Reputation for Confidence*

Laura Gáti† and Amy Handlan‡

February 7, 2026

Abstract

We model how a central bank communicates their noisy forecasts (forward guidance) while taking into account their own uncertainty (confidence) and the public's perception of the bank's uncertainty (*reputation for confidence*). This creates a mismatch between the public and central bank's interpretation of the bank announcement which induces the bank to communicate with partial transparency and deliberate imprecision. Moreover, with higher confidence (lower reputation) announcements are more precise. With text data from internal Fed documents and newspapers, we find communication patterns are largely consistent with the model except the Fed's communication strategy underreacts to reputation compared to the model.

Keywords: communication, forward guidance, reputation, cheap talk, text analysis

JEL codes: E52, E58, C49

*We are grateful for the helpful comments and suggestions from our discussants, James Chapman and Stephanie Ettmeier, and from the many seminar and conference participants. We also thank Manuel Amador, Marco Bassetto, Marek Jarociǹski, and Chris Phelan for their helpful comments. Also, we thank Frank Chiu, Halleluiah Girum, Tobias Müller and Marco Olivari for their research assistance. The views expressed herein are our own and do not necessarily reflect those of the ECB or the Eurosystem. A earlier version of this project was circulated as "Reputation for Competence."

†European Central Bank, Directorate General Research, Sonnemannstraße 20, 60314 Frankfurt am Main, Germany and CEPR. Email: laura_veronika.gati@ecb.europa.eu.

‡Brown University, Department of Economics, 64 Waterman Street, Providence, Rhode Island, USA. Email: amy_handlan@brown.edu.

1 Introduction

> *"As always, there is no guarantee that statement language will be interpreted as intended. We know that. But communicating as much information as we are fairly* ***confident*** *in is desirable."*
>
> - Patrick Harker, Federal Reserve Bank of Philadelphia President, 2019

Monetary policy announcements attract intense attention from media and investors trying to decipher the underlying meaning of the Federal Reserve's (Fed's) guidance. The Fed dedicates time and resources to writing announcements that balance providing detailed information to guide the public's expectations about future policy without implying a strict commitment to past guidance. Meanwhile, how the public responds to the announcement depends on whether they believe the Fed is informed about the economic situation. There is a tension between precise guidance and anticipating the market's reaction which affects the communication strategy for policy announcements.

This paper develops a game-theoretic model to examine how central banks manage communication when they and the market hold different beliefs about the precision of signals regarding policy-relevant shocks. We test the model using text data from and about the Fed. Theoretically, we find that the both bank's own uncertainty (or *confidence*) and the market's views of the bank's uncertainty, which we call the bank's *reputation for confidence*, shape how precisely the bank communicates. While more precise announcements provide stronger forward guidance and shift expectations more, markets may overreact or underreact depending on reputation. The bank balances these tradeoffs by preparing an endogenous set of alternative statement drafts (the *message space*) from which they choose the final policy statement. When the central bank is more confident in their forecasts – or, analogously, there is more agreement between the members of the policy committee – the bank drafts announcements that are more precise. This increase in precision also occurs when the bank's reputation is lower because a lower reputation reduces the market's sensitivity to announcements. To test these dynamics, we construct three text-based measures: (1) the Fed's message space from the set of alternative statements drafted for the Federal Open Market Committee (FOMC) policy meeting, (2) confidence from uncertainty-word use in

transcripts, and (3) reputation based on the Baker et al. (2016) measure of newspaper coverage of Fed uncertainty. The evidence broadly supports the model in that higher confidence or lower reputation are associated with more precise announcements. However, in the data, Fed communication is less responsive to reputation than the model predicts.

In the cheap talk communication game between the central bank and the public, the central bank sets inflation and uses communication policy to stabilize inflation and output in an economy characterized by a simple Phillips curve relationship with inflation-target shocks, similar to the model in Moscarini (2007). Only the central bank observes noisy signals of the shock, capturing the idea that central banks in practice make monetary policy decisions under uncertainty and behind closed doors. We call the precision of the central bank's signal the central bank's *confidence*, so that a more confident central bank forms more accurate forecasts of the shock that hit the economy.[1]

The key novelty of the theoretical setup is that we assume that the public does not observe the central bank's confidence and instead forms beliefs over confidence from economic outcomes through Bayesian learning.[2] We call the public's beliefs over central bank confidence the central bank's *reputation for confidence*. The terminology here reflects the idea that even a central banker with preferences perfectly aligned with those of the public may be operating in a difficult forecasting environment, and the public may have differing views from the central bank on the uncertainty of the environment, just like in Caballero and Simsek (2022).

We conceptualize central bank communication as drafting a set of candidate announcements - which we refer to as the *message space* - from which the bank picks a single message to make public. Importantly, the message space is an equilibrium object where the central bank has no incentive to send any other message than the equilibrium message given the optimal action rules and the signal observed by the central bank. Optimal announcements in the model take the form of *intervals* that contain the central bank's signal of the shock.

[1] We can also think of confidence as standing in for the agreement between various members on the policy committee.

[2] We can microfound this assumption in two ways. First, the central bank cannot credibly communicate their confidence to the public because they are always better off when the public believes the bank is better at forecasting. Second, even if they could communicate it, the expectation survey literature documents that humans have difficulty understanding second moments. Accordingly, we think of confidence as not perfectly seen by the public.

The precision of the announcement is measured as the width of the interval.

The analysis yields two striking results. First, *point revelation* – a situation in which the central bank reveals the exact value of the signal they saw – is possible if and only if reputation equals true confidence. What this is saying is that if the public misperceives how confident the central bank is, then the central bank cannot exactly reveal their information because the information will be misinterpreted. In a world with an *misunderstanding friction*, it is therefore optimal for the central bank to send imprecise messages so that the announcements are misinterpreted in the right way. This finding is a contribution to the central bank communication game literature which historically relies on the central bank wanting to "trick" the public with an inflation surprise (to generate output growth) to obtain coarse communication in equilibrium.[3] We argue that this *misunderstanding friction* better explains central bank communication, rationalizing why central bankers emphasize transparency while at the same time only partially disclosing information.[4] Additionally, we find narrative-evidence from FOMC transcripts which quote policymakers choosing language while struggling with the misunderstanding friction.[5]

Secondly, the model provides predictions for the optimal drafting of alternative statements. In particular, higher confidence or lower reputation lead to more precise messages. Confidence is more strongly associated with the *span* of the alternatives, meaning the distance between the most different alternatives, because a confident central bank may only need to discuss a small range of scenarios, for example those of inflation being slightly below or slightly above the target. A less confident central bank, or one with high disagreement between committee members, will need to be ready to discuss scenarios that are further apart. For example, if the bank is discussing statements where inflation is significantly over-

[3]Inflation surprises refer to models where the central bank targets a positive output gap, and thus tries to trick the public into low inflation expectations so that the bank can secretly create inflation, generating a positive output gap through the Phillips curve.

[4]An example of partial information revelation is that FOMC meeting materials, including staff forecasts, are released with a five-year lag.

[5]One quote from former FRB-Cleveland President Loretta Mester in October 2019 speaks to this challenge: "*But even with clear communications, the inference problem likely will always be an issue, regardless of the particular form the tools take. In particular, it's very hard to communicate the message that we're invoking nontraditional tools because the economy is slow and needs more accommodation, but taking these actions means that 'The future will be bright, so you should spend more today.' Most people are going to stop at 'The economy is slowing and needs more accommodation.' So they are going to hunker down, and they initially may spend less, not more.*"

or undershooting the target at the same meeting, then this corresponds to the bank considering a large span of scenarios. Reputation is instead more strongly associated with the number of alternatives. Adjusting the count of alternative statements for a given span of scenarios is a way for the central bank to dial-in the precision of each statement in light of the misunderstanding friction. In practice, this looks like the Fed entertaining slight variations in wording – with each set of wordings counting as a new candidate announcement – to consider the possible ways the public would respond to each candidate announcement.

To test the model's predictions in the data, we quantify reputation and confidence using the Baker et al. (2016) approach: counting the frequency of uncertainty words. For reputation, we invert the Baker et al. (2016) index of monetary policy uncertainty based on newspapers because we assume that increased uncertainty about monetary policy implies a decrease in the public's view of how confident the central bank is; in other words, a decrease in reputation. For confidence, we use the same methodology of uncertainty word counts, but instead of newspapers, we apply it to the transcripts of the FOMC meetings. This uncertainty count is similarly is inverted so that FOMC meetings with less uncertainty discussions represent meetings where the FOMC is more confident. We validate our newly constructed measure of confidence by showing it accounts for 90 percent of the variation in FOMC member uncertainty responses in the Summary of Economic Projections.

Nevertheless, the novelty of our paper on the empirical front is the measurement of the Fed's message space. For this, we use the *alternative FOMC statements*, which are multiple draft statement versions that form part of the internal FOMC meeting materials, available since 2005 in the Bluebooks and Tealbooks.[6] For each policy meeting, we count the number of alternative statements drafted, obtaining a data analogue of the model's notion of the count of alternatives. Then, using a state-of-the-art BERT large language model approach, we also quantify the distance between the most different alternative statements at each policy meeting, obtaining a data counterpart of the model's span.

We then estimate a baseline OLS regression and conduct an extensive list of robustness checks and alternate specifications. We find the model's predictions are largely borne

[6] Doh et al. (2026) and Handlan (2022) use the alternative statements to aid in identification of monetary policy shocks. While they also use the variation in alternatives, we are the first to use them as an empirical representation of a time-varying message space.

out by the data. Higher confidence and lower reputation are associated with more precise announcements. Confidence is more associated with the span of the alternatives, while reputation more with the count. However, the reputation-count connection is weaker in the data than predicted by the model. While we do not interpret our results as causal, we see the fact that our alternative specifications and robustness checks provide the same broad conclusions as underscoring the robustness of our finding: the Fed drafts their alternative statements taking reputation and confidence into account.

