be relatively straightforward.

Ex post verifiable shocks[5]

Let us start with the case in which the random variables θ and ϵ are *ex post* verifiable by the principal and can thus be written into the contract. Hence, the transfer specified by the general form: $t = T(P, \theta, \epsilon)$. It is clear in this case that the optimal monetary policy rule $P^*(\theta, \epsilon)$ can be enforced through a highly nonlinear contract of the form: any deviation from $P^*(\cdot)$ will be punished severely. There is, however, a simpler and much more revealing way of enforcing the optimal policy. Contrast the *ex ante* and *ex post* first-order conditions, (5) and (6). They differ by a function of θ

$$\mu(\theta) = E_{\epsilon/\theta} W_x(P^*(\theta,\epsilon), P^*(\theta,\epsilon) - P^e(\theta); \theta, \epsilon), \tag{8}$$

where $\mu(\theta)$ is the Lagrange multiplier on the expectations constraint in (4). The central bank will thus find it *ex post* optimal to choose its policy according to (5) if the contract entails an extra marginal cost from inflation exactly equal to $\mu(\theta)$. By (7) it is then easy to verify that the following contract implements the *ex ante* optimal policy:

$$T(p; \theta, \epsilon) = c - k(\theta)p, \tag{9}$$

where c is an arbitrary constant and $k(\theta) = \mu(\theta)$ as defined in (8). The proof is immediate and follows simply by taking the first-order conditions of (7), subject to (9) for a given p^e.

Notice that the optimal contract in (9) does not use up all our degrees of freedom. In particular, we could amend it by making the intercept state dependent: $c(\theta)$. If we set $c(\theta) = k(\theta)p^*$, and let p^* denote the average (across θ and ϵ) of $P^*(\theta, \epsilon)$, we can write the contract on the simple form: $T(p; \theta, \epsilon) = k(\theta)(\bar{p}^* - p)$.

Thus, the optimal contract can be interpreted as a mandate to achieve price stability. The central bank is punished by $\mu(\theta)$ for any percentage point of inflation. Essentially, by punishing *ex post* the central bank for realized inflation, this contract adds a cost that the central bank had "forgotten": the cost of higher expected inflation. According to the optimal contract, the marginal penalty for inflation, $k(\theta)$, must be equal to the *ex ante* marginal cost of expected inflation, $\mu(\theta)$.[6]

Two important remarks are in order. First, the contract is *independent of* ϵ. Thus, the principal does not need to verify the central bank's private information to enforce the optimal policy, even though this information plays an important role in determining what the policy action is (recall from footnote 3 that $P^*_\epsilon \geq 0$ in some standard examples). This result has a simple intuition. Since the inflation bias does not depend on ϵ, it can be corrected simply by adding the correct marginal cost of inflation to the central bank's *ex post* social-welfare function. Note, however, that the correct marginal cost does depend on the realization of θ. By our present assumption this variable is observed by wage-setters, and hence it is arguably easier to verify *ex post* than the realizations of ϵ. Farther on in the paper, we discuss how to design a central-bank contract when θ, though observed by wage-setters, is not verifiable and hence not contractible.

Second, the optimal contract is *linear* in inflation. This result, too, has a simple intuition. The key distortion here is that *ex post* the policymaker does not internalize the effects of his choice on expected inflation, since policy is set taking p^e as given. Of course, expected inflation is linear in realized inflation by construction. Making the performance contract linear in inflation makes the central bank internalize the effect on expected inflation, provided it faces the correct marginal cost. Also, as discussed above, the correct marginal cost is precisely the multiplier $\mu(\theta)$ on the expectation constraint in the *ex ante* problem.

The linearity makes the contract robust to minor variations in the assumptions. In particular, suppose that the policymaker does not control prices directly, but instead controls a policy instrument, q, which in turn affects inflation through a stochastic equation of the form:

$$p = \pi(q; \nu), \tag{10}$$

where ν is a random variable. It is easy to show that the contract in (9) still implements the optimal policy irrespective of whether or not the realization of ν is observed by the central bank when it sets policy — see Appendix 1.

[5] Many of the results in this section were derived by Walsh (1992) in the special case of the usual Barro-Gordon-Rogoff linear-quadratic example without any θ-shocks (in our notation).

[6] It is easy to show that, in the simple linear-quadratic example given in (2), the marginal penalty $k(\theta)$ that enters in (9) reduces to $\lambda\theta$ (see also footnotes 3 and 4).

Discussion

The credibility problem we have studied above is really the starting point for much of the recent literature on monetary policy games. We find it remarkable that the contractual solution to the problem is so simple and that researchers, including ourselves, working in the field have failed to see it. Some of the institutionally-oriented solutions that have been proposed in the literature can indeed be interpreted as adding a transfer function $T(\cdot)$ to the central bank's payoff. However, these solutions—again Walsh (1992) is the exception—all impose particular restrictions on the functional form for $T(\cdot)$. Thus, Rogoff (1985) and the subsequent literature on the "conservative central banker" restricts $T(\cdot)$ to be quadratic in inflation. Lohman (1992) follows the same approach, but adds a fixed cost for dismissing the central bankers and thus makes $T(\cdot)$ costly for the principal. The literature on "escape clauses" in monetary policy (Flood and Isard (1989) and Obstfeld (1991)) basically assumes that $T(\cdot)$ is a step function.

A natural question arises. Do we observe in the real world institutions that enforce an inflation contract like that discussed above? The answer is mixed. On the one hand, such a contract can be interpreted as a general mandate to achieve price stability. The penalty on the central bank for allowing high inflation would then take the form of a general loss of prestige of the institutions, or even as a higher probability that the governor and top officials be dismissed (cf. New Zealand) or that some new piece of legislation harmful to the bank be enacted by the legislature. Under this interpretation, the inflation contract does correspond to aspects of central bank institutions that we observe in the real world.[7]

On the other hand, the contract spelled out above is more specific than just a general mandate for price stability. The penalty suffered by the bank is proportional to inflation. In that sense, the contract is reminiscent to fixing the budget of the bank—or maybe the remuneration for its leadership—in *nominal* terms. However, the constant of proportionality in the penalty depends on the realization of the random variable θ. The optimal inflation contract thus holds the bank precisely accountable for its actions in a way that depends on the state of the world. Such contracts are rarely observed in the real world.[8]

In reality we do observe, however, that central banks often announce precise numerical values for nominal targets, such as intermediate monetary aggregates or exchange rates, and are held accountable for deviating from them. Can we interpret such nominal targets as performing a function similar to the inflation contract illustrated above, namely, of increasing the marginal cost of inflation for the bank? We think the answer is positive. Exchange-rate targets are often imposed on the bank from other political agents, like our inflation contract, and are bound by international agreements. Intermediate monetary targets are instead self-imposed by the bank, but at least in some countries the central-bank prestige would suffer for deviating from them. In both cases, it is clear that such targets provide an anchor for inflationary expectations.

