Optimal Contracts for Central Bankers

By CARL E. WALSH*

This paper adopts a principal–agent framework to determine how a central banker's incentives should be structured to induce the socially optimal policy. In contrast to previous findings using ad hoc targeting rules, the inflation bias of discretionary policy is eliminated and an optimal response to shocks is achieved by the optimal incentive contract, even in the presence of private central-bank information. In the one-period model that has formed the basis for much of the literature on discretionary monetary policy, it is shown that the optimal contract ties the rewards of the central banker to realized inflation. (*JEL* E52, E58)

Finn Kydland and Edward Prescott (1977) and Robert Barro and David Gordon (1983) have shown that, if the monetary authority faces an incentive to expand output above its equilibrium level, discretionary policy has an inflationary bias. At a zero rate of inflation, the marginal benefit of an expansion induced by surprise inflation exceeds the marginal cost of the resulting inflation. When the central bank chooses the rate of inflation to equate marginal costs and benefits, and the public understands that it will do so, the central bank's announcement of a zero-inflation policy will not be credible. The public will expect a positive rate of inflation, and the central bank can do no better than to fulfill those expectations.

A key insight that has motivated the large literature expanding on this analysis is the recognition that central banks respond to the incentives they face.[1] Most existing work has viewed the incentive problem as one involving many principals (the individuals in the economy) and one agent (the central banker). In this framework, the principals may select the agent, but they are unable to specify the objective function of the agent.[2] However, in no country is the institutional framework such that the actual agents in charge of monetary policy are directly chosen by the individual citizens. Instead, citizens in democracies choose a government, and the central-bank head is chosen by the government. Thus, monetary policy involves a multilevel principal–agent problem. Kenneth Rogoff (1985), Susanne Lohmann (1992), and Christopher Waller (1992) focus on the government's choice of

*Department of Economics, University of California, Santa Cruz, CA 95064, and Economic Research, Federal Reserve Bank of San Francisco, P.O. Box 7702, San Francisco, CA 94120. I thank Ramon Moreno, David VanHoose, and two anonymous referees for helpful comments. This research was partially supported by grants from the Center for German and European Studies at UC Berkeley, the Division of Social Sciences at UC Santa Cruz, and the National Science Foundation. Any opinions expressed are those of the author and not necessarily those of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

[1]This literature has analyzed the conduct of monetary policy under such different central-bank objective functions as a social loss function, partisan political preferences, personal preferences of the central-bank head, and targeting rules imposed by (in the context of the United States) Congress or the President. As Torsten Persson and Guido Tabellini (1990 p. 2) describe this recent literature, "From a positive point of view, the theory describes the policymaker's behavior under alternative incentive constraints. From a normative point of view, it suggests how to embed desirable incentive constraints in the existing political and economic institutions, through appropriate institutional reform."

[2]Although see Brendan O'Flaherty (1990).

a central-banker, viewing the government as choosing from a population of potential bankers with differing preferences over inflation and output fluctuations. The government picks the banker whose preferences are such that the resulting conduct of monetary policy maximizes the government's expected utility. In particular, the inflationary bias of discretionary policy can be reduced, at the cost of suboptimal policy responses to aggregate supply shocks, if the central banker minimizes a loss function that puts less weight on output fluctuations than the principal's loss function does.

However, in dealing with a central-bank head, a government can directly affect the central banker's objective function and incentives in many ways. Lohmann (1992) considers the case in which the central banker knows she will be overridden by the government if the economy is subject to a disturbance that is "too" big, and Rogoff (1985 p. 1180) argues that targeting rules might be enforced by making the monetary authority's budget depend on adherence to the rule. Similarly, Michelle Garfinkle and Seonghwan Oh (1993) suggest that a targeting rule might be enforced by legislation punishing the monetary authority if it fails to achieve the target. Such institutional aspects of the central bank's structure and its relationship with the government can be thought of as representing a "contract" between the government and the monetary authority. The conduct of monetary policy is then affected by the contract the government offers to the central bank.

This point of view raises the question of whether there is an optimal contract the government should offer to the central banker. That is, if central banks respond to the incentives they face, with what incentives should they be faced? This is exactly the type of question addressed in the principal–agent literature. In a standard principal–agent problem, one would think of the principal offering to her agent a contract that is designed to affect the agent's choice of action. The principal might structure a contract so that the agent's income depends on the agent's actions, thereby affecting the incentives the agent faces. In the monetary-policy context, the institutional design of the central bank can influence the incentives the bank faces in its conduct of policy. For example, incentives are affected by such facets of design as the appointment and reappointment procedures for members of the policy-making committee, the existence of reporting requirements, and the presence of legislated policy goals. The determination of these aspects of the contract would constitute what Lohmann (1992) has described as the institutional design stage or would correspond to the legislative approach suggested by Matthew Canzoneri (1985) and Garfinkle and Oh (1993) for establishing targeting rules. The institutional structure of the central bank also includes the budgetary procedures that determine how central-bank operations are financed, and Rogoff (1985) has suggested that a government might make the central bank's budget or the income of the head of the central bank contingent on the state of the economy, thereby influencing the incentives the central banker faces in choosing the rate of inflation. As part of its central-bank reform, the New Zealand government actually considered including a financial incentive in the contract for the head of their central bank that would have resulted in a bonus payment if the bank's inflation target were achieved.[3] Legislated aspects of the central bank's structure are costly and time-consuming to change. Governments can plausibly commit to an institutional structure more easily than they can commit to a specific monetary-policy action.

The objective of this paper is to determine how the rewards to the central bank should be structured in order to induce the socially optimal policy (as defined below). That is, the optimal incentive structure is derived directly. This seems a necessary first step in any analysis of policy design, one that must be taken before the issue of implementing optimal policy design can be

[3]As ultimately passed by New Zealand's Parliament, the Reserve Bank Act of 1989 did not include such an incentive (see Walsh, 1994).

usefully addressed. In order to focus on the nature of the incentives with which the central bank should be faced, I assume that both the government and the individuals who might head the bank share the same preferences over inflation and output fluctuations. This could be taken to reflect the outcome of some (unmodeled) appointment process that ensures a similarity of views between the central-bank head and the government. As in the standard model of time-inconsistent monetary policy, both the government and the central bank would like to be able to commit to a low-inflation policy, but both face incentives to support inflationary policies once private agents have entered into wage and price contracts. I examine the incentives the government would like to establish for the central bank at the institutional-design stage in order to eliminate the inflation bias of discretionary policy while still preserving the ability of the central bank to respond with discretion to new information. Thus, the focus is similar to that of the literature on targeting rules in which the government attempts to impose some limits on the central bank's independence in order to reduce a bias toward inflation.[4]

To determine the optimal incentive structure for the central bank, I assume that the government can offer the central-bank head a state-contingent wage contract. Such a contract allows one to derive explicitly the manner in which the bank's incentives should be dependent on the state of the economy. While there are numerous reasons to question the effectiveness and implementability of such employment contracts in the context of monetary-policy determination, a (possibly) state-contingent wage contract for the central banker represents a useful fiction for deriving the optimal incentive structure with which the central bank should be faced and provides a convenient starting point for the analysis of optimal central-bank incentives.[5]

While the use of a state-contingent wage contract is convenient for determining the optimal incentive structure, the specific institutional structure that might implement these incentives is only briefly touched upon in this paper. However, viewing the contract only in the sense of a wage contract is an unnecessarily narrow interpretation of the contract between the government and the central bank. The incentives implied by the optimal contract might be implemented by more traditional institutional structures. For example, the optimal contract is shown to

[4]Implicit in this and other analyses of central-bank design is the assumption that the specialized knowledge required to manage the day-to-day conduct of monetary policy requires the delegation of responsibility to a separate agency, a central bank, equipped with the ability to respond flexibly to new information. Much of the focus associated with the design of a common central bank for Europe and the reform of the Reserve Bank of New Zealand has been on limiting the government's ability to interfere with the conduct of monetary policy. However, in the United States, congressionally imposed targeting and reporting requirements arose from a desire to limit inflationary monetary policy. I assume that *ex ante* the government and the central bank share a desire for a low-inflation policy. Consequently, the government sets an incentive structure for the central bank to ensure that, *ex post*, the central bank still prefers low inflation. The government, *ex post*, prefers higher inflation; therefore, it is desirable to limit the government's influence over the conduct of policy. Thus, the optimal contract derived in Section I allows the central bank complete discretion in the actual determination of policy.

