Seonghwan Oh and
Michelle R. Garfinkel

Seonghwan Oh, an assistant professor of economics at the University of California, Los Angeles, was a visiting scholar at the Federal Reserve Bank of St. Louis. Michelle R. Garfinkel is a senior economist at the Federal Reserve Bank of St. Louis.
Scott Leitz provided research assistance.

Strategic Considerations in Monetary Policy with Private Information: Can Secrecy Be Avoided?

THE FEDERAL RESERVE System has been criticized often for the secrecy that surrounds monetary policy. In particular, many observers have questioned the desirability of the Fed's practice of not disclosing the decisions of the Federal Open Market Committee (FOMC) immediately following its meeting. This criticism has been heightened recently by legislation introduced in the House of Representatives, proposing, among other things, that the Fed release the contents of the FOMC's directives *immediately* after each meeting rather than with a seven-week delay.[1]

The economic rationale behind this proposal is that the Fed's maintained secrecy limits the informational content of prices in financial markets and thereby detracts from the markets' ability to allocate resources efficiently. If, for example, the FOMC voted to maintain its current policy stance but subsequently added reserves to the banking system as a technical and temporary action, market participants might mistakenly interpret such an action as a fundamental change in policy. According to this view, without immediate disclosure of the FOMC's policy directive, confusion about the Fed's intentions can add to the variability of market interest rates.

Those who are skeptical of the value of this legislation argue that immediate disclosure of the FOMC's directive would complicate the implementation of monetary policy.[2] For example, the markets' response to announcements could generate large changes in interest rates that, according to this view, would be excessive and destabilizing.

[1]Lee Hamilton and Byron Dorgan, HR2735-the Federal Reserve Reform Act of 1989. See Hamilton (1989) for a brief discussion of the key changes in the structure of the Fed proposed by this legislation. As discussed by Goodfriend (1986), however, legislation proposed in this spirit is not new.

[2]See, for example, Mooney (1989), Rosenbaum (1989) and Uchitelle (1989). Also, see Goodfriend (1986) for an interesting and useful critique of the arguments made for maintained secrecy at the Fed.

In the context of a relatively simple game-theoretic model of monetary policy, in which the Central Bank would expect to be better off if it had no private information, this article shows why the Central Bank cannot reveal its private information credibly and precisely. The Central Bank might be able to reveal this information partially through imprecise or noisy announcements. From the Central Bank's perspective, however, such announcements are not costless, nor can they remove secrecy from policy perfectly. Hence, the analysis illustrates that, even if the Central Bank perceived monetary policy secrecy as undesirable, fully eliminating it might not be feasible.

STRATEGIC MONETARY POLICY: THE BASIC MODEL

To address issues of secrecy in monetary policy, it is helpful to study a model of monetary policy that specifies the objectives and constraints faced by a Central Bank. Given the particular specification, the model provides a framework for analyzing various strategies for the Central Bank and, in turn, for predicting which strategy is optimal for the Central Bank. The model, a slight variant of Canzoneri (1985), builds on a simple specification of the economy.[3] Output is given by

(1) $y_t = y^n + (p_t - w_t)$,

where y_t, p_t and w_t denote, respectively, the logarithms of output, prices and nominal wages in time t; y^n denotes the log of output that corresponds to the "natural" rate of unemployment. In this model, the natural level of output is the one that would prevail with a steady rate of inflation.

The public attempts to specify wages so as to minimize deviations of output from its natural level. Accordingly, it wants to set $w_t = p_t$. But, in this model, prices are not known at the time wages are set. Hence, wages are set to satisfy

(2) $w_t = p_t^e$,

where p_t^e denotes the public's expectation, as described below, of the log of the price level conditional on information available to the public at the beginning of period t. By combining equations 1 and 2, output can be expressed as follows:

(3) $y_t = y^n + (\pi_t - \pi_t^e)$,

where $\pi_t = p_t - p_{t-1}$ is the actual rate of inflation in time t; $\pi_t^e = p_t^e - p_{t-1}$ denotes the public's expectation of inflation.

Equation 3 captures the notion that the long-run Phillips curve, which is the relationship (trade-off) between inflation and unemployment, is vertical. On average, unemployment and, consequently, output are independent of both expected and actual inflation. In any period, however, unanticipated inflation can create a wedge between output and its natural level. Specifically, the existence of contracts that fix nominal wages for a specific period means that actual output can depart from its natural level if people underestimate or overestimate the future rate of inflation.[4] The effect of unanticipated inflation on output is only temporary. In this model, it lasts only one period. The variance of output implied by equation 3 is simply the variance of the market's inflation forecast error.

The following simple variation of the quantity theory equation describes how prices are determined in each period given monetary policy:

(4) $p_t = m_t - y^n + v_t$,

where v_t denotes an innovation to money demand and m_t denotes the log of the money supply in time t.

Taking the first-difference of equation 4 and rearranging shows how monetary policy affects inflation:

(5) $\pi_t = g_t - \delta_t$,

where $g_t = m_t - m_{t-1}$ is the growth rate of money, the Central Bank's policy instrument, and $\delta_t = v_{t-1} - v_t$ denotes a random disturbance. This disturbance, which is bounded between

[3]The model is intended only to be an illustration, not a complete characterization of the economy. Canzoneri's (1985) model resembles that of Barro and Gordon (1983) except that it provides a role for the Central Bank to react to shocks. As will be evident below, this model does not imply that the first-best policy is a constant money growth rule. Rather, it is a contingent money growth rule. See Cukierman (1986) for a helpful review of this relatively new literature on central bank behavior.

[4]That unanticipated inflation can drive output above its natural level would also be implied by the Lucas-type (1973) supply curve. The important feature of this equation—that output, on average, will be independent of inflation—assumes that the public forms expectations rationally. The assumption that the elasticity of output with respect to unanticipated inflation is equal to one is used to simplify the notation and does not affect the qualitative results discussed below except where noted.

$-D$ and $+D$, is assumed to have a zero unconditional mean and a finite, constant variance, σ_δ^2. As revealed by equation 5, the Central Bank's control over inflation is imperfect; inflation depends not only on monetary policy but on the disturbance to money demand. Thus, equation 5 implies that the public's expectation for inflation in time t equals the difference between its expectation of money growth in time t, g_t^e, and its expectation of δ_t, δ_t^e.