It is then interesting to compare a contract based on an inflation target with one based on an intermediate target. Specifically, let $m = M(q, \eta)$ be the intermediate target where η is a shock (for instance, a velocity shock, if the target is a monetary aggregate) and q is the policy instrument. Suppose further that the central bank does not control prices directly but instead faces a stochastic link between inflation and its policy instrument in the form of equation (10) above. The central bank observes the realization of θ and ϵ, but not of ν and η. Let the contract imposed on the bank be:

$$T(m; \theta, \epsilon) = c - k^m(\theta, \epsilon) m. \tag{11}$$

Faced with such a contract the central bank *ex post* optimality condition is:

$$E_{\nu/\theta,\epsilon}\{(W_p + W_\pi)\pi_q\} = k^m(\theta, \epsilon) E_{\eta/\theta,\epsilon} M_q. \tag{12}$$

To implement the optimal monetary policy just set

$$k^m(\theta, \epsilon) = \frac{\mu(\theta) E_{\nu/\theta,\epsilon}\pi_q}{E_{\eta/\theta,\epsilon} M_q}, \tag{13}$$

where $\mu(\theta)$ is defined as in (A.3) of the Appendix. Thus, a contract based on an intermediate monetary target is equivalent to an inflation contract, in the sense that both implement the optimal monetary policy.

Comparing (7) and (13), however, it is clear that the inflation contract is more direct and simpler to enforce. Through the terms π_q and M_q the constant of proportionality k^m now depends on *both* θ and ϵ. Moreover, if the central bank observes the "velocity shock" η before choosing its policy instrument (or has some information about it), then k^m also depends on η (or whatever else the central bank knows about η). Hence, a contract based on an intermediate monetary target is much more demanding on the principal's information compared to an inflation contract. While the inflation contract needs to be conditional on θ only, the optimal contract based on an intermediate monetary target imposes a penalty on the agent which depends on both θ and ϵ, and possibly even η. The same argument can be shown to be true for a contract that directly targets the policy instrument, q. Generally, the principal finds it easier to monitor the outcome rather than

[7] As discussed in Grilli, Masciandaro, and Tabellini (1991), the central bank statutes of several industrial countries explicitly contain a mandate to achieve price stability.

[8] With the possible exception of New Zealand and Canada, where the central bank and the government jointly announce an inflation target. We return to these countries below.

the policy instrument, because the optimal instrument choice depends on detailed information which may not be available to the principal.

We are thus led to a general conclusion. An inflation contract like (9) minimizes the informational requirement of the principal and thus generally dominates contracts based on intermediate monetary targets or directly on the policy instrument. The intuition for this result is straightforward and related to our discussion on page 62. The conflict of interest between the principal and the agent concerns the inflation bias of the agent. Directly addressing this bias is the simplest and more efficient way to correct the incentive problem.

In the light of this result, the previous remarks about the rarity of explicit inflation targets in the real world are particularly telling. Why do we see exchange-rate targets or monetary targets often imposed (or self-imposed) on central bankers, but rarely see central bankers accountable for the rate of inflation? One reason may have to do with the commitment technology available to the principal. The effect of policy actions on asset prices or the money supply is readily observable. The effect on prices is observable only with substantial delay. It may thus be harder for society to commit to "punishing" a central bank for actions undertaken six months or a year ago. If the central bank deviates from a financial target, the penalty is more immediately related to the policy actions. It may, therefore, be easier to sustain such penalties than in the case of inflation targets.

A second reason may be that contrary to our assumptions the central bank is "risk averse," in the sense that the transfer or penalty t does not enter linearly in the central-bank objective function, (7). In this case the central bank (but not necessarily society) clearly prefers a contract contingent on the money supply or some other easily controllable nominal anchor rather than an inflation target, which it will miss more frequently.

Central bank announcements

Even though it minimizes the information requirements of the principal, the optimal inflation contract in the previous section is still contingent on θ. What is the optimal contract if the principal is uninformed about θ when designing the contract and θ is not verifiable *ex post*? This is the question addressed in this section.[9]

[9] Remember that the variable θ is observed and conditioned upon by wage-setters when forming their expectations. This fact, however, does not imply that θ is also *ex post* verifiable and hence contractible by society. Many observable shocks, such as changes in indirect taxes or in energy prices, affect the costs and benefits of inflation in ways that vary over time and are difficult to quantify in a general and precise way. A trade union which at a particular point in time observes a change in indirect taxes, say, can infer the consequences of that event on inflation in the near future. But it could nevertheless be very hard or impossible for society to legislate *ex ante* how to make the central bank contract contingent in that event. As will be seen below, the contract contingent on the central bank announcement of θ, on the other hand, can be interpreted as a contract which includes an announced inflation target, as well as the realized inflation rate. Legislating such a contract may be a feasible task.

We show that, as long as we retain the assumption that "transfers" to the central bank are not costly to the principal, the *ex ante* optimal monetary policy can still be implemented. However, the contract becomes more complex. The transfer to the central bank now depends both on the realized inflation rate and on the central bank announcement of the variable θ unobserved by the principal. Thus, we have to use those degrees of freedom that were left unexploited in the previous contract (9). Following the literature on the design of regulatory institutions (Baron (1989), Laffont and Tirole (1993)), we show that the optimal contract can be designed to induce truth-telling by the agent. If transfers are costless to the principal, moreover, the agent can always be induced to enact the optimal policy. If transfers are costly, truth-telling remains a property of the equilibrium, but the optimal contract implements a policy of higher inflation than the *ex ante* optimum.