The paper proceeds as follows. First, in Section 2, we provide institutional context for Fed communication. Section 3 lays out the communication game setup. Section 4 analyzes the equilibrium message space, and Section 6 the time series generated by simulating the model. Section 7 introduces our text-based measures and includes the empirical analysis. Section 8 concludes.

1.1 Related Literature

Our paper connects both to the macro theory and to the empirical literature on central bank communication. The theory literature explores the effects of central bank communication policies on macroeconomic expectations and outcomes (Granziera et al., 2025), the interaction between the central bank's and the market's expectations (Caballero and Simsek, 2022; Lorenzoni, 2010; Sastry, 2026) and normative implications for central bank communication (Caballero and Simsek, 2025; Gáti, 2023; Iovino et al., 2022; Ko, 2023; Polis et al., 2025). A line of related papers explore cheap talk games building on the Crawford and Sobel (1982) result that misalignment in preferences leads a sender to send noisy signals to the receiver in the form of a partitioned message space (Amador and Weill, 2010; Bassetto, 2019; Cukierman and Meltzer, 1986; Frankel and Kartik, 2018; Li and Madarász, 2008). This also connects our paper to the classical literature on monetary policy when preferences are misaligned (Barro and Gordon, 1983; Clarida et al., 1999; Rogoff, 1985), with the important distinction that misalignment in preferences is not the key driver of our results. From within these literatures, our paper is particularly related to Moscarini (2007), which forms the point of departure for our model framework.

Our work also relates to papers on reputation-building about preferences in incomplete in-

formation settings (Amador and Phelan, 2021; Kreps and Wilson, 1982; Mailath and Samuelson, 2001; Sobel, 1985), and about the quality of information or ability to process information (Camous and Matveev, 2025; Guembel and Rossetto, 2009; Ottaviani and Sørensen, 2006). Our setup differs from these papers in that our sender (the central bank) does not have an explicit concern for reputation a priori, nor an incentive to trick the public. Reputation-building in our model is a separate, Bayesian parameter learning process (Baley and Veldkamp, 2022; Ghofrani, 2023; King and Lu, 2025) because the central bank does not internalize their effect on tomorrow's reputation.

Our empirical work relates to a growing literature seeking to measure central bank communication and its effects. In particular, papers that use text analysis to measure variation in monetary policy announcements are related. Gáti and Handlan (2025) show that the language in monetary policy announcements has a systematic relationship with the Fed's forecasts of macroeconomic variables. Recent work, like Ahrens et al. (2025); Alexopoulos et al. (2024); Aruoba and Drechsel (2026); Bholat et al. (2015); Cieslak and McMahon (2023); Doh et al. (2026); Granziera et al. (2025); Haldane and McMahon (2018); Handlan (2020, 2022); Hansen et al. (2018); Shapiro and Wilson (2022), use computational text analysis of the Fed's announcements and meeting materials to better understand monetary policy effects via policy announcements. In particular, Doh et al. (2026) and Handlan (2022) use alternative statements that the Fed considers at their policy meetings. Our innovation is to use computational text analysis to discipline a structural model to better understand the strategic pressures on the Fed's communication decision.

2 Background on Federal Reserve Communication

Over the past decades, the Fed has increased the mediums in which they systematically communicate their policy and expectations. In this paper we focus on the post-FOMC-meeting statement as the communication strategy to study for two reasons: first, it is most similar to other policy instruments – as it is voted on by the FOMC in the same fashion as how they vote on setting a new target interest rate – and thus allows us to treat the FOMC as a single decision-maker. Second, there are pre-drafted versions of this statement that the

FOMC chooses from at their policy meeting which will act as the data counterpart to the endogenous message space in our theoretical model.

The FOMC statements convey both the current policy rate and forward guidance, i.e., signals about future policy. Such guidance may be explicit or implicit (e.g., stressing inflation risks suggests future rate hikes). Gáti and Handlan (2025) show that most variation in language reflects the Fed's information set, reinforcing the view that statements are important signals of future policy.

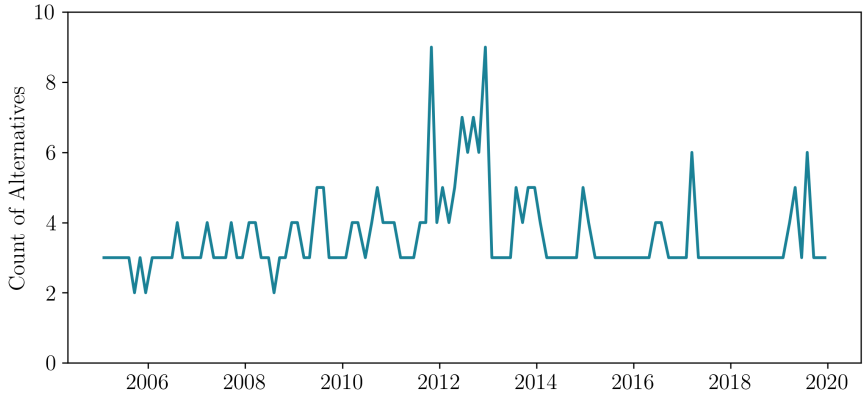
The modeling of this paper is based on the process of how the Fed arrives at the FOMC statement that is released to the public. The Fed staff prepare a set of draft statements, which are revised with input from Governors and regional banks, and finally presented as a set of alternative statements to the FOMC in the policy meeting. These so-called *alternative FOMC statements* are ranked in order of dovishness/hawkishness, and given letter names going up alphabetically from the most dovish to the most hawkish alternative, with "Alternative B" being the baseline alternative statement. At the meeting, FOMC members vote on the statement language usually by selecting one of the statement drafts from the set. Almost always Alternative B is selected.[7] The selected alternative is then released to the public, while the unselected alternatives are kept as internal Fed materials, and only released to the public with a five-year lag.

We now zoom in on the alternative statements themselves, which we retrieve from the FOMC policy meeting materials (Tealbooks/Bluebooks) where they are available since 2005. Over time, the number of alternative statements presented at FOMC meetings has varied considerably. These alternatives differ in language and emphasis, thereby conveying distinct signals about the policy path. For some meetings, the alternatives also embed different rate decisions, though this is not always the case. By construction, every newly worded draft reflects a message regarded by staff or Committee members as substantively different from the others, because if the language was not seen as substantively different then it would not be included.

In Figure 1, we graph the number of alternative statements from 2005 through 2019. We

[7]Sometimes new phrases are added to the statement that were not in the alternative draft, but this is rare.

Figure 1: Alternative Statement Count Over Time



Note: Fed staff members produce the alternative versions of the FOMC post-meeting statement to be included in the FOMC meeting materials. Each set of proposed wordings for the alternative statements is counted as a new alternative.

count each distinct phrasing as a new alternative message to be considered for release.[8] For most meetings, there are three alternatives: Alternative A, Alternative B, and Alternative C. However, we observe between two and nine alternatives over the sample. As noted above, the range of alternative statements is on a spectrum from dovishness to hawkishness. Alternative A is considered more "dovish," concerned with supporting economic growth and employment, and the last letter alternative, usually C or D, is considered more "hawkish," concerned with preventing high inflation. The baseline alternative, Alternative B, which is also often called the "likely" alternative, lies in between the extremes with respect to the policy decision and forward guidance.

Furthermore, the intensive margin across these alternatives also varies. Sometimes the alternative statements are extremely similar, only changing one or two adjectives, and other times they paint entirely different pictures of the economic outlook and forward guidance. In this paper, we model the drafting of alternative statements as the process of how the central bank arrives at the single released policy statement in order to study the central bank's forward-guidance communication decision in the presence of uncertainty.

[8]Sometimes there are multiple language options provided for each alternative. This shows up explicitly with sub-labels – "Alternative A1" vs "Alternative A2" – or more subtly – under "Alternative A" the text includes "... [phrase 1]/[phrase 2]..." – and we include both types in our count of distinct statements.

3 The Communication Game Setup

This section details the model setup. First, we describe the environment and information structure of the game. Then we detail the players' actions and the timing. The communication block of our model builds on the endogenous message space framework of Moscarini (2007) to match the institutional reality at the Fed, where the FOMC drafts multiple potential statements and then releases only one. To isolate the role of this richer communication structure, we keep the rest of the model deliberately simple. For expositional clarity, we first present equilibrium strategies conditional on the central bank announcement, and then characterize the communication stage at the end of this section.

There are two players, a central bank (CB) and a public (P). The economic environment consists of a Phillips curve that relates output (y) to inflation (π) and to the public's inflation expectations (x) as

$$y_t = s(\pi_t - x_t), \tag{1}$$

where $s > 0$ is the slope of the Phillips curve, and $x_t \equiv \mathbb{E}_t^P[\pi_t]$ denotes the inflation expectations of the public. The superscript $i = \{CB, P\}$ on the expectation operator indicates whose information set, $\mathcal{I}_t^i$, the expectations are conditioned on. In other words, $\mathbb{E}_t^i[\cdot] \equiv \mathbb{E}[\cdot|\mathcal{I}_t^i]$. The only shock in this simple environment is an inflation-target shock, which represents policy-relevant shocks in reduced form:[9]

$$\omega_t \sim \mathcal{N}(0, 1). \tag{2}$$

The information structure is as follows. The shock ω_t is not revealed to the players, but the central bank observes a noisy signal of it

$$\theta_t = \omega_t + \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, \sigma^2_{\varepsilon,t-1}), \qquad \sigma^2_{\varepsilon,t-1} \in [0, \infty). \tag{3}$$

[9]This could be a demand shock or preference shock. One could, for example, have a structural demand shock by adding a demand block to the model. For a preference shock, one could model a rotating committee wherein differences in individual preferences aggregate to a committee-level shock in the policy vote. These are examples of extensions that would add expositional complexity without altering the intuition.