Secrecy arises in this model because, in contrast to the public, the Central Bank has a (nontrivial) forecast of the disturbance to money demand.[5] The Central Bank's "private" forecast, $d_t \equiv E_t\{\delta_t\}$, satisfies

(6) $\delta_t = d_t + \varepsilon_t$,

where $E_\tau\{\bullet\}$ denotes the Central Bank's expectation, based on information available to it at the beginning of period τ, before wage contracts are signed. The Central Bank's forecast error, ε_t, has an expected value of zero, a finite variance, σ_ε^2, and no correlation with the Central Bank's forecast. The assumption that the forecast is independent of the forecast error implies $\sigma_\delta^2 = \sigma_d^2 + \sigma_\varepsilon^2$, where σ_d^2 is the variance of the private forecast.

Although this forecast is made just before wages are set, the markets' expectation of δ_t equals zero without any meaningful announcements by the Central Bank. When the Central Bank does not attempt to convey its private information, $\pi_t^e = g_t^e$. The public observes δ_t *after* policy is implemented when π_t is realized. The public, however, cannot infer from that observation what the Central Bank's forecast had been. Similarly, it cannot identify the Central Bank's forecast error. (See figure 1 which summarizes the sequence of events during any period t.) Nevertheless, people understand the Central Bank's objectives as described below and its constraints subject to the unknown disturbance δ_t; they incorporate that understanding into their expectations of money growth and, accordingly, their wage specification.

Figure 1
Sequence of Events in Period t.

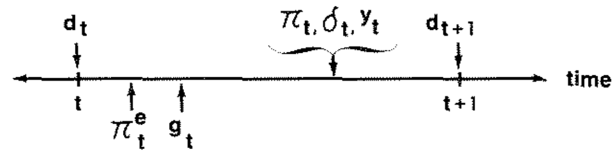

Following Canzoneri (1985), the analysis assumes that the Central Bank has two goals: output and inflation stabilization. Its expected lifetime utility in period t=1 is given by

$$\text{(7a)}\quad U_1 = \sum_{t=1}^{\infty} \beta^{t-1} E_1\{u_t\},\ 0<\beta<1$$

where

$$\text{(7b)}\quad u_t = -(y_t - y^*)^2 - f(\pi_t - \pi^*)^2,\ f>0.$$

β is the Central Bank's discount factor.[6] The parameter f is the weight the Central Bank places on its objective of stabilizing inflation around its target level, π^*, relative to its objective of hitting its target for the log of output, y^*. These targets are given and fixed parameters.

The Central Bank's inflation target need not be zero. But its objective to stabilize inflation is consistent with the public's objective to forecasting future inflation correctly. In other words, by minimizing the variability of inflation, the Central Bank minimizes the variance of the public's inflation forecast error. The Central Bank's

[5] That the public does not have a forecast of δ_t implies $\delta_t^e = 0$, providing that the Central Bank does not communicate to the public its own forecast. Note that it is not crucial that the public has no forecast of the disturbance to money demand. Provided that the Central Bank's forecast is private, the following analysis is relevant. Furthermore, the private information could be in terms of a forecast about a supply shock or the Central Bank's preferences. The qualitative results to follow would not be affected. Also, it should be noted that the present model differs from Canzoneri's (1985) model in that the timing of the forecast here is such that, if the Central Bank released this information, it could be used by the public. The assumed sequence of events, shown in figure 1, is necessary for the analysis of imprecise announcements below.

[6] Note that equation 7 implies that the Central Bank is infinitely-lived. This assumption is only important for the discussion of reputational considerations below. This discussion would be qualitatively the same if, instead, the Central Bank lived only a finite number of periods, T, provided that T is not known with certainty. In this case, β would reflect the Central Bank's chances of survival as well as its time preference. See Grossman and Van Huyck (1988), for example.

objective to stabilize output, however, is consistent with the public's objective to forecast inflation correctly *only* if the Central Bank's target for output equals the natural level. But, in this case, the interesting issues revolving around monetary policy secrecy do not arise.

As in much of this literature, then, the present analysis assumes that $y^* > y^n$. That is to say, the Central Bank prefers output to exceed the public's target. Possible interpretations of this assumption could stem from either social welfare or self-interest considerations.[7] It is only important for the present analysis that the natural output level or the public's target for output be different than the Central Bank's (given) target. This assumption implies that the Central Bank does not have enough instruments to reach its two goals, giving rise to a credibility problem in policy as illustrated below.

Using equations 3, 5 and 7b and noting that the public's expectations for inflation, π^e_t, equals $g^e_t - \delta^e_t$, the Central Bank's utility in period t can be written as

$$(8)\ u_t = -(g_t - g^e_t - \delta_t + \delta^e_t - \Delta^*)^2 - f(g_t - \delta_t - \pi^*)^2,$$

where $\Delta^* = y^* - y^n > 0$ and $\delta^e_t = 0$, without any announcements by the Central Bank about its private forecast. The Central Bank's problem is to choose g_t to maximize the expected value of its lifetime utility, after the markets set wage growth equal to expected inflation, π^e_t. The solution depends on how the Central Bank treats the markets' expectations.

The First-Best Solution

To see why the Central Bank might want to disclose its private information (that is, its forecast of the money demand disturbance), consider the benchmark case wherein the Central Bank recognizes the impact it can have on the markets' expectations and d_t is public information. Furthermore, assume that the Central Bank can make binding commitments to pursue an announced policy. In this case, it chooses g_t, subject to the restriction that expectations are consistent with its policy, to maximize its expected lifetime utility. Because of the stationary (time-independent) nature of the model, this maximization problem reduces to a sequence of one-period problems, in which the Central Bank chooses g_t to maximize its expected one-period utility, shown in equation 8, for each period t.

Given the constraint that $g_t = g^e_t$, creating surprise inflation in an effort to increase output above its natural level is precluded. Rather, the Central Bank commits itself to the following policy:

$$(9)\ \hat{g}_t = \pi^* + d_t,$$

where $\hat{g}_t = g^e_t$ for all t. Note equation 9 implies that, on average, inflation would be equal to the Central Bank's target rate. Because the policy fully accommodates the part of the disturbance to money demand predicted by the Central Bank, $\pi^e_t = \pi^*$ and wage growth is set equal to π^*.[8]

The Central Bank's expected one-period utility in this regime can be found using equations 8

[7]See Barro and Gordon (1983) and Canzoneri (1985) for a discussion of possible social-welfare interpretations of this assumption. These interpretations build on existing distortions in the economy. For example, the existence of large unions that keep real wages too high or the use of income taxes that influence labor decisions depress average output (or the natural level) below the "potential" level (or that level considered desirable from a social-welfare perspective). Although these distortions could be modeled explicitly here, the associated modifications would add unnecessary complexity to the model without providing much insight into the issues at hand. But see Cukierman (1986) for a useful critique of the social-welfare interpretation. Cukierman (1986) also provides an extensive discussion of a political interpretation. For example, although the Central Bank might be an independent institution, it might feel compelled, in order to preserve its existence or independence, to react to signals by the fiscal authority. The fiscal authority might be motivated to stimulate the economy to enhance its chances for re-election.