The mechanism that implements the optimal policy has the following interpretation. Through its announcement of θ, the central bank essentially chooses among a menu of contracts, which are all linear in realized inflation. In the model of this section, a higher realization of θ implies a higher *ex ante* optimal inflation rate. As shown below, the optimal contract associates a higher announcement of θ by the bank with a smaller marginal penalty for inflation (a lower k in (9)). Thus, the central bank announces a target inflation rate, given its private information about the economy. The higher the announced target, the looser the penalty for deviating from it.

Throughout the section, we need two simplifying assumptions. The first is that ϵ-shocks are either *ex post* verifiable by the principal or else are not present. Since the latter assumption minimizes on notation, we simply disregard the ϵ-shocks.[10] Second, we add a fourth assumption about the welfare function:

A 4 $W_{p\theta} > 0, W_{xp} < 0$.

The first part of A4 says that a higher value of θ corresponds to a higher *ex ante* optimal inflation rate (in terms of the previous notation, $P^*_\theta > 0$). With this new assumption, a higher θ amounts to both a reduction in the marginal cost of inflation (by A4) and an increase in the marginal incentive

[10] The results of this section carry through even with nonverifiable ϵ-shocks, provided the central bank announces the realization of θ before having seen the realization of ϵ. Baron (1989) discusses the design of optimal contracts when the agent has private information about more than one parameter. In this case there are few general results and the optimal contract is considerably more complex.

to engage in surprise inflation (by A2). The second part of A4 says that the incentive to engage in inflation surprises decreases at higher inflation rates. This assumption is needed for the second-order conditions of the bank optimum to be satisfied (see Appendix 2).

The optimal contract

Consider the following timing of events. First, the principal commits to a contract that specifies a transfer to the central bank. The transfer is a linear function of inflation and contingent on the central bank announcement of θ, which we denote by $\hat{\theta}$. Specifically, the principal commits to a transfer of the following form:

$$T(p,\hat{\theta}) = c(\hat{\theta} - k(\hat{\theta})p. \tag{14}$$

Next, the realization of θ is observed, and wage-setters form their expectations of inflation. Then the central bank announces $\hat{\theta}$ and finally it chooses the actual inflation rate, taking expectations as given. This timing is illustrated in Figure 1. The alternative timing, where the central bank announces $\hat{\theta}$ before wages are set, is discussed below.

The question we now address is what contract, if any, induces the bank to tell the truth and implement the optimal policy.[11] To answer, we work backwards. First, we find the optimal central-bank behavior given a contract. The bank thus maximizes a function similar to (8) by choice of p and $\hat{\theta}$, subject to (14) and for p^e given. The first-order conditions are:

$$W_p + W_x = k(\hat{\theta}) \tag{15}$$

$$c_{\hat{\theta}} = k_{\hat{\theta}}p, \tag{16}$$

equation (15) immediately implies that the contract induces the bank to choose the *ex ante* optimal inflation rate *if*

$$k(\hat{\theta}) = \mu(\hat{\theta}) = W_x(P^*(\hat{\theta}),0;\hat{\theta}), \tag{17}$$

and *if* the bank tells the truth (i.e., $\hat{\theta} = \theta$). This nails down the optimal choice of the $k(\cdot)$ function for the principal. We are then free to determine the $c(\cdot)$ function so as to induce truth-telling by the bank.

Next consider equation (16), and note that $W_x = k(\hat{\theta})$ at the optimal policy. The central bank finds it optimal to tell the truth if its first-order condition (16) is satisfied at the point $\hat{\theta} = \theta$, for all θ. That is, truth-telling requires:

$$c_{\hat{\theta}}(\theta) = k_{\hat{\theta}}(\theta)P^*(\theta). \tag{18}$$

[11] The results of this section also hold under imperfect monitoring – that is, if the central bank controls prices only up to a random noise, as in Section 3.

Integrating (18) by parts, we obtain the $c(\cdot)$ function that induces truth-telling:

$$c(\hat{\theta}) = k(\hat{\theta})P^*(\hat{\theta}) - \int_{\underline{\theta}}^{\hat{\theta}} k(t)P^*_{\theta}(t)dt \tag{19}$$

where $\underline{\theta}$ is lower support of θ.

Combining (19) with (14) and (17), we thus obtain that the optimal contract is:

$$T(p,\hat{\theta}) = \mu(\hat{\theta})(P^*(\hat{\theta}) - p) - a(\hat{\theta}) \tag{20}$$

where $\mu(\cdot)$ is the function defined in (8) and $a(\hat{\theta}) = \int_{\underline{\theta}}^{\hat{\theta}} \mu(t)P^*_{\theta}(t)dt$. Section 2 of the Appendix verifies that the central bank indeed tells the truth and implements the *ex ante* optimal policy when faced with a contract like (20).

Note that, by A4, $P^*(\cdot)$ is a monotonically increasing function. Hence, announcing $\hat{\theta}$ is equivalent to announcing an *ex ante* optimal inflation rate, given the information available to the bank. Moreover, the Appendix proves that the second-order conditions imply that $\mu(\cdot)$ is a decreasing function, while $a(\hat{\theta})$ is clearly increasing in $\hat{\theta}$. The optimal contract is thus very simple and intuitive. The central bank announces a target inflation rate, P^*. It is punished *ex post* for any upwards deviation from the announced rate (or given a premium for any downward deviation from it). The punishment (or premium) consists of a flat negative transfer $(-a(\hat{\theta}))$ plus a term proportional to the deviation of inflation from target. The higher the announced target, the smaller the constant of proportionality ($\mu_{\hat{\theta}} < 0$), but the larger is the flat negative transfer ($a_{\hat{\theta}} > 0$).

More intuitively, the lower the announced inflation target, the larger the incentive for the bank to impose inflationary surprises (by our assumption A4). Therefore, the bank is monitored more tightly by its principal, and the rewards or punishments for deviations from the inflation target are greater. However, the bank is discouraged from choosing an excessively low inflation target by the flat punishment $-a(\hat{\theta})$. When faced with this contract, the bank reveals all its private information and finds it optimal to behave according to the *ex ante* preferences of society. Note also that the performance contract is still linear in realized inflation with a slope coefficient which, in a revealing equilibrium, coincides with the multiplier on the expectations constraint in the *ex ante* policy problem.