[5]As noted in footnote 4, the government's preferences *ex post* will differ from those of the central bank under a state-contingent contract. If the transfer payment to the central bank represents the employment contract of the central banker, then the legal enforceability of contracts might ensure that the government could credibly commit to the transfer scheme. However, one referee has raised the concern that a government could always make side-payments to undo any incentives created by a legislated wage contract. While there are agencies of government that are designed to make such circumventions more difficult (such as the GAO in the United States), the feasibility of implementing any incentive contract is problematic. This does not reduce the usefulness of determining the optimal incentive structure. In Walsh (1993b), a simple dismissal rule (i.e., fire the central banker if inflation exceeds a critical value) is shown to sustain the optimal policy when the government's ability to offer a state-contingent wage contract is limited. A dismissal threat is credible when the government can choose from a supply of identical potential central bankers.

resemble an inflation-targeting rule, and Rogoff (1985), Canzoneri (1985), and Garfinkle and Oh (1993), have suggested that central banks could be punished for deviating from legislative targeting rules. Walsh (1993b) demonstrates that, in some circumstances, a dismissal rule can substitute for a state-contingent wage contract. Legislated budget procedures, targeting rules, or publicly announced conditions under which the central banker will be fired can be costly to change.[6] Thus, while the formal analysis treats contracts in a fairly narrow sense, a broader interpretation is often possible.

The existing literature on monetary-policy design has not adopted as its starting point the consideration of optimal incentive structures. Instead, the traditional approach has normally analyzed the effects on policy of ad hoc incentive structures, such as targeting rules, imposed on the central bank. As a result, it is never clear to what extent the economic outcomes and policy trade-offs that arise are inherent in the policy problem or simply arise from the suboptimality of the targeting rules chosen for analysis. For example, by deriving the optimal contract, I show that the trade-off between inflationary bias and suboptimal stabilization in Rogoff (1985), Garfinkle and Oh (1993), and others arises because these authors place arbitrary restrictions on the targeting rules they consider. This trade-off disappears under the optimal contract; full credibility and flexibility are simultaneously achieved. By determining whether trade-offs arise from the nature of the model or are artifacts of the targeting rules, the optimal contract provides insight into how the targeting rules might be respecified or how the limitations of the model might appear to make the policy choice too easy by neglecting important aspects of actual policy problems.

In this paper, I consider the optimal contract for the policymaker to offer the central banker in the one-period model that has formed the basis for much of the literature on discretionary monetary policy. It is shown that a contract exists that eliminates the inflation bias of discretionary policy while still ensuring that inflation responds optimally to aggregate supply shocks. In the standard strategic-policy model of Barro and Gordon (1983), a contract that makes the payment to the central banker depend only on inflation eliminates the inflationary bias and leads to an optimal policy in the presence of aggregate demand and supply shocks if the central banker also shares the government's preferences over output and inflation. The intuition is straightforward. In the standard monetary-policy game, the inflationary bias is constant across states of nature; a contract therefore only needs to raise the marginal cost of inflation to the central banker by a constant amount, leaving the central banker free to respond with discretion to economic disturbances. Such a contract can be interpreted as a type of inflation-targeting rule; such rules are normally thought to be suboptimal in the presence of supply shocks.

The impact of imperfect information on the optimal contract is also examined. Canzoneri (1985) has argued that reputational solutions to the time-inconsistency problem will be difficult to sustain if the central bank has private information. I show that central-bank private information of the type considered by Canzoneri is, in fact, irrelevant. Optimal stabilization policy with no inflationary bias can be achieved by a contract that depends only on publicly observable variables. The optimal contract is simple in the sense that it depends only on the actual rate of money growth or inflation. Canzoneri (1985) and Garfinkle and Oh (1993) have emphasized that the presence of information that is private to the central bank makes it difficult to determine whether the central bank is cheating or deviating from the optimal policy, and this leads them to consider various forms of targeting rules. These problems are avoided under the optimal contract.

Optimal contracts are derived when the central bank has private information on a

[6] Under New Zealand's 1989 central-banking law, changes in the target rate of inflation must be published and tabled in Parliament.

forecast of an aggregate supply shock, when the government can only observe a broad monetary aggregate that reflects both central-bank actions and control errors, and when central bankers differ in their ability and can choose among alternative operating procedures that affect the informational content of observable signals of economic disturbances. In all cases, contracts exist that eliminate average inflation while still achieving an optimal policy response to shocks.

I. A Contract Solution to Inflation Bias and Suboptimal Stabilization

Rogoff (1985), Lohmann (1992), and Waller (1992) have studied models of discretionary monetary policy in which the inflationary bias is reduced by appointing as head of the central bank an individual who dislikes inflation more than society does as a whole. From society's point of view, the cost of giving control of monetary policy to such a "conservative" central banker is the lower priority placed on employment or output stabilization by the central bank. In the face of an adverse aggregate supply shock, for instance, the central banker is willing to trade off too much employment reduction to prevent inflation from rising than would be optimal from the perspective of the social welfare function.

Instead of viewing the central bank as headed by a "conservative" whose welfare function differs from that of society, one can interpret Rogoff's (1985) framework as modeling the situation in which the central bank maximizes its budgetary transfer from the government, where this transfer depends on inflation and employment. The problem for the government is to set the optimal weight on inflation relative to employment in this transfer function. Rogoff, in fact, offers this interpretation of his results. In this section, I show that the conflict that arises between inflationary bias and suboptimal stabilization results from the restrictive set of central-bank contracts previously considered.

Rogoff (1985) and Lohmann (1992) carry out their analysis using a framework that is by now standard in this literature. Suppose the government's objective is to minimize a quadratic cost function that depends on the rate of inflation, π, and deviations of real output around a target level, $y - y^*$:

$$V = (y - y^*)^2 + \beta\pi^2. \tag{1}$$

Since I am focusing on the principal–agent relationship between the government and the central banker, and not on that between the public and the government, equation (1) will be viewed interchangeably as reflecting both the government's and society's preferences.[7]

The presence of nominal contracts set at the start of the period leads to an aggregate relationship between output and unexpected inflation of the form

$$y = y^c + \alpha(\pi - \pi^e) + \varepsilon \tag{2}$$

where ε is a mean-zero, serially uncorrelated real aggregate supply shock, π^e is the public's expectation of inflation, and y^c is the equilibrium level of output in the absence of supply shocks or unanticipated inflation. It is assumed that expectations are formed (i.e., nominal wage contracts signed) before ε can be observed but that the monetary authority can set its policy instrument after observing a signal θ about ε. The signal will be taken to be the private information of the central bank and is equal to ε plus a measurement error term ϕ: $\theta = \varepsilon + \phi$, where ε and ϕ are mutually uncorrelated, normally distributed, white-noise processes. The expectation of ε, conditional on observing θ is $s\theta$, where $s = \sigma_\varepsilon^2/(\sigma_\varepsilon^2 + \sigma_\phi^2)$, $0 < s \leq 1$.[8]

[7] In keeping with the Barro and Gordon framework used by Rogoff (1985) and others, it will be assumed that the social welfare function is well defined and represents the preferences of a homogeneous population. When the economy is not characterized by homogeneous preferences, Waller (1992) shows that it may not always be optimal to appoint a conservative central banker.

