Competence Implies Credibility

By GIUSEPPE MOSCARINI*

The (reputation for) competence of a central bank at doing its job makes monetary policy under discretion credible and transparent. Based on its reading of the state of the economy, the central bank announces its policy intentions to the public in a cheap-talk game. The precision of its private signal measures its competence. The fineness of the equilibrium message space measures its credibility and transparency. This is increasing in the competence/inflation bias ratio: the public expects a competent central bank to use its discretion more to pursue its "objective" targets than to surprise expectations and stimulate output. (*JEL* E52, E58)

In the 1990s the Chairman of the Board of Governors of the Federal Reserve System, Alan Greenspan, acquired a near-magical reputation for competence in his conduct of monetary policy. At the same time, US monetary policy became increasingly transparent. In his public briefings and testimony before Congress, Greenspan consistently described the Fed's view of the macroeconomic outlook and his own approach to monetary policy. Minutes of the Federal Open Market Committee's meetings were published with decreasing delay. Investors became quite familiar with Greenspan's reasoning (transparency) and trusted his judgement (competence).[1]

The increasing transparency in the Fed's conduct of monetary policy is part, and certainly a cause, of a worldwide trend (Petra Geraats 2002; Ben Bernanke 2004). Central bankers surveyed by Alan S. Blinder (1999) and economists surveyed by Sandra Waller and Jakob de Haan (2004) answered that credibility is the main asset of a central bank, and the best way to build a reputation for credibility is a track record of honesty, namely "matching deeds to words," and transparency. The European Central Bank (ECB) is no exception to this trend, but its communication with the private sector appeared, at least early on, less effective and credible than the Greenspan Fed. The ECB does not enjoy such a great reputation for competence, due both to its shorter track record and to the more complex task it faces in a newly minted monetary union of heterogeneous economies, lacking fiscal policy coordination. Indeed, the competence of a central bank (CB) should be judged relative to the complexity of the task it faces.

In this paper we argue that the positive correlation, both over time and across central banks, between reputation for competence and transparency in the conduct of monetary policy is not a coincidence. Simply put, (a reputation for) *competence implies credibility and transparency.* A CB that commands the respect of

* PO Box 208268, New Haven, CT 06520-8268 (e-mail: giuseppe.moscarini@.yale.edu; www.econ.yale.edu/faculty1/moscarini.htm). I thank Stephen Morris and Marco Ottaviani for extended and insightful conversations on this topic, Roland Bénabou and William Brainard for specific and very useful comments, audiences at Ente Einaudi, 2003 SED Meetings Paris, Yale University, University of Pennsylvania, ITAM, University of Michigan, UCLA, Columbia University, and Johns Hopkins University, as well as two anonymous referees for many useful suggestions. I am also grateful to the Alfred P. Sloan Foundation and to Yale University for financial support, and to the Ente per gli studi monetari, bancari e finanziari "Luigi Einaudi" for hospitality.

[1] Despite Greenspan's reputation for cryptic wording, during his tenure the Fed conveyed an increasing flow of information. As an example of Fedpeak, in August 2005 the Fed raised the target Federal Funds rate by 25 basis points and wrote: "[...] longer-term inflation expectations remain well contained [...]." A month later, after hurricane Katrina hit and energy prices spiked, the Fed reissued the statement, without the qualifier "well." This apparently minor change of wording was unanimously interpreted as a clearly hawkish signal for the months to come. Subsequent monetary tightening confirmed this interpretation.

the private sector (P) for clarity of vision can more credibly communicate its intentions and motivate its actions vis-a-vis P. Competence ameliorates the credibility problem that originates from an inflation bias, or reputation thereof, through a purely strategic channel.

The intuition behind this argument is as follows. Given P's inflation expectations, CB faces a trade-off between two partially conflicting objectives: choosing the most appropriate policy, given its reading of the state of the economy, and surprising expectations by an appropriate amount to stimulate output.[2] Flexibility and discretion are beneficial to the former objective, but detrimental to the latter. A competent CB trusts its own judgement about what needs to be done, given the state of the economy. So, it uses its discretion more to set policy appropriate to the circumstances than to surprise inflation expectations and to stimulate output. Therefore, it is less inclined to lie. This fact triggers a strategic interaction that leads to credibility *in equilibrium.* Since P believes CB more, CB can manage expectations more effectively, and can tell smaller lies to stimulate output. That is, the more credible it is, the less it needs to exaggerate its claims. In turn, this reinforces its credibility, and so on and so forth. For a central banker, this gain in credibility can be a powerful motive, beyond career concerns à la Bengt Holmström (1982), to acquire and to maintain a reputation for competence.

Following Matthew Canzoneri (1985), we extend a static version of Robert Barro and David Gordon's (1983) classic monetary policy game to allow for CB's private information about the state of the economy. Our main innovation is to make this CB's observation—or "judgment," in Lars Svensson's (2003) terminology—of the state of the economy noisy, and to study the effect of the precision of this noise, that we refer to as CB's *competence,* on policy communication. The more competent CB, the clearer an idea it has about what to do. After observing its private signal, CB sends to P a message about the required inflation target in a "cheap talk" Bayesian game, in the style of Vincent P. Crawford and Joel Sobel's (1982) classic model of partisan advice. P then formulates its rational expectations of inflation (and accordingly sets some prices) based on the information it can extract from CB's announcement. Finally, given P's expectations, CB takes a monetary policy action that ultimately determines the rates of inflation and output. The fineness of the equilibrium message space is a measure of monetary policy's *transparency:* without credible precommunication, the motives for the final monetary policy action remain unexplained and unverifiable to P, thus opaque and open to different interpretations. Thus, we provide one of the first formal definitions of transparency of the motives of monetary policy.

We obtain the following formal results. If and only if CB has an inflation bias, in equilibrium truth-telling is impossible and communication is coarse. CB observes signals of the optimal inflation target from a continuum, but can credibly announce many of them only finitely. Among multiple equilibria of the communication game, one with two messages always exists, where CB can announce only the qualitative direction of monetary policy. In the equilibrium with the most communication, both the fineness of the equilibrium message space and P's ability to forecast CB's future actions increase with competence. That is, *competence implies credibility.* This central result derives from our assumption of conflicting objectives in CB's preferences. In fact, we show that the effect of competence on credibility vanishes with the weight that CB's preferences place on the appropriate monetary policy, the limiting case being that of a CB-sender who has—just like in the classic Crawford and Sobel game—only one biased objective, in this case to surprise expectations of inflation.

The strategic impact of competence on communication can be decomposed in two separate

[2] Albeit central to a vast academic literature, time inconsistency has received an increasingly cold reception by central bankers. Blinder (1998) states that, during his stint at the Fed's Board, the issue of surprising expectations never appeared on the agenda. Similarly, Otmar Issing, as an Executive Board Member of the ECB, has argued that the ECB is not in the business of surprising expectations. These statements, however, do not diminish in any way the importance of time inconsistency, which truly concerns a counterfactual: if for some reason P were to expect a (say) zero inflation rate, would CB want to choose a monetary stimulus that accelerates growth and also generates positive inflation? We contend that the answer is likely to be: Yes! The point is that, *because of this,* inflation is rarely expected to be zero, as rational expectations cannot be fooled (systematically), and there is indeed no point in trying to create further surprises. Therefore, we maintain that time inconsistency from inflation bias remains a significant impediment to the credibility of at least some types of statements by CBs.

effects. The first is the *power of the words* that CB trades off against its inflation bias. Formally, CB's *bias/competence ratio* constrains the equilibrium message space. The more precise is CB's information about the state of the economy, the more weight it puts on doing something about it, relative to surprising P's expectations; the more weight (relative to prior beliefs) P puts on CB's announcement in forming these expectations, the less CB needs to lie to surprise expectations and to stimulate output, and the more credible its announcements. The second effect is the *credibility of likely announcements.* A more competent CB is relatively less likely to observe and to announce the need for extreme monetary policy measures. Thus, it is less credible in the rare instances when extreme measures are in fact called for, but more credible in fine-tuning the frequent small deviations of the inflation target from its long-run mean.

We present an algorithm to compute all equilibria of the communication game, and we provide some numerical examples with realistic parameter values. Technically, Crawford and Sobel's (1982) algorithm does not apply because they confine their analysis to a bounded set of states of nature, while we rely on Gaussian distributions to obtain a natural parameterization of competence. Nonetheless, in our unbounded Gaussian model, also, the equilibrium message space is always finite.

Finally, we address some extensions. Alternative sources of uncertainty yield the same result and intuition. Repeated play of the same game generates a larger equilibrium set. We study sequential equilibria supported by grim trigger: if CB ever lies, P never heeds its announcements again. We show that a sufficiently competent, patient, and unbiased CB can sustain truth-telling in equilibrium. In this sense, competence still implies credibility. Otherwise, as in the static case, equilibria still take the form of a partition that can be constructed from a known algorithm.

This discussion adds a new dimension to the endless debate on rules versus discretion in monetary policy. The world is too complex to be dealt with by a single rule. Commitment to a rule has a cost in terms of flexibility in addressing each new situation with an appropriate response. This loss of flexibility is behind Svensson's (2003) stance against the adoption of Taylor-like *instrument* rules, and in favor of explicit targets. But a more competent CB knows better how to use the freedom afforded by discretion. Thus it has a stronger incentive, and a more credible reason in P's eyes, to reject binding rules. The beneficial effect of CB's competence on flexibility adds to that on transparency, to work against any kind of commitment, and in favor of case-by-case evaluation, followed by a transparent explanation to P of the outcome of the analysis of the state of the economy. Indeed, the Greenspan Fed never adopted any explicit rule.

This principle applies to other monetary policy biases, such as the "stabilization bias" (Richard Clarida, Jordi Galí, and Mark Gertler 1999), and reasons to hide information, such as the "CNBC effect" of monetary policy announcements (Stephen Morris and Hyun Shin 2002), as well as to all kinds of public announcements by policymakers, such as electoral campaign promises. Often, and justifiably so, public opinion discounts policy statements—such as claims on the effects of a proposed tax cut—as possibly motivated by some kind of political (preference) bias. At the same time, people heed such announcements more if they consider their sender authoritative and competent on the issues. Competence of the message-sender enhances the scope for successful communication and coordination, not only because more accurate information is more valuable to its receiver, but also because it soothes the receiver's concerns about the hidden motives behind the sender's message. Indeed, the basic result and intuition likely extend to all communication games where the sender has conflicting objectives that are smooth in actions and states of nature.

Section I reviews the relevant literature, Section II presents the model, Section III characterizes equilibrium communication, Section IV presents the key comparative statics, Section V illustrates some numerical examples, Section VI demonstrates some extensions, and Section VII concludes.

I. Related Literature

The relevance of CB's private information and the consequent signaling role of monetary policy are the subject of an ongoing debate (see Alex Cukierman 1995). Christina Romer and David Romer (2000) present evidence that the Fed's inflation forecasts contain valuable private information, relative to commercial forecasts. In contrast, Blinder (1998) and others argue that the Fed knows privately only its own intentions, because the relevant economic data are publicly and readily available. One may reply that research departments afford CBs a superior ability to process the available data. In our model, this issue is moot: *a CB's private information can be interpreted as its own subjective view of the macroeconomic outlook,* a view that plays a purely strategic role. We assume that P is interested only in correctly anticipating CB's moves, and does not care about the state of the economy per se. Credible communication develops as long as CB believes that its own opinion is informative about economic fundamentals and that P does agree on this, even if, in fact, P does not.

Canzoneri (1985) argues that the standard solutions to time inconsistency are inadequate when CB has private information on a state of the economy. He extends in this direction a static monetary policy game à la Barro and Gordon (1983). The timing of events makes cheap talk communication redundant: the inflation/output trade-off originates from a wage-contracting model, where P sets nominal wages *before* CB observes the state and chooses its policy. But, if the game were repeated, any serial correlation in the state of the economy would create a scope for communication. To keep the analysis of communication tractable, we obtain the same effect by changing the timing of events in the static game.[3] Later, we address the repeated game.

Our approach to cheap talk communication follows Crawford and Sobel's (1982) classic model of partisan advice, and the vast literature generated by this seminal contribution. In their model, a sender observes without noise a state of nature and announces a message to a receiver. The two agents have genuinely incongruent preferences over the state and the receiver's action. In our model, the receiver (P) cares only about the sender's (CB) final action (monetary policy),[4] not about the state of nature that CB observes with noise, and CB has an inflation bias, measuring the degree of time inconsistency. Jeremy Stein (1989) shows, in a different monetary policy game, that time inconsistency induces the sender to lie, just like a genuine preference bias. Our contribution is to identify the beneficial effects of competence on credibility.

The literature on cheap talk advice has not explored the comparative statics effects of changing the sender's competence, formally equivalent to changing the distribution of states of nature, on equilibrium communication.[5] The one exception is Marco Ottaviani (2000). In his model, the sender can observe a uniformly distributed state with a precision that she may choose covertly; so the issue is moral hazard, as opposed to our "adverse selection-reputational" view of precision as competence. The sender just speaks, and the receiver cares only about the objective informational content of the message, in the style of Crawford and Sobel (1982). So, there is no "credibility" in the usual time-inconsistency sense, because there are no deeds (sender's final action) to match to words (message). In contrast, in Roland Bénabou and Guy Laroque (1992), the message sender is an insider trader who has only *one* (biased) objective, to make money out of other traders, but has valuable private information about the fundamental value of the asset. In their Proposition 5 they show that the fewer the noise traders, who do not incorporate information into their decisions, the more "powerful" is the insider's announcements, the stronger his temptation to lie,

[3] Susan Athey, Andrew Atkeson, and Patrick Kehoe (2005) analyze a repeated version of Canzoneri's (1985) game, but assume i.i.d. states, the only (and nongeneric) case precluding communication. Therefore, explicit incentives are needed to extract information from CB. The best equilibrium of the revelation game, society's optimal "dynamic mechanism" for the CB, appropriately balances flexibility with a simple inflation-cap restriction to moderate time inconsistency.

