

Forward Guidance: Communication, Commitment, or Both?

IFS Working Paper W19/20

Marco Bassetto

Forward Guidance: Communication, Commitment, or Both?*

Marco Bassetto[†]

July 25, 2019

Abstract

A policy of forward guidance has been suggested either as a form of commitment (“Odyssean”) or as a way of conveying information to the public (“Delphic”). I analyze the strategic interaction between households and the central bank as a game in which the central bank can send messages to the public independently of its actions. In the absence of private information, the set of equilibrium payoffs is independent of the announcements of the central bank: forward guidance as a pure commitment mechanism is a redundant policy instrument. When private information is present, central bank communication can instead have social value. Forward guidance emerges as a natural communication strategy when the private information in the hands of the central bank concerns its own preferences or beliefs: while forward guidance per se is not a substitute for the central bank’s commitment or credibility, it is an instrument that allows policymakers to leverage their credibility to convey valuable information about their future policy plans. It is in this context that “Odyssean forward guidance” can be understood.

*I thank Gadi Barlevy, Jeffrey Campbell, Martin Cripps, Mariacristina De Nardi, Fernando Duarte, Marco Del Negro, Charles L. Evans, Antonella Ianni, Luigi Iovino, Spencer Krane, Eric Mengus, Benoît Mojon, and Thomas J. Sargent for helpful conversations, and Matthew Easton for valuable research assistance. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

[†]Federal Reserve Bank of Chicago and IFS. Email address: bassetto@nber.org

1 Introduction

When interest rates attained their effective lower bound in the United States, the Euro Area, Great Britain, and Japan, monetary authorities across the world looked for alternative tools to stimulate the economy. The resulting unconventional monetary policy revolved around two main strategies: forward guidance and quantitative easing. In this paper, we will consider the conditions which make forward guidance effective. We define forward guidance as a situation in which central banks provide direct statements about the future path of their policy tools. To some extent, central banks have provided such statements for years, as part of their broader discussion of their view on the underlying conditions of the economy. In fact, Campbell et al. (2012) find evidence that forward rates often reflect ahead of time what would appear as a monetary policy shock in VARs which are purely based on spot-market interest rates and macroeconomic variables. What has been different recently is that announcements have become more explicit,¹ and that they have been tied to a desire to precommit future policy.² Campbell et al. (2012) emphasized this distinction by defining “Odyssean” forward guidance a situation in which monetary authorities make statements with the primary objective of committing their future policy, and “Delphic” forward guidance a situation in which statements about future policy are primarily meant to share with the public any superior information that the central bank may have about the future course of policy. In practice, forward guidance is of course never purely Odyssean nor ever purely Delphic. An important message of this paper is that the “Odyssean” and “Delphic” elements of forward guidance are intrinsically linked: when a central bank has no superior information compared to the public (the purely “Odyssean” case), forward guidance will be shown to be a redundant policy instrument: while expectations about future policy actions matter, an explicit message giving advance notice of those actions is not needed. Furthermore, while pure Delphic forward guidance could be valuable on its own, a central bank will find

¹Campbell et al. (2017, 2019) discuss the evidence about the shifting nature of forward guidance.

²Commitment is very often desirable when a policymaker faces forward-looking agents, as emphasized by Kydland and Prescott (1977). Within the new Keynesian framework, commitment to keep interest rates at zero for longer than would be optimal ex post is particularly valuable; see e.g. Eggertsson and Woodford (2003) and Werning (2011).

it much easier to sustain truthful announcements about its future intentions when these announcements form part of a bigger equilibrium in which central bank credibility is at stake.

The approach that we pursue in this paper is based on the theoretical literature on cheap talk.³ Cheap talk refers to a situation in which a player in a game has the possibility of sending messages that have no direct consequences on the set of future actions available to the players nor on their payoffs. Because of the lack of direct consequences, equilibria in which these messages are ignored are always present, but cheap talk opens the possibility for Pareto superior equilibria to emerge, in which messages reveal some information. When the messenger and the receiver of the message have conflicting objectives, full disclosure of private information will often not be possible.

The current policies of forward guidance map well into the theoretical framework of cheap talk. While central banks around the world have stated their intention not to raise interest rates for extended periods of time, these statements have not *directly* affected their ability to do so. As an example, the Federal Open Market Committee continued to meet eight times a year, and a simple vote on each of these occasions could have led to a rate increase, independently of previous statements. Naturally, such a course of action could have been detrimental in that it would have led to a loss of credibility; this *indirect*, endogenous response of private-sector expectations plays an important role in our analysis. Similarly, the statements per se do not have a direct effect on macroeconomic fundamentals nor on welfare; to the extent that they have been successful, it is because they influenced the private sector's expectations about the future course of policy. Furthermore, as in all interesting applications of cheap talk, the central bank's incentives are likely to be misaligned with those of the public, at least in the short run: as an example, in the throes of a major recession, the monetary authorities would most likely prefer sending optimistic messages and try to prevent expectations from adding to a downward spiral. The temptation to manipulate messages for short-term gains must be tempered with the potential loss of credibility.

In this paper, we show that forward guidance will be a particularly valuable policy tool when two conditions are met: first, the central bank must have some private information;

³See Crawford and Sobel (1982).

second, this private information concerns the central bank's preferences or beliefs. The first point is straightforward: cheap talk is redundant if all the players in a game have symmetric information. The second condition highlights situations in which central bank communication is naturally thought as forward guidance, rather than a more general notion of transparency. As an example, if the central bank has superior information about the current state of the economy, it would be most natural for the central bank to make statements about the state of the economy itself, rather than revealing it indirectly through its intended course of action, which may be a poor proxy for the information that the private sector needs. In such a context, we will show an extreme example in which messages about future policy are *useless*, whereas statements about the state of the economy help in coordinating the private sector.

Forward guidance is a natural message space when the central bank has private information that is only useful to the private sector to improve its forecasts of future policy. In particular, this happens when private information is about a central bank's own preferences and/or beliefs: private agents may not care directly about them, but knowing them helps predicting how the central bank will behave. In this case, forward guidance conveys precisely the information that the private sector needs.

Moving beyond the simple model developed here, the intuition developed in this paper is useful to think about the role of forward guidance for current monetary policy. Is forward guidance a way for the central bank to commit to keep interest rates effectively at the zero lower bound longer than it would otherwise be optimal *ex post*? If it were widely understood and accepted that optimal monetary policy calls for extended periods of zero interest rates, central banks could have simply staked their credibility on their actions, with no need to supplement them with promises. As an example, it is widely understood and accepted that keeping inflation low is desirable; many central banks around the world have not adopted an explicit inflation target, but they still have managed to build their credibility in sustaining low inflation. It is more plausible that forward guidance is needed because the models in which interest rates are optimally kept at zero for an extended period are not universally understood and accepted. In this case, forward guidance could play a role to signal that central banks believe in the prescriptions stemming from these models, and will set their

policy based on these beliefs.

1.1 Related literature

Many recent papers have used forward guidance as a term to directly refer to a policy of keeping interest rates low for an extended period of time. As an example, this is the case in the recent analysis of the “forward guidance puzzle” (Del Negro et al., 2012; McKay et al., 2016; Werning, 2015). Implicitly, these papers have assumed that any announcement of a policy of extended low rates would be fully credible, and conversely also that the lack of such an announcement would imply a different path of the expectations of the private sector. Here, we draw a distinction between *announcements* about future policy and the actual implementation of the announcements. We then study the conditions under which there is an independent role for advance notice about the future path. In common with Angeletos and Lian (2018) and Angeletos and Sastry (2018), forward guidance acts by providing the private sector with information that was previously not common knowledge. In Angeletos and Lian (2018) and Angeletos and Sastry (2018) this is driven by the limited rationality of the private sector: forward guidance acts by coordinating beliefs in a world in which private agents fail to fully capture the strategic interaction among themselves. In my environment, the private sector is instead fully rational, but the central bank has superior information, particularly concerning its own preferences and beliefs.⁴

Several papers have modeled monetary policy statements as (imperfectly) binding commitments. In Bodenstein et al. (2012), preannounced path can only be revised at (exogenous) random times; in King et al. (2008, 2013), central bankers come in two types, one of which is fully committed to carry out its announced policy and one which is free to reoptimize after any announcement.⁵ In contrast, we study here a situation in which it is well understood that statements are not direct constraints on policy, and policymakers must find it in their

⁴Limited rationality on the part of the central bank or the public may be the reason why in practice the central bank’s view of the world is not common knowledge, in which case my paper points to a different reason why limited rationality and communication may matter.

⁵King et al. extend previous work by Barro (1986), who considered direct commitment to policy in the absence of any announcement.

interest to carry through what they previously promised.

Within the context of monetary policy, cheap talk was first analyzed by Stein (1989) in an environment which shares many traits with section 5 in this paper. He focused purely on Delphic announcements, abstracting from credibility and the repeated game aspect of the interaction. He also only focused on asymmetric information about a central bank target, which would not allow to distinguish between statements about future policy and transparency more in general. In a similar context, Moscarini (2007) analyzed the interaction between the precision of a central bank’s information and its ability to credibly communicate it to the public.

Theoretical models of credibility and cheap talk have focused on persistent private information and learning over time (Sobel, 1985; Benabou and Laroque, 1999). Here, the ability to support greater disclosure is not supported by the presence of “honest” types, but by the infinitely-repeated nature of the game, as in Abreu et al. (1986, 1990).

Finally, our model is designed so that full transparency is optimal under commitment. In particular, the assumption of a quadratic loss function is instrumental to achieve this result. Our point is that, even when full transparency is optimal, there may still be no role for forward guidance. With more general objectives, policymakers may find it optimal even from an *ex ante* perspective to withhold part of the information; this issue is analyzed in detail in Jehiel (2015).⁶

2 The model

As in Angeletos and Sastry (2018), I build upon a tractable model where the analysis is particularly transparent. To discuss forward guidance, what is essential is that expectations about *future* policy affect *current* household decisions. It is in this context that we can draw a distinction between announcements about future policy vs. the actual future actions taken

⁶In the case of the new Keynesian model, Fujiwara and Waki (2015) show that it is not optimal for the central bank to reveal information about future cost-push shocks, in that better information about these shocks generates immediate noise without providing any social gain. Fujiwara and Waki (2017) study the conditions under which information disclosure about fiscal policy is optimal in a dynamic stochastic general-equilibrium model.

by the policymaker. To this end, I work within the context of the Barro-Gordon model (Barro and Gordon, 1983). From an expositional perspective, the advantage of the Barro-Gordon model is that inflation is directly determined by the central bank. This sidesteps the thorny issue of the ultimate source of equilibrium determinacy in new Keynesian models, which may be determined by fiscal as well as monetary policy for the reasons emphasized by Sims (1994), Woodford (1994), and Cochrane (2011). Nonetheless, the intuition developed here extends to those models to the extent that the zero lower bound offers a special role to forward guidance and long-run determinacy is not an issue.