[8]In fact, the same outcome would be obtained if the Central Bank's forecast were not known by the public until after wages were set, so that δ^e_t still equaled zero. Because the Central Bank fully accommodates d_t, expected inflation, π^e_t, is independent of d_t in this regime. This is not to say that the Central Bank has no preferences about maintaining the privacy of its forecast. As will become obvious, the Central Bank wants to reveal its private forecast so that it can obtain this outcome. Whether the Central Bank *should* accommodate disturbances to the economy is a matter of controversy. In this model, its motive to react to d_t is compatible with the public's interests. The public prefers the Central Bank to react to its forecast, because such reactions minimize the variance of the public's forecast error. An argument against such a policy, for example, would be that it is destabilizing because the Central Bank's forecasts are inaccurate. As shown below, however, even if its private forecasts are fairly accurate (provided that $\sigma^2_\varepsilon \neq 0$), the Central Bank might not find it desirable to react to its forecast. (Given σ^2_δ, however, the more accurate the forecast, the less likely the Central Bank would be willing to sacrifice flexibility in policy.) The alternative argument against flexibility in policy in this paper builds on the Central Bank's credibility problem.

and 9, with the assumptions that $g^e_t = g_t$ and $\delta^e_t = d_t$:

(10) $E_t\{\hat{u}\} = -(1+f)\sigma^2_\epsilon - \Delta^{*2}$,

for all t, where, as defined previously, σ^2_ϵ denotes the finite variance of the Central Bank's forecast error. It is equal to the variance of inflation and output in this regime. The contingent policy in equation 9 is referred to as the first-best solution since it yields the highest utility to the Central Bank among those policies that are consistent with the public's expectations.

As demonstrated by Kydland and Prescott (1977), however, the policy in equation 9 is not "dynamically consistent." That is, given the public's expectations, the Central Bank has an incentive to deviate from the first-best policy. Specifically, given $\pi^e_t = \pi^*$, the Central Bank would rather implement the following policy:

(11) $g^{CH}_t = \pi^* + d_t + \Delta^*/(1+f)$.

If the Central Bank could create surprise inflation with the policy shown in equation 11, it could augment output above the natural level to approach its target.[9] Such a "cheating" policy would increase the Central Bank's expected one-period utility by $\Delta^{*2}/(1+f)$.

The Myopic Solution

But, even if the Central Bank could break its commitment to follow the first-best policy, cheating would be *impossible* as long as people cannot be fooled. That is, rational people will always anticipate the Central Bank's incentive to cheat, if it cannot make binding commitments.

To consider another solution, one that is more likely to emerge as the equilibrium outcome when the Central Bank has private information, suppose the Central Bank ignores any impact that it could have on the public's expectations. This is not to say that the Central Bank actually fails to understand the impact of its actions on the public's actions that, in turn, influence its own welfare. Rather, given the Central Bank's incentive to cheat, it cannot control the public's expectations directly unless it could somehow be committed to follow an announced policy and to disclose its private information truthfully. Without being able to exploit the dependence of its actions on the public's actions, the Central Bank chooses g_t to maximize its expected one-period utility, shown in equation 8, *as if* it were not trying to influence g^e_t or δ^e_t.

Before the Central Bank sets g_t, the public specifies wage growth equal to its expectations of inflation. Because the public understands the Central Bank's maximization problem, it forms g^e_t by taking an (unconditional) expectation of the Central Bank's first-order condition given by

(12) $-2(g_t - g^e_t - \delta_t + \delta^e_t - \Delta^*) - 2f(g_t - \delta_t - \pi^*) = 0$,

for each t. Even though the Central Bank observes d_t before the public forms its expectations, without any announcements, $\delta^e_t = 0$. Since the public's expectation of g_t equals g^e_t and its expectation of δ_t equals δ^e_t, $g^e_t = \pi^* + \Delta^*/f$.

People recognize the Central Bank's incentive to engineer surprise inflation so as to augment output above its natural level. To protect themselves against a decline in their real wage, then, people specify higher rates of wage growth (equal to g^e_t) than in the first-best solution with commitments. Given that specification, the Central Bank's policy, $\bar{g}_t$, which is referred to here as the "myopic" solution for reasons that will become obvious later, is given by

(13) $\bar{g}_t = \pi^* + d_t + \Delta^*/f$,

for each t. With the myopic policy, the Central Bank fully accommodates its prediction of the money demand disturbance as in the first-best solution. Further, the policy shown in equation 13 validates the public's expectations, implying an average inflation rate equal to $\pi^* + \Delta^*/f$.

When the Central Bank acts as if it were ignoring the impact that it can have on the public's expectations, the best it can do is to follow the policy shown in equation 13. This policy, however, is myopic. Because it essentially ignores the potential benefit of reducing the public's expectations for inflation, it generates an "inflationary bias" for the economy. That is,

[9]The solution in equation 11 is found by substituting $g^e_t - \delta^e_t = \pi^*_t$ into the Central Bank's expected one-period utility function and maximizing that function with respect to g_t. (See the first-order condition below in equation 12.) The Central Bank would follow the same cheating strategy if it had not announced its private information before wages were set. It should be noted that, since such cheating strategies are not consistent with the public's expectations, they are implausible equilibrium strategies and are assumed not to be observed in equilibrium.

inflation, on average, exceeds the Central Bank's target level by Δ^*/f without the benefit of increasing average output above the natural level. It is important to note that the inflationary bias would emerge even if d_t were not private information, as long as the Central Bank did not try to influence the markets' expectations.[10]

The Central Bank's expected one-period utility in this regime can be found by using equations 8 and 13 with $g_t^e = \pi^* + \Delta^*/f$ and $\delta_t^e = 0$:

(14) $E_t\{\bar{u}\} = -(1+f)\sigma_\varepsilon^2 - (1+(1/f))\Delta^{*2}$,

for all t. Because the variance of inflation and output are the same as in the first-best regime, σ_ε^2, the only difference between equations 10 and 14, Δ^{*2}/f, is the Central Bank's one-period disutility of the inflationary bias or, equivalently, the inefficiency of taking the market's expectations as given. Note that the larger Δ^* (which reflects the difference between the Central Bank's and the public's target for output) and the smaller f (the Central Bank's preference for inflation stability relative to output stability), the larger is the inflationary bias.