Does the optimal contract change if we alter the timing in Figure 1, and let the bank announce $\hat{\theta}$ before wages are set? Perhaps surprisingly, the answer is no. The optimal contract still implements the *ex ante* optimal policy, $P^*(\theta)$. Such a policy only depends on θ, not on $\hat{\theta}$. Hence, if the optimal contract is chosen by society and once wage setters have observed the realization of θ, they can exactly forecast inflation. As a result, in equilibrium the central-bank announcement of $\hat{\theta}$ carries no additional information for wage-setters. This is not to say that the central-bank announcement is irrelevant. It matters a lot, but only in so far as it shapes the central-bank incentives as spelled out by the central-bank contract.[12]

Discussion

We find it striking that such a simple and intuitive contract resolves intricate incentive and information problems in monetary policy. Even more striking, perhaps, is the similarity between such a contract and the recent practice in New Zealand and Canada of having the bank announce jointly with the government a target inflation rate, for which the bank is then held accountable.[13]

In most of the recent literature on monetary policy games, announcements of targets for monetary policy do not play any role whatsoever. What matters in that literature is what central banks do, not what they say.[14] The results in this section, by contrast, suggest that monetary policy announcements may be quite important, even though they may carry no essential information for the private sector. The reason is that policy announcements create a standard against which policymakers can subsequently be held accountable. In many countries, the monetary authorities indeed report to the legislature or announce in advance their nominal targets. Such reports or announcements may have important effects on their incentives later on, if, as plausible, their behavior will be judged with reference to their previous announcements.

This discussion suggests a mechanism for overcoming the adverse incentive problems of discretionary monetary policy. Announcements, by their very nature, are often made before private nominal contracts are signed. Hence they are less subject to the well-known inflation bias. If announcements exert effects on policy later on, the credibility of monetary policy can be enhanced. Naturally, for policy announcements to have such effects they must be clearly understood and be nonambiguous. Bernanke and Mishkin (1992) discuss some interesting evidence in this regard. They find that the central banks of those countries that have been most successful in keeping inflation low and stable – Germany, Switzerland and, more recently, Japan – have a practice of announcing clearly understood monetary targets, and have a transparent procedure for revising them or accepting short-run deviations. This contrasts with countries such as the United Kingdom and the United States, where the monetary authorities have often operated with multiple targets for policy and where the procedure for revising these targets has been quite opaque.

Legislation vs. government-imposed targets

In the two previous sections we have seen how carefully-designed contracts can induce the central bank to implement the optimal policy. Even though these contracts are relatively simple, they are still contingent on the realizations of shocks to the economy, or the central bank's announcements of such shocks. In this section we make the perhaps realistic assumption that only simpler contracts, based on the central bank's performance, can be embodied by law in the central bank institution. In the Introduction we identified two interesting routes followed in real-world institutional reforms, a *legislative* route and a *targeting* route. Inspired by this discussion, we study the following specific problem in "second best institution design": the constitutional framer has two options, namely, to set up a central bank that is governed either by strict legislation or by targets imposed by the government.

In the first case, the framer enacts legislation that delegates the task and operating procedures of monetary policy to a very independent central bank, sidestepping to a large extent other government institutions. (We can think about this as a caricature of the German, or maybe the Swiss, case.) It is reasonable to assume that such legislation must be relatively inflexible and thus cannot easily adapt incentives to specific situations or changing circumstances. This arrangement thus has the cost that it may not fully tailor the bank's incentives to the information available when monetary policy is to be carried out.

In the second case, the framer instead enacts legislation that delegates to the central bank's political principal, which may be the executive or the legislature, a great deal of discretionary influence in imposing and enforcing specific simple inflation targets for monetary policy. (We can think about this as the New Zealandian, or maybe the Canadian, case.) This targeting arrangement allows current information to influence the targets, so it would clearly dominate the legislative arrangement with a truly benevolent political principal. However, suppose we make the reasonable extension that the

[12] This point can be seen more precisely as follows. Let $P^e(\hat{\theta},\theta)$ be the expectations function of wage-setters, which in general is allowed to depend on both $\hat{\theta}$ and θ. Under the alternative timing the first-order condition (16) for the optimal announcement is replaced by

$$c_{\hat{\theta}} = W_r P^e_{\hat{\theta}} + k_{\hat{\theta}} P. \qquad (16')$$

The first term enters since the announcement in principle could change wage-setters' expectations. In equilibrium, it must be true that $P^e(\theta,\theta) = P^e(\theta)$ and that $\hat{\theta} = \theta$. Hence, by (15), $P^e(\hat{\theta},\theta)$ must satisfy $W_p(P^e(\theta,\theta),0;\theta) = 0$ for all values of θ, at the point $\hat{\theta} = \theta$. At the point $\hat{\theta} = \theta$, we must therefore have $P^e_{\hat{\theta}} = 0$. But then the term $W_r P^e_{\hat{\theta}}$ falls out of (16'), and we are left with (16).

[13] In the case of Canada, however, it is not clear in what sense the bank is accountable for missing the target.

[14] An exception is Stein (1989).

executive or the legislature also has a private agenda, maybe tied to electoral or partisan objectives. Then, the arrangement has the cost that these private objectives get embodied in the central-bank contract. In the remainder of this section, we show how these two arrangements can be formally illustrated and analyzed within the context of our contractual model.

Central bank ruled by legislation

We return to our original framework, where there are both θ-shocks and ϵ-shocks. The θ-shocks are known by the bank and by private agents, while the ϵ-shocks are known only by the bank. From (3), we know that the *ex ante* second-best policy satisfies:

$$(W_p + W_x) - E_{\epsilon/\theta} W_x = 0. \tag{21}$$

The framer imposes an incomplete linear performance contract, $T(p) = c - kp$, on the central bank. Going through the same steps as in previous sections, we find that the bank sets p in state (θ, ϵ) such that

$$(W_p + W_x) = k \tag{22}$$

and that the optimal choice of k is

$$k = E_{\theta,\epsilon/\theta} W_x. \tag{23}$$

Combining these two expressions, we get

$$(W_p + W_x) - E_{\epsilon/\theta} W_x = (E_{\theta,\epsilon/\theta} W_x - E_{\epsilon/\theta} W_x). \tag{24}$$

The RHS of (24) is thus a measure of the inefficiency generated by the incomplete contract. We see that the central bank's incentives are correct on average. But there is a cost, in the form of a deviation from the second-best policy, because the incomplete contract does not vary the punishments (rewards) for inflation performance, as W_x varies with θ. Obviously, this cost is larger the more the inflation bias fluctuates: that is, the larger is $W_{x\theta}$ and the larger the variance of θ.