Time is infinite and discrete. The economy is populated by a continuum of private agents (“households”) and a government (or central bank). Fundamental uncertainty in the economy is described by a state space Ω , whose generic element is ω . The state ω will contain a realization of two sequences: a sequence of potential output $(y_t^*)_{t=0}^\infty$, which we assume to be contained in $[y^\ell, y^h]$; and a sequence of target inflation $(\pi_t^*)_{t=0}^\infty$, contained in $[\pi^\ell, \pi^h]$.

In addition to a realization of these sequences, the state of nature may contain other variables, such as advance signals that the households or the government may receive about current and future realizations. In general, the state of nature will not be known by either the government or the agents as of time 0, but it will be gradually revealed over time. For now, we generically denote as $\{\mathcal{F}_t\}_{t=0}^\infty$ and $\{\mathcal{G}_t\}_{t=0}^\infty$ the filtration of what is known at the beginning of period t by the households and the government, respectively. Throughout the paper, we will retain the following assumption about \mathcal{F}_t and \mathcal{G}_t :

Assumption 1. *\mathcal{F}_{t+1} is a finer partition than \mathcal{G}_t , that is, any private information that the government may have at the beginning of period t becomes common knowledge at the beginning of period $t + 1$.*

Assumption 1 implies that households eventually learn the same information that the government had; this will make it easier for them to detect government deviations from equilibrium play and will in turn greatly simplify our analysis. In what follows, it is not essential that this information is known by the households with a delay of at most one period; we could assume longer delays, as long as they are finite.⁷

⁷While I conjecture that results extend to the case of persistent private information, in which households may never learn what the government observed, the analysis of this case is considerably more involved.

In addition to fundamental uncertainty, there are sunspots s_t which (without loss of generality) have a i.i.d. uniform distribution⁸ on $[0, 1]^n$, with n arbitrary. s_t is observed by both households and the government at the beginning of period t . The role of sunspots is to provide a source of coordination between households and the government which is independent of any messages that may be sent by the government.

After information has been revealed to households and the government, the sequence of events within each period t unfolds as follows.

First, the government can send a message m_t to the households, out of some set M . “Forward guidance” is a situation in which the message is directly about the inflation level that will be chosen at the end of the period (or the inflation level that will be chosen in future periods). For generality, we allow the possibility that the government randomizes over messages, and denote by \mathcal{M} a σ -algebra with respect to which this randomization is measurable.

Second, households form expectations about inflation π_t and aggregate output y_t , based on the information currently available to them; y_t^e and π_t^e represents the household average expectation. Without loss of generality, we do not consider the possibility that households randomize over their expectations, because their best response is always single valued.⁹

Finally, the government sets inflation $\pi_t \in [\underline{\pi}, \bar{\pi}]$,¹⁰

$$y_t = \theta y_t^* + (1 - \theta)y_t^e + \lambda(\pi_t - \pi_t^e). \quad (1)$$

Here too the government potentially randomizes, in which case its choice is assumed to be

⁸While the assumption of a uniform distribution is without loss of generality, it is important that the sunspot distribution is absolutely continuous, so that the sunspot is not restricted to put probability mass on specific values.

⁹For this reason, in equilibrium, all households will have the same expectation.

¹⁰In order not to deal with corner solutions in some of the proofs below, we assume that this interval is sufficiently wide, and in particular

$$\bar{\pi} > \pi^h + \frac{\lambda k}{\alpha} + \frac{\lambda(1 - \theta)}{\alpha + \lambda^2}(y^h - y^\ell)$$

and

$$\underline{\pi} < \pi^\ell - \frac{\lambda(1 - \theta)}{\alpha + \lambda^2}(y^h - y^\ell).$$

measurable with respect to Borel's σ -algebra $\mathcal{B}([\underline{\pi}, \bar{\pi}])$.

Barro and Gordon's original specification sets $\theta = 1$.¹¹ In their setting of symmetric information, that we will study in Section 3, this makes no difference. However, it is possible that the government has superior information about y_t^* and that a strategic complementarity among private households leads to higher output when private-sector output expectations are more favorable.¹²

The government's¹³ loss function is

$$(1 - \beta)E \sum_{t=0}^{\infty} \beta^t [(y_t - y_t^* - k)^2 + \alpha(\pi_t - \pi_t^*)^2], \quad (2)$$

where E represents the unconditional expectation before any information is revealed and k is a bias in the government's output target. As in Barro and Gordon, the interesting case is when $k \neq 0$ (and usually we think $k > 0$), so that the government has a temptation ex post to resort to unexpected inflation to stimulate the economy.

A household's loss function is

$$(1 - \beta)E \sum_{t=0}^{\infty} \beta^t [(y_t - y_t^e)^2 + (\pi_t - \pi_t^e)^2],$$

which simply means that in an equilibrium households will set their expectations rationally.

While ω describes the sequence of exogenous fundamental shocks, to define an equilibrium we also need to keep track of sunspots and (public) histories of play. As is standard (see e.g. Chari and Kehoe, 1990), to describe the set of possible equilibrium payoffs it is sufficient to

¹¹Barro and Gordon's model is also expressed in terms of unemployment, rather than output, but this makes no difference for the results.

¹²There is a large literature that studies the role of dispersed information in coordinating individual actions, using the global games approach from Morris and Shin (1998). Particularly relevant are the macroeconomic applications appearing in Hellwig (2005), Amador and Weill (2010), Lorenzoni (2010), and Angeletos and La'O (2018). In that literature, government disclosure of information need not be beneficial, in that it may lead the private sector to focus too much on public signals and not enough on private signals; these issues are discussed especially in Morris and Shin (2002) and Angeletos and Pavan (2007, 2009). Here, I abstract from this complication because information is symmetric within the private sector, so government disclosure will be unambiguously beneficial.

¹³We will use "government" and "central bank" interchangeably. There is a single policymaker.

keep track of the history of play by the large agent in the economy (the government).¹⁴ We will thus define a history at the message stage as $h^t := (\{s_s, m_s, \pi_s\}_{s=0}^{t-1}, s_t)$, and a history at the expectations-setting stage as $h^{et} := (\{s_s, m_s, \pi_s\}_{s=0}^{t-1}, (s_t, m_t))$.¹⁵ H^t and H^{et} are the corresponding sets of histories, and \mathcal{H}^t and \mathcal{H}^{et} are the corresponding filtrations.¹⁶

A government strategy σ^g is a $\mathcal{G}_t \times \mathcal{H}^t$ -measurable mapping from (Ω, H^t) into a distribution over (M, \mathcal{M}) and a $\mathcal{G}_t \times \mathcal{H}^{et}$ -measurable mapping from (ω, H^{et}) into a probability distribution over $([\underline{\pi}, \bar{\pi}], B[\underline{\pi}, \bar{\pi}])$.¹⁷ A (symmetric) household strategy is a $\mathcal{F}_t \times \mathcal{H}^{et}$ -measurable mapping σ^p from (Ω, H^{et}) to $[\underline{y}, \bar{y}] \times [y^l, y^h]$.

When the government and the households play a strategy profile σ^g, σ^p , their play induces a probability distribution over outcomes $((\omega, H^\infty), \bigcup_{t=0}^\infty \mathcal{G}_t \times \bigcup_{t=0}^\infty \mathcal{H}^t)$.

Definition 1. A strategy profile (σ^g, σ^p) is a perfect Bayesian equilibrium if:

- (i) Given any (ω, h^t) and given that future play will occur according to (σ^g, σ^p) , any message in the support of $\sigma^g(\omega, h^t)$ is optimal for the government.
- (ii) Given any (ω, h^{et}) and given that future play will occur according to (σ^g, σ^p) , any inflation rate in the support of $\sigma^g(\omega, h^{et})$ is optimal for government.
- (iii) Given any (ω, h^{et}) and σ^g , household expectations are rational:

$$y_t^e = E[y_t | \mathcal{F}_t, \mathcal{H}^{et}; \sigma^g], \quad \pi_t^e = E[\pi_t | \mathcal{F}_t, \mathcal{H}^{et}; \sigma^g]. \quad (3)$$

¹⁴As discussed in Bassetto (2005), this is no longer sufficient to study the set of allocations that can be uniquely implemented under commitment.

¹⁵At time 0, $h^0 = s_0$ and $h^{e0} = (s_0, m_0)$.

¹⁶Formally, $\mathcal{H}^t := \mathcal{B}([0, 1])^{nt} \times \mathcal{M}^t \times \mathcal{B}([\underline{\pi}, \bar{\pi}])^t$, where the superscript on the right-hand side refers to Cartesian power.

¹⁷We assume here that the government does not directly observe y_t^e and π_t^e . When the government has all the information that households have, this makes no difference: within an equilibrium, y_t^e and π_t^e are deterministic functions of what the government knows. This assumption is more relevant when households have superior information, as in the example of section 6. This example is relevant when the government cannot infer the missing information directly from household expectations.

3 Odyssean forward guidance

Consider first the case in which the government has no private information. Then we obtain the following result:

Proposition 1. *Assume that $\mathcal{F}_t = \mathcal{G}_t$. If $M \subseteq \mathbb{R}^a$, then the set of equilibrium government payoffs is the same as that of an economy in which no messages are allowed ($M = \emptyset$).*

Proof. See appendix.

To better illustrate what this proposition does and does not imply, we consider a specific example. Suppose that $k = 0.01$, $\pi_t^* \equiv 0.02$, $\beta = 0.96$, $\alpha = 1$, and $\lambda = 40$. Furthermore, y_t^* is known (along with the entire past history of play) to both the government and the private sector. In this case, the government inflation target is deterministic at 2%. Potential output may be random, but the government always wants to overstimulate the economy by 1%. In a repeated game context, this economy admits many equilibria, independently of the message space M . We focus on two of them.¹⁸

First, suppose $M = [\underline{\pi}, \bar{\pi}]$.¹⁹ The following is an equilibrium strategy profile. If $\pi_t \neq m_t$ never occurred in the past, the government announces $m_t = 0.02$; otherwise, the government can send an arbitrary message (it could be 0.02 again, or anything else, since the households will no longer condition their strategy on the government's reports). Households set $y_t^e = y_t^*$ independently of the past history. If $\pi_t \neq m_t$ never occurred in the past and $m_t \leq 0.43$, households “believe” the government and set $\pi_t^e = m_t$. Otherwise, households set $\pi_t^e = \pi_t^* + (\lambda/\alpha)k = 0.43$. Finally, if $\pi_t \neq m_t$ never occurred in the past and $m_t \leq 0.43$, the government follows through on its announcement and sets $\pi_t = m_t$; otherwise, it sets $\pi_t = 0.43$.

Alternatively, suppose that $M = \emptyset$: no messages can be sent by the government. The following is an equilibrium strategy profile. Households set $y_t^e = y_t^*$ independently of the past history. If $\pi_t \neq 0.02$ never occurred in the past, households set $\pi_t^e = \pi_t^* = 0.02$. Otherwise, households set $\pi_t^e = 0.43$. Finally, if $\pi_t \neq 0.02$ never occurred in the past, the government sets $\pi_t = 0.02$; otherwise, it sets $\pi_t = \pi_t^* + (\lambda/\alpha)k = 0.43$.