The inflationary bias is not easily avoided without the ability to make commitments. The problem stems from the Central Bank's incentive to create surprise inflation. This incentive to cheat, given expectations, ultimately stems from the insufficient number of instruments available to the Central Bank. In the present model, the Central Bank has two objectives with only one instrument. If it had two independent instruments, the Central Bank could achieve both of its goals simultaneously.[11] Alternatively, if the Central Bank "ignored" its goal of output stabilization or f became infinitely large, then the credibility problem would disappear and there would be no inflationary bias in equilibrium.[12] But, with an insufficient number of instruments, the Central Bank's incentive to surprise the public remains, making the first-best policy dynamically inconsistent and not credible, thereby calling into question the feasibility of the first-best solution.

REPUTATIONAL CONSIDERATIONS

If the Central Bank did not possess any private information, then a legislated rule could be imposed to force the Central Bank to follow the first-best policy. Even if it were not feasible to enforce such a rule, the Central Bank could recognize the importance of its "reputation" to eliminate or mitigate the inflationary bias.[13]

To see why its reputation could be important, suppose the Central Bank announces that it will always follow the first-best policy as shown in equation 9. Further, assume that the public always expects the Central Bank to adhere to that policy, provided that it never has cheated in the past by having deviated from the first-best policy. Through its policy actions, then, the Central Bank can maintain a reputation for not deliberately creating surprise inflation.

If, however, the Central Bank were to cheat, then people would expect the Central Bank to continue to cheat in the future. Once having lost its reputation by cheating, the Central Bank is "punished." Anticipating that the Central Bank will continue to cheat in the future because it has done so in the past, people will incorporate an inflationary bias into their wage specification. Given this specification for expectations,

[10]In this regime, as in the first-best outcome, expected inflation is independent of d_t, since d_t is fully accommodated by the myopic policy. Nevertheless, because the presence of private information makes it difficult for the Central Bank to avoid the inflationary bias, as discussed below, it would like to be able to reveal its private forecast truthfully and precisely. See, however, Cukierman and Meltzer (1986) who show that the Central Bank might prefer to maintain the secrecy of its private information when it cannot control the growth of the money perfectly. In their analysis, maintained secrecy about its changing preferences permits the Central Bank to engineer inflation surprises when desired.

[11]Actual policy and expected policy are not independent instruments provided that the public is rational and forward-looking. If it were not, however, the Central Bank would optimally announce $g_t^A = g_t - \Delta^*$, where $g_t = \pi^* + d_t$, so that the Central Bank could systematically fool the public. If the public believed that announcement, the Central Bank's expected one-period utility could increase to $-(1+f)\sigma_\varepsilon^2$.

[12]If the objective function in (7) were interpreted as a social-welfare function, then the analysis above suggests that appointing a "conservative" Central Banker (i.e., one whose concern about pursuing a goal of inflation stability exceeded that of society) would enhance social welfare. See Rogoff (1985) for a detailed discussion of this point. Indeed, this is the thrust of Representative Stephen L. Neal's recently proposed legislation to make price stability the ultimate objective of the Federal Reserve System (H.R. Res. 409). But also see Neumann (1990) who argues that strengthening the independence of the Central Bank could similarly help to avoid the credibility problem in monetary policy without explicitly imposing a goal of price stability on the Central Bank.

[13]See, for example, Barro and Gordon (1983).

the Central Bank can do no better than to follow the myopic policy shown in equation 13 once having cheated. During the "punishment," the outcome would return to the myopic solution that includes the inflationary bias, Δ^*/f.[14]

In some cases, the Central Bank's concern for its reputation can provide the same result as binding commitments when there is no private information. The critical condition is that the expected long-term gain from eliminating the inflationary bias must always exceed the expected short-term gain that could be realized by creating surprise inflation. The long-term gain is simply the present discounted disutility of the inflationary bias, $\frac{\beta}{1-\beta}\left(\frac{\Delta^{*2}}{f}\right)$. The short-term gain is the difference between the expected one-period utility if the Central Bank were to cheat and the expected one-period utility from adhering to the first-best policy, $\Delta^{*2}/(1+f)$ Note that as the Central Bank's discount factor, β, increases (that is, as it cares more about the future), the expected long-term gain from maintained reputation is more likely to exceed the short-term gain from cheating in the current period. Hence, as β increases, the Central Bank's concern for its reputation is more likely to support the first-best outcome.

Even if reputational considerations were not a perfect substitute for binding commitments to achieve the first-best outcome, they could still diminish the magnitude of the equilibrium inflationary bias. As long as the threat of punishment is sufficiently large, the Central Bank will be induced to adhere to the reputational policy that involves a smaller (if not zero) inflationary bias.[15] Hence, in the reputational equilibrium, cheating is never observed.

The presence of private information, however, greatly complicates this situation, influencing the possibilities for cheating. Specifically, because the public does not observe d_t (the Central Bank's private forecast) directly, it can never be certain that the Central Bank has actually implemented the reputational policy that depends on d_t. The public can easily verify that money growth equals the Central Bank's announced reputational policy. But the public cannot be sure that the Central Bank's announcement about d_t is truthful. Indeed, as shown below, the Central Bank has an incentive to misrepresent its private information.

WHY ARE PRECISE ANNOUNCEMENTS NOT FEASIBLE?

The existence of private information weakens the ability of reputational considerations to achieve the efficient outcome. This can be illustrated by showing that it is impossible to force the Central Bank to adhere to the first-best policy, because the Central Bank cannot make credible announcements that precisely reveal its private information.

Suppose that the Central Bank could be forced to adhere to a specified policy, but could not be forced to reveal its private information credibly and precisely.[16] For example, the following rule might be legislated:

---

[14]Making this reputational mechanism effective, in the present model, requires that the Central Bank is infinitely-lived or has a finite but uncertain lifetime, which is consistent with the Central Bank's objective function shown in (7). If the Central Bank were to live a finite and certain number of periods, T, then it would always cheat in the last period, T. But, if the public expects such behavior, the period T outcome would just be the myopic solution. Along this line of reasoning, the solution unravels and the reputational mechanism cannot diminish the inflationary bias below Δ^*/f. Alternatively, if the Central Bank were finitely lived, but its preferences were private information (e.g., the value of the parameter f), the Central Bank could "build" credibility as an inflation-fighter by signaling with monetary policy actions. See, for example, Backus and Driffill (1985).