[8] If one prefers to think of θ as a rational forecast of ε, then s can be set equal to 1 in the subsequent analysis of this and the following section. The measurement-error specification is adopted here because in Section III I will assume that the quality of the signal is affected by the central bank's competency and its choice of operating procedure.

In order to provide an incentive for the policymaker to attempt to create inflation surprises, it must be the case that

$$(3) \qquad k \equiv y^* - y^c > 0.$$

It is standard in this type of model to treat inflation as the central bank's policy instrument. However, since I will later wish to examine contracts that attempt to influence the central bank's choice of operating procedures, it will be useful to distinguish between inflation and the central bank's policy instrument, the latter taken to be the rate of growth of a monetary aggregate directly controlled by the central bank (the monetary base, for example). Denoting this growth rate by m, the inflation rate is given by

$$(4) \qquad \pi = m + v - \gamma\varepsilon$$

where v is interpreted as either a control error or a velocity shock, and the term $\gamma\varepsilon$ allows aggregate supply shocks to have a direct negative impact on inflation. The stochastic shock v is taken to be an exogenous white-noise process whose realization occurs after m is set.

Using (1)–(4), the optimal policy that minimizes expected social loss conditional on θ and subject to the constraint that average inflation be zero is given by

$$(5) \qquad m(\theta) = \left(\gamma - \frac{\alpha}{\alpha^2 + \beta}\right) s\theta \equiv \delta s\theta.$$

The optimal response of the money supply is proportional to the forecast of ε, $s\theta$. If the direct effect of aggregate supply shocks on inflation is zero ($\gamma \equiv 0$), the factor of proportionality is negative; a positive aggregate supply shock raises output, and to stabilize output, the money supply should be reduced. Since this leads to fluctuations in the rate of inflation, the response of money is smaller, in absolute value, the larger is the weight placed on inflation stability in the government's objective function (β). If $\gamma > 0$, a positive aggregate supply shock acts directly to reduce the rate of inflation. If γ is sufficiently large, the desire to stabilize inflation around zero can produce a positive monetary response to the aggregate supply shock.

As is well known, the policy given by equation (5) is time-inconsistent and therefore not credible if implemented either directly by the government or by a central bank whose objective function is given by equation (1). Suppose, however that monetary policy is conducted by an independent central bank, one who shares the government's preferences as in (1), but who also receives a monetary transfer payment from the government. This payment could be thought of either as the direct income of the central banker or as the budget of the central bank. Alternatively, the transfer payment can be viewed more broadly as reflecting legislated performance objectives for the central bank. Let t represent the contract transfer to the central bank, and assume that the central bank's utility is given by

$$(6) \qquad U = t - V.$$

That is, the central bank cares about both the transfer it receives and the social loss generated by inflation and output fluctuations.[9] Equation (6) will arise if the central

[9] This assumption affects the form of the optimal contract but not the set of policies that can be implemented: if a transfer function $t^*(\cdot)$ implements the policy $m^*(\cdot)$ when the central bank's preferences are given by (6), then the transfer function $\tau^* = t^* - V$ implements m^* if the central bank cares only about its transfer. Most central banks seem concerned with the same macroeconomic outcomes that concern politicians, so equation (6) is the most natural assumption. One advantage of the specification in (6) is that the transfer function directly indicates how the government would like the central bank's objectives to differ from its own. In Sections II and III, it will be assumed that the central bank is interested solely in its transfer. As one referee points out, V might be huge compared to any feasible transfer to the central banker. However, if the objective of the government is expressed in terms of the utility of a representative agent, V would be a per capita measure and not the aggregate social loss. While these considerations are critical for analyzing the implementability of an optimal contract, they are not directly relevant at this stage, when the objective is to determine the optimal incentives with which the central banker should be faced.

banker's preferences are separable in social loss and income and the banker is risk-neutral. The central bank sets m to maximize the expected value of U, conditional on the realization of θ. The problem faced by the government (the principal) is to design a transfer function t that induces the central bank to choose $m = m(\theta)$, subject to the requirement that $E(t - V) \geq U_0 = 0$, where U_0 is the central banker's reservation level of utility, normalized to 0 for convenience.[10]

If the government can verify θ *ex post*, there are clearly many contracts that would achieve the desired result. For example, any contract that imposes a large penalty on the central bank if m deviates from $m(\theta)$ will ensure that $m(\theta)$ is chosen. Such knife-edge, or forcing, solutions are of little practical interest as they require that the government specify a complete contingent contract that mandates a value of m for each realization of θ. While the standard Barro-Gordon framework offers no reasons to explain why such contracts might not be feasible, it is generally agreed that the actual difficulty of determining both the possible states of nature *ex ante* and the actual realization of shocks *ex post* make such contracts infeasible to write and impossible to enforce. A major argument, in fact, for the existence of specialized institutions for the conduct of monetary policy must rest on the difficulty of specifying a complete set of rules to follow under all contingencies.

In the present model, the informational structure precludes contracts that are contingent on θ since the signal realization is assumed to be the private information of the central bank. The contract payment to the central bank must depend, therefore, only on the observable variables m, π, and y. As I will show, taking the transfer to be a function solely of either m or π is not restrictive in that the optimal policy can still be supported.

Consider then a transfer function $t(m)$ that makes the government's payment to the central bank contingent on the observed rate of money growth. The transfer function $t(m)$ implements the optimal policy $m(\theta)$ if $m(\theta)$ maximizes $E_\theta(t(m) - V)$ for all θ where E_θ denotes the central bank's expectation conditional on θ. That is, $t(m)$ implements $m(\theta)$ if $m(\theta)$ is the solution to the central banker's problem of maximizing her expected utility.

The first-order condition for the central banker's problem can be solved for $m^{\text{CB}}(\theta)$, the optimal discretionary policy:[11]

$$(7) \quad m^{\text{CB}}(\theta) = \frac{\alpha k}{\beta} + \left(\frac{1}{2\beta}\right) E_\theta(\partial t/\partial m) + \frac{\alpha^2}{2(\alpha^2 + \beta)\beta}\left[E(\partial t/\partial m) - E_\theta(\partial t/\partial m)\right] + \delta s \theta$$

where $E(\cdot)$ denotes the public's expectations. Note that in the absence of the terms involving the transfer function, m^{CB} exhibits the standard inflation bias, $\alpha k/\beta$. The last term in (7) shows that the optimal discretionary policy response to the signal θ is equal to the response under the optimal commitment policy (5). This is important since it implies that the government's objective will be to design a contract that eliminates the inflationary bias while leaving the central bank free to respond with discretion to θ. Setting $m^{\text{CB}}(\theta)$ equal to $m(\theta)$ for all θ

[10]An alternative approach would assume that the government's objective is to minimize the expected loss plus the transfer to the central banker. However, in the present context, the government will always be able to offer a minimum-cost contract. Therefore, assuming that the government minimizes $E(V)$ instead of $E(V + t)$ involves no loss of generality. In Section III, however, central bankers may potentially earn rents, so it will be important to account for the cost of the contract the government offers.