Central bank ruled by government-imposed targets

In this case, we add another agent to our framework: "the government." We assume the following timing. (i) The government observes θ and ϵ and imposes a performance contract $T(p, \epsilon, \theta) = c - k(\theta, \epsilon)p$ on the central bank. Note that this notation represents something different than before. Like before, the government is restricted to writing a simple performance contract, which is linear in p. But the slope coefficient k will now

incorporate the government's information on θ and ϵ, hence the notation $k(\theta, \epsilon)$. (ii) Private agents have rational expectations as usual, so they know $P(k(\theta, \epsilon), \theta, \epsilon)$. However, they only observe θ, not ϵ, so they form expectations $P^e(\theta) = E_{\epsilon/\theta} P(k(\theta, \epsilon), \theta, \epsilon)$. (iii) The central bank sets p, having observed θ, ϵ and $P^e(\theta)$.

The bank's first-order condition is the usual one, namely,

$$(W_p + W_x) = k(\theta, \epsilon). \tag{25}$$

The government has the following objective function

$$E_{\theta,\epsilon}[W(p, p - p^e; \epsilon, \theta) + \alpha G(p, p - p^e; \epsilon, \theta)], \tag{26}$$

where α is a weight and $G(\cdot)$ is a function with properties similar to $W(\cdot)$, which captures the government's private agenda. The government chooses $k(\theta, \epsilon)$ to maximize (26) subject to (25) and the definition of $P^e(\theta)$. Combining the first-order conditions with (25), we get

$$k(\theta, \epsilon) = E_{\epsilon/\theta} W_x - \alpha((G_p + G_x) - E_{\epsilon/\theta} G_x). \tag{27}$$

Note, that the slope of the optimally-chosen performance contract now may depend on ϵ, unlike in Section 3. The reason is that the government, rather than a benevolent constitutional framer, chooses the contract, and the government's private objective $G(\cdot)$ generally depends on ϵ in a different way than society's welfare $W(\cdot)$.

If we combine (25) and (27), we can characterize the equilibrium monetary policy in state (θ, ϵ):

$$(W_p + W_x) - E_{\epsilon/\theta} W_x = -\alpha((G_p + G_x) - E_{\epsilon/\theta} G_x)). \tag{28}$$

It follows from (21) and (28) that the equilibrium policy generally differs from the second-best policy. The deviation is more pronounced the larger is α—the weight on the private agenda and the larger the discrepancy between the functional forms of $W(\cdot)$ and $G(\cdot)$. In other words, equilibrium policy ruled by government-imposed targets is more distant from the second-best: (a) the more important are the government's own goals; and (b) the more they differ from society's goals for monetary policy.

Discussion

The recent literature has stressed the tradeoff between credibility and flexibility in the design of institutions governing monetary policy. The analysis here basically handles that problem by contractual arrangements, assuming that either the framer or the government can commit to a contract with the central bank before private agents form inflationary expectations. Thus, policy may not suffer from any inflation bias, even in the case we have just treated when it is dictated by a government with private objectives. As is clear from equation (32), there is a systematic inflation bias only if the *ex ante* optimal policy that follows from the government's private objective $G(\cdot)$ is systematically higher than society's *ex ante* optimal policy; that is, if for most (ϵ, θ), the expression $G_p + G_\pi - E_{\epsilon/\theta} G_\pi$ is positive at the value of p which sets $W_p + W_\pi - E_{\epsilon/\theta} W_\pi = 0$.

Instead, our analysis points towards a different tradeoff in the design of central bank institutions, namely a tradeoff between *incentives* and *information*. As we have seen, trying to control central-bank policy by legislation alone may lead to informational problems, in that a truly independent bank may not adapt policy to specific circumstances. Giving the government an important role in formulating the targets for monetary policy avoids these informational problems. On the other hand such a mandate may lead to incentive problems if the objectives of the sitting executive and society disagree. This kind of tradeoff is very similar to the one that arises in the regulatory literature on privatization. Shapiro and Willig (1990), in particular, consider the choice between running some production activity in a public company or, alternatively, in a private, regulated company. They end up emphasizing a completely analogous tradeoff.

Suppose a targeting approach is contemplated. An important issue, which we have not addressed, is *who* should impose the inflation target and *to whom* the central bank officials should be held accountable. To put it differently, how should the central bank be ordered relative to other government bodies, the executive and the legislature in particular? Looking across democratic countries, the most common institutional structure is that the central bank is vertically subordinated to the executive, which in parliamentary political systems itself is subordinated to the legislature. The only exceptions to this rule among the parliamentary democracies are Sweden and Finland, where the central banks are instead accountable directly to the parliament and thus parallel to the executive. Further exceptions are the federal (presidential) systems in the United States, Germany, and Switzerland, where the hierarchical relations between government and legislature are more complicated.

Drawing on the analysis in this section, it is not self-evident which arrangement is best. The executive generally has more control over the government bureaucracy and therefore it may have better access to timely information. But the legislature generally embodies more dispersed interests—because it includes the opposition—and therefore it may be less prone to influence policy on the basis of partisan or electoral incentives.

Incomplete contracts

In this section we continue our study of simple contracts, which can only be made contingent on observable central-bank performance and which are restricted to be linear. A new dimension is added to the design of such contracts when the central bank only controls p indirectly and imperfectly via some policy instrument q. An optimal incomplete performance contract will then have to take into account exactly how the central bank's incentive problem covaries with its control over p. Our analysis here draws on the recent paper by Baker (1992).