Here the CB's signal precision is

$$H_{t-1} \equiv (1 + \sigma^2_{\varepsilon,t-1})^{-1} \in (0, 1], \tag{4}$$

which we refer to as the central bank's *confidence*. This captures the idea that monetary policy is generally conducted in an environment of uncertainty, where policymakers try to form an idea of the shocks hitting the economy in real time. Thus one could think of θ_t as the central bank's assessment of the economy formed during the policy meeting at time t, which combines the incoming data with the CB's judgment. This also implies the intuitive interpretation of H_{t-1} as the CB's confidence in their assessment. The wording reflects that confidence determines the width of the CB's confidence intervals around their point forecasts. We index confidence relevant at time t (for meeting t) with $t-1$ to reflect the idea that the central bank's confidence when going into a policy meeting is given by their historical forecasting performance leading up to period t.

Note that a different but complementary interpretation of confidence is that the "confidence interval" it represents actually stands for the diverging views of different FOMC members. In other words, a low degree of confidence can be interpreted as a high degree of disagreement within the committee. For simplicity, we model the central bank as a single decision-maker, and therefore we talk about the confidence of a single "central bank" for most of the paper, but it is important to keep in mind that this modeling choice maps neatly to one with a committee as the decision-maker, with individuals having diverging views. Thus we view the interpretation of confidence as the extent of agreement within the committee as both complementary to our baseline interpretation as well as realistic, and we return to it when discussing our results.

To get some intuition on the CB's confidence, H_{t-1}, note that the information structure implies that the conditional distribution of the shock given the CB's signal is

$$\omega_t|\theta_t \sim \mathcal{N}(H_{t-1}\theta_t, 1 - H_{t-1}). \tag{5}$$

This is saying that, given the signal θ_t, the CB's point forecast of ω_t is $H_{t-1}\theta_t$, with a variance of $1-H_{t-1}$. Thus, if the signal precision is perfect ($H_{t-1} \to 1$), the CB is completely confident

in their assessment of the economy, as seen by the fact that the confidence interval around the point forecast goes to zero. Thus the CB relies heavily on their signal. When the signal is infinitely noisy ($H_{t-1} \to 0$), by contrast, the CB is exceedingly uncertain about the shock after observing the signal, the confidence interval becomes infinitely wide, and so the CB discounts the signal strongly. Thus the CB's confidence is the inverse of the uncertainty surrounding their economic assessment.

To capture the idea of time-varying uncertainty, we assume that with a Poisson arrival rate η, a new H_t gets drawn from a Gamma-distribution with parameters a and d:

$$H_t = \begin{cases} \Gamma(a,d) & \text{if redrawn,} \\ H_{t-1} & \text{otherwise.} \end{cases} \tag{6}$$

Beyond considering exogenous changes to the precision of the CB's signal, we provide an extension in Appendix B where we microfound a time-varying confidence with the economic fundamental, ω_t, having time-varying variance, so that there are times in which the underlying economic conditions are more uncertain, leading to lower CB confidence. We show that depending on the informational assumptions about the volatility of ω_t, the mechanisms of the baseline setup are either exactly the same or qualitatively the same after accounting for a wedge between the CB's communication and the public's inference.

The novelty and contribution of our theory is to assume that the public does not observe confidence. This allows us to introduce our notion of *reputation for confidence*: the public's belief over confidence, which we denote by $\bar{H}_t \equiv \mathbb{E}_t^P[H_t]$. One interpretation of reputation is the public's perception of the central bank's confidence. To borrow Caballero and Simsek (2022)'s terminology, reputation reflects the market's opinion of the the central bank's uncertainty. We call this "reputation" to emphasize that the market can form a favorable opinion of the CB's confidence over time by observing a historically good forecasting performance. The key question we will ask is how the evolution of reputation affects central bank communication.

We assume that reputation evolves through Bayesian parameter learning (Baley and Veldkamp, 2022; Ghofrani, 2023). Note that, from the perspective of the public, the CB's

signal is normally distributed, $\theta_t \sim \mathcal{N}(0, H_{t-1}^{-1})$, so that our notion of reputation corresponds to learning about true confidence, H_{t-1}, the second moment of the signal distribution. We postulate that the public knows that confidence follows a Gamma-distribution, but entertains (potentially) different parameters for it: $H_{t-1} \sim \Gamma(\alpha, \beta)$.[10] Reputation then evolves as P updates beliefs on H_t according to the standard Bayesian parameter learning equations as

$$\bar{H}_t = \frac{\alpha_t}{\beta_t}, \tag{7}$$

where

$$\alpha_t = \alpha_{t-1} + \frac{1}{2}, \tag{8}$$

$$\beta_t = \beta_{t-1} + \frac{(k_t - \mu_k)^2}{2}. \tag{9}$$

Here k_t is the observable the public learns from, and μ_k is its average. As we will see later, the public will observe the CB's forecast of ω_t, $H_{t-1}\theta_t$, at the end of period t, when the learning takes place. But because the public cannot separately observe H_{t-1} and θ_t, the observable that the public learns from is the public's best guess of θ_t, which is formed by cleaning the current observation of $H_{t-1}\theta_t$ by the current best guess for H_{t-1}, reputation. Thus $k_t = \mathbb{E}_t^P[\theta_t | H_{t-1}\theta_t] = H_{t-1}\theta_t / \bar{H}_{t-1}$, and $\mu_k = 0$ (since θ_t is mean zero).

The CB has two policy tools: setting inflation, π_t, and making a public announcement, A_t. We model conventional monetary policy as simply setting inflation directly in order to focus our attention on the intricacies of the communication setup. Turning to the communication policy, note that in the model, as in the data, the central bank will draft a set of alternative statements, from which they will pick one to release to the public. Also as in the data, the public will not see the alternative statements, only the single one that is released.

We refer to the set of candidate statements as the *message space*, and denote it by $\mathcal{M}_t$. The message space will be a key object in our analysis, as it allows us to make statements about the whole set of alternative statements the CB drafts, not just the single announcement

[10] Because the public's prior has the correct functional form and the learning parameters can converge to the truth, the public's beliefs are correctly specified. The belief formation of the public is thus only subject to information frictions, but not to behavioral biases.

that is released to the public. As foreshadowed in the beginning of the section, we postpone the discussion of the message space until the end of the section in order to first give a bird's eye view of the full model. For the model overview, it is sufficient to note that the single announcement the CB releases is about the CB's signal and comes from the set of alternative statements, $A_t(\theta_t) \in \mathcal{M}_t$, in a way that we will describe below. For now, it is sufficient to think of the announcement as providing some information about the CB's signal.

Note that we allow the CB to lie, so that the announcement can contain false information about the signal.[11] However, as spelled out in more detail below, we assume that with some probability γ the public can verify whether the CB told the truth. When the public catches the CB in a lie, they play a grim trigger strategy in which they never believe any message from the CB again. The possibility of the grim trigger results in it being optimal for the CB to always send truthtelling messages.[12]

Conditional on the announcement and reputation, the public forms (rational) expectations of the signal the CB saw

$$\bar{\theta}_t \equiv \mathbb{E}_t^P[\theta_t | A_t, \bar{H}_{t-1}], \tag{10}$$

which they use to form (rational) inflation expectations $x_t(\bar{\theta}_t(A_t, \bar{H}_{t-1}))$.[13] Note that in any given period, the central bank can affect the private sector's inflation expectations by making an announcement because, due to the truthtelling result, any announcement will contribute weakly useful information on π_t via θ_t.

Given inflation expectations $x_t(\bar{\theta}_t(A_t, \bar{H}_{t-1}))$, the central bank sets inflation so that inflation and output are realized. The public can then use the Phillips curve and knowledge of the equilibrium relations to back out the central bank's forecast of ω_t, $H_{t-1}\theta_t$. This brings us

[11] Throughout the paper, we use "truthtelling" to mean that the announcement does not contain false information. But, importantly, truthtelling does not imply full revelation of the central bank's information.

[12] Consider the off-equilibrium behavior where the central bank lies. Until the audit reveals the deception, the public believes the CB's announcement and updates their inflation expectations accordingly. Reputation updates based on the discrepancy between announcement-induced expectations versus actual policy outcomes. Therefore, if the CB lies substantially, this creates a larger gap between the public's expected inflation and realized inflation, which decreases reputation.

[13] In Appendix A.6, we explore how much the public can learn about confidence from the announcement in a "pre-screening" (Cheng and Hsiaw, 2022). We show that A_t contains limited information because it only allows set identification of H_{t-1}. Then we demonstrate that in the majority of cases (at least 70 percent), the identified set contains the public's prior. Thus, assuming that the public only learns about θ_t from A_t is not a restrictive assumption.

to the last within-period stage of the game where reputation is updated. As foreshadowed earlier, we assume that with some probability $\gamma \in [0,1]$, the central bank's signal, θ_t, is revealed to the public. In this case, if the public catches the central bank in a lie (meaning A_t was inconsistent with the revealed θ_t), the public reverts to never believing the CB again (reputation goes to 0 forever). Otherwise, $\bar{H}_t$ updates to H_{t-1}, which the public can now back out from observing both the composite $H_{t-1}\theta_t$ and θ_t separately. If θ_t is not revealed, the public still backs out the central bank's forecast of ω_t, $H_{t-1}\theta_t$. In this case the public uses this information to update reputation through Bayesian learning, Equation 7, using $k_t = H_{t-1}\theta_t/\bar{H}_{t-1}$ as their best guess of θ_t.