[11]It is well known in the principal–agent literature that the first-order approach adopted here may be inappropriate. However, the first-order conditions are both necessary and sufficient as long as $t(\cdot) - V$ is continuous and concave, which will be the case here. For a recent discussion of the first-order approach, see Ian Jewitt (1988).

requires that $t(m)$ satisfy

$$(8)\quad \tfrac{1}{2}E_\theta(\partial t/\partial m) + \frac{\alpha^2}{2(\alpha^2+\beta)\beta}[E(\partial t/\partial m) - E_\theta(\partial t/\partial m)] = -\alpha k \leq 0.$$

The optimal commitment policy $m(\theta)$ can be implemented, therefore, by the transfer function

$$(9)\qquad t(m) = t_0 - 2\alpha km$$

with the constant t_0 set to ensure $E(t - V) = 0$. Note that the transfer is based solely on the observed value of m, the central bank's instrument. Presenting the central banker with the incentive contract (9) achieves the dual objectives of eliminating the inflationary bias while still ensuring optimal stabilization policy in response to the central bank's private information about the aggregate supply shock.

If the government's loss function in (1) took the form $(y - y^*)^2 + \beta(\pi - \pi^*)^2$ where π^* is a target rate of inflation [implicitly equal to zero in (1)], then the transfer function would take the form

$$t = t_0 - 2\alpha k(m - m^*)$$

where $m^* = \pi^*$. In this form, it looks like a targeting rule in which the central bank is penalized for deviations of actual money growth above m^* and rewarded for money growth rates below m^*. Because the penalty is linear in money growth, it raises the marginal cost of monetary expansion by the same amount for all realizations of θ. This does not distort the central bank's response to supply shocks, but it reduces the average rate of money growth under discretion to m^*.

Optimal policy can also be implemented through the use of a transfer function of the form $t_0 - 2\alpha k\pi$ based directly on the realized rate of inflation. Even though the expected rate of inflation conditional on θ will differ from m when $\theta \neq 0$, $E_\theta d\pi/dm = 1$. Consequently, the linear transfer function based on inflation has the same effect on the marginal cost of money growth as does the function based directly on m. The same policy can be achieved regardless of whether the central bank's transfer is based on a money target or an inflation target.

The contract in (9) implements the optimal policy despite the fact that the policy depends on the central bank's private, unverifiable information on θ. It has often been suggested that reputational equilibria might serve to substitute for commitment and act to reduce the resulting inflationary bias. However, Canzoneri (1985) has argued that this would be difficult to enforce If the central bank has private information. The central bank, for example, might justify faster than expected money growth by claiming its private forecast revealed an upward shift in money demand; the public may be unable to determine whether money demand has in fact increased or the central bank is simply attempting to produce a surprise expansion.[12] If the central bank's objective is to minimize the expected social loss function (1), conditional on the observation of θ, and if the public believes that the central bank will follow the optimal rule given by (5), the central bank would actually like to set m equal to $\delta s\theta + \alpha k/(\alpha^2 + \beta)$. To implement such a policy, while attempting to maintain a reputation for following the rule in (5), the central bank will announce that θ is equal to θ^a, where $\theta^a = \theta + \alpha k/\delta s(\alpha^2 + \beta)$. In other words, the central bank has an incentive to misrepresent its observation on θ.

Under the optimal contract, however, this problem does not arise. Since the policy $m(\theta)$ can be implemented by a transfer function based solely on the realized rate of money growth, the government does not need to know θ; the fact that the central bank has private information on its forecast of ε is irrelevant. The reason for this result is straightforward. In Canzoneri's (1985) model, as in all basic versions of the generic

[12] In Canzoneri's (1985) model, the central bank has private information about a shock to money demand, as opposed to the aggregate supply considered here. The same issues, however, arise in either case.

monetary-policy model studied by Barro and Gordon (1983), Rogoff (1985), and others, the inflationary bias of policy is a constant, independent of the realization of θ. The optimal incentive contract acts to raise the marginal cost of inflation to the central bank by the same amount across all states of nature. It is thus independent of θ, and the government gains nothing from either observing θ or from inducing the central bank to reveal θ.

Equation (9) was derived under the assumption that the control error v was realized after the central bank had set m. Since v can also be interpreted as a velocity shock, it is interesting to note that the contract in (9) continues to support the optimal policy if the central bank (but not the public) can observe v before choosing m or if the central bank observes a private signal on v before setting m. The reason is that the central bank will, under discretion, optimally respond to its information about v. Again, the contract only needs to raise the marginal cost of money growth by a constant, independent of the state of nature. The central bank can then be allowed to engage in stabilization based on its private information.

While (9) implements $m(\theta)$ when the central bank cares about the social loss function V, this same outcome is equivalent to the situation in which the central banker is risk-neutral and cares only about the monetary transfer she receives from the government, and the transfer function is $t_0 - 2\alpha k\pi - (y - y^*)^2 - \beta\pi^2$. Rogoff (1985) implicity consider a government that offers the central bank a contract of the form $t' = t_0 - (y - y^*)^2 - b\pi^2$ with b chosen to minimize the government's objective function (1). Since the set of contracts of the form t' is more restrictive in that it imposes the constraint that no linear term in inflation appears, the resulting monetary policy fails to eliminate completely the inflationary bias or to achieve optimal stabilization. By parameterizing contracts solely in terms of the weight placed on the quadratic inflation term, the marginal cost of inflation can be raised to reduce the inflationary bias, but the impact on the marginal cost is made to depend on the level of the rate of inflation, unlike the contract in (9). This distorts the central bank's response to θ and leads to the trade-off emphasized by Rogoff (1985) and Lohmann (1992).[13]

In the standard monetary-policy framework, the inflationary bias of discretionary policy can be eliminated while still achieving an optimal policy response to economic disturbances through the use of a contract between the government and the central bank based only on the observed rate of money growth. The contract resembles a targeting rule. The apparent trade-off between stabilization policy and price stability suggested by previous work arises only when the contract offered to the central banker is arbitrarily restricted.

II. Contracts Based on Performance Measures

The previous section has assumed that the central bank shares the same preferences over inflation and output fluctuations as the government. If monetary policy is implemented by an agent who cares only about her transfer income from the government, the optimal contract based on *ex post* inflation and output is not unique.[14] The previous section showed that the optimal policy could be implemented if the central bank acted to maximize the expected value of $t - V$, where t was given by equation (9). If the central bank cares only about its

[13] The trade-off studied by Garfinkel and Oh (1993) also arises from the form of the targeting rules they consider. That is, the trade-off is due to the nature of the targeting rules analyzed, and not to the fundamental nature of the policy problem.

[14] When the agent cares only about her monetary transfer, it might be thought that the time-inconsistency problem would disappear, since such an agent could be paid a fixed amount to "just follow instructions" and would have no incentive to reoptimize after private agents have entered into wage contracts. However, the assumption that θ is private information to the central banker implies that the government would be unable to verify whether the instructions had in fact been followed. More importantly, this informational structure captures the notion that it is not possible to provide a complete set of contingent instructions to the central banker; some discretion is unavoidable.

expected transfer, then the transfer function $\tau(\pi, y) = t(\pi) - V$, t given by (9), will clearly lead the central bank to choose $m = \delta s\theta$. Because of the time-inconsistency problem, this differs from the standard result in the principal–agent literature with risk-neutrality in which the agent should receive $t_0 - V$, where t_0 is a constant determined by the individual rationality constraint.