We study a situation where there are both θ-shocks and ϵ-shocks. There are no private agendas, only the usual credibility problem. To model imperfect and indirect control, we adopt our previous assumption (10): $p = \pi(q, \nu)$, where q is the monetary policy instrument.

The central bank's objective function is thus

$$E_{\theta,\epsilon,\nu}[W(p, p - p^e; \theta, \epsilon) + T(p)], \tag{29}$$

where p is given by (10) and $T(p)$ is the restricted performance contract.

$$T(p) = c - kp. \tag{30}$$

As before, the bank sets its policy instrument knowing θ and ϵ but not knowing ν. The first-order condition to the bank's maximization problem

$$E_{\nu/\theta,\epsilon}[(W_p + W_\pi)\pi_q] = kE_{\nu/\theta,\epsilon}\pi_q, \tag{31}$$

implicitly defines the optimal value of q as a function

$$q = Q(k, \theta, \epsilon), \tag{32}$$

with derivative $Q_k = E_{\nu/\theta,\epsilon}\pi_q/\Delta$. We assume that Δ—which is a positive number proportional to the second-order condition of the bank's optimization problem—is approximately constant across states.

Expected inflation is thus given by:

$$P^e(\theta) = E_{\epsilon/\theta,(\nu/\theta,\epsilon)}\pi(Q(k, \theta, \epsilon), \nu). \tag{33}$$

The framer's task is to choose k to maximize

$$E_{\theta,\epsilon,\nu}[W(p, p - p^e; \theta, \epsilon)], \tag{34}$$

subject to (10) and (31)–(33). As we demonstrate in the Appendix, the solution is

$$k = E_{\theta,(\epsilon/\theta),(\nu/\theta,\epsilon)}W_\pi + \frac{Cov_\theta[E_{\epsilon/\theta,\nu/\theta,\epsilon}W_\pi, E_{\epsilon/\theta}((E_{\nu/\theta,\epsilon}\pi_q)^2)]}{E_{\theta,\epsilon/\theta}((E_{\nu/\theta,\epsilon}\pi_q)^2)}. \tag{35}$$

To interpret this result, it is perhaps easiest to refer to the first version of the model where p-contracts contingent on θ are indeed feasible. In that case, the slope coefficient of the p-contract (generalizing (7)) is simply:

$$k(\theta) = E_{\epsilon/\theta,\nu/\theta,\epsilon} W_x, \tag{36}$$

so that

$$E_\theta(k(\theta)) = E_{\theta,(\epsilon/\theta),(\nu/\theta,\epsilon)} W_x. \tag{37}$$

Thus, the unconditional (35) contract penalizes (rewards) high (low) inflation more severely *on average* than in the conditional contract if and only if the second term of the RHS of (35) is positive. This term is proportional to the covariance (across θ) between the inflation bias and the volatility of the relation between p and q. However, since Q_k is proportional to $E_{\nu/\theta,\epsilon}\pi_q$, the term $(E_{\nu/\theta,\epsilon}\pi_q)^2$ measures the average, (across ν) sensitivity of p to k (since $\pi_k = \pi_q Q_k$). By A2, when θ is high W_x also tends to be high, so that the incentive problem tends to be serious. If $Cov_\theta[\cdot]$ is positive, this means that p is also particularly sensitive to k in those states. Equation (35) says that the unconditional contract should punish high inflation severely under those circumstances. This makes intuitive sense. With a positive $Cov_\theta[\cdot]$, a high k has particular bite on the central bank's incentives precisely when the incentive problem is especially serious.

Of course, our model is much too stylized to allow us to speak with any confidence about the likely sign of $Cov_\theta[\cdot]$. Nevertheless, the above results and related results we have derived for incomplete performance contracts basesd on the central bank's fulfillment of intermediate targets $M(q,\eta)$—bring home a general point which is reminiscent of a general result in Baker (1992). The optimal incomplete performance contract takes into account how the effects of the agent's actions on the principal's objective covary with the effects on the agent's performance measure.

An interesting issue for further work is to construct more structural models of monetary policy and investigate incomplete performance contracts over alternative intermediate targets. Maybe this can provide a clearer understanding of why rules defined over certain intermediate targets, exchange rates in particular, seem to be enforced much more strictly than rules for other intermediate targets such as monetary aggregates.

Conclusion

In the existing literature on monetary policy credibility, there is a common idea. Delegating policy to an independent central banker who is averse to inflation restores credibility at the expense of the optimal output stabiliza-

tion policy.[15] One of the insights of this paper, which expands on earlier results by Walsh (1992), is that this idea is not generally correct. The optimal central-bank contract, when feasible, does not entail any loss of welfare nor any sacrifice of stabilization policies, even if the central bank has private nonverifiable information. Moreover, if the optimal contract cannot be written, because of transactions or computation costs, writing a suboptimal contract may still not require a loss of stabilization policies. For instance, with a simple linear inflation target not contingent on θ, like in Sections 5 and 6, the central bank carries out an optimal stabilization policy, but inflation is either too high or too low, depending on the realized value of θ. In general, the design of optimal monetary institutions does not seem to entail any tradeoff between credibility and flexibility (inflation and output stabilization). This fiding is consistent with the empirical results in Grilli, Masciandaro and Tabellini (1991) and Alesina and Summers (1991), who find no association between various measures of central bank independence and the variability of output growth (or other real macroeconomic indicators), while they do find that a higher degree of independence is associated with lower average inflation and less variable inflation.

A second general insight of this paper concerns the desirability of inflation targets as opposed to other nominal targets. This presents a puzzle: why do we so rarely see central banks held accountable for the rate of inflation? One answer was suggested earlier. It may be that a commitment to a more readily observable nominal variable, such as a monetary aggregate or the exchange rate, is easier to enforce. A second possible answer is that central bankers would not like to be held accountable for something they do not control tightly, and hence are liable to miss rather often. Overall, we regard this as an important puzzle, with relevant implications for real world institutions.