In both equilibria described above household expectations are set according to a trigger

¹⁸The verification that the two strategy profiles are indeed equilibria is relegated to the appendix.

¹⁹Assume $\underline{\pi} < 0.02$ and $\bar{\pi} > 0.43$.

strategy and the outcome coincides with what would arise under government commitment ($\pi_t = \pi_t^*$ and $y_t = y_t^*$); in both cases, the threat that disciplines the government’s temptation to overstimulate the economy is future reversion to permanently repeating the equilibrium outcome of the static one-shot game.

The first equilibrium resembles what Campbell et al. (2012) call “Odyssean forward guidance:” the government announces future policy, and puts its credibility at stake. If the government fails to deliver on its announcements, it loses its ability to coordinate expectations favorably, and high inflation ensues. This equilibrium shows that Proposition 1 *does not* say that government messages are necessarily irrelevant.

The second equilibrium achieves the same outcome, but without resorting to forward guidance, and is based on the idea that “actions speak louder than words.” Notice that, in the first equilibrium, while the expectations that households form about future policy are essential, the private sector can perfectly forecast what message the government will send. It is thus possible to bypass the message and stake the government’s credibility directly on its *actions*. Pure Odyssean forward guidance is unnecessary. When the private sector and the government share the same information, the temptation that the government faces to renege on its promises is the same whether those promises have been made explicit or left implicit. The same is true in a new Keynesian framework: if it is common knowledge that holding interest rates at zero for an extended period of time is optimal, households could directly form expectations based on this appropriate policy, and the central bank’s credibility would rely on following through with the expected policy, without need of advance messages.

Finally, in the example above, forward guidance is only about the policy that the government will undertake subsequently within the period, but the proposition extends to announcements about policy further into the future. As an example, the following is also an equilibrium. Suppose again that $M = [\underline{\pi}, \bar{\pi}]$. Assume that there is some outstanding message $m_{-1} = 0.02$. If $\pi_t \neq m_{t-1}$ never occurred in the past, the government announces $m_t = 0.02$; otherwise, the government can send an arbitrary message (it could be 0.02 again, or anything else, since the households will no longer condition their strategy on the government’s reports). Households set $y_t^e = y_t^*$ independently of the past history. If $\pi_t \neq m_{t-1}$ never occurred in the past and $m_{t-1} < 0.43$, households “believe” the government and set

$\pi_t^e = m_{t-1}$. Otherwise, households set $\pi_t^e = 0.43$. Finally, if $\pi_t \neq m_{t-1}$ never occurred in the past and $m_{t-1} < 0.43$, the government follows through on its announcement and sets $\pi_t = m_{t-1}$; otherwise, it sets $\pi_t = 0.43$. In this equilibrium, the government announces its inflation plan one period ahead of time. Nonetheless, the equilibrium outcome remains the same, and it coincides with what happens in the trigger-strategy equilibrium with no messages that we described above.

4 Private information about the state of the economy

In this section, the government has superior information about the underlying state of the economy. Let the inflation target π_t^* still be known by both the private sector and the government at the beginning of period t . In contrast, potential output y_t^* is not known at the beginning of period t ; let $F_{y_t^*}(\cdot|\mathcal{F}_t)$ its distribution, conditional on the private-sector information. At the beginning of the period, the government has access to the same information as the private sector, but it also receives a potentially noisy signal \tilde{y}_t . Conditional on the government's superior information,²⁰ the distribution of potential output is thus $F_{y_t^*}(\cdot|\mathcal{G}_t)$.

As a benchmark, suppose that the government could commit to a strategy for its future reports and inflation choices at the beginning of time, before any information is revealed. It is straightforward to prove that the best equilibrium outcome arises when the government commits to report truthfully its information to the private sector and to set $\pi_t \equiv \pi_t^*$. Formally:

Proposition 2. *Let $M = [y^l, y^h]$ be the message space. Suppose that the government commits to a strategy σ^g that reports $m_t = E[y_t^*|\mathcal{G}_t]$ and sets $\pi_t = \pi_t^*$ with probability 1. Let $(y_t, \pi_t, y_t^e, \pi_t^e)_{t=0}^\infty$ be the resulting equilibrium outcome. Then there is no other message space and/or government strategy that generates an equilibrium outcome that strictly dominates $(y_t, \pi_t, y_t^e, \pi_t^e)_{t=0}^\infty$.*

Proof. See appendix.

²⁰Observing \tilde{y}_t implies that the σ -algebra \mathcal{G}_t that represents the government information at time t is finer than the one of the households, \mathcal{F}_t .

Under commitment, the government realizes that it cannot fool the private sector; any deterministic misreporting of its signals would be undone by the private agents, and any garbling of the signals would simply increase the variance of output around potential, which is undesirable.

When the government cannot commit, the typical cheap talk conflict emerges: since incentives are not aligned, the government has a temptation to misreport its signal to induce the households to increase their output. Nonetheless, the ability to send messages is in general valuable, and superior equilibria in which some information is revealed emerge.²¹ This necessarily happens when the government is sufficiently patient, so that information revelation may be supported by trigger strategies in which the future credibility of the government is at stake.

We formalize these points in the following proposition.

Proposition 3. *Fix $\alpha, \lambda, k, \underline{\pi}, \bar{\pi}$, and the stochastic process for $(y_t, \tilde{y}_t, \pi_t^*)_{t=0}^\infty$.*

- (i) *As long as there exists a period t such that $\text{Prob}(E[y_t^*|\mathcal{F}_t] \neq E[y_t^*|\mathcal{G}_t]) > 0$, there exists a value $\bar{\beta} \in [0, 1)$ and a message space M , such that for all $\beta > \bar{\beta}$ the set of equilibrium payoffs attainable in the game in which the government can send messages from M is strictly larger than the corresponding set if no messages are allowed.*
- (ii) *The expansion of the set includes higher equilibrium payoffs than what can be supported without messages. If in addition $\bar{\pi}$ is sufficiently high, $\bar{\beta}$ in part (i) can be chosen so that the set of payoffs for the game with messages does not include any payoffs that are worse than those achievable in equilibria of the game without messages.*

Proof. See appendix.

Proposition 3 is not unambiguously optimistic. First, as is always the case in games with cheap talk, there will be equilibria in which the government “babbles,” sending the same message independently of its information, and households in turn disregard the government message, defeating any attempt to convey extra information. Perhaps even more disappointingly, there are instances in which allowing for a nontrivial message space will create the

²¹We evaluate welfare from the perspective of the government, but the same is true for the households, whose expectations become more precise.

possibility of equilibria whose welfare is *worse* than the worst possible equilibrium under no messages. Abreu et al. (1986, 1990) show that there often is a link between the payoff of the best and the worst equilibrium: the worst equilibrium represents a threat that can be used to support the best equilibrium, and the best equilibrium can be used as a reward for the government to be willing to endure the worst punishment. The ability to send messages offers a way for the government to better coordinate the private sector, reducing the volatility of output; paradoxically, by increasing the payoff in the best equilibrium, this ability opens the door for the worst equilibrium to become worse. The last part of the theorem proves that this does not happen if the maximal level of attainable inflation is sufficiently high and the government is sufficiently patient. In this case, the “punishment stage” of the worst equilibrium does not last a single period and the continuation of the worst equilibrium is not the best equilibrium. Then, the government cannot do worse than what happens when households expect it to babble and ignore its messages. This case is reassuring: government communication may lead to better equilibrium payoffs, but not to worse ones.²²

From here on, we take the optimistic view that the economy coordinates on superior outcomes, in which case the government messages are unambiguously helpful. However, the next question is whether these messages take the form of forward guidance versus a generic need for “transparency.” A transparent government discloses (truthfully) a variety of information that is not publicly available. This information takes the form of forward guidance if it concerns the future path of policy. The example at hand is designed to be particularly stark. For β sufficiently high, the Folk theorem implies that the best equilibrium outcome coincides with the outcome under commitment described in proposition 2. In this equilibrium, inflation is always at its target value π_t^* and it does not depend on the private information available to the central bank. As a consequence, a message reporting future policy would not allow the households to infer the information that they need to form the appropriate expectations: while communication is potentially valuable, it is not about future policy.

²²Of course, we are silent on the process by which households and the government coordinate to one among many equilibria. If the introduction of government communication leads to coordination toward equilibria with a worse payoff, the ability to send messages might still be harmful.

The example above is clearly extreme. In richer environments, optimal government policy depends on the realization of the shocks about which the government has superior knowledge. However, even in this case, reporting future policy, as opposed to the underlying information that rationalizes the policy choice, is at best an indirect way to convey the information that the households need to form their expectations. For the sake of concreteness, consider an example in which, for some reason, the government objective function is²³

$$(1 - \beta)E \sum_{t=0}^{\infty} \beta^t \{ (y_t - y_t^* - k)^2 + \alpha[\pi_t - \pi_t^* - f(y_t^*)]^2 \}, \quad (4)$$

so that optimal inflation under commitment and full transparency will be $\pi_t = \pi_t^* + E[f(y_t^*)|\mathcal{G}_t]$. In this case, sending a message about future policy π_t will reveal some information to the private sector about what the government observed through \tilde{y}_t . However, even in this case, unless f is affine, knowing $E[f(y_t^*)|\mathcal{G}_t]$ need not be the same as knowing $E[y_t^*|\mathcal{G}_t]$, which the households need. Announcing future policy is an imperfect and roundabout way of announcing the underlying information that the private sector requires to properly coordinate. Appendix B expands on this point.

The next two sections consider two cases in which the government's private information is not about the underlying state of the economy, but rather about its objective or its beliefs. In this case, sending messages about future policy is a natural way to convey the information that households need to make their decisions.

5 Private information about government's objective

Assume now that the government and the private sector have symmetric information about y_t^* , but the government has private information about π_t^* . Without loss of generality, let the government know π_t^* perfectly (the only role of π_t^* is to act as the government's preferred inflation rate). At the beginning of period t , conditional on the information available to the private sector, π_t^* has a distribution $F_{\pi_t^*}(\cdot|\mathcal{F}_t)$. As in the previous section, we rule

²³We do not consider microfoundations for this example. It is simply meant as an illustration of a situation in which optimal government policy depends on the underlying information about the exogenous state of the economy.

out persistent private information by assuming that π_t^* becomes common knowledge at the beginning of period $t+1$.²⁴ The following propositions are the counterparty to Propositions 2 and 3 in the previous section.

Proposition 4. *Let $M = [\pi^\ell, \pi^h]$ be the message space. Suppose that the government commits to a strategy σ^g that reports $m_t = \pi_t$ and sets $\pi_t = \pi_t^*$ with probability 1. Let $(y_t, \pi_t, y_t^e, \pi_t^e)_{t=0}^\infty$ be the resulting equilibrium outcome. Then there is no other message space and/or government strategy that generates an equilibrium outcome that strictly dominates $(y_t, \pi_t, y_t^e, \pi_t^e)_{t=0}^\infty$.*

Proof. See appendix.