The transfer function $\tau(\pi,y)$ can be thought of as a performance-based incentive mechanism; the central banker's reward is based on the realization of output and inflation. To demonstrate that τ does in fact lead to the optimal policy, consider the general quadratic performance contract based on inflation and output:

$$(10)\quad \tau(\pi, y) = b_0 + b_1\pi + b_2\pi^2 + b_3(y - y^*) + b_4(y - y^*)^2 + b_5\pi(y - y^*).$$

Using equations (2)–(4) and denoting the public's expectations of m by $E(m)$, the first-order condition for the maximization of $E_\theta\tau(\pi, y)$ can be written as

$$(11)\quad (b_1 + b_3\alpha) + 2(b_2 + \alpha^2 b_4 + \alpha b_5)(m - \gamma s\theta) + (2\alpha b_4 + b_5)[s\theta - k + \alpha E(m)] = 0.$$

Solving for m under the assumption of rational expectations yields

$$(12)\quad m^* = \frac{-(b_1 + \alpha b_3) + (2\alpha b_4 + b_5)k}{2b_2 + \alpha b_5} - \frac{(2\alpha b_4 + b_5)s\theta}{2(b_2 + \alpha^2 b_4 + \alpha b_5)} + \gamma s\theta.$$

Now, m^* will equal $m(\theta) = \delta s\theta$ if and only if the parameters of the contract jointly satisfy

$$(13)\quad \frac{\alpha b_4 + 0.5b_5}{b_2 + \alpha^2 b_4 + \alpha b_5} = \frac{\alpha}{\alpha^2 + \beta}$$

$$(14)\quad -(b_1 + \alpha b_3) + (2\alpha b_4 + b_5)k = 0.$$

It is clear that the transfer function contains redundant parameters. Since the first-order condition is unaffected by arbitrary normalization, one can set $b_4 = -1$. Setting $b_3 = b_5 \equiv 0$, the optimal values for the remaining parameters are

$$(15)\quad b_1 = 2\alpha b_4 k = -2\alpha k$$

$$(16)\quad b^2 = \beta b_4 = -\beta.$$

Thus,

$$(17)\quad \tau(\pi,y) = b_0 - 2\alpha k\pi - \beta\pi^2 - (y - y^*)^2 = t(\pi) - V$$

which verifies the claim that $t(\pi) - V$ supports the optimal policy.

Equations (13) and (14) imply that the transfer function that achieves the optimal money-supply rule is not unique. For example, a rule based on $(y - y^*)$, $(y - y^*)^2$, and $\pi(y - y^*)$ with coefficients $-2\alpha^2 k/(\alpha^2 - \beta)$, -1, and $-2\alpha\beta/(\alpha^2 - \beta)$ will also implement $m = \delta s\theta$. In all cases, the optimal performance contract must involve output; a contract expressed solely in terms of inflation fails to achieve the optimal policy.[15]

The transfer function (17) makes the central banker's payoff a function of both inflation and output. Most recent discussions of central-bank design have focused on the establishment of monetary-policy goals expressed solely in terms of the rate of inflation. For example, New Zealand's Reserve Bank Act of 1989 makes inflation control

[15] That is, stabilizing the rate of inflation in the face of aggregate supply shocks is suboptimal when preferences are given by (1), as is well known.

the sole objective of the central bank.[16] Price stability is also the only macroeconomic goal proposed for the new European Central Bank. Such contracts are socially optimal only if the central banker's utility is also made to reflect society's preference over output fluctuations, either because the central banker directly cares about V or because her contract depends on output.

Contracts based on (17) may be difficult to implement because they depend on y^*, the government's target level of real output. This may be private information to the government, and therefore it may not be feasible to condition the transfer to the central banker on y^*. Suppose then that the transfer function is based only on the realized rate of inflation and takes the form

$$\tau(\pi, y) = b_0 + b_1\pi + b_2\pi^2. \tag{18}$$

The central bank will maximize the expected value of (18) by setting $m = \gamma s\theta - b_1/2b_2$. The deflationary bias is eliminated if $b_1 = 0$. In this case, $m = \gamma s\theta \neq \delta s\theta$ unless there are no supply shocks ($\varepsilon \equiv \theta \equiv 0$), money has no effect on real output ($\alpha \equiv 0$), or the government places no weight on output stabilization ($\beta = \infty$). Otherwise, inflation-targeting rules lead to suboptimal policy (see Rogoff, 1985).

III. Imperfect Monitoring, Competency, and Effort by the Central Bank[17]

The preceding sections have ignored two aspects of monetary-policy implementation that are critically important in considering the design of the incentives the central bank should face. First, most central banks have a variety of policy instruments available to them. The Federal Reserve, for example, uses open-market operations, the discount rate, nonprice administration of the discount window, reserve requirements, and other tools to affect the economy. The multidimensionality of the set of policy instruments makes it difficult to develop unambiguous indicators of the stance of monetary policy.[18] Even if θ were observable by the government, it may be impossible for the government to specify a policy action in response to each realization of θ and then to verify that the central bank has actually implemented that policy. Second, the informational content of financial-market and macroeconomic variables can depend on the actions of the central bank. For example, the central bank's choice of operating procedure can affect the information about future policy contained in interest-rate movements (Michael Dotsey and Robert King, 1983; V. Vance Roley and Walsh, 1985). More generally, the optimal response to the signal θ will depend on the "quality" of the signal, and that can depend on the central bank's conduct of policy in the sense of its choice of operating procedure and the efficiency with which it carries out its policy.

To capture these considerations, two changes to the basic model can be made. First, in contrast to the previous two sections, assume that the signal θ is publicly observable, but the central bank's setting of its instrument is not. Specifically, the government can observe $\mu \equiv m + v$; consequently, the government cannot separate the central bank's setting of its instrument m from the random control error v. The signal θ might correspond to a variable such as a short-term interest rate; m to the central bank's policy stance reflected in open-market operations, reserve-requirement setting, and discount-rate policy; and μ to a broad monetary aggregate like M2.

Second, it is assumed that candidates to head the central bank differ in their competency and that the distribution of the measurement error ϕ is affected both by the

[16] The Act also requires the Governor of the Reserve Bank to sign an agreement with the government establishing a target rate of inflation (currently 0–2 percent) and a date on which the target rate will be achieved. Failure to meet this target can then provide grounds for the government to dismiss the Reserve Bank Governor.

[17] The model of this section draws heavily on Jean-Jacques Laffont and Jean Tirole (1986).

[18] For recent examples of the difficulty of developing indicators of monetary policy, see Ben Bernanke and Alan Blinder (1992) and Lawrence Christiano et al. (1993).

central bank's competency and by its implementation of policy. Measurement errors might be affected by the central bank's choice of operating procedures (which could affect the information content of financial variables as in Dotsey and King [1983]), by its ability to forecast, or by its management of other policy instruments (such as reserve requirements, discount-window borrowing, regulatory oversight, etc.). Improving policy implementation, while raising the quality of the signal, imposes resource costs on the central bank. These costs might arise from increased data-collection requirements, more intensive monitoring of financial-market developments, or greater staff resources devoted to forecasting.[19]

A central banker of type a who makes effort e is assumed to experience a measurement error of $\phi = (a - e)\omega$, where ω is an exogenous, mean-zero serially uncorrelated disturbance with variance σ_ω^2, $a \in [\underline{a}, \bar{a}]$, $e \in [\underline{e}, \bar{e}]$ and $\underline{a} \geq \bar{e}$ so that $a - e \geq 0$ for all a and e. Neither the central banker's competency nor her "effort" in reducing control errors will be assumed to be observable by the government. In this framework, central bankers with low a's are "better."[20] For a given signal realization θ, the optimal setting of the money supply is still given by $\delta s \theta$, but s is now a decreasing function of $a - e$:[21]

$$s = \sigma_\varepsilon^2 \Big/ \left(\sigma_\varepsilon^2 + \sigma_\phi^2\right)$$

$$= \sigma_\varepsilon^2 \Big/ \left[\sigma_\varepsilon^2 + (a-e)^2 \sigma_\omega^2\right].$$

It will be assumed that σ_ε^2 and σ_ω^2 are known to the government. The expected value of the social loss function V will depend on σ_ϕ^2; a reduction in σ_ϕ^2, that is, an improvement in the quality of the signal θ, will reduce the expected social loss: $\partial E(V)/\partial \sigma_\phi^2 \equiv V_\phi \geq 0$.