The third general insight of this paper is the importance of monetary policy announcements. If transparent and clearly understood, policy announcements create a reference point against which to judge subsequent central-bank behavior. As such, they affect incentives for later policy choices and can therefore play an important role in giving credibility to monetary policy. In general, any target which makes the bank clearly accountable for its behavior can play this desirable role. But the results of this paper suggest that the accountability should concern the policy outcomes directly, and that some sanction should be imposed on a central bank for missing the specified targets, or some premium given for hitting them. Again, it is somewhat of a puzzle that so few central banks are accountable in this way. In this regard, the draft treaty for European Monetary Union goes in the wrong direction

[15] The idea was first put forward by Rogoff (1985). Since then it has appeared in many studies of the topic, such as Lohman (1992), Flood and Isard (1989), Cukierman (1992), Persson and Tabellini (1990), Alesina and Grilli (1991), to point out a few

in that it leaves the European Central Bank completely unaccountable to its political principals. As Goodhart (1992) noted, the European Central Bank statute is no surprise in this regard, since it was drafted by the likely future governors of the Bank itself.

Whatever the value of these general insights, the results are still far less specific than we would like them to be. A basic question is exactly how to translate the various contracts that we have derived into real-world institutions. One possibility is, of course, to make a literal interpretation: $T(\cdot)$ in the model then becomes performance-based remuneration for the officials of the central bank. But we have also argued that less literal interpretations are possible. If the central-bank leadership attaches a fixed cost to losing their job, and their probability of losing it increases in the deviation from the policy target, this works pretty much as a performance contract of the type that we have studied. Perhaps these interpretations fit better with a government-imposed inflation target. Yet another interpretation relies on career concerns. Suppose that the central bank enjoys a great deal of independence and has a clearly-stated goal of price stability. In such a situation, the rest of society—and financial markets in particular—may judge the competence of the bank's leadership according to how well they manage to meet their inflation targets. Central-bank officials who worry about their future career or about their status and prestige may then perceive a link between their success on the job and their future payoff, which is similar to $T(\cdot)$ in the model.[16]

More generally, subtle threats and rewards can be imposed on the central bank in many ways. As is clear from the empirical literature on central-bank independence, central-banking institutions across the world differ in many dimensions. A challenging task for further work is to link the incentive approach in this paper with the institutional facts on appointment procedures, budget control, etc. in that literature.

We find the recent reforms of the New Zealand central bank particularly interesting. The very explicit Policy Targets Agreement, signed by the Governor and by the New Zealand Finance Minister in 1990, is a document which comes close to a performance-based contract of the type we have studied in the paper. This document includes a specific target range for the inflation rate (0–2%) in terms of the CPI. It also includes explicit contingency statements: the inflation rate is allowed to deviate from the target in response to specific severe shocks.

The Governor of the Reserve Bank is indeed accountable for achieving the inflation target in that he may be dismissed, or his contract may not be renewed should he fail to meet the target. Furthermore, the Bank's revenue has been fixed in nominal terms for a five-year period. (See Archer (1992) for an interesting account of the New Zealand legislation, with particular emphasis on these kinds of incentive issues.)

Another issue, which is often raised in discussions about central banking, is whether there is any rationale for the secretive practices that characterize many central banks. Our results point in somewhat different directions. On the one hand, we argued that the informational requirements on the optimal contract were less severe in the case of ε-shocks than in the case of θ-shocks. Since the former were private information of the bank, this seems to speak for secrecy. On the other hand, we argued that it was important for supporting the optimal policy that policy announcements be clear and transparent. This seems to speak against secrecy. Surely these are not the only arguments. But we believe the kind of contractual/regulatory approach that we have pursued in this paper, with its focus on private information and incentives, is inherently the right way of posing the secrecy question.

Throughout the paper we have mentioned a number of other open issues. Are sanctions on central bankers genuinely harder to enforce when tied to inflation targts rather than to intermediate targets such as money-growth rates or exchange rates? How important are announcements of policy targets in committing central banks to a desirable course of action? Who should be the political principal of the central bank? Can an incomplete contracting approach help us understand why exchange-rate targets seem to be more seriously enforced, at least in many small countries, than money-growth targets? We believe the approach that we have pursued in this paper is a promising way to go when seeking answers to these and other questions regarding monetary policy institutions. More research should follow.

[16] Tirole (1992) discusses and sketches formally how career concerns may shape incentives and affect the behavior of government officials.

Appendix

Optimal inflation contract with imperfect monitoring

Here we illustrate only the case where ν is *not* observed by the policymaker when policy is set. Suppose that inflation is given as in (10). The *ex ante* optimal policy must then satisfy the first-order condition:

$$E_{\nu/\theta,\epsilon}\{(W_p + W_x)\pi_q\} = E_{\nu/\theta,\epsilon}\pi_q \cdot E_{\epsilon/\theta}E_{\nu/\theta,\epsilon}W_x. \qquad (A.1)$$

An agent facing a contract like (9) sets q to satisfy the following (*ex post*) first-order condition:

$$E_{\nu/\theta,\epsilon}\{(W_p + W_x)\pi_q\} = \mu(\theta)E_{\nu/\theta,\epsilon}\pi_q. \qquad (A.2)$$

Hence, setting:

$$\mu(\theta) = E_{\epsilon/\theta}E_{\nu/\theta,\epsilon}W_x \qquad (A.3)$$

yields the *ex ante* optimal policy.

If ν is observed by the policymaker, the proof is similar except that the term E_ν drops out everywhere.

Verification that (20) induces the central bank to set $\hat{\theta} = \theta, p = P^*(\theta)$

The central bank solves:

$$\max_{p,\hat{\theta}} W(p, p - p^e; \theta) + \mu(\hat{\theta})(P^*(\hat{\theta}) - p) - \int_{\underline{\theta}}^{\hat{\theta}} \mu(t)P_\theta^*(t)dt. \qquad (A.4)$$

The first-order conditions with respect to p and $\hat{\theta}$ are, respectively:

$$W_p + W_x - \mu(\hat{\theta}) = 0 \qquad (A.5)$$

$$\mu_\theta(P^*(\hat{\theta}) - p) + \mu(\hat{\theta})P_\theta^*(\hat{\theta}) - \mu(\hat{\theta})P_\theta^*(\hat{\theta}) = 0. \qquad (A.6)$$

By $(A.4)$, the optimal policy $p = P^*(\theta)$ is chosen if $\mu(\hat{\theta}) = W_x$ (cf. (17)). This, together with $(A.6)$, implies that $\hat{\theta} = \theta$. Consider the derivative of $(A.6)$ with respect to $\hat{\theta}$. It is:

$$\mu_\theta P_\theta^*. \qquad (A.7)$$

The second-order conditions for a central-bank optimum with respect to $\hat{\theta}$ require $(A.7)$ to be negative. By A4, $P_\theta^* > 0$. Hence $\mu_\theta < 0$, as claimed in the text. By the definition of $\mu(\theta)$ in (17), this inequality is satisfied only if $W_{px} < 0$ (cf. A4 again).