Proposition 5. *Fix $\alpha, \lambda, k, \underline{\pi}, \bar{\pi}$, and the stochastic process for $(y_t, \tilde{y}_t, \pi_t^*)_{t=0}^\infty$.*

- (i) *As long as there exists a period t such that $\text{Prob}(E[\pi_t^* | \mathcal{F}_t] \neq \pi_t^*) > 0$, there exists a value $\bar{\beta} \in [0, 1)$ and a message space M , such that for all $\beta > \bar{\beta}$ the set of equilibrium payoffs attainable in the game in which the government can send messages from M is strictly larger than the corresponding set if no messages are allowed.*
- (ii) *The expansion of the set above includes higher equilibrium payoffs than what can be supported without messages. If in addition $\bar{\pi}$ is sufficiently high, $\bar{\beta}$ can be chosen so that the set of payoffs for the game with messages does not include any payoffs that are worse than those achievable in equilibria of the game without messages.*

Proof. The proof follows from the reformulation in Appendix A.2 and the proof of Proposition 2.

In this case, households have all the information about the underlying state of the economy that they need to make decisions, given government policy. While the government could

²⁴That π_t^* becomes common knowledge is once again assumed because to simplify the arguments. However, for reaping the benefits of repeated interaction, it is important that at least *some* additional information about π_t^* will become available to the private sector after the government policy choice. When this is not the case, Athey et al. (2005) show that the optimal mechanism would involve a static provision of government incentives. Even in that environment, cheap talk, which is ruled out in their paper, could still be valuable. Within a new Keynesian model, this issue is revisited in Waki et al. (2018).

report its underlying information that leads it to prefer π_t^* , this is more information than necessary: all they need is to know the policy choice π_t that the government will take.²⁵ In other words, what the private sector needs is precisely forward guidance about monetary policy.

5.1 A Numerical Example

To better understand the interplay of credibility and disclosure, we consider a numerical example. Our model is too stylized to afford a plausible calibration and parameters are chosen with the aim of illustrating the range of possibilities that emerge when the CB provides forward guidance about its actions in a setting of repeated interaction, rather than providing a quantitative assessment of the model. Nonetheless, a robust feature of the model is that, when the CB credibility on inflation is at stake, even a very impatient central bank will in general be led to communicate its future plans truthfully.

In our first example we set $\lambda = 2$, implying that an extra 1% unexpected inflation increases output by 2%, $\alpha = 4$ (that is, the CB puts 4 times as much weight on inflation when both inflation and output are at target),²⁶ $k = 0.05$ (if it could be done at no cost, the CB would want to set output 5% above its equilibrium value), $\pi^\ell = -0.1$ and $\pi^h = 0.2$ (the CB could generate at most a 10% deflation and a 20% inflation),²⁷ the inflation target in each period (which is private information to the CB) is 1.5%, 1.75%, or 2.5% with equal probability ex ante, and it is i.i.d. over time.²⁸

Figure 1 plots inflation in the three states, as a function of the discount factor.²⁹ Figure 2 plots instead the unconditional expected value of inflation, along with the expected value of its target (1.92%), and highlights the conditions that lead the CB to partially succumb to

²⁵Under commitment, $\pi_t = \pi_t^*$, but this need not be the case when the government acts under discretion.

²⁶Notice, however, that output is below the CB target in equilibrium.

²⁷The rationale for tight bounds is that it would take time for the CB to generate hyperdeflation/hyperinflation.

²⁸We restrict our analysis to pure-strategy equilibria. Also, since y_t^* is common knowledge, we can normalize it to zero without loss of generality, and the equilibria are independent of the value of θ .

²⁹We only plot the range in which all the interesting action occurs. The lines are constant for greater values of β and they keep growing in a straight line as $\beta \rightarrow 0$.




Figure 1: Inflation in the best equilibrium, as a function of the discount factor β . Each line corresponds to a different realization of the target: 1.5% (lowest), 1.75% (middle), and 2.5% (highest).

its inflationary bias inherent in the Barro-Gordon framework.

The figures are most intuitively read from right to left, and it is in this way that we will discuss the results. In this example, the best equilibrium features full CB transparency and inflation always at target as long as $\beta > .157$. This equilibrium is supported by private-sector expectations which trigger high inflation if ever the CB is found to be misreporting its target or deviating from it. As β drops below $.157$, when the CB target is at the intermediate value of 1.75%, the temptation to report 1.5% and fool the private sector by choosing higher inflation becomes too strong, and a trade-off emerges between low inflation and truthful revelation. For $\beta \in [0.154, .157]$, the best equilibrium maintains full disclosure, but it reduces the temptation to deviate by prescribing higher inflation when the CB target is 1.5%. This is shown in Figure 1, where the bottom level of inflation moves above the target level. In this range, the CB could equally well announce its target or its actual intended policy. Since its intended policy sometimes deviates from target, communicating the future policy (“forward guidance”) will offer a simpler way to focus the expectations of the private sector and will thus be the more natural communication.

At $\beta = 0.154$, the cost of the inflationary bias is greater than the coordination benefit from




Figure 2: Unconditional expectation of inflation in the best equilibrium, compared to the unconditional expectation of its target.

full revelation of the CB target. Consequently, for $\beta \in [0.145, 0.154]$ the best equilibrium features only partial revelation: the CB truthfully announces its target when it is the at the highest level (2.5%), and otherwise announces that its target is one of the two bottom ones. While expected inflation jumps almost all the way back to its target, a small discrepancy persists, and it grows as β shrinks: this time, the incentive constraint of the highest-target CB is at play, and implementing lower inflation in the two bottom states would provide too great a temptation for the CB with the high target to abandon the equilibrium path. As the cutoff threshold of $\beta = 0.154$ is crossed, the level of inflation for both of the two bottom types jumps down (going from right to left). This happens for two different reasons. First, for the middle type (the CB with a 1.75% target), the drop occurs because private-sector expectations jump down. Since households can no longer tell apart the two bottom types, their inflation expectations are based on a mix of the two policies. The optimal policy from the perspective of the CB is a linear combination of its target and private expectations: hence, when expectations jump down, the optimal policy also decreases. Second, for the lowest type (the CB with a 1.5% target), the drop is caused by the lesser need for an inflationary bias to keep the incentives of the middle type in check. This effect is partially offset by the fact that private-sector expectations jump up, a mirror image of the discussion

of the middle type above. In this range, forward guidance is still valuable, but it cannot be too precise: the CB would have to provide an advance notice of a *range* of values, rather than a numerical target. Too precise guidance would not be credible.

Below $\beta = 0.145$, providing any credible forward guidance is too costly, and the best equilibrium no longer features any informative message. The temptation for the lowest type to defect is now reduced by the ability to pool with the other types and benefit from a correspondingly lower inflation expectation, and this allows the CB to again avoid any inflationary bias, at the cost of greater variability of output, determined by the fact that households know even less about future monetary policy. As we cross $\beta = 0.145$ from the left, inflation in the best equilibrium drops for all 3 types, for the same reasons described in the two bullets above: in the case of the low and middle type, this is driven by the removal of the need for an inflationary bias, and in the case of the high type by the lower expectations it faces. Finally, below $\beta = 0.125$, CB impatience is such that even with no CB communication it is impossible to avoid giving into some of the inflation bias.

As we discussed in commenting the theorems above, forward guidance can be beneficial, but in principle it could also be a disadvantage, if household expectations coordinate on a perverse equilibrium. In this example, the worst possible equilibrium is worse under forward guidance when $\beta \in [0.145, 0.384]$. In this range, the worst equilibrium when no messages are allowed is characterized by a single period of very high inflation, followed by a return to the best equilibrium. When the best equilibrium becomes better, owing to the possibility of sending informative messages, the CB is willing to endure a worse temporary loss to obtain the best continuation as a reward, and the net effect is a higher total loss. In the “punishment stage” of the worst equilibrium, no messages are sent, so it is not that sending messages is directly making things worse, but rather it is the *anticipation of informative messages in the future* that has indirect adverse consequences. Notice that the lower cutoff is precisely the point at which the best equilibrium stops featuring informative messages: for lower values of β , the best equilibrium is the same with or without messages, and the same is then necessarily the case for the worst equilibrium as well. In contrast, for $\beta > 0.384$, the worst equilibrium features inflation at its upper bound for at least one period, and an interior continuation value. With an interior continuation, changes to the payoff of the best

equilibrium become irrelevant.

In the example above, the repeated nature of the interaction is necessary for any communication to take place: at $\beta = 0$, only babbling is possible.³⁰ This is not necessarily the case: when the target levels are spread further apart, the temptation to misreport may be more limited and partial or full revelation may happen even in a purely static equilibrium (corresponding to $\beta = 0$). As an example, for $\beta = 0$ the best equilibrium features partial revelation if we change π^* to be 0, 5% or 12% with equal i.i.d. probability, using the same parameters as above except for $\alpha = 40$ and setting the upper bound for inflation at 50%.

6 Private information about government's beliefs

In Section 4 the government has superior information about the underlying state of the economy. Here, we assume instead that the households receive a perfect signal of y_t^* at the beginning of period t , while the government only observes a noisy signal \tilde{y}_t , so that it faces nontrivial uncertainty described by the conditional distribution $F_{y_t^*}(\cdot|\mathcal{G}_t)$. Even though households know y_t^* perfectly, the government still has private information, about its own imperfect beliefs.³¹

When government preferences are given by (2), optimal government policy under commitment requires the government to set $\pi_t = \pi_t^*$ unconditionally, and households do not need any report from the government to achieve perfect foresight, setting $y_t^e = y_t^*$ and $\pi_t^e = \pi_t^*$.

However, suppose instead that the loss function is once again distorted as in (4). We then obtain:

Proposition 6. *Let $M = [\pi^l, \pi^h]$ be the message space. Suppose that the government commits to a strategy σ^g that reports $m_t = \pi_t$ and sets $\pi_t = \pi_t^* + E[f(y_t^*)|\mathcal{G}_t]$ with probability*

³⁰Suboptimal equilibria with full revelation but higher inflation exist for $\beta > 0.025$. Suboptimal equilibria with partial revelation and higher inflation exist for $\beta > 0.021$. Below these thresholds, no pure equilibria with informative messages exist.

³¹Formally, it would be incorrect to say that households have superior information compared to the government. If households unambiguously had superior information, \mathcal{F}_t would be a finer σ -algebra than \mathcal{G}_t ; but in this case, households would know what the government knows. In our example, instead, households do not know what the government observed.

1. Let $(y_t, \pi_t, y_t^e, \pi_t^e)_{t=0}^\infty$ be the resulting equilibrium outcome. Then there is no other message space and/or government strategy that generates an equilibrium outcome that strictly dominates $(y_t, \pi_t, y_t^e, \pi_t^e)_{t=0}^\infty$.

Proof. See appendix.

As in Section 5, the government's report could be about the information that led it to choose its policy actions; in this case, this would be the signal \tilde{y}_t that it received. However, this information is redundant: the only reason households need to know the imperfect signal observed by the government is to form expectations about government policy. Once again, in this setting forward guidance is a natural message space for the government to communicate the information that can coordinate households towards desirable equilibria. As in Sections 4 and 5, the benefits of transparency carry over to equilibria of the game in which the government has no commitment, provided the government is sufficiently patient.