Higher effort imposes costs on the central bank. These may be pecuniary (larger staffs, more frequent data collection) or measured directly in terms of utility (greater time invested in monitoring financial markets). Costs are assumed to be quadratic in e, so that the *ex post* utility of the central banker is $t - C(e) = t - \eta(e - \underline{e})^2$. Note that I have now assumed that the central banker's utility does not depend directly on social welfare.[22] Unlike the situations considered in previous sections, the government can no longer guarantee that contract costs will be minimized. Instead, "better" central bankers will earn rents. Thus, the objective of the principal will be to offer a contract to the central banker that minimizes expected social loss plus transfer, $E(V + t)$.

A. *Perfect Information on a*

Consider first the case in which there is no uncertainty about the central banker's type: a is known. In this case, the sequence of events evolves according to the following pattern. The government offers a contract to the central banker, after which the central banker chooses an effort level e while the public forms expectations about the rate of inflation and enters into wage contracts. The signal θ on the aggregate supply shock is then observed, after which m is set. Finally, the control error v is realized, and output and actual inflation are determined.

In this environment, the government's problem is to minimize $E(V + t)$ subject to $E(m) = 0$ and the individual rationality constraint $E(t - C(e)) \geq 0$, where the expectations are taken over the joint distribution of ε, θ, and v.[23] The optimal rule for m is again given by $m(\theta) = \delta s \theta$, although s is

[19] For an analysis of the effects of costs on the choice of operating procedure, see David VanHoose (1992).

[20] That is, the expected utility of both the government and the central banker can be shown to be decreasing in a, so central bankers with low a's earn higher expected utility and are preferred by the government.

[21] That is, once a and e are set and θ is then observed, the policy problem is exactly that posed in Section I, so that $m = \delta s \theta$ is the optimal response.

[22] This assumption is not restrictive since the case analyzed is equivalent to one in which the central banker maximizes $T - V - C(e)$ for $T = t + V$.

[23] I continue to assume that the central banker's alternative utility level is normalized to 0 so that the participation constraint takes the form $E(t - C(e)) \geq 0$.

now a function of $a-e$. Substituting $m(\theta)$ into the definition of V, the first-order condition for the optimal level of effort can be written

$$(19)\quad -V_{\phi}(a-e)\sigma_{\omega}^{2}+\eta(e-\underline{e})=0.$$

Solving for the optimal e when central-bank type is known yields

$$(20)\quad e^{*}(a)=\psi a+(1-\psi)\underline{e}$$

$$\psi=\frac{V_{\phi}\sigma_{\omega}^{2}}{V_{\phi}\sigma_{\omega}^{2}+\eta}.$$

From (20) it follows that $0\leq\partial e^{*}(a)/\partial a\leq 1$, so $(a-e)$ is increasing in a. It is optimal to have central bankers with less ability exert more effort, but not enough to completely compensate for their lower ability. As a result, social welfare is decreasing in a.

B. *Imperfect Information on the Central Banker's Type*

Now suppose that neither the government nor the public can observe the central banker's type. I restrict consideration to truthful mechanisms under which the central bank reveals its type, chooses an effort level, observes the signal θ, and sets m. It receives a transfer $\tau(a,\theta,\pi)$ and utility $\tau(a,\theta,\pi)-C(e)$. Denote by $e(a)$ the effort level the government wishes to induce from a central banker of type a, and let $s(a)=\sigma_{\varepsilon}^{2}/\{\sigma_{\varepsilon}^{2}+[a-e(a)]^{2}\sigma_{\omega}^{2}\}$ be the value of s if the central banker of type a chooses $e(a)$. As before, the socially optimal value of m, conditional on θ, equals $m(a,\theta)=\delta s(a)\theta$. I restrict attention to subgame-perfect equilibria.

Let a^{a} denote the central bank's announced type and let $e(a,a^{\mathrm{a}})$ be the effort level that maximizes the expected utility of a central banker of type a who announces her type as a^{a}. In equilibrium $a^{\mathrm{a}}=a$, and $e=e(a,a^{\mathrm{a}})$ must maximize $E(\tau(a^{\mathrm{a}},\theta,\pi)-C(e))$ for a central banker of type a.

The properties of the optimal transfer function can be derived following the analysis of Laffont and Tirole (1986); details are contained in an earlier version of the present paper and are available upon request. The requirement that the transfer function must induce truthful revelation of the central banker's type imposes restrictions on the effort levels the government will be able to elicit. In particular, the first-order condition needed to ensure that $a^{\mathrm{a}}=a$ maximizes the central banker's utility can be shown to require that

$$E\big(\tau_{a}(a,\theta,\pi)-C'(e(a))e'(a)\big)$$
$$=-C'(e(a))=-2\eta[e(a)-\underline{e}]$$

while the second-order condition is satisfied if and only if $e'(a)\equiv\partial e(a)/\partial a\leq 1$. This second-order condition implies that any implementable effort function must ensure that $a-e$ is increasing in a; therefore, the measurement-error variance $(a-e)^{2}\sigma_{\omega}^{2}$ increases with a. Better central bankers (i.e., those with lower a's) produce smaller forecast errors. Otherwise, a "good" central banker could mimic a "poor" central banker at a lower disutility of effort. Since $s(a)$ is decreasing in $a-e$, "better" central bankers respond more to a given signal than do "poor" central bankers. It can also be shown that central-bank utility is decreasing in a; better central bankers gain higher expected utility.

The government's objective is to minimize expected loss plus transfer. Assuming that the government has a uniform prior over central-bank types, and using the fact that the expected transfer equals $E(\tau(a,\theta,\pi))=u+C(e)$ where u is the central banker's utility, the government's problem can be expressed as

$$(21)\quad \min E\int_{\underline{a}}^{\bar{a}}(V+u+C)\,da$$

subject to

$$E(m)=0 \qquad u\geq 0$$
$$u_{a}=-2\eta[e(a)-\underline{e}] \qquad e'(a)\leq 1.$$

The expectations are taken with respect to the distributions of ε and v.

The effort level that it is optimal for the government to induce, $e(a)$, is given as the solution to

$$(22)\quad \eta(e-\underline{e}) = V_{\phi}(a-e)\sigma_{\omega}^{2} + \eta(\underline{a}-a).$$

The level of effort that solves (22) satisfies the second-order condition for a minimum of the government's objective function. In this case, (22) implies that $e'(a) \leq 1$, so that the second-order condition for the central bank's decision problem is also satisfied.

Recall that $e^{*}(a)$ given by (20) is the optimal effort level when central-banker type is known. Equations (20) and (22) can be used to show that

$$(23)\quad e(a) - e^{*}(a) = \frac{-\eta(a-\underline{a})}{V_{\phi}\sigma_{\omega}^{2} + \eta} \leq 0.$$

Therefore, $e(a) \leq e^{*}(a)$ for all a. Less investment in improving forecasts is induced when the central banker's ability is private information. Since $e(a) \leq e^{*}(a)$, $a - e(a) \geq a - e^{*}(a)$, so for $a > \underline{a}$, s is smaller when a cannot be observed. As a result, the optimal response to a given signal is smaller because the quality of the signal is lower. However, for the best central banker, the one for whom $a = \underline{a}$, $e(\underline{a}) = e^{*}(\underline{a})$. The least efficient central banker (the $\bar{a}$ banker) receives an expected transfer of $C(e(\bar{a}))$ that just compensates her for the cost of choosing an effort level of $e(\bar{a})$. More efficient central bankers earn rents.