The assumption that θ_t is revealed with probability γ is important because it ensures that there are states of the world in which the public can verify whether the central bank tells the truth. The threat of the grim trigger, in turn, keeps the central bank from lying, allowing us to focus on truthtelling equilibria. For the rest of the paper, we will therefore let γ be infinitely small but nonzero, and restrict the analysis to equilibria where the CB never lies and θ_t is never actually revealed.

The public's loss is a quadratic in current inflation, so that the expected loss is:

$$\mathcal{L}^P(x_t, \bar{H}_{t-1}, A_t) = \mathbb{E}_t^P[L^P(x_t, \pi_t)|\bar{H}_{t-1}, A_t] = \mathbb{E}_t^P[(x_t - \pi_t)^2], \tag{11}$$

while the period expected loss of the central bank is given by:

$$\begin{aligned}\mathcal{L}^{CB}(A_t, \pi_t, \theta_t, H_{t-1}, \bar{H}_{t-1}) =& \mathbb{E}^{CB}[L^{CB}(y_t, \pi_t, \omega_t)|\theta_t, H_{t-1}, \bar{H}_{t-1}] \\ =& \mathbb{E}^{CB}[(y_t - b)^2 + \lambda(\pi_t - \pi^* - \omega_t)^2].\end{aligned} \tag{12}$$

Here $\pi^* > 0$ is the inflation target, $\lambda > 0$ is the weight on inflation stabilization in the loss function and $b \geq 0$ is an output bias that introduces an incentive for the CB to induce surprise inflation.[14]

Before finally turning to the details of the communication block of the model, we conclude

[14] A large literature in the wake of Crawford and Sobel (1982) and Barro and Gordon (1983) has focused on misaligned preferences between the CB and the public as a mechanism that leads to coarse communication in equilibrium. Here, instead, the mere presence of evolving reputation will enable coarse equilibrium communication, and we can set $b = 0$.

the bird's eye view of the model by describing the timing of the game. In the beginning of period t, reputation is inherited from period $t-1$, while with probability η, confidence H_{t-1} is redrawn from $\Gamma(a,b)$, and with probability $1-\eta$ it stays the same as in last period. Given confidence H_{t-1} and reputation $\bar{H}_{t-1}$, the message space is created, i.e. the CB drafts the alternative statements. The shock ω_t is drawn, and the CB gets a noisy signal of it, θ_t. Given their measurement of the state of the economy (θ_t), the CB picks an equilibrium announcement A_t about θ_t from the set of alternatives, and releases the announcement to the public, keeping the unreleased alternatives confidential, so the public never gets to see them. The public uses the released announcement to form their inflation expectations x_t, taking into account the CB's reputation for confidence in making inferences from the announcement. Then the CB chooses inflation, and thus output y_t is realized through the Phillips curve. Given the realizations of the observables, the public backs out the CB's forecast of ω_t, $H_{t-1}\theta_t$. With probability γ, θ_t is revealed, so the public verifies whether the CB told the truth or not. If yes, $\bar{H}_t = H_{t-1}$, otherwise the public reverts to never believing any announcement from the CB for the rest of the game. With probability $1-\gamma$, θ_t is not revealed, and the public uses the backed-out forecast $H_{t-1}\theta_t$ to update reputation.

Now we turn to the details of the CB's communication policy. Having seen the full game in terms of loss functions and actions of the players, it is clear that the CB wants to send an announcement that induces an inflation expectation of the public, x_t, that allows the CB to pick inflation in a way to minimize its loss in Equation 12 through the expectation's effect in the Phillips curve.

We make the following assumptions on the announcement $A_t(\theta_t)$. First, it needs to be about θ_t. This assumption is motivated by the interpretation of θ_t as representing the CB's assessment of the economy formed during the policy meeting at time t by combining the available economic data at that time with the judgment of the CB. Thus, since FOMC statements are issued about the decisions and economic assessments of a particular policy meeting, our model needs to capture the announcement being about the decisions and assessments of that meeting.[15]

[15]For clarity, we focus on the case of a single message about θ_t. With an additional message about confidence (H_{t-1}), the model becomes more involved because since the CB has an incentive to misrepresent H_{t-1}, any perfectly revealing announcement on H_{t-1} would be disregarded by the public. Hence, for an-

Second, while we do not rule out the CB providing false information on θ_t, lying behavior does not occur in the equilibria we focus on because of the threat of the grim trigger that is invoked if the CB is caught in a lie. Thus, in the equilibria we consider, it will be true that $\theta_t \in A_t(\theta_t)$. Third, the announcement has to belong to the set of real numbers, but we make no assumption on the properties of the set. Restricting to the reals is just a technical assumption that is motivated both by the just mentioned result that $\theta_t \in A_t(\theta_t)$ and the fact that θ_t itself is on the real line. Thus it simplifies the exposition not to deal with imaginary numbers, just like the FOMC statements do not deal with them.

To make these assumptions concrete, suppose that $\theta_t = 3.1$. Our assumptions on $A_t(\theta_t)$ make many announcements possible. For example, the point $A_t = 3.1$ is a permissible announcement. So is the closed interval $A_t = [0, 5]$, or the open and non-convex set $A_t = 2 \bigcup (2.9, 4)$, or even the full real line $A_t = \mathbb{R}$. However, $A_t = a + bi,\ a, b \in \mathbb{R},\ i^2 = -1$, is not a permissible message. Note that the messages $A_t = 3$, $A_t = [4, 5]$ or $A_t = \mathbb{R}^-$ are all permissible messages, but will not occur in the equilibria that we consider because they do not satisfy the result that $\theta_t \in A_t(\theta_t)$. In other words, the latter example announcements provide incorrect information on θ_t, and as such do not constitute truthtelling behavior.

Above, we defined the message space $\mathcal{M}_t$ loosely as the set of candidate messages. Now, having specified the assumptions on the announcement, we can make that more concrete. Since the announcement maps a particular θ_t to a set on the real line, we can define the message space as the set of all announcements $A_t(\theta_t)$ **for all** $\theta_t \in \mathbb{R}$. Thus, the message space $\mathcal{M}_t$ is either $\mathbb{R}$ itself, or a partition of $\mathbb{R}$.

Crucially, we will allow the message space to be an equilibrium object. More precisely, our equilibrium definition will involve a construction of the message space, given confidence and reputation, that minimizes the central bank's loss (Equation 12). This allows us to interpret the message space as the communication policy of the central bank: for the current levels of confidence and reputation, and for any possible outcome of the upcoming policy meeting (θ_t), what would be the optimal the FOMC statement? It is in this sense that the message space can be interpreted as the set of alternative statements that the Fed staff draft

nouncements on H_{t-1} to matter in equilibrium, they would have to be coarse. That coarseness would still induce a wedge between H_{t-1} and $\bar{H}_{t-1}$ and thus the mechanism of the θ_t-message that we focus on would still go through. Thus we opt for the simpler framework.

before the policy meeting, each one drafted to be optimal for a specific potential outcome of the policy meeting. As we will see in the next section, the fact that the message space is an equilibrium outcome will result in a lot of structure on $\mathcal{M}_t$ that will be useful for our analysis.

As a last technical note on the message space: because the public does not see the message space, we need to place some assumptions on how they can conceive of announcements as a basis for forming expectations. For the purpose of this discussion, we will denote the message space for a given confidence and reputation as $\mathcal{M}_{H,\bar{H}}$, noting that it is not necessarily unique. We assume the public knows the message space will be supported by $\mathbb{R}$, and thus, because they do not know H_{t-1} they will entertain $\mathcal{M}_{\bar{H}}$, the combined set of possible messages given reputation:

$$\mathcal{M}_{\bar{H}} \equiv \underset{H}{\cup} \mathcal{M}_{H,\bar{H}} \subseteq \mathbb{R}. \tag{13}$$

The CB sees the precise $\mathcal{M}_{H,\bar{H}}$ and selects an announcement $A(\theta) \in \mathcal{M}_{H,\bar{H}} \subseteq \mathbb{R}$. If an announcement is chosen outside of $\mathcal{M}_{\bar{H}}$, the public will automatically interpret it as the closest message within the set $\mathcal{M}_{\bar{H}}$.[16] It is important to note here that because the public only sees one message from the space, they are limited in what they can infer indirectly from communication. In particular, any given message interval A_t could be part of multiple $\mathcal{M}_{H,\bar{H}}$, so the public will not be able to perfectly back out H_{t-1} from a single message. For simplicity of notation, for the rest of the paper, we will use the notation $\mathcal{M}_t$ for the message space at time t as a function of H_{t-1} and $\bar{H}_{t-1}$.

We proceed by analyzing the communication problem in two stages. First, we characterize the equilibrium message space, and investigate how an exogenous reputation affects communication in one period of the game. As above, we zoom in on period t, taking the prior on reputation, $\bar{H}_{t-1}$, as given. We then consider the time series the model generates when in each period, given evolving reputation and confidence, a new equilibrium message space arises.

For tractability we do not allow the CB to anticipate how their announcements affect their reputation in future periods, beyond communicating in a way to avoid the grim trigger.

[16] An effectively equivalent assumption would be to assume that the public interprets any message outside of $\mathcal{M}_{\bar{H}}$ as babbling and thus ignores it.