Proposition 7. Fix $\alpha, \lambda, k, \underline{\pi}, \bar{\pi}$, and the stochastic process for $(y_t, \tilde{y}_t, \pi_t^*)_{t=0}^\infty$.

- (i) As long as there exists a period t such that $\text{Prob}(E[y_t^* | \mathcal{G}_t] \neq y_t^*) > 0$, there exists a value $\bar{\beta} \in [0, 1)$ and a message space M , such that for all $\beta > \bar{\beta}$ the set of equilibrium payoffs attainable in the game in which the government can send messages from M is strictly larger than the corresponding set if no messages are allowed.
- (ii) The expansion of the set above includes higher equilibrium payoffs than what can be supported without messages. If in addition $\bar{\pi}$ is sufficiently high, $\bar{\beta}$ can be chosen so that the set of payoffs for the game with messages does not include any payoffs that are worse than those achievable in equilibria of the game without messages.

Proof. The proof follows from the reformulation in Appendix C and the proof of Proposition 2.

In the simple model analyzed here, the central bank only has imperfect information about potential output in the economy. In practice, there are many unknown aspects that affect the optimal choice of monetary policy: uncertainty about the transmission mechanism, about the sectors which are most affected, about financial intermediation interlinks... Optimal policy is informed by the models that the CB adopts to account for all of these complicated

interactions, and a primary source of asymmetric information about beliefs is precisely the class of models that the CB believes to be an accurate representation of the economy. If the central bank revealed all the details of these models, households and firms may find it difficult to absorb all the information, which is in any case of limited interest for them beyond forecasting policy actions. It is in this context that forward guidance emerges as a simpler way of summarizing relevant information. Statements about keeping interest rates low for extended periods of time were not per se a form of commitment. Rather, central banks could use such statements to explain that, according to the models that they believe in, extended period of future low rates would be beneficial for the current state of the economy. CBs then would have to rely on their existing credibility to persuade private agents that they would carry through their announced policy even once the beneficial effect of keeping rates low was in the past. These announcements were important because it is not a settled matter that extended periods of low rates are necessarily beneficial (see e.g. the discussion of neo-Fisherian views in Cochrane (2018) or Uribe (2017, 2018)). Relying purely on credibility that the central bank would take the “right actions” would have been insufficient, because the private sector might not know what the “right actions” are, and especially what actions the central bank deems “right.” This is the key difference between this environment and the case of Section 3.

7 Conclusion and Discussion

Forward guidance is unnecessary if it’s purely a commitment tool. When the government and the private sector have symmetric information, actions speak louder than words, and the same outcomes can be achieved when the government stakes its credibility on its actions directly, with no need to have interim messages. When the government has private information, optimal disclosure may call for transparency, but not necessarily forward guidance; in many instances, households learn more if the government discloses the actual information, rather than its future policy plans. We identified two circumstances in which a policy of forward guidance (that is, sending messages about future policy) helps in achieving the best possible equilibria: these arise when the primary source of asymmetric information concerns

directly the preferences or the beliefs of the policymaker. Duarte (2019) presents other examples in which this is the case. What all of these examples have in common is that the private sector has no independent value for the underlying information of the central bank, other than to forecast the future course of monetary policy.

How much information a central bank can credibly disclose depends on its credibility. In the context of repeated interaction, credibility can be better established if the policymaker is sufficiently patient, and the interaction is sufficiently frequent. In New Keynesian models, forward guidance is particularly beneficial in times in which the interest rate is at the zero lower bound, which was an infrequent event in the past. This may make it more difficult to credibly deploy forward guidance. At the same time, central banks rely on their credibility not just to promise long period of low rates, but for a variety of other policies: in normal times, they rely on it to keep inflation anchored and to avoid giving in to the temptation to overstimulate the economy.³² The model presented here is silent on equilibrium selection; nonetheless, central bank credibility spills over across different policies, the rare occurrence of periods of interest rates at the zero lower bound is not necessarily a challenge for information revelation through forward guidance.

References

- Abreu, D., Pearce, D., Stacchetti, E., 1986. Optimal cartel equilibria with imperfect monitoring. *Journal of Economic Theory* 39, 251–269.
- Abreu, D., Pearce, D., Stacchetti, E., 1990. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58, 1041–1063.
- Amador, M., Weill, P.O., 2010. Learning from prices: Public communication and welfare. *Journal of Political Economy* 118, 866–907.
- Angeletos, G.M., La’O, J., 2018. Optimal monetary policy with information frictions. *Journal of Political Economy* Forthcoming.

³²This may be important in political-economy models, where partisan pressures might emerge, particularly around election periods.

- Angeletos, G.M., Lian, C., 2018. Forward guidance without common knowledge. *American Economic Review* 108, 2477–2512.
- Angeletos, G.M., Pavan, A., 2007. Efficient use of information and social value of information. *Econometrica* 75, 1103–1142.
- Angeletos, G.M., Pavan, A., 2009. Policy with dispersed information. *Journal of the European Economic Association* 7, 1–50.
- Angeletos, G.M., Sastry, K.A., 2018. Managing expectations without rational expectations. Working Paper 25404. NBER.
- Athey, S., Atkeson, A., Kehoe, P.J., 2005. The optimal degree of discretion in monetary policy. *Econometrica* 73, 1431–1475.
- Barro, R.J., 1986. A model of monetary policy with incomplete information. *Journal of Monetary Economics* 17, 3–20.
- Barro, R.J., Gordon, D.B., 1983. A positive theory of monetary policy in a natural rate model. *Journal of Political Economy* 91, 589–610.
- Bassetto, M., 2005. Equilibrium and government commitment. *Journal of Economic Theory* 124, 79–105.
- Benabou, R., Laroque, G., 1999. Using privileged information to manipulate markets: Insiders, gurus, and credibility. *The Quarterly Journal of Economics* 107, 921–958.
- Bodenstein, M., Hebden, J., Nunes, R., 2012. Imperfect credibility and the zero lower bound. *Journal of Monetary Economics* 59, 135–149.
- Campbell, J.R., Evans, C.L., Fisher, J.D., Justiniano, A., 2012. Macroeconomic effects of federal reserve forward guidance. *Brookings Papers on Economic Activity* 2012, 1–54.
- Campbell, J.R., Ferroni, F., Fisher, J.D., Melosi, L., 2019. The limits of forward guidance. *Journal of Monetary Economics* , forthcoming.

- Campbell, J.R., Fisher, J.D., Justiniano, A., Melosi, L., 2017. Forward guidance and macroeconomic outcomes since the financial crisis. *NBER Macroeconomics Annual* 31, 283–357.
- Chari, V.V., Kehoe, P.J., 1990. Sustainable plans. *Journal of Political Economy* 98, 783–801.
- Cochrane, J.H., 2011. Determinacy and identification with Taylor rules. *Journal of Political Economy* 119, 565–615.
- Cochrane, J.H., 2018. Michelson-Morley, Fisher, and Occam: The radical implications of stable quiet inflation at the zero bound. *NBER Macroeconomics Annual* 32, 113–226.
- Crawford, V., Sobel, J., 1982. Strategic information transmission. *Econometrica* 75, 1431–1451.
- Del Negro, M., Giannoni, M., Patterson, C., 2012. The forward guidance puzzle. Staff Report 574. Federal Reserve Bank of New York.
- Duarte, F., 2019. Discussion of forward guidance: Communication, commitment, or both? *Journal of Monetary Economics*, forthcoming.
- Eggertsson, G.B., Woodford, M., 2003. The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity* 2003, 139–211.
- Fujiwara, I., Waki, Y., 2015. Private News and Monetary Policy. Working Paper 238. Federal Reserve Bank of Dallas, Globalization and Monetary Policy Institute.
- Fujiwara, I., Waki, Y., 2017. Fiscal Forward Guidance. Discussion paper 17-E-087. RIETI.
- Hellwig, C., 2005. Heterogeneous information and the welfare effects of public information disclosures. <http://www.econ.ucla.edu/people/papers/Hellwig/Hellwig283.pdf>. Mimeo, UCLA.
- Jehiel, P., 2015. On transparency in organizations. *Review of Economic Studies* 82, 736–761.
- King, R.G., Lu, Y.K., Pastén, E.S., 2008. Managing expectations. *Journal of Money, Credit and Banking* 40, 1625–1666.

- King, R.G., Lu, Y.K., Pastén, E.S., 2013. Policy design with private sector skepticism in the textbook new Keynesian model. http://ihome.ust.hk/~yanglu/KLP_ME2.pdf. Mimeo, Hong Kong University of Science and Technology.
- Kydland, F.E., Prescott, E.C., 1977. Rules rather than discretion: the inconsistency of optimal plans. *Journal of Political Economy* 85, 473–491.
- Lorenzoni, G., 2010. Optimal monetary policy with uncertain fundamentals and dispersed information. *Review of Economic Studies* 77, 305–338.
- McKay, A., Nakamura, E., Steinsson, J., 2016. The power of forward guidance revisited. *American Economic Review* 106, 3133–3158.
- Morris, S., Shin, H.S., 1998. Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review* 88, 587–597.
- Morris, S., Shin, H.S., 2002. Social value of public information. *American Economic Review* 92, 1521–1534.
- Moscarini, G., 2007. Competence implies credibility. *American Economic Review* 97, 37–63.
- Sims, C.A., 1994. A simple model for study of the determination of the price level and the interaction of monetary and fiscal policy. *Economic Theory* 4, 381–399.
- Sobel, J., 1985. A theory of credibility. *Review of Economic Studies* 52, 557–573.
- Stein, J.C., 1989. Cheap talk and the Fed: A theory of imprecise policy announcements. *American Economic Review* 79, 32–42.
- Uribe, M., 2017. The neo-Fisher effect in the United States and Japan. Working Paper 23977. NBER.
- Uribe, M., 2018. The neo-Fisher effect: econometric evidence from empirical and optimizing models. Working Paper 25089. NBER.
- Waki, Y., Dennis, R., Fujiwara, I., 2018. The optimal degree of monetary-discretion in a new Keynesian model with private information. *Theoretical Economics* 13, 1319–1368.

Werning, I., 2011. Managing a liquidity trap: monetary and fiscal policy. Working Paper 17344. NBER.

Werning, I., 2015. Incomplete markets and aggregate demand. Working Paper 21448. NBER.

Woodford, M., 1994. Monetary policy and price level determinacy in a cash-in-advance economy. *Economic Theory* 4, 345–380.

A Appendix

Provided below are omitted proofs and relevant extensions.

A.1 Proof of proposition 1

Let V be the set of equilibrium government payoffs for an economy with a message space M . First, we show that V is nonempty, since there is always at least an equilibrium associated with playing in period s the static equilibrium that would arise in a one-period game.