[4] M3 is a good example of a variable that the private sector paid attention to mostly because the Bundesbank did (and the ECB still does), so it helped forecast its intentions, not because it was necessarily of any real relevance to the economy.

[5] David Austen-Smith (1990) is the seminal contribution on partisan advice with noise in the observation of the payoff-relevant state of nature. Roland Bénabou and Marco Battaglini (2003) allow for multiple noisy experts. Neither article considers the comparative statics effects of noise on equilibrium communication.

and the lower his credibility. This is the opposite of our result and originates from the single objective of the sender. In our context, CB must balance *two* potentially conflicting objectives: stimulating output by surprising expectations, and doing the right thing given the state of the economy.

From a technical viewpoint, most of the literature on partisan advice made two convenient assumptions, which we show to have substantive consequences. First is a bounded state space, which fails to guarantee existence of an equilibrium with communication, while in our unbounded setting a two-message equilibrium always exists. Second, in solved examples, is an uninformative (uniform) prior, which eliminates by construction the effect of competence on the distribution of private signals, thus on the "credibility of likely announcements."

Two other roles of competence in strategic communication have received attention. The first is the concern that an advisor may have for his reputation for competence (Morris 1997; Ottaviani and Peter Sorensen 2006), which distorts cheap talk advice independently of bias. The second is the coordination of expectations. Morris and Shin (2002) study the effects of public information when, as in Michael Woodford (2003), private information is dispersed among private agents, whose actions are strategic complements. In a monetary policy context, CB's actions signal its valuable information about the state of the economy, but also coordinate the beliefs and actions of private agents, who then place too much weight on CB's policy/announcements. The social marginal benefit of CB's competence is initially negative, because a modest increase in the accuracy of CB's information about fundamentals amplifies the undesirable coordination. Similarly, in our model competence is the power of CB's words in P's expectations. P weighs these words against its prior beliefs on the state of nature, because it has no private information of its own. Therefore, competence is always beneficial, because it works against time inconsistency, the impediment to communication.

II. The Model

Our two players are a central bank (CB) and the private sector (P). The economy is described by an output-gap version of the natural-rate Phillips curve:

$$y = s(\pi - x), \tag{1}$$

where y is the growth rate of real output (or of minus the output gap), π is the rate of inflation, x is P's rational expectation of inflation conditional on its information set $\mathbb{I}_P$,

$$x = \mathbb{E}[\pi | \mathbb{I}_P],$$

and $s > 0$ measures the sensitivity of output to inflation forecast errors, namely the degree of nominal rigidity in the economy.

CB controls π to minimize a quadratic loss function

$$L(y, \pi) = \mathbb{E}[(y - b)^2 + \lambda(\pi - \pi^* - \omega)^2 | \mathbb{I}_{CB}],$$

where b is desired output growth rate, π^* is the average level of desired inflation, ω is a shock to desired inflation, $\lambda > 0$ is the relative weight on inflation, and the expectation is conditioned on CB's information $\mathbb{I}_{CB}$. P has possibly a different information set, and its action is to formulate the expectation of inflation x and to set some nominal prices (e.g., wages) accordingly.[6]

When all parameters and shock realizations are public information, as is well known, this game has a unique Nash equilibrium outcome with $y = 0$ and

$$\pi = x = \pi^* + \frac{s}{\lambda} b + \omega. \tag{2}$$

This outcome is inefficient if CB is biased, namely if CB attempts to attain a growth rate b different (we assume larger) than the one $y = 0$ dictated by correct expectations. In fact, in this

[6] As shown by Canzoneri (1985), ω can represent several shocks affecting the economy, such as money demand shocks, or the effectiveness of monetary policy at that particular juncture. In Section VI we append ω to the supply curve (1) to capture a cost shock. The formal analysis and results are essentially the same.

case the equilibrium loss is positive, $b^2(1 + s^2/\lambda) > 0$, while it would be 0 if CB could commit to the Stackelberg inflation rate $\pi = x = \pi^* + \omega$. This is a simple illustration of the time-inconsistency or inflation-bias problem which has preoccupied monetary economists in the last decades.

We assume, instead, that ω is uncertain and drawn from

$$\omega \sim \mathcal{N}(0, 1),$$

where the mean and variance of the state ω are normalized to 0 and 1 without loss in generality. CB privately observes an informative but noisy signal of ω:

$$\theta = \omega + \varepsilon \qquad \text{where} \qquad \varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2).$$

Bayesian updating implies

$$\omega | \theta \sim \mathcal{N}(H\theta, 1 - H),$$

where we introduce the convenient notation

$$H \equiv \frac{1}{1 + \sigma_\varepsilon^2} \in [0, 1]$$

to denote CB's "competence." This measures CB's ability, *relative to the difficulty of the problem* as parameterized by $\sigma_\omega^2 = 1$, to correctly observe or interpret the state of the economy ω. For future reference, notice that the unconditional (from the point of view of the uninformed P) distribution of CB's private signal is

$$\theta = \omega + \varepsilon \sim \mathcal{N}(0, 1 + \sigma_\varepsilon^2) = \mathcal{N}(0, 1/H).$$

The timing of the game is as follows:

1. CB observes θ and announces to P a message A (silence is a message);
2. P formulates rational expectations of inflation x conditional on A;
3. CB chooses the optimal inflation rate $\pi(\theta, A)$, which depends on both its private information θ and its announcement A.

The private signal θ captures either private information that CB has about the state of the economy, or just CB's (in Svensson's (2003) terminology) "judgement" of the state of the economy. The economy is hit by different shocks, as summarized by ω, and the more precise CB's observations and/or interpretation of these shocks, namely the lower the variance σ_ε^2 and the higher the competence H, the more precise an idea CB has of what to do about it.[7]

To illustrate the trade-off faced by CB in this game, consider the commitment solution. CB chooses an inflation rate subject to the constraint of not surprising expectations:

$$\pi^C(\theta) = \arg\min_\pi \{(-b)^2 + \lambda \mathbb{E}[(\pi - \pi^* - \omega)^2 | \theta]\}$$

$$= \pi^* + H\theta.$$

When CB cannot commit, it has an incentive to surprise expectations. As we will see shortly, the "optimal surprise" (best response to expectations x) is

$$\pi^{BR}(\theta, x) = \frac{s^2 x + \lambda \pi^C(\theta) + sb}{s^2 + \lambda},$$

which reduces to the commitment solution if $b = 0$ and $x = \pi^{BR}$. Clearly, information is valuable, and CB would like to coordinate actual and expected inflation on $\pi^C(\theta)$, which depends on the signal θ that it received about the state of the economy. But, if P believes that CB will play the commitment solution, $x = \pi^C(\theta)$, and CB is biased, then

$$\pi^{BR}(\theta, \pi^C) = \pi^C + \frac{sb}{s^2 + \lambda} > \pi^C.$$

[7] It is immaterial whether P believes θ to be genuinely informative of the state of the economy ω, or θ just represents CB's own view of the economy, which might be totally uninformative about ω from P's viewpoint. The role of θ for P is purely strategic; P cares about what CB believes, even if P has no confidence whatsoever in CB's competence, because P is concerned only with anticipating CB's moves. Of course, CB has to believe that P believes θ to be informative, even if P really does not. This is another difference from Morris and Shin (2002), where public announcements play both an allocative role, and a signaling role of *genuinely* valuable information about fundamentals.

CB wants to surprise expectations by inflating the economy. Thus, CB faces a dilemma. On the one hand, it wants to do "the right thing," by setting policy to a level $\pi^C(\theta)$ appropriate to the perceived state of the economy $H\theta$, and make sure that expectations are coordinated accordingly. On the other hand, it wants to induce low expectations to surprise them with high inflation and stimulate output. The higher competence H, the more weight CB puts on its view of what the economy needs and the less on surprising expectations. In the following sections, we develop this analysis formally.

III. Communication and Transparency of Monetary Policy

In this section we characterize the set of equilibria of the communication game. Without loss of generality, we can restrict attention to messages that are subsets of the real line such that announcing the set A means saying that $\theta \in A$. A *Perfect Bayesian Nash Equilibrium with Communication* (PBEC) is a measurable partition $\mathcal{A}$ of $\mathbb{R}$ and a set of beliefs about θ with the following properties. For all $A \in \mathcal{A}$, if CB announces A to P, P "believes" it, namely P updates its beliefs about θ in a Bayesian fashion from the prior $\mathcal{N}(0, 1/H)$ to a posterior $\Pr(\theta | \theta \in A)$. If CB makes any other announcement not in $\mathcal{A}$, P assimilates it to some message $A' \in \mathcal{A}$. After announcing A, CB chooses the optimal inflation rate π. Anticipating the consequences of any announcement, a CB that privately observes $\theta \in A$ has no incentives to send any other message than A.

When $\mathcal{A} = \mathbb{R}$, no communication takes place: there always exists a fully nonrevealing equilibrium, equivalent to "babbling," where P believes nothing that CB says. When every $A \in \mathcal{A}$ is a singleton, we have full communication, or truth-telling.

A. *Equilibrium for a Given Policy Announcement*

We find equilibria of this game by backward induction. First, for every CB's private observation θ ("type"), we fix the announcement A and find the equilibrium inflation rate $\pi(\theta, A)$ of the policy subgame. Second, we find the equilibrium announcement.

In any PBEC, after CB announces a message A, P "believes it." Formally, P reformulates an expectation of inflation

$$x(A) = \mathbb{E}[\pi | \theta \in A].$$

Given its type θ and previous announcement A, CB solves

$$\min_{\pi}\{L(A, \pi | \theta) = \mathbb{E}[(s\pi - sx(A) - b)^2 + \lambda(\pi - \pi^* - \omega)^2 | \theta]\}.$$

The necessary and sufficient first-order condition for an optimal inflation rate yields a unique best response to P's expectations $x(A)$:

$$(3) \qquad \pi^{BR}(\theta, A) = \frac{\lambda\pi^* + sb + s^2x(A) + \lambda H\theta}{s^2 + \lambda}.$$

P's rational expectation of inflation (best response) satisfies

$$x(A) = \mathbb{E}[\pi^{BR}(\theta, A) | \theta \in A] = \frac{\lambda\pi^* + sb + s^2x(A) + \lambda H\bar{\theta}(A)}{s^2 + \lambda},$$

where

$$\bar{\theta}(A) = \mathbb{E}[\theta | \theta \in A] = \frac{\int_A \theta e^{-H(\theta^2/2)}\, d\theta}{\int_A e^{-H(\theta^2/2)}\, d\theta}$$

is P's expectation of θ given the believed announcement (that $\theta \in$) A. Solving for $x(A)$:

$$(4) \qquad x(A) = \pi^* + \frac{sb}{\lambda} + H\bar{\theta}(A).$$

Notice that the effect of the inflation bias b on equilibrium expected inflation x is amplified by s and is dampened by λ: the larger s, the more responsive is output to inflation surprises, and the stronger the temptation for CB to inflate the economy, while the higher λ, the more costly to

CB is (off-target) inflation. Notice also that the equilibrium expected inflation $x(A)$ from (4) with incomplete information equals the one of the complete information game, (2), where the estimated inflation target $H\bar{\theta}(A)$ given CB's announcement replaces the true inflation target ω.

Replacing P's best response $x(A)$ from (4) into CB's best response (3) and collecting terms gives an equilibrium inflation rate

$$\pi(\theta, A) = \pi^* + \frac{sb}{\lambda} + H\frac{s^2\bar{\theta}(A) + \lambda\theta}{s^2 + \lambda}.$$

In turn, replacing both $x(A)$ and $\pi(\theta, A)$ into the supply curve (1) gives an equilibrium growth rate of output

$$y(\theta, A) = sH\frac{s^2\bar{\theta}(A) + \lambda\theta}{s^2 + \lambda} - sH\bar{\theta}(A),$$

$$= s\lambda H\frac{\theta - \bar{\theta}(A)}{s^2 + \lambda}.$$

Notice that, if the effect of CB's announcement on P's beliefs about the state of the economy, $\bar{\theta}(A)$, is higher than the signal θ that CB truly observed, then P expects too high inflation and this depresses real activity.