Thus our game belongs to the cheap talk tradition, where reputation (the public's perception of CB confidence) constrains the central bank only by affecting announcement effectiveness each period. The CB does not internalize how their current actions influence the constraints they will face in the future. We justify this assumption by noting that it has no bearing on the stage game results, which is the main focus of the paper. The main role for the repeated game is in disciplining the CB's messaging decisions to reflect the type of coarse truthtelling we see in the time series data. We leave the full dynamic game, which would allow for reputation-building strategies, for future work.

4 Equilibrium Communication Policy

Our equilibrium concept is a Perfect Bayesian Equilibrium with Communication (PBEC), as spelled out below. Because of the threat of the grim trigger in the case where θ_t is revealed and the CB is caught in a lie, we focus on PBECs in which truthtelling always happens, and therefore the public always believes the central bank's announcements. This behavior is summarized as the incentive compatibility constraint.

Definition 1 *(Perfect Bayesian Equilibrium with Communication (PBEC))*
Given reputation and confidence, a PBEC is π_t, x_t a message space $\mathcal{M}_t$, and an announcement A_t such that

- $\pi_t = \textit{argmin}\ \mathcal{L}^{CB}\left(\pi_t, x_t(A_t, \bar{H}_{t-1}), \theta_t, H_{t-1}\right)$,
- $x_t = \textit{argmin}\ \mathcal{L}^{P}(x_t, \bar{H}_{t-1}, A_t)$,
- $\mathcal{M}_t$ *is* $\mathbb{R}$ *or a partition of* $\mathbb{R}$,
- $\mathcal{M}_t$ *induces CB to make an announcement* $A_t \in \mathcal{M}_t$ *such that* $\theta_t \in A_t$,
  - *P believes the announcement so* $\bar{\theta}_t = \mathbb{E}^P[\theta_t | \theta_t \in A_t, \bar{H}_{t-1}]$, *and*
  - *CB prefers* A_t *that induces* $\bar{\theta}_t$ *to any alternative* A'_t *that induces* $\bar{\theta}'_t$, *such that*

$$\mathcal{L}^{CB}\left(\cdot, \bar{\theta}_t(A_t)\right) \leq \mathcal{L}^{CB}\left(\cdot, \bar{\theta}'_t(A'_t)\right) \qquad \textit{(Incentive Compatibility)}$$

Because we do not allow the CB to internalize the effect of their actions on reputation, from the perspective of the CB, there are no endogenous state variables here. Therefore

both players are solving a static problem, and we can solve the stage game backwards by first computing the best-response inflation rate that minimizes the expected period loss of Equation 12

$$\pi_t^{BR} = \frac{1}{s^2+\lambda}(s^2 x_t + \lambda\pi^* + \lambda H_{t-1}\theta_t + sb). \tag{14}$$

Postulating rational expectations on the part of the public and conditioning on a central bank announcement $A_t(\theta_t)$, this gives us the best-response inflation expectations[17]

$$x_t^{BR} = \pi^* + \bar{H}_{t-1}\bar{\theta}_t + \frac{sb}{\lambda}. \tag{15}$$

Given the two best-responses, equilibrium inflation is

$$\pi_t^{eqb} = \pi^* + \frac{sb}{\lambda} + \frac{s^2\bar{H}_{t-1}\bar{\theta}_t + \lambda H_{t-1}\theta_t}{s^2+\lambda}, \tag{16}$$

and equilibrium output is given by the Phillips curve in Equation 1.

Before deriving the equilibrium message space, we establish the following lemmas.

Lemma 1 *(Message Space as a Sequence of Non-Overlapping Intervals)*
Any equilibrium message space $\mathcal{M}_t$ is either $\mathbb{R}$ or a partition of $\mathbb{R}$ with non-overlapping intervals, such that all candidate statements are intervals.
Proof: See Appendix A.2.

The interpretation of this lemma is that optimal communication policies will partition the real line into distinct intervals. Hence, for any realization of θ_t, the chosen announcement will be of the form "we saw a number between θ^k and θ^{k+1}", where the interval borders $\theta^k, \theta^{k+1} \in \mathbb{R}$ form an increasing sequence.

This key property of equilibrium communication arises because the equilibrium definition requires all candidate announcements to be optimal for a given θ_t. Since announcements are about the random variable θ_t, which in the public's perception is distributed according to $\mathcal{N}(0, \bar{H}_{t-1})$, the public's expectation of θ_t given an announcement A_t, $\bar{\theta}_t(A_t)$, is the integral of a normal density over the announcement space. On the real line, such a space is an interval.

[17] The public's expectation of the bank's conditional forecast is $\mathbb{E}^P[H_{t-1}\theta_t] = \bar{H}_{t-1}\bar{\theta}_t$. (Proof in Appendix A.1.)

The intervals cannot have gaps (non-convexity) because those could result in conditional expectations $\bar{\theta}_t(A_t)$ falling outside the interval, violating truthtelling. Lastly, the result that the announcement intervals are non-overlapping comes from the fact that the CB's loss function is non-constant in $\bar{\theta}_t(A_t)$. Therefore indifference regions between two announcements can at most be a point, as the integral of a point has zero mass. Thus the intervals do not overlap.

Given this understanding of the lemma, we make two important observations. First, the width of announcement intervals captures whether announcements are fine or coarse because the integral of a more narrow interval is an expectation $\bar{\theta}_t$ that is closer to the actual θ_t on average than the integral of a wider interval. For example, $A_t(\theta_t) = \theta_t$ is the finest possible announcement, a situation we will refer to as *point revelation*, while $A_t(\theta_t) = \mathbb{R}$ is the coarsest.

Second, the widths of the candidate announcements in the message space will be equilibrium outcomes because the message space itself is an equilibrium outcome. Furthermore, the equilibrium alternative statements do not need to have equal lengths. There is similarly no restriction on the number of alternative statements, which will also be determined in equilibrium.

Therefore, while most papers on communication games concentrate on whether the sender tells the truth or not, given a fixed message space, our approach with an endogenous message space allows varying levels of reputation and confidence to yield message spaces with varying number and width of announcement intervals. Thus we can study how reputation and confidence affect the design of the central bank's communication policy and the precision of communication. Again, this modeling choice reflects what we see in reality: the Fed does not consider lying to the public, but rather communicating more or less precisely depending on their confidence and how they anticipate the public to react.

Lemma 2 *(Message Space Independent of Signal)*
The message space $\mathcal{M}_t$ is determined by confidence H_{t-1} and reputation $\bar{H}_{t-1}$, and the parameters of the model b, s and λ, and is independent of the signal θ_t.
Proof: See Appendix A.3.

The economic intuition behind Lemma 2 is straight-forward. Since our equilibrium concept implies that the message space specifies the optimal announcement for every possible realization of the signal, the message space is independent from what signal actually materializes. This is why the equilibrium message space is the model's notion of the set of alternative statements and at the same time also corresponds to optimal communication policy. Once the signal is realized, the CB picks the corresponding statement from the set of alternative statements, releases the single statement to the public, and keeps the rest of the alternative statements as internal materials.

Now we provide an outline of the derivation of the equilibrium message space, suppressing the details to Section 4 and Appendix A.4. We start by computing best responses for inflation expectations x_t and inflation π_t, taking communication as given, obtaining Equation 15 and Equation 14. Combining the two, we get Equation 16. Substituting π_t^{eqb} and x_t^{BR} into the period loss of Equation 12, we obtain the CB's loss function expressed only in terms of the expectations induced by the CB's announcement.

We proceed by imposing incentive compatibility from the definition of equilibrium. In other words, we require that the equilibrium message space should be such that each possible equilibrium announcement is loss-minimizing for the CB. This leads to a system of equations that characterize the relationship between the borders of each interval and the private sector expectations induced by those intervals (Equation 3 in Appendix A.4). Solving for the equilibrium message space amounts to solving this system of equations.

Concretely, if the solution splits the real line into K intervals, then it contains $K-1$ interval borders which form an increasing sequence. Solving for the message space means solving for this increasing sequence of $K-1$ cutoffs such that each interval satisfies Equation 3 in Appendix A.4.[18] We do this numerically using a shooting algorithm, described in detail in Section 4 and Appendix C. In a nutshell, the algorithm searches for the equilibrium cutoffs by zeroing out the residuals from the equation system, subject to the cutoffs being in increasing order.

It is important to note that this not only results in equilibrium intervals of varying widths, so that the width of every single interval is an equilibrium object and potentially

[18]Note that we use the terminology of "cutoffs" and "interval borders" interchangeably.

different, but also of varying number. In other words, the number of intervals, K, is also an equilibrium outcome.

Solving for the equilibrium message space, we arrive at our first result.

Proposition 1 *(Point Revelation)*
A point revealing equilibrium (where the CB sends the singleton message $\bar{\theta}_t = \theta_t$) exists if and only if $b = 0$, and $\bar{H}_{t-1} = H_{t-1}$ or $\theta_t = 0$.
Proof: See Appendix A.5.

Proposition 1 gives conditions for when the CB point reveals their signal, i.e. makes an announcement that exactly corresponds to the number θ_t the CB observed. The proposition is saying that point revelation is only possible if there is no output bias and either reputation equals true confidence, or the CB's signal corresponds to the unconditional expectation of the shock, zero.[19]

The requirement that $b = 0$ is familiar from the literature on communication games which emphasizes the importance of aligning the CB's preferences with the public in order to get rid of incentives for the CB to trick the public (Barro and Gordon, 1983; Clarida et al., 1999; Crawford and Sobel, 1982; Rogoff, 1985). When trying to rationalize coarse communication in equilibrium, papers either modify the central bank's loss function to include an explicit concern for reputation (Ottaviani and Sørensen, 2006), or set $b > 0$ as in Moscarini (2007).