In general, the static equilibrium is given by the following:³³

$$\pi_t = \frac{\lambda}{\alpha}k + \frac{\alpha}{\alpha + \lambda^2} \sum_{i=0}^{\infty} \left(\frac{\lambda^2}{\alpha + \lambda^2} \right)^i \hat{E}_t^i(\pi_t^*) + \frac{\lambda(1-\theta)}{\alpha + \lambda^2} \sum_{i=0}^{\infty} \left(\frac{\lambda^2}{\alpha + \lambda^2} \right)^i \hat{E}_t^i \left[E(y_t^*|\mathcal{G}_t) - \hat{E}(y_t^*) \right] \quad (5)$$

and

$$y_t = \theta y_t^* + (1-\theta)E(y_t^*|\mathcal{F}_t) + \lambda(\pi_t - E[\pi_t|\mathcal{F}_t]), \quad (6)$$

where, given any random variable X , $\hat{E}_t(X) := E[E(X|\mathcal{F}_t)|\mathcal{G}_t]$, and the i power in the sum indicates repeated application of the \hat{E}_t operator.

Let \hat{V} be the set of equilibrium payoffs for an economy where the message space is empty.

We first prove that $\hat{V} \subseteq V$. This is straightforward, because in cheap talk games babbling is always an equilibrium. Specifically, let $\hat{\sigma}^g, \hat{\sigma}^p$ be an equilibrium of the game with no messages. Let m^0 be an arbitrary message from the set M . We construct an equilibrium of the game with messages as follows.

First,

$$\sigma^g(\omega, h^t) = m^0 \quad \forall h^t;$$

³³Equation (5) is quite involved when it is not possible to establish that either the government or the households have superior information, since in this case higher-order beliefs matter explicitly. In our examples, either $\mathcal{G}_t \subseteq \mathcal{F}_t$ or $\mathcal{F}_t \subset \mathcal{G}_t$. In the former case, equation (5) simplifies to $\pi_t = (\lambda/\alpha)k + \pi_t^*$ and in the latter to

$$\pi_t = \frac{\lambda}{\alpha}k + \frac{\alpha}{\alpha + \lambda^2}\pi_t^* + \frac{\lambda^2}{\alpha + \lambda^2}E(\pi_t^*|\mathcal{F}_t) + \frac{\lambda(1-\theta)}{\alpha + \lambda^2} [E(y_t^*|\mathcal{G}_t) - E(y_t^*|\mathcal{F}_t)].$$

then,

$$\sigma^p(\omega, h^{et}) = \hat{\sigma}^p(\omega, \hat{h}^{et}) \quad \forall h^t,$$

where \hat{h}^{et} is the history that shares the same values of $(\pi_s, \pi_s^e)_{s=0}^{t-1}$ as h^{et} ; and finally,

$$\sigma^g(\omega, h^{et}) = \hat{\sigma}^g(\omega, \hat{h}^{et}) \quad \forall h^t,$$

where \hat{h}^{et} and h^{et} are related as above.

In this equilibrium, household expectations and government inflation choices are independent of the messages being sent, so the government is indifferent over which message to send, and m^0 is (weakly) optimal. Households do not learn any information from government messages, both because the government and the households share the same information sets and because the government message is independent of its information. By construction, the probability distribution over future outcomes for inflation and inflation expectations is the same in the two games, given any choice of π_t for the government. Hence, the choices that were optimal for $\hat{\sigma}^g$ remain optimal under σ^g . Similarly, since the probability distribution over future outcomes remains the same, household expectations remain rational.

Consider now the converse: $V \subseteq \hat{V}$. Let (σ^g, σ^p) be an equilibrium that attains the value $v \in V$ for a game with messages. Let the sunspot in the original economy with messages be $s_t \in [0, 1]^n$. In the economy without messages, households and the government will coordinate based on a sunspot $\hat{s}_t \in [0, 1]^{n+q}$: the additional q dimensions of the sunspot will replace government messages as a source of coordination.³⁴

In each period t , and for (almost) each information set in \mathcal{G}_t , we create a mapping η_t from

³⁴The power of the theorem is that, while government messages may be correlated with the fundamentals of the economy and/or the past history of play, sunspots are not by construction: since messages play at most a pure coordination role and convey no extra information, there is no need for a strategic player to be the sender. The theorem could easily be extended to more general message spaces than \mathbb{R}^q , but this requires identifying the equivalent of the mapping η below, which is difficult to characterize at the current level of generality.

$\hat{s}_t \in [0, 1]^{n+q}$ to $(s_t, m_t) \in [0, 1]^n \times \mathbb{R}^q$:

$$\left\{ \begin{array}{ll} \eta_{it}(\hat{s}_t) = s_{it} & i = 1, \dots, n \\ \eta_{n+1t}(\hat{s}_t) = \inf_v (v : \text{Prob}(m_{1t} \leq v) \geq \hat{s}_{n+1t}) & \\ \eta_{n+it}(\hat{s}_t) = \inf_v (v : \text{Prob}[m_{it} \leq v | (m_{1t}, \dots, m_{i-1t}) = & i = 2, \dots, q. \\ (\eta_{n+1t}(\hat{s}_t), \dots, \eta_{n+i-1t}(\hat{s}_t))] \geq \hat{s}_{n+it}) & \end{array} \right. \quad (7)$$

Next, define recursively a mapping η from histories of the game without messages at the expectation-setting stage to the game with messages at the same stage as follows:

$$\begin{aligned} \eta(\hat{s}_0) &:= \eta_0(s_0, m_0), \\ \eta(\hat{h}^{et-1}, (\pi_{t-1}, \hat{s}_t)) &:= (h^{et-1}, \pi_{t-1}, \eta_t(\hat{s}_t)). \end{aligned}$$

Define the following strategy profiles for the government and the households in the game without messages:

$$\begin{aligned} \hat{\sigma}^g(\hat{h}^{et}) &:= \sigma^g(\eta(\hat{h}^{et})), \\ \hat{\sigma}^p(\hat{h}^{et}) &:= \sigma^p(\eta(\hat{h}^{et})). \end{aligned}$$

By the construction of the mapping η , the probability distribution over $(\omega, \{\pi_t, y_t\}_{t=0}^\infty)$ remains the same, so the government attains the same payoff when $(\hat{\sigma}^g, \hat{\sigma}^p)$ is played in the game without messages as it does according to (σ^g, σ^p) in the game with messages. Similarly, the distribution over future outcomes in the two games remain the same starting from arbitrary histories \hat{h}^{et} and $\eta(\hat{h}^{et})$, whether these histories do or do not occur on the path of play of the corresponding strategy profiles. Hence, if $\sigma^g(\eta(\hat{h}^{et}))$ is an optimal distribution of inflation for the government at h^{et} in the game with messages, then the same distribution $\hat{\sigma}^g(\hat{h}^{et})$ is optimal in the game without messages. Similarly, if $\sigma^p(\eta(\hat{h}^{et}))$ represents rational expectations for the households about (π_t, y_t) , then the same expectations $\hat{\sigma}^p(\hat{h}^{et})$ are also rational in the game without messages. QED.

A.2 A convenient reformulation of the environment

The propositions of Sections 4 and 5 share a common structure. In order not to repeat steps unnecessarily, it is convenient to rewrite the problem by defining

$$z_t := -y_t^*(1 - \theta) + \lambda\pi_t, \quad (8)$$

$$z_t^* := -y_t^*(1 - \theta) + \lambda\pi_t^*, \quad (9)$$

and

$$z_t^e = -y_t^e(1 - \theta) + \lambda\pi_t^e \quad (10)$$

With these definitions, from equations (1) and (2) the loss function for the government becomes

$$(1 - \beta)E \sum_{t=0}^{\infty} \beta^t \left[(z_t - z_t^e - k)^2 + \frac{\alpha}{\lambda^2} (z_t - z_t^*)^2 \right] \quad (11)$$

Let \mathcal{F}_t^m be the filtration that represents the household information after receiving any messages that the CB may send in period t ; this includes what they directly observe (\mathcal{F}_t), but it is finer if the CB provides informative messages. Taking the expectation of equation (1) conditional on \mathcal{F}_t^m , we obtain $E(y_t|\mathcal{F}_t^m) = E(y_t^*|\mathcal{F}_t^m)$ and thus

$$y_t = \theta y_t^* + (1 - \theta)E(y_t^*|\mathcal{F}_t^m) + \lambda[\pi_t - E(\pi_t|\mathcal{F}_t^m)] \quad (12)$$

In Sections 4 and 5 the government has superior information compared to the households. The objective function can then be rewritten as

$$(1 - \beta)E \sum_{t=0}^{\infty} \beta^t \left[\text{Var}(z_t|\mathcal{G}_t) + (E(z_t|\mathcal{G}_t) - E(z_t|\mathcal{F}_t^m) - k)^2 + \frac{\alpha}{\lambda^2} (E(z_t|\mathcal{G}_t) - z_t^*)^2 \right]. \quad (13)$$

The initial variance term is independent of the actions of the government. By choice of π_t , the government can instead control $E(z_t|\mathcal{G}_t)$. This specification has the advantage of being the same whether the private information of the government is about y_t^* , as in Section 4, or about π_t^* , as in Section 5; it is this symmetry that we exploit to condense proofs. Of course, the economic interpretation of the required information disclosure is different in the two cases, which is at the heart of the arguments in this paper.

In the proofs that follow, we save on notation by assuming that y_t^* is known to the CB, so that we can omit the variance term and replace $E(z_t|\mathcal{G}_t)$ with z_t ;³⁵ this is purely for notational simplicity, as the proofs go through identically in the more general case.³⁶

³⁵Assuming that π_t^* is known to the CB is without loss of generality, since π_t^* is its inflation target.

³⁶At the cost of additional complexity, similar propositions and proofs could also be derived for the more-general case in which both the government and households have information not known to the other party. For simplicity, we limit ourselves to the case in which one party has superior information.

A.3 Proof of proposition 2

After reformulating the problem as in Section A.2 and using the law of iterated expectations, we obtain

$$E \left[(z_t - E(z_t | \mathcal{F}_t^m) - k)^2 + \frac{\alpha}{\lambda^2} (z_t - z_t^*)^2 \right] = E \left\{ E \left[(z_t - E(z_t | \mathcal{F}_t^m) - k)^2 + \frac{\alpha}{\lambda^2} (z_t - z_t^*)^2 | \mathcal{F}_t \right] \right\}, \quad (14)$$

and

$$E \left[(z_t - E(z_t | \mathcal{F}_t^m) - k)^2 + \frac{\alpha}{\lambda^2} (z_t - z_t^*)^2 | \mathcal{F}_t \right] = k^2 + E \left\{ [z_t - E(z_t | \mathcal{F}_t^m)]^2 | \mathcal{F}_t \right\} + \frac{\alpha}{\lambda^2} E[(z_t - z_t^*)^2 | \mathcal{F}_t] \geq k^2 \quad (15)$$

The lower bound in equation (15) is exogenous to government policy, and it is attained when the government strategy sets $z_t = z_t^*$ (equivalently, $\pi_t = \pi_t^*$) and provides a message $m_t = z_t^*$, which ensures $E(z_t^* | \mathcal{F}_t^m) = z_t^*$. In the context of Section 4, a message about z_t^* is equivalent to a message about y_t^* , since households know π_t^* . QED.