Finally, conditional on the private realization of the optimal inflation target θ and on the message A that CB announces about θ, the expected loss in equilibrium is

$$\begin{aligned}
&\mathcal{L}(A|\theta)\\
&= L(y(\theta, A), \pi(\theta, A))\\
&= \mathbb{E}\Bigg[\left(\frac{s\lambda H(\theta - \bar{\theta}(A))}{s^2 + \lambda} - b\right)^2\\
&\quad + \lambda\left(\frac{sb}{\lambda} + s^2H\frac{\theta - \bar{\theta}(A)}{s^2 + \lambda} + H\theta - \omega\right)^2\Bigg|\theta\Bigg]\\
&= s^2H^2\lambda\left[\frac{\theta - \bar{\theta}(A)}{s^2 + \lambda}\right]^2\\
&\quad - 2bsH[\theta - \bar{\theta}(A)]\\
&\quad + b^2(1 + s^2/\lambda) + \lambda(1 - H),
\end{aligned}$$

where in the last line we use $\mathbb{E}[(H\theta - \omega)^2|\theta] = \text{Var}[\omega|\theta] = 1 - H$ and collect terms. This function measures the incentives to announce any message A for a CB that observed a signal θ of the state of the economy and expects to be believed by P.

We notice a key fact. Since only P's expectation of θ, given a message $A \in \mathcal{A}$, matters for payoffs, *we can replace messages A with numbers* $\bar{\theta} = \mathbb{E}[\theta|\theta \in A]$, with the interpretation that saying "$\bar{\theta}$" means "θ belongs to the set A where θ has mean $\bar{\theta}$." This fact enormously simplifies the analysis. So we will write

$$(5)\quad \mathcal{L}(\bar{\theta}|\theta) = s^2H^2\lambda\frac{(\theta - \bar{\theta})^2}{s^2 + \lambda} - 2bsH(\theta - \bar{\theta}) + b^2(1 + s^2/\lambda) + \lambda(1 - H).$$

B. *Incentive-Compatible Policy Announcements*

The Incentive Compatibility (IC) constraint that PBEC announcements in $\mathcal{A}$ must satisfy is that, when observing θ, CB prefers to report message $\bar{\theta} = \mathbb{E}[\theta|\theta \in A]$ rather than any other message in $\mathcal{A}$. Formally,

$$\mathcal{L}(\bar{\theta}|\theta) \leq \mathcal{L}(\bar{\theta}'|\theta)$$

for every $A, A' \in \mathcal{A}$, every $\theta \in A$, with $\bar{\theta} = \mathbb{E}_\theta[\tau|\tau \in A]$, and $\bar{\theta}' = \mathbb{E}_\theta[\tau|\tau \in A']$.

Using (5) and simplifying terms, this is

$$s^2H^2\lambda\frac{(\theta - \bar{\theta})^2}{s^2 + \lambda} - 2bsH(\theta - \bar{\theta}) \leq s^2H^2\lambda\frac{(\theta - \bar{\theta}')^2}{s^2 + \lambda} - 2bsH(\theta - \bar{\theta}').$$

Dividing through by $H^2s^2\lambda/(s^2 + \lambda) > 0$ and rearranging, we finally obtain a quite simple inequality:

$$(6)\quad (\bar{\theta} - \bar{\theta}')(\bar{\theta} + \bar{\theta}' - 2\theta + q) \leq 0,$$

where

(7) $$q \equiv 2\,\frac{b}{H}\left(\frac{s}{\lambda} + \frac{1}{s}\right)$$

is a normalized *bias/competence ratio.*

This composite parameter plays a critical role: it increases with CB's inflation bias b and decreases with CB's competence H and with CB's conservatism (weight on inflation) λ. Intuitively, CB's incentives to engineer a high unexpected inflation increase in its bias b and decrease in its cost of inflation λ. The dependence of q on the sensitivity of output to inflation surprises s is ambiguous, as

(8) $$\frac{\partial q}{\partial s} = \frac{2b}{\lambda H}\left(1 - \frac{\lambda}{s^2}\right).$$

This is due to two countervailing effects of nominal rigidity s on CB's incentives to inflate the economy. A large s makes inflation surprises very "productive" in terms of output gap. Therefore, when s is large, on one hand P expects CB to be prone to deception; on the other CB does not need to deceive P by much to attain the desired output growth b. Which of the two effects dominates depends on conservatism λ. If $s \le \sqrt{\lambda}$, then P knows that CB prefers moderate inflation surprises, because inflation is more important than output to CB; so the second effect dominates and q is decreasing in s. Both s and λ are scale-free parameters. As we will argue in Section V, s is likely to be significantly below 1, while modern central banks are generally believed to care more about inflation than about real activity ($\lambda > 1$), so (8) is generally negative.

We can manipulate (6) to obtain an IC constraint in simpler form: for every $A, A', A'' \in \mathcal{A}$, with $\bar{\theta} = \mathbb{E}_\theta[\tau|\tau \in A]$, $\bar{\theta}' = \mathbb{E}_\theta[\tau|\tau \in A']$, $\bar{\theta}'' = \mathbb{E}_\theta[\tau|\tau \in A'']$, and every $\theta \in A$,

(9) $$\begin{cases} \bar{\theta}' \le 2\theta - \bar{\theta} - q & \text{when} \quad \bar{\theta}' < \bar{\theta} \\ \bar{\theta}'' \ge 2\theta - \bar{\theta} - q & \text{when} \quad \bar{\theta}'' > \bar{\theta} \end{cases}.$$

The intuition behind these incentives constraints is simple. CB would like P to expect low inflation and to surprise P to boost growth closer to the bliss rate b. Thus, when observing θ, CB prefers not to tell a lie $\bar{\theta}'$ (another message that P would believe, as part of the equilibrium partition) only if this lie $\bar{\theta}'$ is so far from θ relative to the candidate equilibrium announcement $\bar{\theta}$ that saying $\bar{\theta}'$ and being believed will excessively mislead inflation expectations, and weigh on the inflation term of the social loss function.

C. *Equilibrium Characterization*

We now assume that a PBEC exists, and we illustrate some of its properties. In the next subsection we will tackle existence. First, we confirm a well-known fact.

PROPOSITION 1 (Communication is coarse): *In any PBEC, truth-telling ($\bar{\theta} = \theta$ for all θ) is impossible if and only if $b > 0$.*

PROOF:

Under truth-telling, CB reveals its type θ credibly, so P believes any pointwise announcement, i.e., $\bar{\theta}$ can be any real number. But then for $\theta = \bar{\theta}$, the first line in (9) reads $0 \ge \bar{\theta}' - \theta + q$ for all real numbers $\bar{\theta}' < \theta = \bar{\theta}$, which is clearly impossible if and only if $b > 0$, namely $q > 0$. The converse is straightforward.

Second, it is immediate from (9) that if type θ prefers message $\bar{\theta}$ to $\bar{\theta}' < \bar{\theta}$, then so do all types $\hat{\theta} > \theta$. Conversely, if θ prefers message $\bar{\theta}$ over $\bar{\theta}'' > \bar{\theta}$, then so do all $\hat{\theta} < \theta$. It follows:

LEMMA 1: *In any PBEC, all CB's types θ who send the same message belong to a connected set, either an interval or a (half-)line.*

In light of these results, a PBEC is a collection of intervals $[\theta_k, \theta_{k+1}) \subseteq \mathbb{R}$ such that: (a) all CB's types $\theta \in [\theta_k, \theta_{k+1})$ send the message

(10) $$\bar{\theta}_k = \mathbb{E}[\theta|\theta_k \le \theta < \theta_{k+1}],$$

and (b) if inflation expectations are based on $\bar{\theta}_k$, then no other message will give a smaller expected loss, conditional on θ.

To satisfy all Incentive Compatibility (IC) constraints, it suffices to impose that each boundary type CB's θ_k be indifferent between the message $\bar{\theta}_k > \theta_k$, that it and its right-neighbors $\theta_k + \varepsilon$ for some $\varepsilon > 0$ are supposed to send, and the message $\bar{\theta}_{k-1} < \theta_k$, that its left-neighbors $\theta_k - \varepsilon$ are supposed to send. Because, then, from (9), all $\theta > \theta_k$ strictly prefer $\bar{\theta}_k$ to $\bar{\theta}_{k-1}$, so their IC constraint is satisfied when comparing their equilibrium message $\bar{\theta}_k$ to the largest smaller message $\bar{\theta}_{k-1}$, and *a fortiori* to all smaller messages $\bar{\theta}_{k-1-j}$. At the same time, again from (9), all $\theta < \theta_k$ strictly prefer $\bar{\theta}_{k-1}$ to $\bar{\theta}_k$, so their IC constraint is satisfied when comparing their equilibrium message $\bar{\theta}_{k-1}$ to the smallest larger message $\bar{\theta}_k$, and *a fortiori* to all larger messages $\bar{\theta}_{k+j}$. When this holds for all k, all IC constraints are automatically satisfied, strictly so for θ in the interior of each interval, and weakly so for boundary types.

From (9), the required indifference condition for the boundary type θ_{k+1} between the two nearby messages $\bar{\theta}_k$ and $\bar{\theta}_{k+1}$ can be written as follows:

$$\bar{\theta}_{k+1} = 2\theta_{k+1} - q - \bar{\theta}_k. \tag{11}$$

LEMMA 2 (Form of the equilibrium message space): *A PBEC is a partition of the real line into intervals* $\{[\theta_k, \theta_{k-1})\}$, *where* $\{\theta_k\}$ *is an increasing, possibly doubly infinite, sequence which solves* (10) *and* (11).

From (11), the critical parameter q, the normalized bias/competence ratio, bounds the fineness of the equilibrium message space. First, suppose that CB observes θ and in equilibrium announces $\bar{\theta}_k > \theta$. For every lower message $\bar{\theta}_{k-1} < \bar{\theta}_k$ that is part of the equilibrium, we deduce from (11):

$$\bar{\theta}_k > \theta \geq \theta_k = \frac{\bar{\theta}_k + \bar{\theta}_{k-1} + q}{2}.$$

Rearranging, messages have to be at least q apart: $\bar{\theta}_k - \bar{\theta}_{k-1} > q$. Using this inequality one step forward (at $k+1$), (11) again yields $q < \bar{\theta}_{k+1} - \bar{\theta}_k = 2(\theta_{k+1} - \bar{\theta}_k) - q$, which implies the sharper restriction

$$\theta_{k+1} - \bar{\theta}_k > q. \tag{12}$$

The upper bound of each interval must be at least q larger than the mean of θ on that interval. Since $\theta_k < \bar{\theta}_k$, this also implies that intervals must have a width of at least q: $\theta_{k+1} - \theta_k > q$.

If we interpret the fineness of the equilibrium message space as monetary policy's transparency, we obtain:

PROPOSITION 2 (Maximum transparency of monetary policy): *In any PBEC, messages sent by CB to P about the optimal inflation rate are intervals* $[\theta_k, \theta_{k+1})$ *whose boundaries and means with respect to the prior beliefs* $\bar{\theta}_k = \mathbb{E}[\theta|\theta_k \leq \theta < \theta_{k+1}]$ *are spaced more than q apart, namely* $\theta_k - \theta_{k+1} > q$ *and* $\bar{\theta}_{k+1} - \bar{\theta}_k > q$. *Therefore, the normalized bias/competence ratio q constrains the transparency of monetary policy.*

Next, we show that the sequence $\{\theta_k\}$ is bounded both below and above. For this we need an auxiliary:

LEMMA 3: *For every* $H \in (0, 1]$, *the function*

$$f(t|H) = t - \frac{\int_{-\infty}^{t} \theta e^{-H(\theta^2/2)}\, d\theta}{\int_{-\infty}^{t} e^{-H(\theta^2/2)}\, d\theta}$$

is increasing in t, with $\lim_{t\to-\infty} f(t|H) = 0$, $\lim_{t\to+\infty} f(t|H) = \infty$, *and the function*

$$g(t|H) = \frac{\int_{t}^{\infty} \theta e^{-H(\theta^2/2)}\, d\theta}{\int_{t}^{\infty} e^{-H(\theta^2/2)}\, d\theta} - t$$

is decreasing in t, with $\lim_{t\to-\infty} g(t|H) = \infty$, $\lim_{t\to+\infty} g(t|H) = 0$.

PROOF:

By a change of variable $u = t/\sqrt{H}$, it is easy to verify that $f(t|H) = f(u|1)/\sqrt{H}$ and $g(t|H) = g(u|1)/\sqrt{H}$. Given the definition of u, for any given $H > 0$, it suffices to prove the claims for $f(u|1)$ and $g(u|1)$. These follow from l'Hôpital's rule.

LEMMA 4 (Finite message space): *Any PBEC is a finite partition of the real line into K intervals* $\{(-\infty, \theta_1), [\theta_1, \theta_2), [\theta_2, \theta_3), \ldots, [\theta_{K-1}, \infty)\}$ *where* $\{\theta_k\}_{k=1}^{K-1}$ *is a strictly increasing and finite sequence satisfying* (10) *and* (11).

PROOF:
See Appendix.

D. *Two-Message Equilibrium*

Having characterized any equilibrium of the communication game, we now explore its existence. First, we show that a two-message equilibrium always exists, so CB is always able to credibly communicate something. This stands in contrast to Crawford and Sobel's (1982) results, where the bias b has to be small enough for communication to occur, due to their assumption of a bounded state space.

In a two-message equilibrium, a CB observing $\theta < \theta_1$ for some real number θ_1 announces $\bar{\theta}_0 = \mathbb{E}[\theta|\theta < \theta_1]$, or "The economy needs low inflation," while for all $\theta \geq \theta_1$ CB announces $\bar{\theta}_1 = \mathbb{E}[\theta|\theta \geq \theta_1]$, or "The economy needs high inflation." The only IC constraint to be satisfied with equality by the cutoff type θ_1 is $\bar{\theta}_1 = 2\theta_1 - q - \bar{\theta}_0$, so

$$\phi(\theta_1) \equiv f(\theta_1) - g(\theta_1) = q.$$

By Lemma 3, ϕ is continuous, strictly increasing and onto with $\phi(0) = 0$. By the Mean Value Theorem, a solution $\theta_1^{(2)}$ to this equation always exists and is strictly positive. The superscript "(2)" refers to the number of messages sent in this equilibrium. Since $g(\cdot) > 0$, the last equation also implies the required $f(\theta_1^{(2)}) > q$.