The solution to the government's problem is given by $\{e(a), m(a,\theta)\}$ where $e(a)$ is the solution to (22) and $m(a,\theta) = \delta s(a)\theta$ for δ defined in (5). Of more interest is the form of the transfer function $\tau(a,\theta,\pi)$ that will implement this policy. The transfer function must be such that a central banker of type a finds it optimal to choose $e(a)$ and, observing θ, to set m equal to $\delta s(a)\theta$. The Appendix proves that $\{e(a), m(a,\theta)\}$ is implemented by the transfer function given by

$$(24)\quad \tau(a,\theta,\pi) = A(a) - K(a)[\pi - \pi(a,\theta)]^{2}$$

where $\pi(a,\theta) = m(a,\theta) - \gamma s(a)\theta$. A and K depend on the central banker's announced type but not on θ, and their values are given in the Appendix.

The government can implement the effort level $e(a)$ and induce the optimal setting for the monetary aggregate, $m(\theta)$, by making the central banker's transfer payment depend on the realized rate of inflation. Equation (24) is similar to an inflation-targeting rule for the central bank of the type often analyzed, with the exception that the parameters of the rule (A and K) are chosen based on the central bank's announcement. The transfer is based on a comparison of the actual rate of inflation with $\pi(a,\theta)$, the expected rate of inflation based on the central bank's announced type and the signal on the aggregate supply shock. Equation (24) then is actually an inflation-targeting rule with a reporting requirement. The central bank is penalized for (squared) deviations of actual inflation relative to the inflation rate expected on the basis of the signal θ. Using the definition of $A(a)$ given in the Appendix, the transfer can be shown to depend on

$$-K(a)\Big\{[\pi - \pi(a,\theta)]^{2} - \gamma^{2}s(a)[a-e(a)]^{2}\sigma_{\omega}^{2} - \sigma_{v}^{2}\Big\}.$$

The actual squared deviation of π from $\pi(a,\theta)$ is compared to its expectation under the optimal policy $\gamma^{2}s(a)[a-e(a)]^{2}\sigma_{\omega}^{2} + \sigma_{v}^{2}$. It is interesting to note that the weight given to squared inflation deviations is independent of the signal θ. The parameter $K(a)$ only affects the effort level chosen by the central bank, not its implementation of stabilization policy. The value of θ does play a role, however, in determining the target rate of inflation π.

Rogoff (1985), Canzoneri (1985), and Garfinkel and Oh (1993) find that inflation targeting results in a trade-off between reducing the inflationary bias of discretionary policy and responding optimally to aggregate supply shocks. No such trade-off is faced with the transfer function in (24) even though it closely resembles an inflation-targeting rule.[24] The difference is due to the fact that in the present framework, and unlike the framework utilized by Rogoff (1985), the government is able to implement its desired money rule $\delta s\theta$ once the signal θ is observed. Thus, (24) really incorporates both inflation targeting and a contingent money rule, and the optimal policy can be implemented even when the government is unable to determine whether movements in the monetary aggregate it observes are due to movements in the central bank's instrument or to random control errors.

The optimal contract given by (24) looks like an inflation-targeting rule with a reporting requirement. Such rules are often proposed as solutions to the inflationary-bias problem associated with discretionary monetary policy. Equation (24) and the expressions for $A(a)$ and $K(a)$ given in the Appendix suggest, however, that the informational requirements of the optimal rule may be severe.

IV. Summary

The approach adopted in this paper differs from that normally used in analyzing the design of monetary policy. Here, the central bank is viewed as the agent of the government, one who attempts to maximize an objective function that depends, in part, on contingent transfers from the government. The optimal transfer function from the government's perspective has been derived for a version of the standard Barro-Gordon model of discretionary policy in which the central bank has private information and for a version in which central bankers differ in their abilities and in the effort they devote to implementing policy.

When the central bank cares about its transfer income and about social welfare, the optimal contract resembles an inflation-targeting rule, even though such rules are generally suboptimal in the presence of supply shocks. More generally, when the central bank cares only about its transfer income, the optimal contract takes the form of a contingent inflation-targeting rule in which the target is contingent on the observed signal on aggregate supply disturbances. The contracting approach suggests that targeting rules of the type often observed may serve to eliminate inflation bias while still preserving the advantages of discretion.[25]

While the focus in this paper has been normative, it is natural to ask why state-contingent wage contracts do not appear to be observed in practice if they solve the problem of time-inconsistent optimal monetary policy. Since a wage contract was used as a tool for deriving incentives, one needs to examine existing central-banking structures to see if they give rise to incentives that mimic the effects of the optimal contracts derived here. The optimal transfer function of Section III, for example, resembles an inflation-targeting rule with a reporting requirement, suggesting that such rules may generate the appropriate incentives for the central bank. Making the central banker's employment contingent on achieving prespecified inflation targets, as is the case under the New Zealand Reserve Bank Act of 1989, may serve the same purpose. An interesting line of future research will be to determine the ways in which specific institutional structures observed in practice affect the incentives facing the makers of monetary policy.

The type of state-contingent wages studied here serves as a useful means by which to determine the nature of the incentives with which the central bank should be faced. In contrast with the traditional approach in which the central bank is subject to an ad hoc targeting rule, the contrasting approach more clearly highlights the nature of the policy problems and associated trade-offs that are inherent in the environment of the

[24]Stabilization policy is optimal, given s. However, when a and e are private information, the central bank's choice leads to a value of $a-e$ that is too large and a signal that is too noisy.

[25]Persson and Tabellini (1993) build on an earlier version of this paper to extend the contracting approach to central-bank design in a number of interesting directions. See also Walsh (1993a, b).

model, as opposed to those that arise from the use of suboptimal targeting rules. By deriving the optimal contract, I have shown that the difficulties of achieving both optimal policy responses to new information and zero average inflation, particularly in the presence of private central-bank information, are not inherent in the basic framework employed in much of the recent literature on strategic monetary policy. Instead, apparent trade-offs among price stability and stabilization, flexibility, and credibility, may simply reflect the suboptimal incentives the central bank faces under ad hoc targeting rules.

APPENDIX

This appendix shows that the transfer function given by equation (24) implements the effort function $e(a)$ and the money supply rule $m(a,\theta)$ when both a and m are private information to the central bank. A and K in (24) depend on the central banker's announced type, but not on θ, and their values are given by

$$A(a) = B(a) + K(a)(\sigma_v^2 + \gamma^2 s(a)[a - e(a)]^2\sigma_\omega^2) \tag{A1}$$

$$K(a) = \frac{C'(e(a))}{2\gamma^2 s(a)[a - e(a)]\sigma_\omega^2} = \frac{\eta[e(a) - \underline{e}]}{2\gamma^2 s(a)[a - e(a)]\sigma_\omega^2} \tag{A2}$$

and

$$B(a) = \int_a^{\bar{a}} C'(e(z))[1 - e'(z)]dz \geq 0 \tag{A3}$$

so that $B(\bar{a}) = 0$; the least-efficient central banker earns no rent.[26]

The decisions faced by the central banker form a two-stage process. In the first stage, the central banker announces her type and chooses an effort level. Effort is chosen prior to obtaining any information on the current realizations of the random disturbances (θ, ε, ω, and v). In the second stage, the central banker observes the signal θ and sets the money supply. If m cannot be observed by the government, the central bank's contract cannot specify actions contingent on θ unless those actions are utility-maximizing for the central bank.