A.4 Proof of proposition 3

- (i) Proposition 2 proved that the commitment outcome when messages are allowed requires the government to disclose truthfully its information. Since there exists a period t such that $\text{Prob}(E[z_t^* | \mathcal{F}_t] \neq E[z_t^* | \mathcal{G}_t]) > 0$, there is at least one period t and a set of states of nature of positive probability in which the government has superior information and this disclosure reveals information to the households: this disclosure cannot be replicated without messages, so the commitment outcome when messages are allowed strictly dominates the commitment outcome when no messages are allowed.

We now use the standard logic of the Folk theorem to prove that there exists a value $\bar{\beta}$ such that, if $\beta > \bar{\beta}$, then the commitment outcome with messages is an equilibrium outcome of the game (with messages) under discretion. In Proposition 2, we have shown that the commitment outcome achieves a payoff of k^2 .

We look for an equilibrium in which the commitment outcome is supported by the threat of permanent reversion to the equilibrium of the static one-shot game in which

the government babbles, which we now compute. In this equilibrium, household expectations are independent of past history and of the message sent by the government. Since government messages carry no consequences, we can assume without loss of generality that the government always reports some default message m_0 . In setting inflation, the government takes into account that its action has no consequences on the future, and it thus solves

$$\min_{z_t} E \left[(z_t - z_t^e - k)^2 + \frac{\alpha}{\lambda^2} (z_t - z_t^*)^2 | \mathcal{G}_t \right],$$

taking z_t^e as given. The first-order condition yields

$$z_t = \frac{\lambda^2}{\alpha + \lambda^2} (z_t^e + k) + \frac{\alpha}{\alpha + \lambda^2} z_t^* \quad (16)$$

Taking expected value of (16) conditional on private-sector information and imposing the equilibrium conditions $z_t^e = E(z_t | \mathcal{F}_t)$, we obtain

$$E(z_t | \mathcal{F}_t) = z_t^* + \frac{\lambda^2 k}{\alpha}.$$

Substituting this into (16) we solve for the equilibrium level of z_t :³⁷

$$z_t = z_t^* + \frac{\lambda^2 k}{\alpha} - \frac{\lambda^2}{\alpha + \lambda^2} [z_t^* - E(z_t^* | \mathcal{F}_t)].$$

The payoff of following the static babbling equilibrium from period s into the indefinite future is thus

$$\begin{aligned} & (1 - \beta) E \sum_{t=0}^{\infty} \beta^t \left[(z_t - E(z_t | \mathcal{F}_t) - k)^2 + \frac{\alpha}{\lambda^2} (z_t - z_t^*)^2 \right] = \\ & k^2 (1 + \lambda^2 / \alpha) + (1 - \beta) \frac{\alpha}{\alpha + \lambda^2} \sum_{t=s}^{\infty} \beta^{t-s} E \left[(y_t^* - E(y_t^* | \mathcal{F}_t))^2 | \mathcal{G}_s \right] \geq \quad (17) \\ & k^2 (1 + \lambda^2 / \alpha). \end{aligned}$$

Comparing equation (17) with the loss under commitment, which we established to be k^2 , a government which chooses to deviate from the commitment play in period s will

³⁷Note that our assumptions in footnote 10 ensure that the solution for inflation above is interior. Furthermore, there are no other solutions in which (16) may hold as an inequality due to inflation being sometimes at a corner: this is because $\partial \pi / \partial \pi_t^e \in [0, 1)$ according to (16) when a corner may be binding, so that, starting from $\pi_t^e = \pi_t^* + \frac{\lambda k}{\alpha}$, moving in either direction would be incompatible with a fixed point.

lose at least $k^2\lambda^2/\alpha$ from period $s + 1$ onwards. The best that the government can achieve through such a deviation is to set the loss in period s to zero, which yields a one-period gain of k^2 . We can thus choose

$$\bar{\beta} = \frac{\alpha}{\alpha + \lambda^2} : \quad (18)$$

for any value of $\beta \geq \bar{\beta}$, the one-shot gain of deviating from the commitment outcome is smaller than the loss from subsequent reversion to the static babbling equilibrium. Hence, for such values of β , the commitment outcome is an equilibrium outcome even without commitment; this outcome is better than the best outcome that can be attained without messages, completing the proof.

(ii) We first establish a couple of properties which will be used in the remainder of the proof.

1. Given any period t and any government information set $G_t \in \mathcal{G}_t$, the set of expected payoffs of continuation equilibria from period $t + 1$ onwards is convex. To prove this, consider two continuation equilibria with payoff L_{t+1}^1 and L_{t+1}^2 . Let these equilibria sunspot processes $\{s_s\} \in \mathbb{R}^n$ (assuming the same n is without loss of generality, since the equilibrium play can simply ignore some of the components of s_t). To attain a continuation payoff $qL_{t+1}^1 + (1 - q)L_{t+1}^2$, we construct a continuation equilibrium based on a sunspot process $\{\hat{s}_s\} \in \mathbb{R}^{n+1}$. Specifically, whenever $s_{n+1,t+1} \leq q$, the continuation equilibrium from period $t + 1$ coincides with the equilibrium attaining L_{t+1}^1 , whereas otherwise it coincides with the equilibrium L_{t+1}^2 .
2. In part (i) of the proof, we established a threshold $\bar{\beta}$ that ensured that, in the game with messages, the commitment outcome can be sustained as an equilibrium outcome of the game with discretion using the threat of reversion to the static babbling equilibrium. We now show that the same threshold is also sufficient to sustain the commitment outcome under the same threat even in the game without messages. Simple algebra proves that the commitment outcome of the

game without messages is given by the following policy:

$$z_t = \frac{\lambda^2}{\alpha + \lambda^2} z_t^* + \frac{\alpha}{\alpha + \lambda^2} E(z_t^*).$$

It follows that, in the game without messages, the difference in payoffs between the commitment outcome and repeating the static equilibrium outcome is exactly $k^2\lambda^2/\alpha^2$, the value that we used as a lower bound for the game in which messages are allowed. The one-time gain from deviating is bounded above by $[k+(1-\theta)(y^h - y^\ell)]^2$, and so any value $\beta > [k+(1-\theta)(y^h - y^\ell)]^2 / (k^2\lambda^2/\alpha^2 + [k+(1-\theta)(y^h - y^\ell)]^2)$ will once again ensure that the one-time benefit of deviation from the commitment outcome is smaller than the subsequent loss.³⁸

We now reason by contradiction. Let (σ_w^g, σ_w^p) be an equilibrium of the game with messages that attains a loss L_w higher than the supremum of the losses attainable in the game without messages, which we denote \hat{L} . Since there is no link in the game across different information sets of the private sector, there necessarily is a period-0 information set $F_0 \in \mathcal{F}_0$ for which the loss according to the equilibrium (σ_w^g, σ_w^p) is worse than the conditional supremum of the losses attainable in the game without messages; we establish a contradiction for this information set. Let $L_1(F_0), F_0 \in \mathcal{F}_0$ be the expected loss attained by the equilibrium (σ_w^g, σ_w^p) from period 1 onwards, conditional on F_0 . We distinguish two cases, and rule out each of them in turn:

1. Consider first the case in which $L_1(F_0)$ is smaller than the lowest attainable conditional expected loss from period 1 in the game without messages. This is impossible, provided β is sufficiently high. Specifically, the period-0 loss is bounded above by³⁹

$$L_{\text{worst}} := (1 - \beta)[(y^\ell - y^h - k)^2 + (\bar{\pi} - \pi^\ell)^2]. \quad (19)$$

³⁸When asymmetric information is about π_t^* rather than y_t^* , the one-time gain from deviating is bounded above by $[k + \lambda(\pi^h - \pi^\ell)]^2$, and the threshold for β is adjusted accordingly.

³⁹The worst-case could in principle be either $\underline{\pi}$ or $\bar{\pi}$, but we assume that $\bar{\pi}$ is sufficiently large, so that high inflation is worst. The proof goes through with appropriate replacements if $\underline{\pi}$ is the worst-case scenario. In a model with deeper microfoundations, the zero lower bound imposes a limit on *expected* deflation, which implies that high inflation would be the worst-case scenario.

For $\beta > \bar{\beta}$, the commitment outcome of the game without messages is sustainable from period 1, so the loss from then on must be smaller than such outcome. However, when

$$\beta > \max\left\{\frac{L_{\text{worst}}}{L_{\text{worst}} + \frac{k\lambda^2}{\alpha^2}}, \bar{\beta}\right\},$$

the loss from receiving L_{worst} in period 0 and the commitment outcome of the game without messages from period 1 is smaller than the loss from playing the repeated static babbling equilibrium. Since the repeated static babbling equilibrium is attainable in the game without messages, this contradicts the premise that L_w is higher than the supremum of the losses in the game without messages.

2. Suppose instead that $L_1(F_0)$ is not smaller than the lowest attainable conditional expected loss from period 1 in the game without messages.

To find the supremum of the losses in the game with messages as of period 0 and conditional on F_0 , we follow the logic of Abreu et al. (1986, 1990). Define σ_0^g a government strategy for period 0 (a mapping from \mathcal{G}_0 to a probability distribution over messages m_0 and inflation rates π_0) and let $E^{\sigma_0^g}$ denote expectations conditional on the fact that the government will follow the strategy σ_0^g in period 0. The supremum of the loss will solve the following problem:

$$\sup_{\sigma_0^g, L_1} (1 - \beta) E^{\sigma_0^g} \left\{ \left[z_0 - E^{\sigma_0^g}(z_0|F_0, m_0) - k \right]^2 + \frac{\alpha}{\lambda^2} (z_0 - z_0^*)^2 | F_0 \right\} + \beta E(L_1(G_0)|F_0) \quad (20)$$

subject to the incentive-compatibility constraint⁴⁰

$$\begin{aligned} (1 - \beta) \left[z_0 - E^{\sigma_0^g}(z_0|F_0, m_0) - k \right]^2 + \frac{\alpha}{\lambda^2} (z_0 - z_0^*)^2 + \beta L_1(G_0) &\leq \\ (1 - \beta) \inf_{\tilde{m}_0, \tilde{z}_0} \left\{ \left[\tilde{z}_0 - E^{\sigma_0^g}(z_0|F_0, \tilde{m}_0) - k \right]^2 + \frac{\alpha}{\lambda^2} (\tilde{z}_0 - z_0^*)^2 \right\} + \beta \sup \mathcal{L}(G_0) &\end{aligned} \quad (21)$$

and subject to $\pi_0 \in [\underline{\pi}, \bar{\pi}]$ and $L_1(G_0) \in \mathcal{L}(G_0)$, where $\mathcal{L}(G_0)$ is a convex set that represents all possible values from continuation equilibria from period 1 onwards

⁴⁰If the supremum of $\mathcal{L}(G_0)$ cannot be attained, then the incentive-compatibility constraint should hold as an inequality. However, it is straightforward to show that the theorem of the maximum implies that the overall problem is continuous in $\sup \mathcal{L}(G_0)$ (endowed with the L^1 norm), so that the solution of the supremum in (20) will be the same if equality is included in (21).

conditional on the *government's* information as of period 0. In terms of z_0 , the restriction on π_0 is equivalent to

$$z_0 \in [\underline{z}, \bar{z}] := [-y_0^*(1 - \theta) + \lambda \underline{\pi}, -y_0^*(1 - \theta) + \lambda \bar{\pi}].$$

In the problem above, we first observe that equation (21) must be binding almost surely: if this were not the case on a set of values of G_0 of positive measure and $L_1(G_0) < \sup \mathcal{L}(G_0)$ on those values, an equilibrium with a higher loss could be attained by selecting a continuation equilibrium from period 1 that corresponds to a higher value of $L_1(G_0)$; if instead $L_1(G_0) = \sup \mathcal{L}(G_0)$ almost surely, then by the definition of its terms (21) could not hold as a strict inequality.