Since $\theta_1^{(2)} > 0$ and the normal distribution of θ is symmetric around its mean zero, it follows that $|\bar{\theta}_0^{(2)}| < |\bar{\theta}_1^{(2)}|$, i.e., $\bar{\theta}_0^{(2)}$ is closer to zero than $\bar{\theta}_1^{(2)}$. When CB announces that low inflation is needed, P believes this but sets a relatively high expectation $\pi^* + H\bar{\theta}_0^{(2)}$, below but not far from π^*. When CB says that high inflation is needed, P's expectation goes way above π^*, to $\pi^* + H\bar{\theta}_1^{(2)}$. Preference bias skews credible communication. The higher q, the higher and the more asymmetric are the two equilibrium messages.

Since $\theta_1^{(2)} > 0$, $\Pr(\bar{\theta}_0^{(2)}) = \Pr(\theta \leq \theta_1^{(2)}) > 1/2$. That is, somewhat paradoxically, CB is ex ante more likely to credibly announce that low inflation ($\pi^* + H\bar{\theta}_0^{(2)} < \pi^*$) is appropriate, even if the ex ante chance of high and low inflation is even. CB may even credibly announce below-average inflation ($\pi^* + H\bar{\theta}_0^{(2)}$) when some moderately above-average inflation is truly required by $\theta \in (0, \theta_1^{(2)})$. *So an inflation-biased CB is relatively likely to speak like a conservative one—and get away with it!* Credibility forces CB to be conservative in its announcements. The alternative two-message equilibrium, with $\bar{\theta}_0^{(2)}$ very low and negative and $\bar{\theta}_1^{(2)}$ moderately high and positive, would not be feasible, because the temptation to say $\bar{\theta}_0^{(2)}$ and be believed would be very strong. The more competent CB is, the less conservative it can be.

E. *Equilibrium Existence and Construction*

Our equilibrium characterization suggests a recursive algorithm to construct any equilibrium. Finding the equilibrium with two messages is a trivial nonlinear equation problem in one variable, the cutoff $\theta_1^{(2)}$. Any equilibrium with more than two messages is a finite sequence solving a second-order difference equation, but whose initial value is unknown and whose final value must satisfy a known restriction. Finding this sequence requires a "shooting" method.

First, define the implicit function $h(\theta_1, \bar{\theta}_1|H)$ by

$$\frac{\int_{\theta_1}^{h(\theta_1,\bar{\theta}_1|H)} \theta e^{-H(\theta^2/2)}\, d\theta}{\int_{\theta_1}^{h(\theta_1,\bar{\theta}_1|H)} e^{-H(\theta^2/2)}\, d\theta} = \bar{\theta}_1.$$

Second, notice that from Lemma 4, for $k = 1$ the inequality (12) reads

$$\theta_1 - \mathbb{E}[\theta|\theta < \theta_1] = \theta_1 - \bar{\theta}_0 = f(\theta_1) > q,$$

or $\theta_1 > \theta_1^* \equiv f^{-1}(q)$. This provides a lower bound to the very first (and possibly only) element of the equilibrium sequence $\{\theta_k\}_{k=1}^{K-1}$. The algorithm starts from $\theta_1 = \theta_1^*$, and computes recursively $\bar{\theta}_0 = \mathbb{E}[\theta|\theta < \theta_1]$, then $\bar{\theta}_1 = 2\theta_1 -$

$q - \bar{\theta}_0$, then $\theta_2 = h(\theta_1, \bar{\theta}_1|H)$, then $\bar{\theta}_2 = 2\theta_2 - q - \bar{\theta}_1$, and so forth. At each step we check whether the required order is respected, i.e., $\theta_k < \bar{\theta}_k < \theta_{k+1}$; if either of these two inequalities fails, we go back to increase θ_1 and restart. The algorithm converges at step k such that $\bar{\theta}_k = \mathbb{E}[\theta|\theta > \theta_k]$. If $\bar{\theta}_k < \mathbb{E}[\theta|\theta > \theta_k]$, we go on; if $\bar{\theta}_k > \mathbb{E}[\theta|\theta > \theta_k]$, we have "overshot" and need to restart from a higher θ_1. When we converge, we have found one equilibrium with K messages. Next, we restart the algorithm from $\theta_1 = \theta_1^{(K)} + \varepsilon$ and look for other equilibria. When θ_1 reaches the two-message equilibrium threshold $\theta_1^{(2)}$, we have exhausted the equilibrium set.

This algorithm can be used to find all equilibria of the game for a given set of parameters q, H. Because the parameter space is bidimensional, the search is cumbersome. To make it a simple unidimensional search, we study communication in a different space. Suppose that CB, having observed θ, announces a message about the magnitude

$$\sqrt{H}\,\theta \sim \sqrt{H} \cdot \mathcal{N}(0, H^{-1}) = \mathcal{N}(0, 1).$$

Clearly, the exact form of communication is not important, and the equilibrium partition in this space is dictated by $\{\sqrt{H}\,\theta_k^{(K)}\}_{k=1}^{K-1}$, as we are "changing language." In particular, given parameter values, the set of equilibria is the same, up to the rescaling.

Under this renormalization, the "composite" normalized bias is

$$\bar{q} = \sqrt{H}\,q = \frac{2b(s^2 + \lambda)}{\lambda\sqrt{H}\,s}. \tag{13}$$

Just like q in (7), $\bar{q}$ is decreasing in competence H and conservatism λ and increasing in inflation bias b and in nominal rigidity s for $s \leq \sqrt{\lambda}$. The distribution of the message is a standard normal, invariant to changes in competence and in any other parameter. Therefore, we have collapsed all the parameters of the model into the renormalized bias $\bar{q}$, and we can solve numerically for the entire set of equilibria when $\bar{q}$ ranges over the positive reals, stopping when $\bar{q}$ is so large that only two-message equilibria remain. This clearly provides a complete characterization of the equilibrium set, and we can prove general result numerically.[8]

After this rescaling, we implement the algorithm described earlier in a simple GAUSS program. First, we find that there cannot be a three-message equilibrium whenever $\bar{q} > 0.37$. Given the properties of the function ϕ established earlier, for any $\bar{q} > 0.37$, only the two-message equilibrium exists (it always does). Therefore our characterization of the comparative statics effects of changes in parameters is complete when iterating over values of $\bar{q}$ in a bounded set $[0, 0.37]$.

We obtain, and illustrate later numerically, the following characterization. For every set of parameter values, there exists a maximum number of messages (at least two) that can be supported in a PBEC, and there exists a PBEC with every integer number of messages up to that upper bound. We state the results in terms of the original message space.

PROPOSITION 3 (Equilibrium communication of monetary policy): *Every PBEC takes the following form. For every $\bar{q} \in (0, \infty)$ defined in (7), there exists a finite integer $N(\bar{q}) \geq 2$ and, for each integer $K \in \{1, \dots N(\bar{q})\}$, a finite, strictly increasing sequence of $K + 1$ extended reals, $\{\theta_k^{(K)}\}_{k=0}^{K}$, $\theta_0^{(K)} = -\infty = -\theta_K^{(K)}$, which defines a sequence of K adjacent intervals $[\theta_k^{(K)}, \theta_{k+1}^{(K)})$ with conditional expectations $\bar{\theta}_k^{(K)} = \mathbb{E}[\theta|\theta_k^{(K)} \leq \theta < \theta_{k+1}^{(K)}]$, $k = 0, 1, \dots K - 1$, with the following properties. After privately observing a signal realization $\theta \in [\theta_k^{(K)}, \theta_{k+1}^{(K)})$ of the state ω, CB sends to P the message $\bar{\theta}_k^{(K)}$. Therefore, there always exist both a PBEC with one (uninformative) message $K = 1$, or "babbling," and a PBEC with two messages, $N(\bar{q}) = 2 \Rightarrow K = 2$, where CB credibly announces whether required inflation exceeds a unique cutoff dictated by $\theta_1^{(2)} = \phi^{-1}(\bar{q}/\sqrt{H})$, increasing in $\bar{q}$. For every $\bar{q} \in (0, \infty)$ such that $N(\bar{q}) \geq 3$, a nonempty set, there exist PBEC with K messages for every $K = 1, 2 \dots N(\bar{q})$. In any*

[8] There is also a third message space of interest. The CB announces the expected inflation target $H\theta \sim H \cdot \mathcal{N}(0, H^{-1}) = \mathcal{N}(0, H)$. In this space the normalized bias is independent of competence, and the effect of competence is entirely loaded on the distribution of the message. A competent CB is less constrained by the prior, thus more able to manipulate P's posterior beliefs.

PBEC, both the messages $\bar{\theta}_k^{(K)}$ and the interval boundaries $\theta_k^{(K)}$ are spaced at least $q = \bar{q}/\sqrt{H}$ apart. The lowest finite threshold $\theta_1^{(K)}$ strictly exceeds $f^{-1}(q)$, which is increasing in q. The subsequent thresholds $\theta_k^{(K)}$, $k = 2, 3, \dots K$, solve the second-order nonlinear difference equation (11) *with terminal condition $\theta_K^{(K)} = \infty$.*

F. *The Value of Equilibrium Communication*

How much does equilibrium communication contribute to welfare? Recall the loss function from (5). In any PBEC, the unconditional expected loss before seeing θ is

$$\mathcal{L} = \mathcal{L}_T + \frac{s^2H^2\lambda}{s^2+\lambda}\,\mathbb{E}[(\theta-\bar{\theta})^2 - q(\theta-\bar{\theta})],$$

where the expectation is taken over the outcome of the PBEC, and we use the shorthand

$$\mathcal{L}_T \equiv b^2\left(1+\frac{s^2}{\lambda}\right) + \lambda(1-H) = \mathbb{E}[\mathcal{L}(\theta|\theta)]$$

for the part of the indirect loss that does not depend on announcements, namely, the loss $\mathbb{E}[\mathcal{L}(\theta|\theta)]$ when CB always tells the truth and P believes it. In this case, no communication-strategy issue remains, and the only role of competence H is to enhance the value of flexibility and to enable a better monetary policy. Indeed, $\mathcal{L}_T$ declines in H.

After some manipulations, it is easy to show that, for any equilibrium with K messages:

$$\mathcal{L} = \mathcal{L}_T + \frac{s^2H^2\lambda}{s^2+\lambda}\sum_{k=0}^{K-1}\int_{\theta_k^{(K)}}^{\theta_{k+1}^{(K)}}(\theta-\bar{\theta}_k^{(K)})^2\,\frac{e^{-(\theta^2/2)H}}{\sqrt{2\pi H}}\,d\theta,$$

where the sum of integrals is $\mathbb{E}[\mathrm{Var}[\theta|\theta_k^{(K)} \le \theta < \theta_{k+1}^{(K)}|q, H]]$, the ex ante unconditional expectation of the ex post variance of θ in P's beliefs after CB communicates according to the PBEC. This quantity ranges from 0 when truth-telling is an equilibrium to $\mathrm{Var}[\theta] = 1/H$ when no communication take places in equilibrium. So $1 - H\mathbb{E}[\mathrm{Var}[\theta|\theta_k^{(K)} \le \theta < \theta_{k+1}^{(K)}|q, H]]$ is the percentage reduction in the variance due to communication, a measure of credibility.

IV. Competence Implies Credibility

The central result of this paper is the comparative statics effect of competence H on the structure of the equilibrium. The competence of the informed party, CB, has a beneficial effect on its ability to communicate information to P. Notice that, in the original message space where CB announces its private signal θ, competence H formally plays two roles: it reduces the normalized bias q, and it determines the probability distribution of the private signal $\theta \sim \mathcal{N}(0, H^{-1})$, thus the function $h(\cdot|H)$ defined earlier to construct the equilibrium message space.

The first strategic effect of competence is the *power of the words*. Competence H is the weight that P puts on CB's credible announcement $\bar{\theta}$ in forming an expectation of inflation. In fact, inflation expectations $x(A)$ contain a term $H\bar{\theta}$ which is just P's expectation, conditional on the announcement $\bar{\theta}$ of CB's updated inflation target $H\theta$: namely, from (4),

$$x(A) - \pi^* - \frac{s}{\lambda}b = H\bar{\theta} = \mathbb{E}[H\theta|A, \mathbb{I}_P]$$
$$= \mathbb{E}[\mathbb{E}[\omega|\theta, \mathbb{I}_{CB}]|A, \mathbb{I}_P].$$

If P knows that CB deems its own information θ to be accurate, then P also knows that CB will put a large H weight on the private signal θ, relative to the weight $1-H$ that CB puts on the prior expectation of ω, normalized to 0. The reason for this strategic effect is simple. In CB's best response (the monetary policy function (3)), competence H affects the weight that CB puts on its own private information, as opposed to inflation expectations and to its inflation preferences $\lambda\pi^* + sb$.