[15]Suppose, for example, that the Central Bank announces $g_t = k + \pi^* + d_t$, where k is the average inflation in excess of the optimal rate ($0 \leq k \leq \Delta^*/f$) and d_t is public information. (Note that when $k=0$, this policy is simply the first-best one and when $k=\Delta^*/f$ the policy is the myopic one.) Provided that $k<\Delta^*/f$, the temptation for the Central Bank to cheat, given by $(\Delta^2 + f^2k^2 - 2f\Delta^*k)/(1+f)$, will be positive but will decrease as k increases. The general condition for reputational considerations to work is that this temptation be less than the expected gain to maintained reputation given by

$$\frac{\beta}{1-\beta}\left(\frac{\Delta^{*2}}{f} - fk^2\right),$$

which is also positive as long as $k<\Delta^*/f$. This gain also decreases as k increases. Even if the expected present discounted gain from maintained cooperation is smaller than the Central Bank's temptation to cheat for $k=0$, the reputational equilibrium inflationary bias, k, can be less than Δ^*/f, if the temptation decreases faster than the expected present discounted gain as $k<\Delta^*/f$ increases.

[16]That there is no separate mechanism to force the Central Bank to reveal its private information might seem puzzling. For example, in the United States, Congress or the Administration could set up an agency to monitor the Central Bank's activities and take part in formulating monetary policy, whereby the private forecasts can be revealed to the public. Why such an arrangement is not adopted is beyond the scope of this analysis.

(15) $\hat{g}_t^L = \pi^* + d_t^A$,

for all t, where d_t^A denotes the Central Bank's announcement of its private forecast. If that announcement were believed by the public, the public would form the following expectations: $g_t^e = \pi^* + d_t^A$ and $\pi_t^e = \pi^*$. With these expectations, before setting its policy in period t, the Central Bank would announce optimally

(16) $d_t^A = d_t + \Delta^*/(1+f)$.

If the public were to believe the Central Bank's announcement, the Central Bank would be able to disguise its cheating policy (shown in equation 11) as the first-best policy by overstating the value of its forecast.[17] In this case, the Central Bank could drive output above its natural level by $\Delta^*/(1+f)$.

But, as in the case of simple cheating, the Central Bank's incentive to lie, which also fundamentally stems from its incentive to create surprise inflation, will be fully recognized; as a result, no one will believe the announcement. Given that the public cannot determine with certainty whether or not $d_t^A = d_t$, it can do no better than to protect itself from surprise inflation by setting wage growth equal to $\pi^* + \Delta^*/f$.[18]

Because the Central Bank's forecast is private information, a legislative approach depending on that information is not effective in achieving a better outcome than the myopic solution.[19] Similarly, the Central Bank's private information obscures the relevance of reputational considerations to improve upon the myopic outcome. Although people can see whether the Central Bank has implemented its announced policy—for example, the policy shown in equation 15—they cannot verify that its announcement truly reflects the value of its private forecast (that is, $d_t^A = d_t$), unless the forecast were always perfect (that is, $\varepsilon_t = 0$ for all t). Hence, the public cannot evaluate the Central Bank's reputation based on past policy actions.

A CONSTANT MONEY GROWTH RULE AND THE ROLE OF NOISY ANNOUNCEMENTS

Although the Central Bank cannot make credible announcements that precisely state its private information, it can make announcements that have some informational content. In a recent study, Stein (1989) applies the work of Crawford and Sobel (1982) to show that, through noisy announcements or "cheap talk," the Central Bank can reveal its private information partially. In his application, where the Central Bank's private information concerns its objective for the target exchange rate, Stein illustrates how the Central Bank can make announcements of a range in which its target falls. Because the announcement does not state the exact value of the Central Bank's target, it is a noisy announcement. These announcements are a costless form of communication in that no resources are used in making them. But the announcements are credible because the Central Bank would incur an implicit cost if it were to lie. This cost is sufficiently large to induce the Central Bank to reveal its private information truthfully, though not precisely.

An application of Crawford and Sobel's (1982) analysis to the present model, however, shows that noisy announcements might not be as "cheap" as Stein's (1989) analysis would suggest.

[17]This can be seen by substituting equation 16 into equation 15. To verify that equation 16 is the optimal announcement, substitute equation 14 into the Central Bank's one-period utility function (8) and choose d_t^A to maximize the expected value of (8) subject to the public's expectations $g_t^e - \delta_t^e = \pi^*$. The Central Bank would lie in the same manner if it were not necessary to make its announcement until after the policy was implemented.

[18]To see this, note if the Central Bank were to act on its incentive to create surprise inflation given the public's expectations, it would set its policy optimally to satisfy the first-order condition in equation 12. Rearranging equation 12 and using $E_t(\delta_t) = d_t$, one can verify the following:

$$g_t = \frac{g_t^e - \delta_t^e}{1+f} + \frac{f\pi^* + \Delta^*}{1+f} + d_t.$$

Noting that $g_t - d_t$ equals the Central Bank's expectation for inflation given d_t, $E_t(\pi_t)$, and $g_t^e - \delta_t^e = \pi_t^e$, the expression above implies that

$$E_t(\pi_t) = \frac{\pi_t^e + f\pi^* + \Delta^*}{1+f}.$$

Since $E_t(\pi_t) > \pi_t^e$ for $\pi_t^e \leq \pi^* + \Delta^*/f$, the Central Bank always has an incentive to create surprise inflation unless the public incorporates the inflationary bias Δ^*/f into its wage specification.

[19]Garfinkel and Oh (1990a) have shown how a legislative approach that is independent of the Central Bank's private forecast can achieve a better outcome than the myopic solution studied above. With a multi-period (N periods) average targeting procedure, requiring $\sum_t^N g_t = N\pi^*$, the Central Bank can diminish the magnitude of the inflationary bias that emerges in equilibrium. This procedure is not efficient, however, in that it necessarily limits the Central Bank's flexibility to stabilize output and inflation. Nevertheless, it can permit more flexibility than a strict constant money growth rule.

In contrast to the present model, Stein's model implies that, if it were possible to force the Central Bank to reveal its exchange rate target truthfully and precisely, then the first-best outcome could be obtained. Accordingly, noisy announcements alone can easily achieve a better outcome than no announcements or complete secrecy.

The credibility problem in monetary policy in the present framework, however, is slightly more complicated. As indicated above, even if it were possible to make the Central Bank reveal its private forecast truthfully and precisely, imposing an additional restriction on policy either through a legislative rule or reputational considerations would be necessary to ensure that the Central Bank follow the first-best policy. That is, even if the public's expectations, g_t^e and π_t^e, included information about d_t, the Central Bank would have an incentive to surprise the public (according to equation 12) unless π_t^e also were to incorporate the inflationary bias, Δ^*/f.[20] But the Central Bank has no motivation to reveal d_t if it cannot reduce or eliminate the inflationary bias in doing so. Similarly, the Central Bank's incentive to create surprise inflation would not disappear if it were to make noisy announcements about its private forecast and could contaminate those announcements.