Second, we prove that $L_1(G_0) \geq \inf \mathcal{L}(G_0)$ is not binding when

$$\beta > \max \left\{ \frac{L_{\text{worst}}}{L_{\text{worst}} + \frac{k\lambda^2}{\alpha^2}}, \bar{\beta} \right\}.$$

This follows the proof of step 1:

$$\left[z_0 - E^{\sigma_0^g}(z_0 | F_0, m_0) - k \right]^2 + \frac{\alpha}{\lambda^2} (z_t - z_t^*)^2 \leq L_{\text{worst}}$$

and

$$\inf_{\tilde{m}_0, \tilde{z}_0} \left\{ \left[z_0 - E^{\sigma_0^g}(z_0 | F_0, \tilde{m}_0) - k \right]^2 + \frac{\alpha}{\lambda^2} (\tilde{z}_0 - z_0^*)^2 \right\} \geq 0;$$

furthermore, for the given range of β , $\sup \mathcal{L}(G_0) - \inf \mathcal{L}(G_0) \geq k^2 \lambda^2 / \alpha$, since the commitment outcome is sustainable and the worst equilibrium outcome is at least as bad as the repeated static babbling equilibrium. It then follows that (21) would not hold as an equality if $L_1(G_0) = \inf \mathcal{L}(G_0)$, contradicting what we already proved.

Having proved that $\inf \mathcal{L}(G_0) \leq L_1(G_0) \leq \sup \mathcal{L}(G_0)$ are not binding constraints, while (21) is, the plan that maximizes (20) subject to (21) is equivalent to the plan that solves

$$\sup_{\sigma_0^g} \inf_{\tilde{m}_0, \tilde{z}_0} \left\{ E \left[\left[\tilde{z}_0 - E^{\sigma_0^g}(z_0 | F_0, \tilde{m}_0) - k \right]^2 + \frac{\alpha}{\lambda^2} (\tilde{z}_0 - z_0^*)^2 \right] | F_0 \right\}, \quad (22)$$

with $L_1(G_0)$ determined as a residual so as to ensure that (21) holds with equality.

To solve the problem in (22), we first compute the infimum with respect to \tilde{z}_0 , for given \tilde{m}_0 and σ_0^g . The first-order condition yields

$$\tilde{z}_0 - E^{\sigma_0^g}(z_0|F_0, \tilde{m}_0) - k + \frac{\alpha}{\lambda^2}(\tilde{z}_0 - z_0^*) = 0,$$

which (given the assumptions on $\underline{\pi}$ and $\bar{\pi}$ in footnote 10) yields an interior solution

$$\tilde{z}_0 = \frac{\lambda^2}{\alpha + \lambda^2} \left[E^{\sigma_0^g}(z_0|F_0, \tilde{m}_0) - k \right] + \frac{\alpha}{\alpha + \lambda^2} z_0^*. \quad (23)$$

Substituting (23) into (22) and simplifying the algebra, the problem becomes

$$\sup_{\sigma_0^g} E^{\sigma_0^g} \left\{ \inf_{\tilde{m}_0} \frac{\alpha}{\alpha + \lambda^2} \left[z_0^* - E^{\sigma_0^g}(z_0|F_0, \tilde{m}_0) - k \right]^2 | F_0 \right\}. \quad (24)$$

We have

$$\begin{aligned} E^{\sigma_0^g} \left\{ \inf_{\tilde{m}_0} \left[z_0^* - E^{\sigma_0^g}(z_0|F_0, \tilde{m}_0) - k \right]^2 | F_0 \right\} &\leq E^{\sigma_0^g} \left\{ \left[z_0^* - E^{\sigma_0^g}(z_0|F_0, m_0) - k \right]^2 | F_0 \right\} = \\ E^{\sigma_0^g} \left\{ \left[z_0^* - E^{\sigma_0^g}(z_0^*|F_0, m_0) \right]^2 + \left[E^{\sigma_0^g}(z_0^*|F_0, m_0) - E^{\sigma_0^g}(z_0|F_0, m_0) - k \right]^2 | F_0 \right\} &\leq \\ E \left\{ \left[z_0^* - E(z_0^*|F_0) \right]^2 + \left[-k + \lambda(\pi_0^* - E^{\sigma_0^g}(\pi_0|F_0, m_0)) \right]^2 | F_0 \right\} &\leq \\ E \left\{ \left[z_0^* - E(z_0^*|F_0) \right]^2 + \left[-k + \lambda(\pi_0^* - \bar{\pi}) \right]^2 | F_0 \right\}. & \end{aligned} \quad (25)$$

The last line holds provided $\bar{\pi}$ is sufficiently high, given $\alpha, \lambda, k, \underline{\pi}$, and the stochastic process for $(y_t, \tilde{y}_t, \pi_t^*)_{t=0}^\infty$ (but, importantly, independently of β , which can thus be set sufficiently high in turn as to ensure that the bounds of the continuation payoff are not binding).

If the CB plays a strategy that sets $\pi_0 = \bar{\pi}$ with probability 1 and babbles (i.e., it sends a message that is independent of its information), the resulting payoff attains exactly the last line of equation (25): hence, this is the strategy that achieves the supremum in problem (24). We thus conclude that a strategy that achieves the supremum of the losses in period 0 involves babbling,⁴¹ so that the ability or inability of the CB to send messages in period 0 has no consequences

⁴¹If $\mathcal{L}(G_0)$ is a closed set, the supremum is attained and the strategy just described is the unique strategy that attains the worst payoff

for the upper bound of the set of equilibrium payoffs. Hence, if the supremum of the equilibrium payoffs in the game with messages is higher than in the game without messages, this must be because the continuation payoff $E[L_1(G_0)]$ is worse than the worst possible equilibrium payoff in the game without messages. But we can repeat the same steps as of period 1 and reach the same conclusion: that the worst equilibrium payoff as of period 1 is the same whether messages are or are not allowed in period 1, so that, if $E[L_1(G_0)]$ is worse than the worst possible equilibrium in the game without messages, it must be because there is some information set F_1 such that the continuation payoff $E[L_2(G_1)|F_1]$ is worse than the worst possible equilibrium payoff in the game without messages. We can proceed iteratively and prove that the ability of sending messages up to any arbitrary period t does not affect the worst possible equilibrium payoff. This establishes a contradiction: if the upper bound of the equilibrium payoffs in the game with messages is strictly worse than in the game without messages, then there must be a period t and an information set F_t in which the government sends informative messages. QED.

A.5 Proof of proposition 4

This proof mimics closely that of Proposition 2, by exploiting the reformulation in Appendix A.2. The only difference is that, in this case, since the optimal policy under commitment features $\pi_t = \pi_t^*$, it is equivalent for the CB to send messages about the target π_t^* or the actual policy π_t .⁴²

⁴²Messages about π_t^* are in turn equivalent to messages about z_t^* , since in this context y_t^* is known to the private sector.

B An Economy in which the Policy Target varies with the State of the Economy

In the introduction, I have discussed forward guidance as a situation in which the CB sends messages about future policy. As is well known, in a cheap-talk environment, the message space can be abstract and there are many alternative methods to convey the same information. To give examples, the CB could make announcements in English, Spanish, or Urdu; or it could simply state “blue” or “red:” and in equilibrium people would know that “blue” and “red” mean that inflation will be 2% or 3%. In the baseline case of Section 4, the central bank can equally well send messages about z_t^* or y_t^* , since there is a one-to-one relationship between the two.⁴³ This extends to the case in which the CB does not know y_t^* with certainty: there remains a one-to-one mapping between $E(z_t^*|\mathcal{G}_t)$ and $E(y_t^*|\mathcal{G}_t)$.

Consider now an economy in which the government loss function is described by equation (4). By redefining

$$z_t^* := -y_t^*(1 - \theta) + \lambda[\pi_t^* + f(y_t^*)],$$

we can recover the same reformulation as in Appendix A.2, and we can thus rederive the theorems of Section 4. The optimal policy under commitment requires once again full disclosure of $E(z_t^*|\mathcal{G}_t)$, along with setting $z_t = z_t^*$ (that is, $\pi_t = \pi_t^* + f(y_t^*)$). Consider first the case in which y_t^* is known to the CB at time t . In this case, sending a message about π_t allows the CB to achieve the full-commitment outcome only if f is monotone, in which case

$$E(y_t^*|\mathcal{F}_t, \pi_t^* + f(y_t^*)) = y_t^*$$

Whenever

$$Prob \{f(y_t^1) = f(y_t^2) \wedge y_t^1 \neq y_t^2 | \mathcal{F}_t\} > 0$$

we will have

$$E(y_t^*|\mathcal{F}_t, \pi_t^* + f(y_t^*)) \neq y_t^*,$$

⁴³In Section 4, π_t^* is common knowledge in period t .

so that revealing only information about inflation will be insufficient to achieve the same payoff as the commitment outcome with full disclosure.

More in general, when the CB does not know y_t^* , the optimal ex-ante inflation is $\pi_t^* + E(f(y_t^*)|\mathcal{G}_t)$. Whenever

$$\text{Prob} \{f(y_t^1) = f(y_t^2) \wedge E(y_t^1|\mathcal{G}_t) \neq E(y_t^2|\mathcal{G}_t)|\mathcal{F}_t\} > 0,$$

we will have

$$E(y_t^*|\mathcal{F}_t, \pi_t^* + f(y_t^*)) \neq E(y_t^*|\mathcal{G}_t),$$

and once again sending messages about inflation only is insufficient to deliver the outcome which the CB can attain through full disclosure and setting inflation at its target.

C Reformulating the Economy of Section 6

For this case, define

$$z_t := \lambda\pi_t, \tag{26}$$

$$z_t^* := \lambda(\pi_t^* + E[f(y_t^*)|\mathcal{G}_t]), \tag{27}$$

and

$$z_t^e = \lambda\pi_t^e. \tag{28}$$

With these definitions, the objective maps once more into (11), which allows us to prove the same theorems as in the other cases.

C.1 Proof of proposition 6

This proof mimics closely that of Proposition 2, by exploiting the reformulation of equations (26)-(28). Since the optimal policy under commitment features $\pi_t = \pi_t^* + E(f(y_t^*|\mathcal{G}_t))$, it is equivalent for the CB to send messages about the signal \tilde{y}_t (which determines $E(f(y_t^*|\mathcal{G}_t))$) or the actual policy π_t .