Since P cares only about anticipating CB's moves, when formulating the expectation of $H\theta$, P will use the same value of H as CB. Knowing that P adopts the competence weight H, CB knows that P's inflation forecast error $\pi - x$ and the resulting output gap y depend on $H(\theta - \bar{\theta})$. The larger is H, the smaller the "lie" $(\theta - \bar{\theta})$ that CB needs to tell to align the output gap to the desired rate b. At the same time, CB must choose inflation as close to $\pi^* + H\theta$ as possible to respond to the economy's needs. So the larger is CB's competence H, the smaller

CB's incentives to claim a very different θ than the one it truly observed, and the finer the communication structure that P trusts. The interplay between the two conflicting objectives, output (which requires an inflation surprise) and inflation (which requires an appropriate inflation rate), generates this effect. Setting the right inflation rate for the state of the economy *and* surprising P optimally is impossible because P has a prior it relies upon, so CB must trade off the two. The more competent CB is, the more weight it puts on the first objective, and the more P can believe CB.

If CB had just one biased objective, as in Crawford and Sobel's (1982) original game, then this effect would disappear. This occurs when $\lambda = 0$, so CB's objective is just $L(y, \pi) = \mathbb{E}[(y - b)^2|\mathbb{I}_{CB}]$. Then, in (3), CB puts no weight on its own private information (the weight on θ is λH). Therefore, a change in competence has no effect on CB's actions, and P has no more nor fewer reasons to believe what a more competent CB says. It is easy to show that as $\lambda \searrow 0$, the normalized bias q explodes, the equilibrium message partition changes smoothly, and the effect of competence on credibility disappears continuously, with no discontinuity at $\lambda = 0$. In fact, for a sufficiently small but positive λ, just like for a sufficiently small but positive H, the normalized bias q is so high that the only nontrivial equilibrium has two messages. As λ declines further and approaches zero, the expected value of communication with two messages declines (the unique cutoff $\theta_1^{(2)}$ goes to $-\infty$), and so does the comparative statics effect of competence (cf. (7)). When λ reaches zero and the normalized bias q reaches infinity, we have the limiting case of no communication (babbling is the unique equilibrium), where competence is irrelevant.

The second strategic effect of competence is the *credibility of likely announcements*. The more competent is CB, the more likely it is to observe a true θ near its (zero) average. So, P can believe announcements by CB that are very far from the prior expectation (zero) only if they are made by a larger set of CB types, because such announcements are unlikely to be truthful. This implies that the equilibrium partition is coarser in the tails, and by the same token finer in the middle near zero, where the private signal θ is likely to be from P's viewpoint. Again, this is all irrelevant if $\lambda = 0$ and CB has no conflicting objectives. This discussion suggests that the effects of competence on credibility hold in any Bayesian communication game, provided that the sender has conflicting objectives and is uncertain about one of the two.

To establish formally the effects of changes in competence H on the equilibrium sequences $\{\theta_k\}_{k=0}^{K}$, $K = 1, 2, \ldots N(q, H)$, we again use the normalization that collapses all parameters into the rescaled bias $\bar{q}$. An increase in competence is equivalent to a decrease in $\bar{q}$. We vary $\bar{q}$ and compute numerically all equilibria. We obtain the main result of this paper:

PROPOSITION 4 (Competence implies credibility): *The maximum number of PBEC messages is a finite integer $N(\bar{q}) \geq 2$, nonincreasing in the composite bias $\bar{q}$ in* (13), *thus nonincreasing in the inflation bias b and nondecreasing in competence H and conservatism λ, with* $\lim_{\bar{q}\to 0} N(\bar{q}) = \infty$. *For a given $\bar{q}$ (parameter set), comparing across the multiple equilibria with* 1, 2, ... $N(\bar{q})$ *messages, the minimum expected loss and the minimum percentage reduction in P's belief variance due to communication are both attained by the equilibrium with the finest partition $K = N(\bar{q})$. These minimized values are increasing in $\bar{q}$.*

Equilibria are Pareto ranked by the number of messages. So, even if the game always has multiple equilibria (at least babbling and two-message, but usually more), it is natural to focus on the equilibrium with the finest partition of $N(\bar{q})$ messages. A simple forward induction argument may be used to select this Pareto-dominant equilibrium. CB has to announce a message that can only be part of this best equilibrium. If P plays along, CB has indeed an incentive to announce that message, and P to believe it, so the announcement and the Pareto-dominant equilibrium are self-enforcing. Because in this equilibrium the number of elements and the associated "variance" measure of credibility are increasing in H, we conclude that competence implies credibility.

V. Numerical Examples

The numerical solution of the game allows us to prove Proposition 4 and to illustrate the comparative statics effects of changes in bias, competence, and conservatism on the extent of communication. For comparative statics purposes, for each parameter configuration, we focus on the Pareto-dominant equilibrium with the finest message space, which maximizes the scope for communication. We derive the results with the algorithm for the rescaled message space, but we present them in the original message space.

Since we normalized $\sigma_\omega = 1$, we express quantities in percentage points. For example, we assume an average desired inflation rate $\pi^* = 2$. This implies

$$\Pr(0 = \pi^* - 2\sigma_\omega \leq \pi^* + \omega \leq \pi^* + 2\sigma_\omega = 4) = 0.99,$$

that is, the true inflation target is almost always between 0 and 4 percent.

A key parameter is s, the degree of nominal rigidity in the economy, and the response of output (growth) to a unit impulse to surprise inflation. Although the debate about this number is open, to allow a large scope for communication we set it at a relatively low $s = 0.25$: a 1-percent unexpected inflation raises output growth by 0.25 percent. From (8), this implies that the normalized bias/competence ratio q is decreasing in the nominal rigidity of the economy, s, as long as the relative weight on inflation in CB's preferences (conservatism) λ is smaller than 16. It seems plausible that any CB, even those publicly committed to price stability only, would care about real activity at least this much.

Consistent with the claims made by central bankers against their own time inconsistency, we allow for a strong conservatism $\lambda = 5$ and a tiny inflation bias, $b = 0.0002$. This implies a normalized bias

$$q = \frac{2 \cdot 0.02 \cdot \left(5 + \frac{1}{16}\right)}{H 5 \frac{1}{4}} \simeq \frac{1}{6H}.$$

We study the probability distribution of the equilibrium messages $\bar{\theta}$, for a range of values of the competence parameter $H = (1 + \sigma_\varepsilon^2)^{-1}$. By the Bayesian nature of the equilibrium and the martingale property of posterior beliefs, communication is on average unbiased:

$$\mathbb{E}[\bar{\theta}] = \sum_{k=0}^{K-1} \bar{\theta}_k^{(K)} \Pr(\theta_k^{(K)} \leq \theta < \theta_{k+1}^{(K)})$$

$$= \sum_{k=0}^{K-1} \int_{\theta_k^{(K)}}^{\theta_{k+1}^{(K)}} \theta \frac{e^{-(\theta^2/2)H}}{\sqrt{2\pi H}} d\theta$$

$$= \int_{-\infty}^{+\infty} \theta \frac{e^{-(\theta^2/2)H}}{\sqrt{2\pi H}} d\theta = 0.$$

Figure 1 reveals a few more, interesting facts.

First, even for the tiny inflation bias chosen, two hundredths of 1 percent, communication is fairly coarse. The number of credible messages never exceeds 6, as opposed to the uncountable support of θ. In fact, an inflation bias exceeding 1 results in $q > 2$, even for the perfectly competent CB ($H = 1$), and this can be shown to prevent any communication other than the simple two-message equilibrium. Therefore, any meaningful communication between CB and P requires that P perceive almost no meaningful bias in CB's preferences. Even if the temptation to surprise inflation is very modest and presents itself very occasionally (when the observed θ is extreme), it still hampers communication most of the time, especially for incompetent CBs.

Second, for every value of the competence parameter H, the message distribution is skewed above its average value of 2. This is a consequence of the inflation bias, which makes announcements of low inflation incredible, even when CB truly observes a signal of low required inflation θ. CB cannot announce a very low inflation target because P would not believe it.

Third, as communication is unbiased on average, to compensate for skewness CB must be relatively likely to announce below-average inflation. In order to offset the perception of bias, CB, especially if not highly competent, must speak as a moderately conservative one most of the time. This result has been proven in general for the two-message equilibrium.

Fourth, as CB's competence H rises and the normalized bias/competence ratio q declines, the maximum number of messages that can be sent by CB to P in equilibrium rises from 2 to 6.

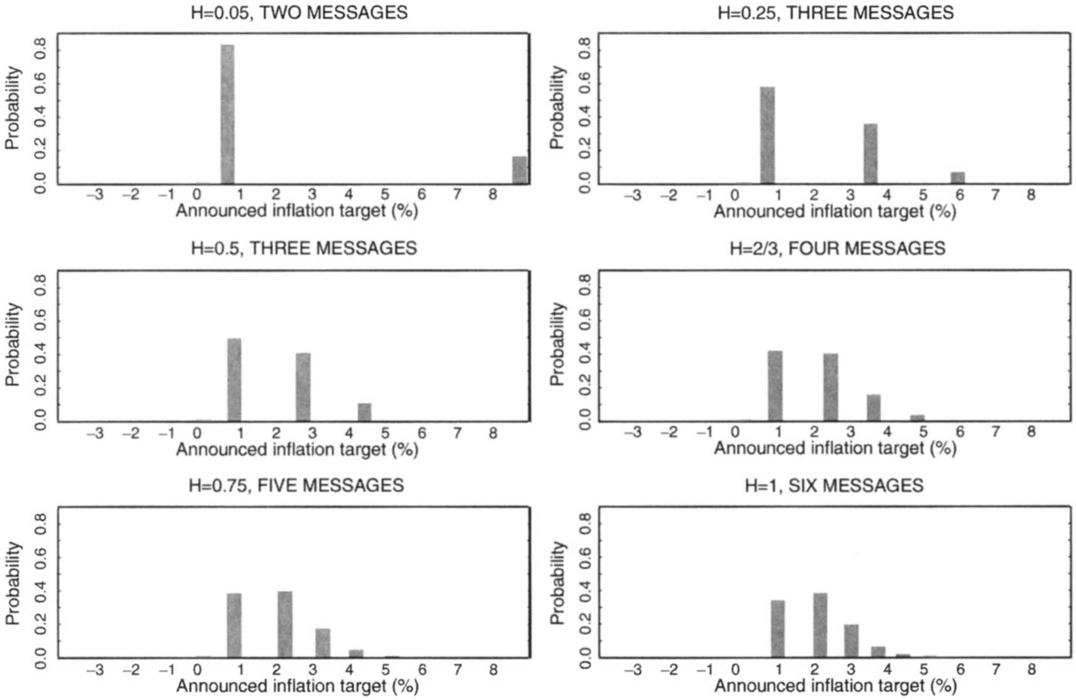

FIGURE 1

Notes: Examples of equilibrium (finest) message space for different values of competence H. In each panel the private signal θ is $\mathcal{N}(0, 1/H)$. Each panel illustrates the reports $\bar{\theta}$ of θ that the CB can credibly announce to P in equilibrium, and the unconditional probability of each announcement. The unconditional mean is 2.

At the same time, the distribution of messages becomes gradually more symmetric and concentrated around the mean, approaching the distribution $\mathcal{N}(0, H^{-1})$ of the private signal θ that CB truly observes. These phenomena are an illustration of the beneficial effect of competence on credibility.

Figure 2 quantifies the gains from communication in a scale-free metric, invariant to affine transformations of payoffs. We plot the percentage reduction in the variance of the private signal θ from the unconditional value H^{-1} to the value $\mathbb{E}[\text{Var}[\theta|\theta_k \leq \theta < \theta_{k+1}]|q, H]$ attained by the best possible equilibrium communication $\{\theta_k\}_{k=1}^{N(\sqrt{H}q)}$. This is a rough measure of the amount of information that CB can fruitfully pass to P. We see that this magnitude is substantial, in relative terms, and increasing in competence.

VI. Extensions

We now discuss the robustness of our results to extensions and variations of the model.[9] We show that, depending on the setup, competence either enhances or, in the least interesting cases, has no effect on credibility. We speculate on the circumstances under which the opposite result, that competence is detrimental to credibility, may emerge, although we could not find such an example, so this remains an open question. Our earlier discussion suggests that the basic intuition behind the result extends to all communication games where the sender has conflicting objectives represented by smooth payoff functions.

[9] I thank two anonymous referees for suggesting the analysis of the first three subsections.

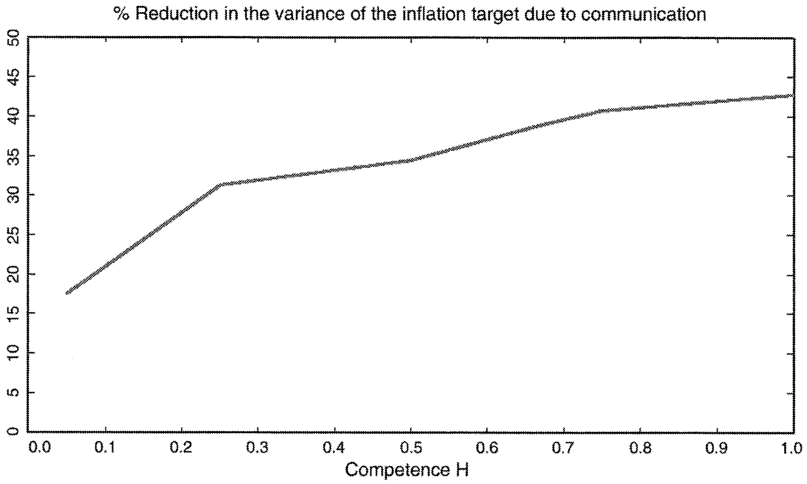

FIGURE 2. % REDUCTION IN THE VARIANCE OF THE INFLATION TARGET DUE TO COMMUNICATION

A. *Alternative Sources of Uncertainty*

Our model of CB's private information follows a tradition in the literature which specifies uncertainty in the desired inflation rate. This is a parsimonious but equivalent representation for several shocks, such as to the demand for money. This specification easily accommodates uncertainty about the mean inflation target π^*. Indeed, only the sum $\pi^* + \omega$ matters for equilibrium. We now consider alternative sources of uncertainty.