A Constant Money Growth Rule

Because of this incentive to surprise the market with inflation, limiting the degree of flexibility permitted in monetary policy is necessary to ensure that the announcements contain some information while allowing the Central Bank to avoid the inflationary bias. In other words, a rule for monetary policy must be imposed to "tie" the hands of the Central Bank. As indicated above, for this constraint to be effective, the rule must be independent of the private information.[21] For example, legislation could require

(17) $\tilde{g}_t^L = \pi^*$.

Although this constant money growth rule eliminates the inflationary bias, it precludes any (otherwise desirable) reactions to the part of money demand disturbances predicted by the Central Bank.[22] As such, this rule produces a higher variance of the public's forecast error for inflation and, hence, a higher variance of output than in both the first-best and myopic regimes.

The Central Bank's expected utility under this regime without any announcements is given by

(18) $E_t\{\tilde{u}\} = -(1+f)\sigma_d^2 - \Delta^{*2}$,

for all t. Expected utility in this regime will exceed that under the myopic regime only if $\Delta^{*2}/f > (1+f)\sigma_d^2$.

This condition underscores the Central Bank's trade-off between eliminating the inflationary bias and eliminating flexibility in monetary policy with the constant money growth rule. The larger is the inflationary bias that emerges in the myopic outcome (that is, the smaller f and/or the larger Δ^*), the more likely this condition will be satisfied. The Central Bank is less likely, however, to prefer a constant money growth rule over the myopic policy the larger the variance of the component of the money demand disturbance predicted by the Central Bank, σ_d^2, which captures the expected benefit of being able to react to d_t. Because the legislated rule in equation 17 does not permit the Central Bank to react to its forecast of the disturbance to money demand to stabilize inflation, the variance of inflation and output increase to σ_d^2. Nevertheless, if the possible benefits of maintained flexibility are not too large (that is, if σ_d^2 is small), the Central Bank might prefer to be constrained not to react to its private forecast to avoid the inflationary bias.

It is important to note that, even with this rule, the Central Bank still would not precisely reveal its forecast. In particular, given $g_t = \pi^*$, the Central Bank would want to overstate the value of its forecast according to

(19) $d_t^A = d_t + \Delta^*$.

[20] See footnote 18.

[21] Whether it is possible to enforce a legislated rule is beyond the scope of this paper. Of course, reputational considerations might be able to support the same rule. To simplify the discussion, the analysis assumes that it is possible to enforce a legislated rule that does not depend on the Central Bank's private information.

[22] If there were another shock, say, in the supply equation, and the Central Bank's information about this shock were not private, then the legislated rule could provide flexibility to react to this shock. Moreover, not all flexibility needs to be removed from policy in this model. The constant money growth rule is not the only way to tie the hands of the monetary authority to make the announcements meaningful. The imposition of a multi-period average targeting rule that permits some flexibility would also work; however, with this constraint, the inflationary bias would not be eliminated totally. See Garfinkel and Oh (1990a).

Equation 19 illustrates again that the credibility problem of monetary policy is not easily resolved in the presence of private information. But if d_t were not private information, the Central Bank's expected one-period utility with a constant money growth rule would be

$$(20)\quad E_t\{\tilde{u}\} = -f\sigma_d^2 - \Delta^{*2}.$$

Hence, the Central Bank would prefer to disclose d_t under a constant money growth rule even though it cannot do so precisely.

Noisy Announcements

By making noisy announcements about its forecast, the Central Bank could enhance its own welfare under the rule. Given that it must follow the rule in equation 17, the Central Bank cannot actively pursue its goal to stabilize inflation and output by reacting to d_t. Making noisy announcements, as an alternative policy tool, permits the Central Bank to pursue its goal of stabilizing output. Specifically, the Central Bank could partly influence expectations by announcing a range in which its forecast falls, thereby reducing the variance of the public's inflation forecast error and, in turn, reducing the variance of output.

To take a concrete example, suppose that the Central Bank announces that d_t lies either between $-D$ and a or between a and D.[23] For any announcement to contain some information about d_t, the Central Bank must perceive that lying is costly. The cost, however, cannot be directly imposed by the market upon observing δ_t because, as mentioned earlier, the market cannot infer the true value of d_t from that observation. Rather, the cost of lying about d_t is implicitly contained in how such a lie would affect the market's expectations about d_t.

Suppose the Central Bank were to announce that d_t fell in the higher range, [a,D]. Given that announcement and the money growth rule shown in equation 17, the market forms an expectation about future inflation. This expectation would equal the Central Bank's target rate of inflation, π^*, minus the expected value of d_t given that it lies somewhere between a and D. Call this conditional expectation d^h. On the other hand, if the Central Bank announced that d_t fell in the lower range, [$-$D,a], the market would expect a higher inflation rate equal to the difference between π^* and the expected value of d_t given that it falls somewhere between $-$D and a. Call this conditional expectation d^L.

If d_t is greater than $-$D but less than a, then the Central Bank's expected one-period utility by announcing $d_t \in [-D,a]$ must be greater than or equal to that by claiming $d_t \in [a,D]$ for the former announcement to be credible. That is,

$$(21)\quad -E_t\{(d^L - d_t - \varepsilon_t - \Delta^*)^2\} \geq -E_t\{(d^h - d_t - \varepsilon_t - \Delta^*)^2\}.$$

The inequality in equation 21 would be reversed if d_t were greater than a and less than D. Finally, if $d_t = a$, then the Central Bank must be indifferent between announcing the higher and lower ranges.

This last condition can be used to determine the dividing point of the distribution of d_t, a, such that for all possible values of d_t, the Central Bank's announcement is credible. The determination of the dividing point from that condition ensures that the Central Bank will not act on its motive to lie about the range in which d_t falls. For example, when d_t is in the lower range, the Central Bank will not announce that d_t is in the upper range. If it did so, the public's inflationary expectations would fall by a sufficiently large amount that, in turn, drives output too far from the Central Bank's output target and, hence, renders lying undesirable.

By making noisy announcements about its private forecast while adhering to the constant money growth rule, the Central Bank can enhance its expected utility above what it would be when it simply follows the rule. This is not to say that the Central Bank will always choose to make noisy announcements. As illustrated with a more specific example in the appendix, the Central Bank would prefer to maintain full discretion and secrecy, the more it cares about inflation stability, the less the difference between its and the public's output goals, and the more accurate the private forecast.[24]

The basic intuition here is essentially the same as that used when discussing the merits of a

[23]See the appendix for a more detailed example. Also see Garfinkel and Oh (1990b).