Here, instead, $b = 0$ is not sufficient for point revelation. The first of the two possible additional requirements, $\theta_t = 0$, is trivial. It simply says that if $b = 0$, so that the CB has no incentive to induce an inflation surprise, and if the CB thinks that no shock is hitting the economy, then they anticipate not to have to take any inflation-stabilizing action. Point revealing the information that no policy action will be taken leaves the economy ex ante in the steady state, as private expectations remain at their steady state value. Clearly, this is not an interesting case in practice.

The crucial novel result here is the other possible requirement for point revelation, namely that $\bar{H}_{t-1} = H_{t-1}$. The interpretation is that if the CB thinks that shocks are hitting the

[19] Because of equilibrium multiplicity, the existence of point revealing equilibria does not close out the existence of other equilibria that involve coarse communication.

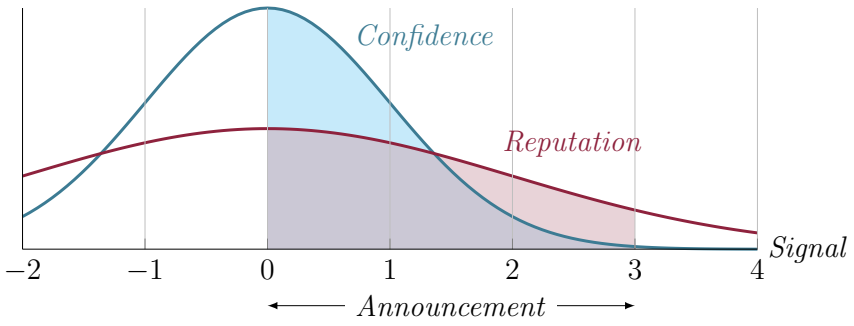
economy ($\theta_t \neq 0$), then not only does $b = 0$ need to hold, but the private sector also needs to believe that the CB is exactly as confident in the shocks as they truly are. Conversely, if the public is mistaken in their assessment of the CB's confidence, then even without an inflation surprise motive, point revelation cannot happen in equilibrium. This result helps rationalize why central banks in practice may not reveal all their information to the public even without an inflation surprise motive.

The intuition behind this result is that if the private sector misperceives the central bank's confidence, then they either overreact (if reputation is higher than actual confidence), or underreact (if reputation is lower than confidence) to a given announcement. Therefore the central bank cannot point reveal their signal, but needs to adjust their announcement to reflect the over- or underreaction of the public. Varying the width of the announcement interval is a way of doing so.

Figure 2 illustrates the key tension that confidence and reputation mean for the central bank's communication policy. The blue bell curve represents the distribution that the central bank believes the signal has (confidence), while the red bell curve is the market's belief of the θ_t-distribution (reputation). Both curves are centered at 0 because both players (correctly) think that θ_t is mean zero. The only difference between the curves is their variances: the tightness of the central bank's beliefs is governed by confidence, while that of the market by reputation.

Now suppose, for example, that the central bank announces "we saw a signal between 0 and 3," as depicted on the figure. The way the CB intends for this announcement to be understood is to integrate the blue bell curve under the interval of the announcement, because the information provided by the announcement is that a normally distributed variable with

Figure 2: Reputation Guides the Interpretation of Announcements



a particular variance (confidence) fell into the interval [0, 3]. Because the market disagrees with the signal variance, they will instead integrate the red bell curve and "misunderstand" the announcement.

This helps explain the idea in Proposition 1 that if reputation is not equal confidence, then the market will not interpret the announcement in the way the central bank meant it. Thus, optimal central bank communication needs to take into account not only the bank's own confidence in their measurement of the economy, but also how their announcements will be interpreted in light of the market's view. In this way, we can think of confidence as guiding what the central bank wants to *say*, and reputation as guiding what the bank wants markets to *understand*. This *misunderstanding friction* is the key driver of equilibrium communication in the model.

5 Model Predictions

Now we turn to exploring how the optimal communication policy handles the tension between what the central bank wants to say and what they want the markets to understand. To do this, we solve for the equilibrium message space numerically for different values of confidence and reputation.

Looking at the numerical message spaces requires more details about the numerical solution procedure, the shooting algorithm.[20] As introduced in the previous section, the idea of the algorithm is to guess a sequence of interval cutoffs, denoted $\{\theta^k\}_{k=1}^{K-1}$, for a given number of alternatives K, compute the associated conditional expectations for each alternative k, which we denote $\bar{\theta}^k$, and adjust the cutoffs to minimize errors to incentive compatibility conditions, subject to the constraint that the cutoff and conditional expectations sequences are increasing sequences. We start the algorithm with an initial guess for the equilibrium number of intervals, K_0, set $K = K_0$, and perform the search for the equilibrium cutoffs.[21] If an equilibrium is found, we stop. If no equilibrium is found, we decrease K by 1. We keep going until $K = 2$, which is the most coarse message space. If no equilibrium message space

[20] For an extensive description of the algorithm, see Appendix C.

[21] We selected $K_0 = 12$ by experimentation. One way to calibrate it would be to set it equal to the largest number of alternative FOMC statements the Fed has entertained, which is nine.

is found, we set $K = 2$, because one can show theoretically that a two-message equilibrium always exists (see Appendix A.7).[22]

There are two things to note here. First, there is a multiplicity of equilibria. For a particular set of parameter values, there are different coarse equilibrium message spaces, ranging from $K = 2$ to some initially unknown $K = K^{max} \in \mathbb{N}^+$. In what follows, from among the multiple coarse equilibrium message spaces, we will focus on the finest message space; that is, the highest finite K. Second, even when the conditions of Proposition 1 are met, and an infinitely fine equilibrium message space exists, we will still focus on the finest coarse one in order to make comparison between equilibria easier.

Table 1: Calibration

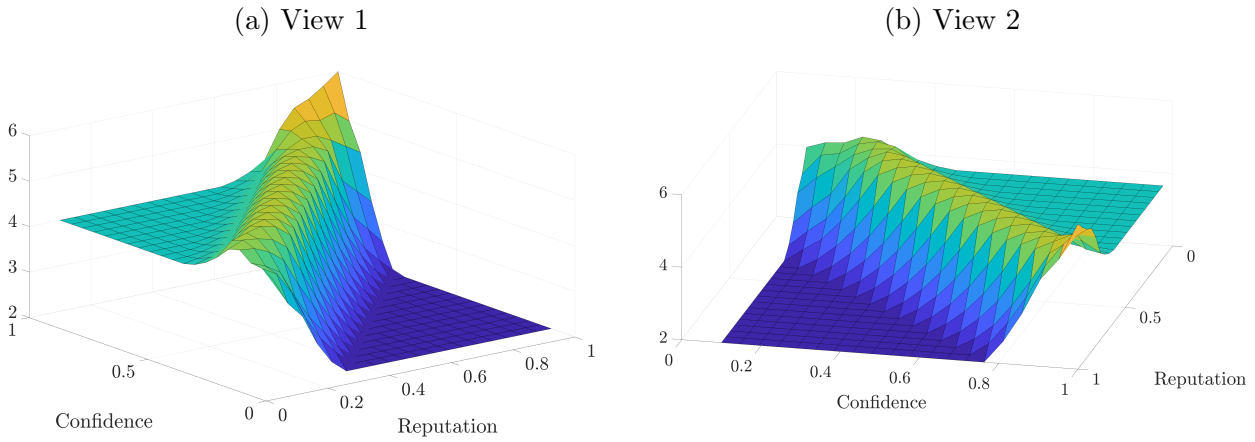
Parameter	Value	Target
π^*	2	Federal Reserve inflation target
s	12	Expected price duration of two quarters
λ	20	Rotemberg and Woodford (1997)
b	(0,0.02)	Moscarini (2007)
a	3.2	Mean H of 0.8 and variance of 0.2
d	0.25	Mean H of 0.8 and variance of 0.2

We use the calibration outlined in Table 1. π^* is set to correspond to the Fed's inflation target. λ is set to match the ratio of weights on inflation versus output stabilization. Following Rotemberg and Woodford (1997)'s estimate of a 0.05 weight on output stabilization when setting the weight on inflation to 1, this implies a 20:1 ratio. Lastly, we set s by observing that we can rewrite our Phillips curve of Equation 1 to express inflation as a function of output and inflation expectations, implying that $1/s$ corresponds to the New Keynesian slope of the Phillips curve. Setting the latter to match an expected price duration of two quarters, we get a number of 0.0850, which we invert to get 11.7647, and then round up.[23]

[22] Two is also the minimum number of alternatives we see in the Fed data.

[23] The slope of the Phillips curve in the standard New Keynesian model is $\kappa = \frac{\omega+1/\sigma}{1+\omega\theta}\frac{(1-\alpha)(1-\alpha\beta)}{\alpha}$, where the parameters $(\omega, \sigma, \theta, \beta, \alpha)$ are the elasticity of marginal cost to output (including the Frisch elasticity), the intertemporal elasticity of substitution, the price elasticity of demand, the discount factor and the probability that a firm is stuck with their price, respectively. We set α to match an expected price duration of two quarters, and the other parameters are set to the values of Chari et al. (2000), which is the calibration followed in Woodford (2003).

Figure 3: Count of Messages



Note: The figure shows a smoothed representation of the equilibrium number of alternative statements for varying levels of confidence (H_{t-1}) and reputation ($\bar{H}_{t-1}$). For other versions of these plots, see Appendix D.