Suppose that the unknown state of nature ω is a shock that affects output

$$y = s(\pi - x) + \omega. \tag{14}$$

Equilibrium output y differs from ω only because of inflation surprises. As before, ω is Gaussian and CB privately observes a signal $\theta = \omega + \varepsilon$ of the shock ω contaminated with Gaussian noise ε. Now CB's competence consists of its ability to estimate this shock ω. The natural interpretation of ω is potential output, realized when inflation expectations are correct.

There are two natural specifications of CB's objective. First, suppose that CB wants to keep actual output above the natural rate by an amount b, so that

$$\pi(x) = \arg\min_{\pi} \mathbb{E}[(y - \omega - b)^2 + \lambda(\pi - \pi^*)^2 | \mathbb{I}_{CB}] = \frac{s(x + b) + \lambda\pi^*}{s + \lambda},$$

where the second equality follows after replacing output y from (14) and solving the minimization problem. In this case, uncertainty over the natural rate of output ω does not matter for equilibrium play: whatever potential output ω happens to be, CB wants to overshoot it by a known b. To do so, CB sets a best response that depends only on P's expectations, not on CB's private information. This is a standard Barro-Gordon game. Competence plays no role, because CB does not care about learning the state ω. Therefore, this specification is inappropriate to capture the idea that CB has private information that is relevant to its decisions.

In the second specification, CB does not know potential output and sets a target b for actual output:

$$\pi(x) = \arg\min_{\pi} \mathbb{E}[(y - b)^2 + \lambda(\pi - \pi^*)^2 | \mathbb{I}_{CB}]$$

$$= \arg\min_{\pi} \mathbb{E}[(s\pi - sx + \omega - b)^2 + \lambda(\pi - \pi^*)^2 | \mathbb{I}_{CB}].$$

The interpretation is that the natural rate of output is on average $\mathbb{E}[\omega] = 0$ and CB wants to exceed it by b. In this case, uncertainty about potential output matters. This case is formally equivalent to CB having private but imperfect information about its own inflation bias. For instance, CB may not know exactly the degree of monopoly distortions in the economy, as captured by ω. Or, a new (E)CB may be unsure about the preferences of its committee members.

Since this case is of particular interest, we go through the same steps as before to show that, and why the main results still obtain. Given P's inflation expectation $x(A) = \mathbb{E}[\pi|\theta \in A]$, CB's private information θ and its previous announcement A, CB solves

$$\min_{\pi}\{L(A, \pi|\theta) = \mathbb{E}[(s\pi - sx(A) + \omega - b)^2 + \lambda(\pi - \pi^*)^2|\theta]\}.$$

The unique best response is

$$\pi^{BR}(\theta, x(A)) = \frac{\lambda\pi^* + s^2x(A) + s(b - H\theta)}{\lambda + s^2}.$$

As before, competence H is the weight that CB places, when setting its policy, on its own private information θ about the state of the economy. The only formal, but substantively irrelevant, difference is that now a high potential output shock ω calls for a noninflationary response, i.e., θ enters negatively, because a high potential output ω (as suggested by a high θ) already stimulates output above the desired level b, calling for monetary tightening.

P's rational expectation of inflation is[10]

$$\begin{aligned} x(A) &= \mathbb{E}[\pi^{BR}(\theta, x)|\theta \in A] \\ &= \frac{\lambda\pi^* + sb + s^2x(A) - sH\bar{\theta}(A)}{s^2 + \lambda} \\ &= \pi^* + \frac{s}{\lambda}[b - H\bar{\theta}(A)]. \end{aligned}$$

Replacing this expression into CB's best-response inflation rate $\pi^{BR}(\theta, x(A))$, we obtain an equilibrium inflation rate given CB's private signal θ and announcement A:

$$\pi(\theta, A) = \pi^* + \frac{s}{\lambda}\left[b - H\frac{s^2\bar{\theta}(A) + \lambda\theta}{s^2 + \lambda}\right].$$

Subtracting expectations from this expression and rearranging terms, the equilibrium inflation forecast error

$$\pi(\theta, A) - x(A) = \frac{s}{s^2 + \lambda}H[\bar{\theta}(A) - \theta]$$

is positive if and only if P's expectations about the supply side shock ω, following CB's announcement, are too optimistic from CB's viewpoint. Using this expression in the aggregate supply curve gives an equilibrium (growth rate of) output

$$\begin{aligned} y(\theta, A) &= s[\pi(\theta, A) - x(A)] + \omega \\ &= \frac{s^2H}{s^2 + \lambda}[\bar{\theta}(A) - \theta] + \omega. \end{aligned}$$

Finally, conditional on the private realization of the optimal inflation target θ and on the message A that CB announces about θ, the expected loss in equilibrium is

$$\begin{aligned} \mathcal{L}(A|\theta) &= \mathbb{E}[(y(\theta, A) - b)^2 + \lambda(\pi(\theta, A) \\ &\quad - \pi^*)^2|\theta] = 1 - H \\ &\quad + \frac{[(\lambda + s^2)b - H(s^2\bar{\theta}(A) + \lambda\theta)]^2}{\lambda(\lambda + s^2)}, \end{aligned}$$

where we add and subtract $H\theta$ from the first

[10] Notice that, if we redefine $\hat{b} = b - H\theta$, then this CB behaves like one that has, and knows for sure (but privately), an inflation bias equal to $\hat{b}$. From P's viewpoint, this inflation bias is uncertain with $\hat{b} = b - H\theta \sim N(b, H)$. Therefore, in this case the "competence parameter" H measures the uncertainty of P's beliefs about CB's inflation bias. If $H = 0$, then P knows exactly how biased CB is, and no residual uncertainty remains. In fact, no announcement can be believed. The higher H, the more uncertain is P about the motives for CB's announcements, and the more P pays attention.

squared expression, then use $\mathbb{E}[(\omega - H\theta)|\theta] = 0$ and $\mathbb{E}[(\omega - H\theta)^2|\theta] = \text{Var}[\omega|\theta] = 1 - H$.

Now let

$$\hat{\theta} = -\frac{s^2}{\lambda}\bar{\theta}(A),$$

which is simply a "change of language" or message space, which does not depend on bias b or competence H. Rather than announcing A or, equivalently, $\bar{\theta} = \mathbb{E}[\theta|\theta \in A]$, CB announces a message $\hat{\theta}$, without changing the content of communication. Then we can rewrite the loss from announcement A as follows:

$$\mathcal{L}(A|\theta) = 1 - H + b^2\left(1 + \frac{s^2}{\lambda}\right) - 2bH(\theta - \hat{\theta}) + \frac{H^2\lambda(\theta - \hat{\theta})^2}{\lambda + s^2}.$$

The IC constraint that PBEC announcements in $\mathcal{A}$ must satisfy requires that, when observing $\theta \in A$, CB prefers to report message $\hat{\theta} = -(s^2/\lambda) \cdot \mathbb{E}_\theta[\tau|\tau \in A]$ rather than any other message in $\mathcal{A}$. Formally,

$$\mathcal{L}(\hat{\theta}|\theta) \leq \mathcal{L}(\hat{\theta}'|\theta).$$

Replacing the expressions above and recalling the definition of normalized bias/competence ratio q from (7), this is equivalent to

$$(\hat{\theta} - \hat{\theta}')(\hat{\theta} + \hat{\theta}' - 2\theta + sq) \leq 0.$$

From here, the analysis proceeds as in the previous case, with the only (irrelevant) difference that the normalized bias is sq rather than q, and the message space has been reformulated accordingly. Our main results go through. In particular, competence implies credibility (Proposition 4).

The intuition behind the main result also remains the same as in the baseline model. CB cares about surprising inflation expectations to stimulate output, but also about inflation per se, with weight λ. In the baseline model, CB is not sure about the latter objective. In this version of the model, it is not sure about how expansionary a given inflation rate is, because output is also driven by a supply-side shock ω. If CB did not care about inflation, $\lambda = 0$, then the game would be a standard Barro-Gordon game with just one biased objective, with the only added twist that the bias is $b - H\theta$ rather than b (cf. the best response inflation rate when $\lambda = 0$). Since θ has zero mean, P expects $b - H\theta$ to be equal to b on average. From the point of view of P, this is exactly like the perfect information situation. This game has no equilibrium, as both expected inflation and actual inflation would be pushed to infinity (cf. (2.2) with $\lambda = 0$). Whatever CB observes, on average P expects it to try to inflate the economy. Lack of CB's concerns for the costs of inflation leads to no equilibrium. Notice that the change of language is meaningful as long as $\lambda > 0$. When inflation matters ($\lambda > 0$) and its costs discipline monetary policy, then P can believe something of what CB says. But then the actual realization of private information θ matters, and again competence H measures the weight that CB places on θ in the output part of its objective. A very competent CB that observes a very high θ thinks that the supply shock ω is high and output needs no further monetary stimulus, so CB is less inclined to give in to temptation from the inflation b. Thus, its announcement of a low inflation rate is more credible than one made by an incompetent CB.

Two possible specifications of uncertainty remain, concerning the parameters λ and s. Preferences for inflation λ are plausibly CB's private information, but it is hard to think of reasons why a CB's governing body would be persistently uncertain about its own λ. The degree of nominal rigidity s, instead, is a structural parameter of the model, and not a preference parameter of CB. Thus, it is reasonable to think that CB has some private view/signal about s, but is not certain about its value. However, unlike shocks to desired inflation and to the natural rate of output, it is plausible that s is slowly changing, and thus quickly learned by P. We do not pursue this possibility further, because our starting point is the competence of CB in evaluating *fresh new information* every period, the type of uncertainty that makes discretionary policy valuable in the first place. As the dynamic analysis developed later suggests, the competence of CB matters most when the state of nature that CB estimates is sufficiently variable over time to preserve asymmetry of information.

B. *Alternative Specifications of Central Bank Competence and Preferences*

So far, we have defined the "competence" of CB to be its ability to accurately read the state of the economy and to set the optimal policy target accordingly. Alternatively, we can think of "competence" as the ability to actually implement a chosen target. Formally, fix the noise in the observation θ of the state ω, thus in the choice of the policy target π. Suppose that CB cannot directly control inflation, but rather an intermediate instrument m, such as the Federal Funds overnight rate, which affects inflation with some noise: $\pi = m + \zeta$, where $\zeta \sim \mathcal{N}(0, 1/\hat{H})$. Then $\hat{H}$ captures this second meaning of ability. It is clear that the structure of equilibria is the same as in the game analyzed in this paper, but changes in "competence" $\hat{H}$ have no impact on equilibrium credibility. Hence, the beneficial effects of competence on credibility are restricted to the meaning given in this paper to the term "competence." This raises a semantic issue: which is the most appropriate notion of CB's "competence?" This question has a substantive counterpart: what is the hardest part of a CB's job—choosing the appropriate course of action, or implementing it? The current practice of monetary policy and statements by its practitioners suggests that, consistently with the view taken in this paper, CB's main concern is the former task, formulating the optimal policy, albeit considerable resources are also devoted to implement it correctly and effectively.

The quadratic loss function specification widely adopted in Barro-Gordon policy games, and often justified as a local Taylor approximation to a more general loss function, is important for the results. In particular, the marginal cost of deviating from truth-telling is initially negligible and increases slowly. Suppose instead that CB is very risk-averse concerning large deviations of inflation and output from a reasonable range: both deflation and excessive inflation have ominous consequences. For example, CB engages in "macro-risk management." If the acceptable range is common knowledge and CB privately formulates a target for inflation or output within that range, then our intuition will stand. An incompetent CB risks making big mistakes when trying to fine-tune the economy, so it cannot be trusted when announcing a target, especially near the boundaries of the range. A competent CB would be more comfortable taking such risks, so one would expect it to be more credible in those circumstances. If, instead, CB has private information only about the acceptable range for inflation and output, rather than about numerical values of the targets, then a less competent CB will be more conservative when trying to manipulate output, for fear of making the fatal big mistake. In this quite different scenario, competence is potentially detrimental to credibility, although it remains an open question whether other effects might be sufficiently strong to preserve the basic result.

C. *Dynamics*

What happens when the game is repeated over and over? Reputational effects from repeated play can sustain a host of equilibria and lessen incentive constraints in the presence of private information. As mentioned earlier, the kind of CB private information that we are interested in concerns shocks that are not too persistent, because those are the types of shocks that require flexibility to be addressed.