[24]Again, see Garfinkel and Oh (1990b). Their analysis produces a somewhat surprising result: under the conditions that noisy announcements are more likely to be preferred by the Central Bank, the credibility problem is more severe so that, at the same time, these announcements cannot be particularly informative.

simple constant money growth rule over those of the myopic policy. The presence of private information forces the Central Bank to face a new trade-off between removing the inflationary bias and limiting flexibility in policy. But the money growth rule with noisy announcements is more likely to dominate the myopic policy than the rule by itself. Although both output and inflation will have a greater variance in the regime with noisy announcements than in the myopic regime, the variance of output will be smaller in this regime than when the Central Bank simply follows a constant money growth rule. The elimination of the inflationary bias possible with the constant money growth rule, combined with the slight reduction in the variance of output possible with noisy announcements, provide the main benefits that would make abandoning the myopic policy—that is, maintaining complete secrecy with full discretion—desirable from the Central Bank's perspective.

CONCLUDING REMARKS

This article has examined the possibility of fully or at least partially removing secrecy in monetary policy. In the context of a model in which the Central Bank has an incentive to create surprise inflation, the Central Bank would like to reveal its private information, whereby it could easily avoid an inflationary bias. The Central Bank's private information combined with its incentive to surprise individuals gives rise to a credibility problem in monetary policy that is nearly impossible to resolve. Neither reputational considerations nor binding commitments to force the Central Bank to adhere to the first-best policy are effective in improving upon the myopic solution if the public never directly observes the Central Bank's private information.

Although the Central Bank cannot make precise announcements, it can make announcements that partially reveal its private information. By announcing a range in which its forecast falls and adhering to the constant money growth rule, the Central Bank can avoid the inflationary bias and influence the market's expectations in a discrete way to lower output variability below that generated by a simple constant money growth rule alone. Nevertheless, some secrecy remains.

Moreover, the Central Bank might prefer to maintain complete secrecy. Unlike Stein's (1989) result that there is always room for improvement with noisy announcements, in the context of the more general model developed here, noisy announcements require constraints on flexibility that can be permitted in the conduct of monetary policy—for example, a legislated constant money growth rule. The constraints are costly if they preclude desirable reactions to disturbances in the economy.

More generally, the analysis suggests that legislation requiring the Fed to disclose the FOMC's decisions immediately after its meeting might be of little value. If the Central Bank has private information about the economy that influences its decisions and has an incentive to surprise the public, it will not release this information truthfully and precisely. The Central Bank's incentive to misrepresent its private information detracts from the value of any information it releases.

That noisy announcements can work in enhancing the efficiency of monetary policy only under restrictive conditions prompts a general but more fundamental conclusion. In the presence of private information, the Central Bank faces a trade-off between higher-than-desired average inflation and limited flexibility. Without eliminating the ultimate source of the credibility problem—namely, that the Central Bank has too few tools to achieve its ultimate goals—this consequence of the strategic considerations of monetary policy is not easily avoided.

REFERENCES

Backus, David, and John Driffill. "Inflation and Reputation," *American Economic Review* (June 1985), pp. 530-38.

Barro, Robert J., and David B. Gordon. "Rules, Discretion and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics* (July 1983), pp. 101-21.

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

Crawford, Vincent P., and Joel Sobel. "Strategic Information Transmission," *Econometrica* (November 1982), pp. 1431-51.

Cukierman, Alex. "Central Bank Behavior and Credibility: Some Recent Theoretical Developments," this *Review* (May 1986), pp. 5-17.

Cukierman, Alex, and Allan H. Meltzer. "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information," *Econometrica* (September 1986), pp. 1099-128.

Garfinkel, Michelle R., and Seonghwan Oh. "Strategic Discipline in Monetary Policy with Private Information: Optimal Targeting Periods," UCLA Working Paper No. 584 (January 1990a).

________. "When and How Much To Talk: Credibility and Flexibility in Monetary Policy With Private Information," Federal Reserve Bank of St. Louis Working Paper No. 90-004 (June 1990b).

Goodfriend, Marvin. "Monetary Mystique: Secrecy and Central Banking," *Journal of Monetary Economics* (January 1986), pp. 63-92.

Grossman, Herschel I., and John B. Van Huyck. "Sovereign Debt as a Contingent Claim: Excusable Default, Repudiation, and Reputation," *American Economic Review* (December 1988), pp. 1088-97.

Hamilton, Lee H. "Regulating the Federal Reserve Board," *Christian Science Monitor*, October 20, 1989.

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

Lucas, Robert E., Jr. "Some International Evidence on Output-Inflation Tradeoffs," *American Economic Review* (June 1973), pp. 326-34.

Mooney, Richard E. "Don't Fiddle with the Fed," *New York Times*, October 3, 1989.

Neumann, Manfred J.M. "Precommittment to Rules in Monetary Policy: A Comment," in Michael T. Belongia, ed., *Monetary Policy on the 75th Anniversary of the Federal Reserve System*, Proceedings of the Fourteenth Annual Economic Policy Conference of the Federal Reserve Bank of St. Louis (Kluwer Academic Publishers, forthcoming).

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

Rosenbaum, David E. "Little Chance Seen for Bills on Fed Rein," *New York Times*, October 10, 1989.

Stein, Jeremy C. "Cheap Talk and the Fed: A Theory of Imprecise Policy Announcements," *American Economic Review* (March 1989), pp. 32-42.

Uchitelle, Louis. "Moves On in Congress to Lift Secrecy at the Federal Reserve," *New York Times*, August 24, 1989.

Appendix
Why Not Complete Secrecy?—An Example

This appendix illustrates with a simple example how noisy announcements work and under what conditions the Central Bank would prefer to use them rather than not reveal anything about its private forecast. For simplicity in what follows, suppose that δ_t has a uniform distribution bounded by $-D$ and D.[1] Consider the simplest example where there is only one dividing point, a_1, over that distribution.[2] Then, given an announcement by the Central Bank, say, that d_t falls in the lower range, $[-D,a_1]$, and (17) in the main text, the public will form expectations according to the following:

$$(A1)\ \pi^e(-D,\ a_1) = \pi^* - \frac{-D+a_1}{2}.$$

With this influence on the public's expectations, it is important to ensure that the Central Bank will announce the correct range. For example, if $d_t \varepsilon [-D,a_1]$, the Central Bank should not announce $d_t \varepsilon [a_1,D]$. To guarantee that the Central Bank will not misrepresent the range in which d_t falls, it must always be indifferent between announcing the ranges, $[-D,a_1]$ and $[a_1,D]$ when $d_t = a_1$.

Formally, this condition, called the "arbitrage condition," is written as

$$(A2)\ E_t\{\tilde{u}([-D,a_1],\ d_t)\} = E_t\{\tilde{u}([a_1,D],\ d_t)\},$$

or equivalently,

$$-E_t\{(\frac{-D+a_1}{2} - a_1 - \varepsilon_t - \Delta^*)^2\} =$$

$$-E_t\{(\frac{D+a_1}{2} - a_1 - \varepsilon_t - \Delta^*)^2\},$$

where $d_t = a_1$. For this condition to be satisfied, a_1 must equal $-2\Delta^*$.