Consider the bank's problem once a^a has been announced, the value of e has been determined, and θ has been observed. The central banker's decision variable is the setting for m. The value of m solves[27]

$$\max_m E_\theta \tau(a,\theta,\pi). \tag{A4}$$

Using (24), the first-order condition is

$$-2K(a^a)E_\theta[\pi - \pi(a^a,\theta)] = 0 \tag{A5}$$

or $E_\theta \pi = \pi(a^a,\theta)$. Using equation (4) and the definition of $\pi(a^a,\theta)$, it follows that

$$m = \delta s(a^a)\theta - \gamma[s(a^a) - s(a)]\theta \tag{A6}$$

so that $m = m(a,\theta) = \delta s(a)\theta$ if the central bank announced its type truthfully.

In the first stage, the central bank of type a announces that it is of type a^a and chooses an effort level e. Consider $\{a^a, e\}$ in the concealment set $\{a^a, e(a,a^a)\}$ where $e(a,a^a) = e(a^a) + a - a^a$.[28] Using equation (24) and (A1), the central banker's expected utility from choosing a member of this set can be written as

$$B(a^a) - \gamma^2 K(a^a)\{s(a-e)^2 - s(a^a)[a^a - e(a^a)]^2\}\sigma_\omega^2 - C(e). \tag{A7}$$

[26] If the central banker's reservation level of utility is $U_0 > 0$, then $B(\bar{a}) = U_0$.

[27] Central-bank utility is $\tau - C(e)$, but since e is a fixed parameter when the decision to set m is made, this last term is dropped for convenience.

[28] If the central bank chooses to announce a^a and undertake effort $e(a,a^a)$, such a deviation cannot be detected by the government (see Laffont and Tirole, 1986).

$$\text{(A9)}\quad B'(a^{a}) - K'(a^{a})\gamma^{2}\left(s(a)(a-e)^{2} - s(a^{a})[a^{a} - e(a^{a})]^{2}\right)\sigma_{\omega}^{2}$$
$$-2K(a^{a})\gamma^{2}\{[a - e(a,a^{a})]s(a)[1 - e'(a^{a})]$$
$$-[a - e(a,a^{a})]^{2}[\partial s/\partial(a-e)][e'(a^{a}) - 1]\}\sigma_{\omega}^{2}$$
$$+2K(a^{a})\gamma^{2}\{[a^{a} - e(a^{a})]s(a^{a})[1 - e'(a^{a})]$$
$$-[a^{a} - e(a^{a})]^{2}[\partial s/\partial(a-e)][e'(a^{a}) - 1]\}\sigma_{\omega}^{2}$$
$$-C'(e(a,a^{a}))[e'(a^{a}) - 1] = 0$$

Maximizing with respect to e, the first-order condition is $2\gamma^{2}K(a^{a})s^{2}(a-e)\sigma_{\omega}^{2} - C'(e) = 0$, which implies, using the definition of $K(a^{a})$ from equation (A2),

$$\text{(A8)}\quad \frac{e(a^{a}) - \underline{e}}{s(a^{a})[a^{a} - e(a^{a})]}s(a)(a-e) = (e - \underline{e}).$$

Maximizing (A7) with respect to a^{a} and using the definition of $e(a, a^{a})$ yields equation (A9), above. However, (A9) is satisfied for $a^{a} = a$ as long as

$$B'(a) = C'(e(a))[e'(a) - 1]$$

which just defines $B(a)$. Thus, if this condition holds, (24) induces the truthful revelation of a. Substituting $a^{a} = a$ into (A8) yields

$$\text{(A10)}\quad \frac{e(a) - \underline{e}}{a - e(a)} = \frac{e - \underline{e}}{a - e}.$$

The right side of (A10) is monotonically increasing in e, while the left side is constant (given a). Therefore, $e = e(a)$ is the unique solution to (A10).[29] Since $a^{a} = a$, equation (A6) implies that $m = \delta s(a)\theta$.

[29]This would hold for any convex cost function $C(e)$.

REFERENCES

Barro, Robert J. and Gordon, David B. "A Positive Theory of Monetary Policy in a Natural Rate Model." *Journal of Political Economy*, August 1983, *91*(4), pp. 589–610.

Bernanke, Ben and Blinder, Alan. "The Federal Funds Rate and the Channels of Monetary Transmission." *American Economic Review*, September 1992, *82*(4), pp. 901–21.

Canzoneri, Matthew B. "Monetary Policy Games and the Role of Private Information." *American Economic Review*, December 1985, *75*(5), pp. 1056–70.

Christiano, Lawrence; Eichenbaum, Martin and Evans, Charles. "The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds." Unpublished manuscript presented at the Canadian Macroeconomics Study Group Meeting, Victoria, BC, November 1993.

Dotsey, Michael and King, Robert G. "Monetary Instruments and Policy Rules in a Rational Expectations Environment." *Journal of Monetary Economics*, September 1983, *12*(3), pp. 357–82.

Garfinkel, Michelle R. and Oh, Seonghwan. "Stategic Discipline in Monetary Policy with Private Information: Optimal Targeting Horizons." *American Economic Review*, March 1993, *83*(1), pp. 99–117.

Jewitt, Ian. "Justifying the First-Order Approach to Principal–Agent Problems." *Econometrica*, September 1988, *55*(5), pp. 1177–90.

Kydland, Finn E. and Prescott, Edward C. "Rules Rather Than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy*, June 1977, *85*(3), pp. 473–91.

Laffont, Jean-Jacques and Tirole, Jean. "Using Cost Observation to Regulate Firms." *Journal of Political Economy*, June 1986, *94*(3) pp. 614–41.

Lohmann, Susanne. "Optimal Commitment in Monetary Policy: Credibility versus Flexibility." *American Economic Review*, March 1992, *82*(1), pp. 273–86.

O'Flaherty, Brendan. "The Care and Handling of Monetary Authorities." *Economics and Politics*, March 1990, *2*(1), pp. 25–44.

Persson, Torsten and Tabellini, Guido. *Macroeconomic policy, credibility and politics.* Chur: Harwood, 1990.

______. "Designing Institutions for Monetary Stability." *Carnegie-Rochester Conference Series on Public Policy*, December 1993, *39*, pp. 53–84.

Rogoff, Kenneth. "The Optimal Degree of Commitment to an Intermediate Monetary Target." *Quarterly Journal of Economics*, November 1985, *100*(4), pp. 1169–89.

Roley, V. Vance and Walsh, Carl E. "Monetary Policy Regimes, Expected Inflation, and the Response of Interest Rates to Money Announcements." *Quarterly Journal of Economics*, Supplement 1985, *100*, pp. 1011–39.

VanHoose, David D. "Optimal Choice of Monetary Policy Instruments When Policymaking Is Costly." Mimeo, University of Alabama, June 1992.

Waller, Christopher J. "The Choice of a Conservative Central Banker in a Multisector Economy." *American Economic Review*, September 1992, *82*(4), pp. 1006–12.

Walsh, Carl E. "Optimal Contracts for Independent Central Banks: Private Information, Performance Measures and Reappointment." Paper prepared for the CEPR/FRBSF Conference on Macroeconomic Stabilization Policy: Lessons for the Future, Stanford, CA, March 1993a.

______. "When Should Central Bankers Be Fired?" Mimeo, University of California, Santa Cruz, September 1993b.

______. "Is New Zealand's Reserve Bank Act of 1989 an Optimal Central Bank Contract?" Mimeo, University of California, Santa Cruz, January 1994.