The literature on repeated games with evolving private information is thin. The closest reference is Athey, Atkeson, and Kehoe (2005), who study a repeated version of this monetary policy game. As already mentioned, in order to preclude any role for communication, they assume that states ω are i.i.d. over time, and expectations are formed each period by P before CB privately observes the signal θ, which is equal to ω; i.e., CB is perfectly competent. Through a dynamic mechanism design approach, they characterize the best equilibrium of the infinitely repeated game, which they implement via an optimal inflation cap: CB is prohibited from setting inflation above some optimal mandated level. Our analysis differs from theirs in two respects: we are interested in the role of communication, and we assume away the ability of P to restrict CB's actions through binding rules.

There exists no analysis of repeated partisan advice. A few authors have analyzed repeated rounds of cheap talk communication ("conversation") before playing a one-shot game (Robert Aumann and Sergiu Hart 2003; Vijay Krishna and John Morgan 2004). The idea of a two-sided conversation does not seem too relevant in a monetary policy context, where CB faces a P comprising millions of people. Here we are interested in a different kind of situation, where each period the state of the economy changes, CB has some private information about it, and a new stage game is played.

To simplify the analysis and to gain some insight, we restrict attention to the case of i.i.d. draws for the state of the economy ω across periods, as in Athey, Atkeson, and Kehoe (2005), but we maintain our timing of actions in the stage game to preserve a role for communication: P sets expectations after CB observes θ. We also assume that P observes without error the inflation rate chosen by CB and, ex post, CB's realized signal θ.[11]

A complete characterization of the equilibrium set of such a complex game goes well beyond the scope of this article. As in the static case, we focus on equilibria where the amount of communication is maximized. A standard way to support this favorable outcome is grim trigger. Since "babbling" (no communication) is always a PBEC of the one-shot game, "babbling forever" is also a PBEC of the repeated game. Hence, P can credibly threaten CB to revert to babbling forever, that is, to never trust CB again, if CB ever lies. We now show that this threat can support truth-telling for a sufficiently competent CB.

If CB always tells the truth and P believes whatever CB says, the flow loss is

$$\mathcal{L}_T = \mathcal{L}(\theta|\theta) = b^2\left(1 + \frac{s^2}{\lambda}\right) + \lambda(1 - H).$$

In the repeated babbling equilibrium where no announcement is believed, $\bar{\theta} = 0$ always, and the expected flow loss for a CB that observed a signal θ is

$$\mathcal{L}_B = \mathbb{E}[\mathcal{L}(0|\theta)] = \mathbb{E}\left[\frac{s^2H^2\lambda}{s^2 + \lambda}\,\theta^2\right] - \mathbb{E}[2bsH\theta]$$
$$+ \mathcal{L}_T = \frac{s^2H^2\lambda}{s^2 + \lambda}\,\mathrm{Var}[\theta]$$
$$+ \mathcal{L}_T = \frac{s^2H\lambda}{s^2 + \lambda} + \mathcal{L}_T.$$

Importantly, a more competent CB enjoys a better outcome from the babbling equilibrium, thanks to its ability to manage the economy:

$$\frac{d\mathcal{L}_B}{dH} = \frac{s^2\lambda}{s^2 + \lambda} - \lambda = -\frac{\lambda^2}{s^2 + \lambda} < 0.$$

Therefore, the threat of grim trigger by P is less punishing for a more competent CB. This implies that a more competent CB's incentives are not obviously in favor of truth-telling.

Finally, the one-shot gain from deviating for one period from truth-telling, when P believes whatever CB says, is given by the best possible lie. After choosing the optimal lie and plugging it back:

$$\min_{\bar{\theta}} \mathcal{L}(\bar{\theta}|\theta) = \frac{s^2H^2\lambda}{s^2 + \lambda}\left(b\,\frac{s^2 + \lambda}{s\lambda H}\right)^2$$
$$- 2bsHb\,\frac{s^2 + \lambda}{s\lambda H} + \mathcal{L}_T$$
$$= \mathcal{L}_T - \frac{s^2 + \lambda}{\lambda}\,b^2 = \mathcal{L}_D,$$

which is independent of θ and is also smaller than the flow loss in the truth-telling equilibrium, $\mathcal{L}_T$, by an amount independent of competence H.

[11] We could also assume that P never observes θ. CB would never deviate by first telling a "lie" (a message not associated to the draw θ in the prescribed equilibrium) and then choosing an inflation rate that is consistent with the lie. This announcement-action sequence would preserve equilibrium play in the future, because deeds would be matched to words and P would not detect the lie. But it would also increase the short-run loss relative to equilibrium play. In fact, CB would force itself to take the wrong action just to preserve credibility, and it can achieve this goal more cheaply by not lying and taking the appropriate monetary action (talk is cheap, inflation is not). Hence, we restrict attention to deviations from equilibrium announcements, followed by the static best inflation rate given P's expectations. This always allows P to detect lies, as if P truly observed the θ realization.

Truth-telling is supported as a Perfect Bayesian Nash Equilibrium of the repeated game if the present discounted value of the loss from truth-telling is smaller than the flow loss from deviating and surprising a trusting P and then being punished with reversion to babbling forever:

$$\frac{\mathcal{L}_T}{1-\beta} \leq \mathcal{L}_D + \beta \frac{\mathcal{L}_B}{1-\beta},$$

where $\beta \in (0, 1)$ denotes the discount factor. Replacing the expressions above and simplifying, this condition is equivalent to

$$(15) \qquad \left(b\,\frac{s^2 + \lambda}{s\lambda}\right)^2 \leq \frac{\beta}{1-\beta} H,$$

which is more likely to hold, the more competent, unbiased, and patient CB is. In this sense, a more competent CB is more credible, because it is more likely to be able to support a truth-telling equilibrium through grim trigger.

The Appendix investigates the structure of equilibria when (15) fails. We show that equilibria with communication supported by grim trigger take again a partitional form, just like in the static case, and we provide an algorithm to compute such equilibria and study the comparative statics properties of changes in competence H when this is too small to support truth-telling.

D. *Alternative Impediments to Credibility*

The beneficial impact of (a reputation for) competence on credibility that we have discussed in this paper may apply beyond the inflation bias to other sources of time inconsistency. A formal analysis is beyond the scope of this paper. But two important examples come to mind and are worth discussing.

First, the "stabilization bias" (Clarida, Galí, and Gertler 1999): when an adverse cost shock hits the economy, CB would like to commit to raise future nominal interest rates, in order to subdue current inflation expectation. Once this goal has been attained, however, the promised rise in interest rates is undesirable, thus time-inconsistent. If CB has some superior information about either the cost shock or just its own interpretation of the effects, then it will have an incentive to downplay the inflation risks, or to announce/signal a tightening bias for the medium run. But P should not believe CB's announcements and stated intentions, due to the aforementioned time inconsistency. Then our mechanism comes into play.

Second, the "CNBC effect" (Morris and Shin 2002): when private agents have a coordination motive that depends on an unobserved state of the economy, and valuable private information on this state is dispersed across the population, then a public announcement that becomes common knowledge serves as a focal point for coordination of expectations and actions. Private agents place a suboptimal weight on their own private information, in favor of the commonly known public information, because they expect the other agents to do the same and to coordinate their actions accordingly. As a consequence, if the public announcement is imprecise relative to private information, too much of the latter is ignored in formulating expectations and actions, so the public announcement is socially harmful on average.

Morris and Shin do not allow the public announcement to be withheld, but in our context this possibility becomes natural. In their linear Gaussian model, in many respects similar to ours, CB would always choose a bang-bang solution, either full communication or no communication, depending on the precision of CB's information (see Christian Hellwig 2002 for a similar conclusion in a monetary general equilibrium model). This particular effect, however, unlike the general idea behind the CNBC effect, appears to be model-specific. It is not too difficult to imagine a different setup where a relatively incompetent CB, one that is moderately confident in its information, would add noise to its announcements and engage in "opaque" communication, or be intentionally vague, again in the style of the cheap talk equilibrium analyzed in this paper. Again, competence and transparency would go hand in hand, at least over some range.

E. *Reputation for Competence and Reputation for Conservatism*

Canice Prendergast and Lars Stole (1996) first analyzed the effects of a concern for a reputation for competence on noncheap actions, and Morris (1997) on cheap talk, while Morris (2001) focuses on the effects of a reputation for bias on cheap talk. Suppose that both CB's competence and its bias are private information. Therefore, CB's decision to commit or not to a rule (to resist proposed incentive schemes or delegation) may convey information about its competence *and* bias, and have an impact on P's expectations about policy and the credibility of announcements. CB's choice of a discretionary policy may be interpreted by P not as a sign of bias, or lack thereof, but rather as an expression of competence and of the value attributed to flexibility. Discretion is not necessarily punished by P with high inflationary expectations, as it usually is in the literature on time-inconsistent monetary policy, because CB has many motives to choose discretion.

In this light, the ECB's rigidity, relative to the Fed, can be due either to the need to convince the private sector that the ECB did not inherit the historical bias of the Bundesbank, or due to its "incompetence" relative to the formidable task that it faces. Conversely, the Fed's activism and ability to shape expectations through communication may be explained by the reputation for competence acquired during Greenspan's tenure, which resolves any doubts of inflation-biased discretionary policy.

F. *Central Bank's Over/Underconfidence*

Finally, relax the assumption of common prior beliefs. Suppose that CB believes it is more competent than P thinks it is. That is, CB is relatively overconfident in its own ability, and the two players agree to disagree on the key parameter H. As documented by an ample psychology literature, human beings have an almost innate tendency to overestimate their own ability. Clearly, *CB's overconfidence has a beneficial effect on communication and welfare.* The reason is that CB is benevolent, and its incentives to truthful communication depend on its perceived degree of competence. A confident CB is self-disciplined, even if it thinks that P incorrectly (in CB's view) disagrees with this assessment, it acts paternalistically, and it chooses a transparent approach. In turn, P believes CB because it knows that CB believes it is right.

Conversely, suppose that competence is CB's private information, which is discovered by P over time from monetary policy performance. If CB gets lucky early on, then it enjoys what it knows to be an excessive reputation for competence, and becomes relatively underconfident. Nonetheless, due to the beneficial effects of (reputation for) competence on credibility, CB has no reason to reveal its incompetence. This result stands in contrast to that in Morris and Shin (2002), where a benevolent policymaker who knows it possesses valuable but imprecise information would prefer not to announce it publicly. When both bias and the CNBC effect are present, CB would have to trade off the two effects in relaying to P its confidence in its own forecasts.

VII. Conclusions

Communication and transparency are taking center stage in the analysis and in the practice of monetary policy. We analyze a simple setting where a central bank has an incentive to communicate to the private sector its intentions, based on its own reading of the state of the economy. Communication finds an impediment in a classic inflation bias, giving rise to time inconsistency. We uncover two separate beneficial effects of the competence of a CB at doing its job on the credibility and transparency of its monetary policy.

Our findings suggest a possible explanation for the common wisdom that a CB's governing body should speak with one voice. Dissenting and conflicting opinions leaked by its members to the public are likely to convey the impression of confusion and/or diverging views. In terms of our model, this amounts to "incompetence": CB appears to lack a clear and focused idea of what needs to be done, even if individual members of the board think they do.[12] This public perception undermines CB's credibility and ability to "manage" markets' expectations. Knowing that its words carry less weight in expectations, CB exag-

[12] De Haan and David-Jan Jensen (2004) find that statements by different (groups of) European central bankers on the interest rate, inflation, and economic growth have indeed been contradictory, although they have been converging over time.

gerates its announcements, making things even worse. This may be an equilibrium outcome of our game, without resorting to some form of bounded rationality in markets, only if CB's competence is its private information, so that it cares about its reputation for competence. Exploring this environment is a natural direction of future research.

Another open question is under which circumstances, if any, the opposite result holds, namely, competence is detrimental to credibility. While this is an interesting possibility, Section VI shows that, within our class of monetary policy games, the beneficial effects of competence on credibility are robust. Indeed, in this respect, there is nothing special about central banking, and the analysis supports the same intuition in any communication game where the sender has conflicting objectives and is uncertain about the nonbiased part of his preferences. An important example is an electoral campaign, where each candidate has an ideological bias, but also a concern for setting a correct social policy. An extreme but very competent candidate may prevail over a centrist incompetent opponent. More concretely, consider a candidate for president, who is leading in the polls and has a known preference for strong employment protection, but believes that the economy in its current state cannot tolerate rigid labor markets. If he is perceived to be very competent on labor issues, he will be able to make a credible campaign promise to introduce no new employment protection. Otherwise, whatever he says, mass layoffs will occur when his election appears likely.

Appendix

PROOF OF LEMMA 4:

First, we show that the sequence $\{\theta_k\}$ is bounded below. By contradiction: for every N finite there exists $k > -\infty$ such that $N \in [\theta_{k-1}, \theta_k)$ for some $\theta_{k-1} > -\infty$. By definition of PBEC and $\bar{\theta}_{k-1} = \mathbb{E}[\theta|\theta_{k-1} \leq \theta < \theta_k]$, this means that

$$\mathbb{E}[\theta|\theta_{k-1} \leq \theta < \theta_k] = 2\theta_k - q - \mathbb{E}[\theta|\theta_k \leq \theta < \theta_{k+1}] > \mathbb{E}[\theta|\theta < \theta_k]$$

because θ has full support on the real line. Rearranging,

$$\theta_k - \mathbb{E}[\theta|\theta < \theta_k] = f(\theta_k|H) > q + \mathbb{E}[\theta|\theta_k \leq \theta < \theta_{k+1}] - \theta_k > q > 0.$$

Since this must be true for every N and every $\theta_{k-1} \leq N$ unbounded below, this contradicts Lemma 3.