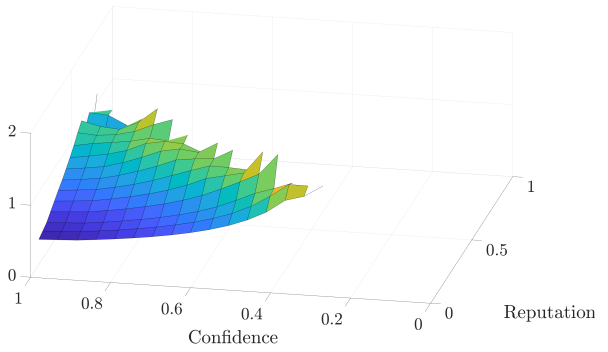
Our baseline value for the output bias b is zero in order to prevent preference misalignment from driving the results. We also include results for the case with $b = 0.02$ (Moscarini, 2007) in Appendix D. Lastly, the parameters of the Gamma-distribution that H_t is drawn from are set to induce a mean and variance for H of 0.8 and 0.2 respectively.

We start our analysis by introducing two quantitative notions of announcement precision from the model that will be useful both for characterizing the equilibrium message space and for taking the model predictions to the data. The first is the count of equilibrium intervals, K, which is the model's notion of the number of alternative FOMC statements the Fed drafts. The second is a notion of distance which we call "span of messages," and which we define as the distance between the outmost cutoffs. For example, in an equilibrium with $K = 4$ messages, there are three interior cutoffs: $\theta^1, \theta^2, \theta^3$. Then the count is $K = 4$, and the span is $\theta^3 - \theta^1$. Clearly, a large span means a wider range of possible outcomes, while a larger count given a fixed span means more narrow intervals.

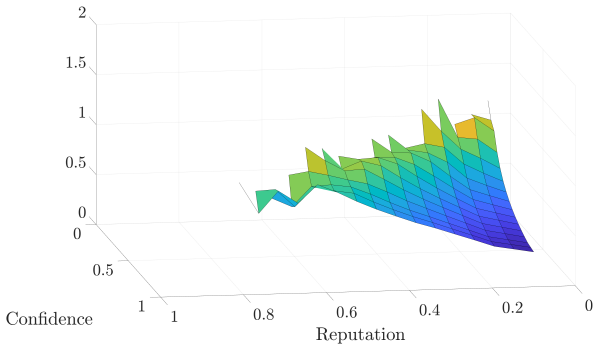
Figure 3 and Figure 4 show 3D plots comparing how the count and the span vary for different values of confidence and reputation. Higher z-axis values are indicated by lighter colors, while darker colors mean low z-axis values. Figure 4 shows an interpolation of the direct model output for the span, and Figure 3 has a smoothed representation of the model output for the count. Appendix D shows raw model outputs along with alternative specifications.

Figure 4: Span

(a) View 1



(b) View 2



Note: The figure shows the equilibrium span, defined as the distance between the two outmost cutoffs ($\theta^{K-1} - \theta^1$), for varying levels of confidence (H_{t-1}) and reputation ($\bar{H}_{t-1}$). Note that only the region in which confidence is weakly above reputation ($H_{t-1} \geq \bar{H}_{t-1}$) is shown because the span is not defined in a two-message equilibrium. The model output has been smoothed using a moving average with a window length of 5 along both dimensions. For other versions, see Appendix D.

Starting with Figure 3, we see that the variation in the count is mainly attributable to whether we are above or below the diagonal. Therefore, it is useful to define $h \equiv \bar{H}/H$ as the "reputation-over-confidence ratio." We see that when this is sufficiently above one, the CB always just drafts two alternatives, while when it is sufficiently below one, the CB drafts four. Close to the diagonal, the count increases.

Thus the reputation-over-confidence ratio splits the count into broadly three regimes. First, when reputation is much above confidence, there are only two alternatives, for example, a high-inflation and a low-inflation alternative. Second, when reputation is much below confidence, there are four alternatives, such as a very-low-inflation, a mildly-low-inflation, a mildly-high-inflation and a very-high-inflation alternative. Leaning on the intuition from Figure 2, we can interpret this in the light of the misunderstanding friction emphasized in the discussion to Proposition 1. When reputation is above confidence, the market's expectations overreact to a given announcement, which is why in this case the CB drafts fewer, more vague announcements. When instead reputation is lower than confidence, the market's expectations underreact to a given announcement, and so the CB drafts a higher number of more specific alternative statements. Third, the diagonal represents the scenario when Proposition 1 holds. Thus, as $h \to 1$, the count goes up starkly, approaching point revelation,

keeping in mind that, in these figures, exact point revelation is excluded by definition.

The model's predictions for the derivatives of the count with respect to reputation or confidence are specific to the three regimes. When h is far from 1, the derivatives are predicted to be zero, because for a very low reputation-over-confidence ratio, the equilibrium message space always contains four messages, while for a very high ratio, it always consists of just two. The derivatives are only well defined when h is close but not equal to 1. In that region, when reputation is above confidence, a higher count is associated with lower reputation and higher confidence. When reputation is below confidence, this reverses: a higher count is associated with higher reputation and lower confidence. These predictions will allow us to tease out from the data which reputation-over-confidence regime the Fed is in.

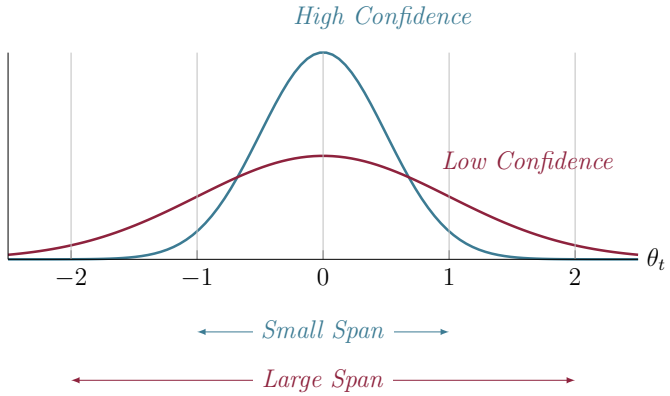
Turning to Figure 4, we immediately notice that when reputation is above confidence, the span is not defined, leading to the empty area on the figure "above" the diagonal. This is because, as can be seen on Figure 3, in the $h > 1$ region, the count is generally two, so there is only a single cutoff, and thus no span.

The model's predictions for the effect of reputation and confidence on the span are very clear in the $h < 1$ region where the span is defined. A higher confidence is associated with a lower span, while a higher reputation with a higher span. For reputation, the intuition continues to be the misunderstanding friction discussed on Figure 2. A too low reputation, which is the case here since we are in the $h < 1$ regime, means that the market does not respond to announcements as much as the CB would like them to. Therefore the CB sends more precise announcements, and a lower span allows the bank to send announcements more focused on specific ranges of outcomes, thus increasing message precision.

As for higher confidence lowering the span, Figure 5 illustrates the intuition on an example. If the CB is highly confident that the shock was, for example, between [-1, 1] (teal line), or analogously, if there is strong agreement in the committee, then the CB staff do not need to draft alternatives that discuss shock values of say -2 or 2. If the CB is not so confident (red line), or there is disagreement in the committee, then shock values of -2 or 2 are not unlikely enough for the CB not to draft up what policy would be in the event that such a shock materialized. Analogously, drafting up alternatives that concern themselves

with the extreme outcomes 2 or -2 may be a way of getting committee members with very diverging views to sign off on the bank's communication during the policy meeting. This echoes the words of Phildelphia Fed President Patrick Harker from the January 31, 2018 FOMC meeting transcript, where he suggests that with "elevated uncertainty regarding the appropriate path for the policy rate, [...] I can easily envision outcomes calling for a steeper pace of tightening. But I can also envision outcomes that would call for a prolonged pause."

Figure 5: Confidence and Span



Thus the count and span of the alternative statements are two ways to finetune message precision. For a given span, a higher count increases message precision. For a given count, a lower span increases message precision. With these insights, we can succinctly summarize the model's predictions on the effects of reputation and confidence on the equilibrium alternative statements as follows:

1. High confidence or low reputation result in more precise messages (more narrow equilibrium intervals) in the region where the span is defined.

2. Confidence primarily operates through the span, but not exclusively.

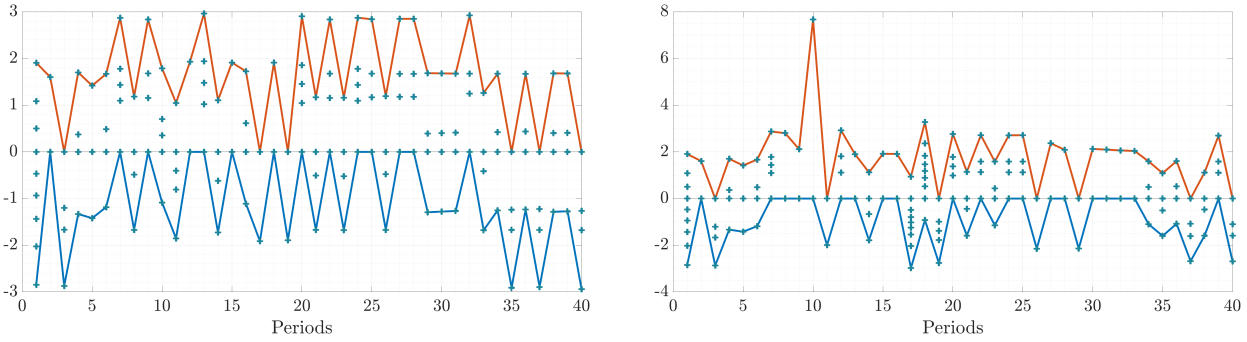
3. Reputation primarily operates through the count, but not exclusively.