The basic idea here is that, given that the Central Bank must follow the constant money growth rule, its incentive to lie depends on its

[1]Hence, the probability that $\delta_t = \bar{\delta}$, where $\bar{\delta}$ is any possible realization of δ_t, is the same for any value of $\bar{\delta}$: 1/2D. The distributions of d_t and ε_t are not specified here. They need only be independent random variables with zero means and finite variances that sum to the variance of δ_t. See Crawford and Sobel (1982) for a more general analysis of the noisy announcement equilibrium.

[2]See Garfinkel and Oh (1990b) for a derivation of a more general noisy announcement equilibrium of size n in this framework. (In this particular example, with $n = 2$, $a_0 = -D$ and $a_2 = D$.)

forecast. The dividing point a_1, determined from the arbitrage condition, implies that if the Central Bank were to overstate the value of its forecast, when $d_t < a_1$, it would have to do so by an amount so large that it is too costly to lie.

Note that the dividing point is such that for $d_t < a_1$ the announcement is more precise—that is, informative. More generally, when there are n steps, the subintervals become longer as they move away from the lower bound. For example, consider when there are two dividing points, a_1 and a_2. In this case, the arbitrage condition requires $a_1 = -D/3 - 4\Delta^*$ and $a_2 = D/3 - 4\Delta^*$. The length of the first interval from $-D$ to a_1 equals $2D/3 - 4\Delta^*$; the length of the next interval equals $2D/3$; and, the length of the last interval equals $2D/3 + 4\Delta^*$. When the disturbance is smaller (closer to $-D$), the Central Bank's incentive to overstate the value of the forecast is smaller.

Although a constant money growth rule is not first-best in that it does not permit (otherwise desirable) reactions to the Central Bank's forecasts of money demand disturbances, it does eliminate the inflationary bias. When the Central Bank also makes noisy announcements, it can enhance its expected welfare above that with a simple constant money growth rule. With only one dividing point, its expected one-period utility is given by:

$$\text{(A3)}\quad \tilde{u} = -(6\,\Delta^{*2} + D^2(f + \tfrac{1}{4}))/3,$$

which is always greater than the Central Bank's utility when it simply follows a constant money growth rule, provided that $\Delta^* < D/2$.[3] Note that this condition will be satisfied by the requirement that $a_1 > -D$. More generally, noisy announcements with any number of dividing points (greater than or equal to 1) will always be better than a simple constant money growth rule provided that the first step is greater than $-D$.[4]

In addition, the Central Bank's utility under this regime can be greater than that under the myopic regime. In the present example, this condition is given by

$$\text{(A4)}\quad D^2(\alpha^2(1+4f) - 3(1-\alpha))/12 < \Delta^2(1-f)/f,$$

where $\alpha^2 = \sigma_d^2/\sigma_\delta^2$ with $0 < \alpha < 1$.[5] The parameter α captures the degree of accuracy of the Central Bank's forecast. As α approaches 1, the Central Bank's forecast is generally more accurate.

The condition in (A4) is weaker than that for a strict rule to dominate the myopic policy. Nevertheless, this condition is quite strong, reflecting the idea that, although the inflationary bias can be avoided, the resulting loss of flexibility in this regime can be costly. In fact, when the monetary authority's forecast is extremely accurate (that is, α approaches 1), a sufficient condition for the myopic policy to dominate the constant money growth rule with noisy announcements and one dividing point, a_1, is simply that $f > 1$. If the Central Bank cares more about inflation stability than about output stability (and its forecast is extremely accurate), then it will not prefer noisy announcements, with a_1, over the myopic policy. When α is close to 1, it can be shown that, given that $f > 1$, noisy announcements with any number of partitions will not be desired by the Central Bank.[6]

Nevertheless, noisy announcements might enhance the Central Bank's utility if $f < 1$. Even if a strict constant money growth rule without

[3]See Garfinkel and Oh (1990b), who show that, for a general noisy announcement equilibrium of size n, the Central Bank's expected one-period utility is given by $\tilde{u}_t = -(\Delta^{*2}(n^2+2) + D^2(f + 1/n^2))/3$. Under the specifications for the distribution of δ_t, the one-period expected utility for the Central Bank is $-(1+f)\,D^2/3 - \Delta^{*2}$ when it follows a simple constant money growth rule. This can be easily verified by either using the above expression for expected utility with $n = 1$ or by using equation 18 and noting that the variance of a random variable which has a uniform distribution bounded by x_1 and x_2 is given by $(x_2 - x_1)^2/12$.

[4]This no-nonsense condition is automatically satisfied by the requirement that the partition equilibrium of size n is feasible. See Garfinkel and Oh (1990b).

[5]In the myopic regime, the Central Bank's one-period expected utility is $-(1+f)(1-\alpha^2)D^2/3 - (1+(1/f))\Delta^{*2}$ since, by the definition of α; $1 - \alpha^2 = \sigma_\epsilon^2/\sigma_\delta^2$.

[6]See Garfinkel and Oh (1990b). The intuition here, as discussed in the main text, follows simply from the tradeoff between the benefits of reducing the inflationary bias and the benefits of maintained flexibility. Assuming that α is sufficiently close to 1, the larger is f, the smaller is the inflationary bias that emerges in the myopic regime and the smaller is the benefit of avoiding the inefficiency of that bias relative to the expected costs of not reacting to money demand disturbances. In the case that the elasticity of output with respect to unanticipated inflation were not equal to 1, the sufficient condition for the Central Bank to prefer the myopic policy is that f be greater than the square of that elasticity. The smaller that elasticity, the greater the likelihood of the Fed preferring the myopic policy. For example, if the elasticity were equal to 1/2, then $f > 1/4$ would imply that the myopic policy dominates the constant money growth rule with noisy announcements.

any announcements does not dominate the myopic policy, there can be room for improvement with noisy announcements and the strict rule. In the case of one dividing point, there can be room for improvement provided that $f < 1/2$ even when the private forecasts are extremely accurate (that is, α is close to one). More generally, the condition in equation A4 implies that noisy announcements are more likely to be preferred over no announcements with full flexibility in monetary policy, the less accurate the Central Bank's forecast (when there is a smaller desire for flexibility in monetary policy). Further, the larger the difference between the output goals of the Central Bank and the public and the smaller the Central Bank's relative preferences for inflation stability, the Central Bank is less likely to prefer complete secrecy over noisy announcements.