Second, we show that the sequence $\{\theta_k\}$ is bounded above. Using $\bar{\theta}_k = 2\theta_k - q - \bar{\theta}_{k-1} \geq \theta_k$, $\mathbb{E}[\theta|\theta > \theta_k] \geq \bar{\theta}_k$, and rearranging,

$$\mathbb{E}[\theta|\theta > \theta_k] - \theta_k = g(\theta_k|H) \geq \bar{\theta}_k - \theta_k = \theta_k - \bar{\theta}_{k-1} - q \geq 0,$$

where the last inequality follows from $\bar{\theta}_k \geq \theta_k$. By contradiction, suppose θ_k grows unbounded with k. Then by Lemma 3, $\lim_{k\to\infty} g(\theta_k|H) = 0$, so for every $\varepsilon > 0$ there exists K_ε such that $g(\theta_k|H) < \varepsilon$ for all $k > K_\varepsilon$, and therefore

$$\varepsilon > \theta_k - \bar{\theta}_{k-1} - q \geq 0,$$

which in turn implies

$$\text{(A1)} \qquad \lim_{k\to\infty} (\theta_k - \bar{\theta}_{k-1}) = q.$$

It follows

$$\lim_{k\to\infty} [\theta_k - \bar{\theta}_{k-1} - (\theta_{k+1} - \bar{\theta}_k)] = 0.$$

So for every $\delta > 0$, there is K_δ such that for all $k > K_\delta$

$$\theta_{k+1} - \theta_k < \bar{\theta}_k - \bar{\theta}_{k-1} + \delta = 2(\theta_k - \bar{\theta}_{k-1}) + \delta - q,$$

where in the last equality we used (11). If we take $k > \max\langle K_\delta, K_\varepsilon\rangle$,

$$\theta_{k+1} - \theta_k < 2(q + \varepsilon) + \delta - q = q + 2\varepsilon + \delta.$$

This fact, along with $\theta_{k+1} - \theta_k > q$ from Proposition 2, finally proves

$$\text{(A2)} \qquad \lim_{k\to\infty} (\theta_k - \theta_{k-1}) = q.$$

So we have established that we can always satisfy

$$\theta_{k-1} + q - \varepsilon < \bar{\theta}_{k-1} < \theta_k < \theta_{k-1} + q + \delta.$$

It follows

$$\theta_{k-1} + q - \varepsilon < \bar{\theta}_{k-1} = \mathbb{E}[\theta | \theta_{k-1} \leq \theta < \theta_k]$$

$$< \mathbb{E}[\theta | \theta_{k-1} \leq \theta < \theta_{k-1} + q + \delta]$$

$$= \frac{\int_{\theta_{k-1}}^{\theta_{k-1}+q+\delta} \theta e^{-H(\theta^2/2)}\, d\theta}{\int_{\theta_{k-1}}^{\theta_{k-1}+q+\delta} e^{-H(\theta^2/2)}\, d\theta}.$$

But notice that for $\varepsilon = \delta = 0$, the leftmost term is strictly larger than the rightmost term, so this inequality must be violated for some ε, δ small enough, and $k > \max\langle K_\delta, K_\varepsilon \rangle$, the desired contradiction.

Dynamic Analysis when (15) Fails.—When condition (15) fails, we can construct equilibria with partial communication, supported by the threat of grim trigger to "babbling forever" (flow loss $\mathcal{L}_D$) after P detects a lie by CB "outside the equilibrium" partition (message space).

First, we show that equilibria are again partitional also in this dynamic setting. Given any equilibrium, whether partitional or not, the associated payoff when observing θ and announcing $\bar{\theta}$ is

$$\mathcal{L}(\bar{\theta}|\theta) + \frac{\beta}{1-\beta} \mathbb{E}[\mathcal{L}(\bar{\theta}|\theta)] = \frac{s^2H^2\lambda}{\lambda + s^2} (\theta - \bar{\theta})^2$$

$$- 2bsH(\theta - \bar{\theta}) + \mathcal{L}_T + \frac{\beta}{1-\beta} \mathbb{E}\left[\frac{s^2H^2\lambda}{\lambda + s^2}\right.$$

$$\left.\times (\theta - \bar{\theta})^2 - 2bsH(\theta - \bar{\theta}) + \mathcal{L}_T\right] = \frac{s^2H^2\lambda}{\lambda + s^2}$$

$$\times (\theta - \bar{\theta})^2 - 2bsH(\theta - \bar{\theta}) + \frac{s^2H^2\lambda}{\lambda + s^2} \frac{\beta}{1-\beta}$$

$$\times \mathbb{E}[(\theta - \bar{\theta})^2] + \frac{\mathcal{L}_T}{1-\beta},$$

where $\mathbb{E}[\mathcal{L}(\bar{\theta}|\theta)]$ is the expected flow loss from equilibrium play, and $\mathbb{E}[(\theta - \bar{\theta})^2]$ is the expected variance of θ conditioning on equilibrium communication. The incentive constraint, comparing equilibrium play to a deviation to another of the equilibrium messages followed by babbling forever, is

$$\mathcal{L}(\bar{\theta}|\theta) + \frac{\beta}{1-\beta} \mathbb{E}[\mathcal{L}(\bar{\theta}|\theta)] \leq \mathcal{L}(\bar{\theta}'|\theta)$$

$$+ \frac{\beta}{1-\beta} \mathcal{L}_B ;$$

$$\frac{s^2H^2\lambda}{\lambda + s^2} (\theta - \bar{\theta})^2 - 2bsH(\theta - \bar{\theta})$$

$$+ \frac{\beta}{1-\beta} \frac{s^2H^2\lambda}{\lambda + s^2} \mathbb{E}[(\theta - \bar{\theta})^2] + \frac{\mathcal{L}_T}{1-\beta}$$

$$\leq \frac{s^2H^2\lambda}{\lambda + s^2} (\theta - \bar{\theta}')^2 - 2sbH(\theta - \bar{\theta}') + \mathcal{L}_T$$

$$+ \frac{\beta}{1-\beta} \left(\frac{s^2H\lambda}{s^2 + \lambda} + \mathcal{L}_T \right);$$

rearranging terms,

$$\frac{s^2H^2\lambda}{\lambda + s^2} [\bar{\theta}^2 - \bar{\theta}'^2 - 2\theta(\bar{\theta} - \bar{\theta}')] + 2bsH$$

$$\times (\bar{\theta} - \bar{\theta}') + \frac{s^2H^2\lambda}{\lambda + s^2} \frac{\beta}{1-\beta} \mathbb{E}[(\theta - \bar{\theta})^2]$$

$$\leq \frac{\beta}{1-\beta} \frac{s^2H\lambda}{s^2 + \lambda};$$

$$(\bar{\theta} - \bar{\theta}')(\bar{\theta} + \bar{\theta}' - 2\theta + q)$$

$$\leq \frac{\beta}{1-\beta} \left\{ \frac{1}{H} - \mathbb{E}[(\theta - \bar{\theta})^2] \right\}.$$

The payoff from a deviation (the left-hand side) must fall short of the damage (the right-hand side), which is the decrease in the variance of the signal θ from the prior $1/H = \mathbb{E}[\theta^2]$ to the posterior after communication $\mathbb{E}[(\theta - \bar{\theta})^2] \leq 1/H$. This latter inequality is the same condition as in the static game (cf. equation (6)), except that the zero on the right-hand side is replaced by a positive magnitude.

Let

$$\hat{q} \equiv q - \frac{\beta}{1-\beta}\frac{1}{\bar{\theta}-\bar{\theta}'}\left(\frac{1}{H} - \mathbb{E}[(\theta-\bar{\theta})^2]\right).$$

When $\bar{\theta}' < \bar{\theta}$, namely, when comparing the equilibrium message with a downward lie, this is equivalent to reduce the normalized bias, as $\hat{q} < q$, and the inequality becomes

$$\bar{\theta}' \leq 2\theta - \bar{\theta} - \hat{q},$$

which is less binding than the analogous condition in the static case $\bar{\theta}' \leq 2\theta - \bar{\theta} - q$. That is, lower equilibrium messages $\bar{\theta}'$ can be closer to the message $\bar{\theta}$ prescribed for type θ. So CB is more credible because the lie is followed by a punishment. When $\bar{\theta}' > \bar{\theta}$, when comparing the equilibrium message with an upward lie, this is equivalent to increase the normalized bias, as $\hat{q} > q$, and the inequality becomes

$$\bar{\theta}' \geq 2\theta - \bar{\theta} - \hat{q},$$

which is less binding than the analogous condition in the static case $\bar{\theta}' \geq 2\theta - \bar{\theta} - q$. In either direction, the equilibrium takes the form of a partition of the state space, where the interval's bounds solve the two conditions above with equality, and messages are as usual the conditional means of θ inside those intervals:

$$\theta_k = \frac{\bar{\theta}_{k-1} + \bar{\theta}_k + \hat{q}}{2};$$

$$\bar{\theta}_k = \sqrt{\frac{H}{2\pi}}\int_{\theta_{k-1}}^{\theta_k} \theta e^{-(H\theta^2/2)}\, d\theta.$$

This double recursion allows us to construct all equilibria, similarly to the static case, and to verify the comparative statics effects of changes in H. For given parameter values, the best equilibrium's partition is finer than in the static case. Given the tail properties of the normal distribution, there always exists a two-message equilibrium.

REFERENCES

Athey, Susan, Andrew Atkeson, and Patrick J. Kehoe. 2005. "The Optimal Degree of Discretion in Monetary Policy." *Econometrica*, 73(5): 1431–75.

Aumann, Robert J., and Sergiu Hart. 2003. "Long Cheap Talk." *Econometrica*, 71(6): 1619–60.

Austen-Smith, David. 1990. "Information Transmission in Debate." *American Journal of Political Science.* 31(1):124–152.

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

Battaglini, Marco, and Roland Bénabou. 2003. "Trust, Coordination, and the Industrial Organization of Political Activism." *Journal of the European Economic Association*, 1(4): 851–89.

Bénabou, Roland, and Guy Laroque. 1992. "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility." *Quarterly Journal of Economics*, 107(3): 921–58.

Bernanke, Ben S. 2004. "Fedspeak." Paper presented at the One Hundred Sixteenth Annual Meeting of the Allied Social Science Associations, San Diego, CA.

Blinder, Alan S. 1998. *Central Banking in Theory and Practice* (Lionel Robbins Lectures). Cambridge and London: MIT Press.

Blinder, Alan S. 1999. "Central Bank Credibility: Why Do We Care? How Do We Build It?" National Bureau of Economic Research Working Paper 7161.

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

Clarida, Richard, Jordi Galí, and Mark Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature*, 37(4): 1661–1707.

Crawford, Vincent P., and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica*, 50(6): 1431–51.

Cukierman, Alex. 1995. *Central Bank Strategy, Credibility, and Independence: Theory and Evidence*. Cambridge, MA, and London: MIT Press.

de Hann, Jakob, and David-Jan Jansen. 2004.

"Look Who's Talking: ECB Communication during the First Years of EMU." CESifo Working Paper 1263.

Geraats, Petra M. 2002. "Central Bank Transparency." *Economic Journal*, 112(483): F532–65.

Hellwig, Christian. 2002. "Public Announcements, Adjustment Delays, and the Business Cycle." Unpublished.

Holmström, Bengt. 1982. " Managerial Incentive Problems: A Dynamic Perspective." In *Essays in Economics and Management in Honor of Lars Wahlbeck.* Helsinki: Swedish School of Economics.

Krishna, Vijay, and John Morgan. 2004. "The Art of Conversation: Eliciting Information from Experts through Multi-Stage Communication." *Journal of Economic Theory*, 117(2): 147–79.

Morris, Stephen. 1997. "The Reputational Cost of Truthful Information Transmission." Unpublished.

Morris, Stephen. 2001. "Political Correctness." *Journal of Political Economy*, 109(2): 231–65.

Morris, Stephen, and Hyun Song Shin. 2002. "Social Value of Public Information." *American Economic Review*, 92(5): 1521–34.

Ottaviani, Marco. 2000. "The Economics of Advice." Unpublished.

Ottaviani, Marco, and Peter N. Sorensen. 2006. "Reputational Cheap Talk." *RAND Journal of Economics,* 37(1):155–175.

Prendergast, Canice, and Lars Stole. 1996. "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning." *Journal of Political Economy*, 104(6): 1105–34.

Romer, Christina D., and David H. Romer. 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review*, 90(3): 429–57.

Stein, Jeremy C. 1989. "Cheap Talk and the Fed: A Theory of Imprecise Policy Announcements." *American Economic Review*, 79(1): 32–42.

Svensson, Lars E. O. 2003. "What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules." *Journal of Economic Literature*, 41(2): 426–77.

Waller, Sandra, and Jakob de Haan. 2004. "Credibility and Transparency of Central Banks: New Results Based on IFO's *World Economic Survey.*" CESifo Working Paper 1199.

Woodford, Michael. 2003. "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, ed. Philippe Aghion, Ronald Frydman, Joseph Stiglitz, and Michael Woodford, 25–58. Princeton and Oxford: Princeton University Press.