Derivation of (32)

The framer maximizes the expression

$$E_{\theta,\epsilon,\nu}W(p, p - p^e; \theta, \epsilon), \qquad (A.8)$$

where p is given by (10), q by (30), and $P^e(\theta)$ by (31), subject to the bank's first-order condition (31). The first-order condition with respect to k is:

$$E_{\theta,\epsilon/\theta}[E_{\nu/\theta,\epsilon}(W_p + W_x) \cdot \pi_q Q_k] - E_\theta[E_{\epsilon/\theta,(\nu/\theta,\epsilon)}W_x \cdot \pi_q Q_k] = 0. \qquad (A.9)$$

Using the bank's first-order condition (31) and the result in the text that $Q_k = E_{\nu/\theta,\epsilon}\pi_q/\Delta$, $(A.8)$ can be rewritten

$$k \cdot E_{\theta,(\epsilon/\theta)}((E_{\nu/\theta,\epsilon}\pi_q)^2) = E_\theta[E_{\epsilon/\theta,(\nu/\theta,\epsilon)}W_x \cdot E_{\epsilon/\theta}((E_{\nu/\theta,\epsilon}\pi_q)^2)]. \qquad (A.10)$$

From $(A.9)$ we can solve for k. Using the definition of covariance to rewrite the solution, we obtain the expression in (35).

References

Alesina, A., (1987). Macroeconomic Policy in a Two-Party System as a Repeated Game. *Quarterly Journal of Economics*, **102**: 651 78.

Alesina, A. and Grilli, V., (1991). The European Central Bank: Reshaping Monetary Politics in Europe. Chapter 3 in *Establishing a Central Bank: Issues in Europe and Lessons from the U.S.*, (eds.) M. Canzoneri, V. Grilli, and P. Masson. London: Cambridge University Press.

Alesina, A. and Summers, L., (1991). Central Bank Independence and Macroeconomic Performance: Some Comparative Evidence. Forthcoming in *Journal of Money Credit and Banking*.

Archer, D., (1992). Organizing a Central Bank to Control Inflation: The Case of New Zealand. Mimeo, Reserve Bank of New Zealand.

Baker, G., (1992). Incentive Contracts and Performance Measurement. *Journal of Political Economy*, **100**: 598 614.

Baron, D., (1989). Design of Regulatory Mechanisms and Institutions. Chapter 24 in *Handbook of Industrial Organization*, (eds.) R. Schmalensee and R. Willig. Amsterdam: North Holland.

Barro, R., (1983). Inflationary Finance under Discretion and Rules. *Canadian Journal of Economics*, **16**: 1 25.

Barro, R. and Gordon, D., (1983). Rules, Discretion and Reputation in a Model of Monetary Policy. *Journal of Monetary Economics*, **12**: 101–22.

Bernanke, B. and Mishkin, F., (1992). Central Bank Behavior and the Strategy of Monetary Policy: Observations from Six Industrialized Countries. *NBER Macroeconomics Annual*. Cambridge: MIT Press.

Calvo, G., (1978). On the Time-Consistency of Optimal Policy in a Monetary Economy. *Econometrica*, **46**: 647 61.

Canzoneri, M., (1985). Monetary Policy Games and the Role of Private Information. *American Economic Review*, **75**: 1056 70.

Cukierman, A., (1992). *Central Bank Strategy, Credibility and Independence: Theory and Evidence*. Cambridge: MIT Press.

Cukierman, A. and Meltzer, A., (1986). A Theory of Ambiguity, Credibility and Inflation Under Discretion and Asymmetric Information. *Econometrica*, **53**: 1099 1128.

Flood, R. and Isard, P., (1989). Monetary Policy Strategies. *IMF Staff Papers*, **36**: 612–32.

Giovannini, A., (1992). Central Banking in Monetary Union: Reflections on the Proposed Statute of the European Central Bank. *Carnegie-Rochester Conference Series*, **38**: 191 229.

Goodhart, C., (1992). The Statute of the European Central Bank. Mimeo, London School of Economics.

Green, E. and Porter, R., (1984). Non-Cooperative Collusion Under Imperfect Price Information. *Econometrica*, **51**: 87-100.

Grilli, V., Masciandaro, D., and Tabellini, G., (1991). Institutions and Policies. *Economic Policy*, **13**: 342 91.

Kydland, F. and Prescott, E., (1977). Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, **85**: 473 90.

Laffont, J-J. and Tirole, J., (1993). *A Theory of Incentives in Regulation and Procurement*. Cambridge: MIT Press.

Lohman, S., (1992). The Optimal Degree of Commitment: Credibility versus Flexibility. *American Economic Review*, **82**: 273–86.

Obstfeld, M., (1991). Destabilizing Effects of Exchange-Rate Escape Clauses. Mimeo, University of California at Berkeley.

Persson, T. and Tabellini, G., (1990). *Macroeconomic Policy, Credibility and Politics*. London: Harwood Economic Publishers.

Rogoff, K., (1985). The Optimal Degree of Commitment to an Intermediate Monetary Target. *Quarterly Journal of Economics*, **100**: 1169 90.

Rogoff, K. and Sibert, A., (1988). Elections and Macroeconomic Policy Cycles. *Review of Economic Studies*, **55**: 1 16.

Shapiro, C. and Willig, R., (1990). Economic Rationales for the Scope of Privatization. Mimeo, Princeton University.

Stein, J., (1989). Cheap Talk and the Fed: A Theory of Imprecise Policy Announcements. *American Economic Review*, 79: 32-42.

Tirole, J., (1992). The Internal Organization of Government. Mimeo, MIT.

Walsh, C., (1992). Optimal Contracts for Central Bankers. Mimeo, University of California at Santa Cruz.