6 Simulation

In this section we simulate the model to explore how the equilibrium message space evolves over time and to lead up to our regression analysis. We continue working with the calibration in Table 1, focusing on the case of zero output bias, and providing the positive output bias case in Appendix D, along with other alternative specifications. We set initial confidence to $H_0 = 0.9$, a fairly high value. We then simulate a sequence of shocks, and initializing reputation at true confidence ($\bar{H}_0 = H_0$), we solve for the equilibrium message space in every period.

Figure 6: Message Spaces Over Time

(a) H constant, $b = 0$ (b) H evolving, $b = 0$



Note: The figure shows the evolution of the equilibrium message space over time. The teal crosses designate the alternative statements in every period. The blue line connects the lowest candidate messages across the periods, and the red line connects the highest candidate messages. The alternatives are normalized so that the chosen alternative lies on the zero line.

Figure 6 shows how the message space evolves for an economy subject to the same sequence of shocks. The teal crosses represent what we have called $\bar{\theta}_t^k$, the public's expectation induced by the alternative k, the model's quantitative notion of each alternative statement. Adding up the teal crosses in each period t thus represents the count. We also use the normalization that the chosen message lies on the zero line, analogously to the baseline alternative in the data. The red and blue lines connect the highest and lowest alternative messages in every period, respectively, so that the distance between the red and blue lines in every period t is that period's span.

There are two main takeaways from Figure 6. First, as reputation (and, on the right panel, confidence) evolves, this generates a time-varying message space with a varying count of alternative statements and a varying span as well. Second, in the left panel, we hold H fixed so that $H_t = H_0 \ \forall t$. This underscores the fact that the time-varying nature of the message space does not primarily come from a time-varying H_t. The intuition behind this is that even if the CB did not have time-varying uncertainty about the economy, the fact that the market has a time-varying perception of the CB's uncertainty means that the misunderstanding friction takes a new form in every period. Therefore, the CB needs to adjust the message space in every period to correct for the time-varying misunderstanding.

We provide additional results in Appendix H, where we regress the message space measures on confidence and reputation measures from the simulated data. We find coefficients consistent with the analysis in Section 5. In the next section, we will bring the model to data on Fed communication to see how its predictions hold up empirically.

7 Empirical Analysis

In this section, we study how the FOMC's message space in the data varies with respect to measures for confidence and reputation. As explained in detail below, we use the alternative FOMC statement drafts from the Tealbooks to find data analogues to the model's notions of message space count and span: the count is the number of alternative statements and the span comes from the text-similarity of alternatives. Alongside these, we create "uncertainty" measures à la Baker et al. (2016) from newspaper articles on monetary policy and FOMC meeting transcripts which inversely approximate our model notions of reputation ($\bar{H}$) and confidence (H), respectively. Having these data counterparts, we regress our count and span measures on the measures for confidence and reputation.

7.1 Message Space in Practice

To make the idea of similarity or difference between alternative announcements more concrete, we will take the most extreme alternative statements from the 2008-12 and 2014-04 meetings as an example. The full statement texts of these examples are included in

Appendix F.

The 2008-12 meeting featured a lot of uncertainty, which is reflected in striking differences between the announcement texts. The most hawkish statement not only contradicts the most dovish statement in terms of its assessment of inflation and the real economy, but also features a different policy decision, suggesting that the target rate should be kept at one percent. In contrast, the most dovish alternative is not even satisfied with suggesting that the target rate should be lowered, but suggests that since the economy is so weak, specifying a target rate is pointless. We can also see that when the alternatives feature different policy recommendations, they also back those up with differences in the assessment of the economy. This is thus an example of alternative statements being really far in meaning, and one way to see that is by noting the strong differences in wording.

The 2014-04 alternative statements, instead, are very close together. Across the board, they share the policy decision of keeping the target rate at 0-1/4 percent. Their assessment of the economy is also very similar, with only minor differences. One of those is word order: for example, while the most dovish alternative suggests that growth slowed, but is picking up, the hawkish alternative reverts the word order to first indicate that growth is picking up, only then to admit that there are "slowing factors." The second is the choice of slightly different qualifiers when describing the same phenomenon. To go on with the growth example, where the dovish statement simply says that growth slowed, the hawkish alternative instead talks about "slowing factors," which, moreover are also described with the word "fade," which is absent from the dovish statement. The 2014-04 alternatives are thus much closer semantically than the 2008-12 ones, and it is this variation that we will exploit to quantify our model's notion of the span of messages.

We quantify this notion of semantic similarity or distance between alternative statements using a frontier language model. More precisely, we use a BERT sentence encoder which takes FOMC statement text as input, and outputs a 768-dimension vector, called an embedding, that encodes the meaning behind the joint occurrence of words.[24] Since we now have the alternative statements represented as vectors of equal length, we can then measure the

[24]We use a general BERT model to embed FOMC statements, since these are written in standard English for public consumption. Standardizing all data series removes the potential average biases in the encoding with the general language model. Appendix G provides further details on text processing and analysis.

differences between statements by measuring the differences between the statement vectors using Euclidean distance.[25] Summing these up, we get a measure of the total distance of the Fed's message space spanned at a given meeting. In other words, we have a measure of the span.

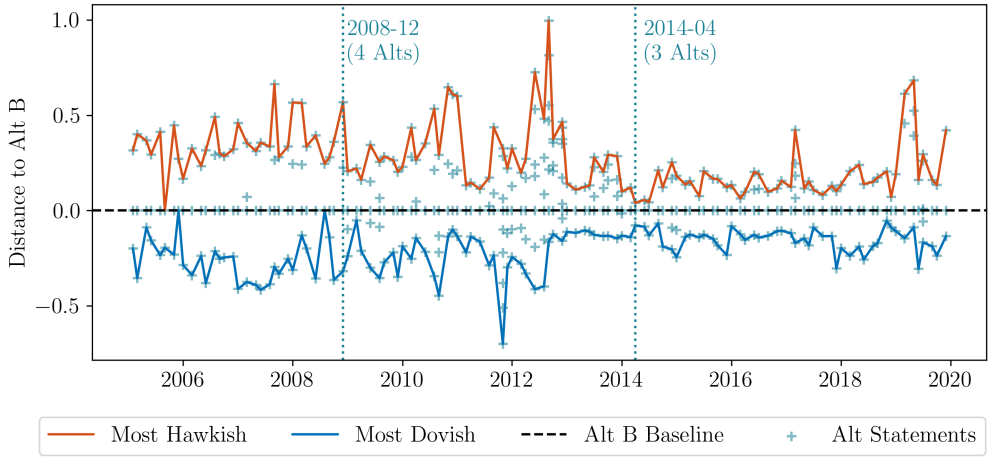
Thus we are able to measure the distance between messages in the data. In the model, the width of a message captures the precision of that message because the announcement width is the range of possible values of the signal θ_t. Therefore we assume that the distance measure in the data also captures a notion of announcement precision: the greater distance between alternatives, the less precise they are. To test whether this assumption makes sense, we perform a statistical test whether the message space satisfies a triangle inequality property: neighboring messages (A and B, or B and C) should be more similar by our measure than the extreme messages (the most dovish A and the most hawkish C). We find that 97 percent of our sample have message spaces that pass this test and present results in Appendix G. This is evidence that the one-dimensional distance measure captures our conceptual notion of message space span as envisioned in our model, and that the data support the interpretation of the model that smaller distances between messages correspond to more precise messages.

Figure 7 plots the alternative messages in relation to each other, normalized around the baseline alternative (Alt. B). The orange line shows the distance between the most hawkish alternative (Alt. C, or the last available letter) and Alt. B, and the blue line shows the distance between the most dovish alternative (Alt. A) and Alt B. This ordering, from hawkish to dovish, is taken directly from the Tealbooks and is assigned by Fed policymakers and staff. On the graph, we stack alternatives-distances such that hawkishness, or greater preference for price stability and higher interest rates, is considered more positive. As the orange and blue lines deviate more from zero, this represents that the Fed is considering a wider range of messages at that meeting, in other words, a larger span. Concretely, the span comes from adding up the relative distances to alternative B and is thus the distance between the orange and blue lines.[26]

[25] Another common measure of similarity between texts in the text analysis literature is the notion of cosine similarity. We provide robustness using 1-cosine similarity in Appendix G.

[26] In Appendix G, we provide robustness analysis for this measure where we simply directly measure the dissimilarity between the extreme alternatives (the first A and the last C), and we find the same qualitative results.

Figure 7: Distance Between Alternative Statements Over Time



Note: The graph shows the span between alternative messages at each FOMC meeting, centered around that meeting's baseline statement, Alt. B (black dashed line). We use a pretrained sentence BERT language model to generate text embedding vectors for comparison. Distance between statement embeddings is measured as Euclidean distance. The most hawkish statement (orange line) is the last letter (usually Alt C) in the Tealbook/Bluebook. The most dovish statement (blue line) is the first, Alt A. Teal crosses, +, represent the other alternative statements for that meeting.

In Figure 7, the number of alternatives is captured as the teal crosses, +, plotted between the hawkish and dovish contours. As in the plot of the simulated data from the model, Figure 6, the sum of the teal crosses at a given meeting is the count of alternatives at that meeting, and the distance between the orange and blue lines is the span of alternatives at that meeting. Looking at the data analogues of the model objects, we can already already see interesting patterns of Fed communication: the count and the span of alternative statements seem weakly correlated, but exhibit distinct variation. Accordingly, we will control for the span when the count is the regressand, and vice versa.

7.2 Reputation, Confidence, and Uncertainty

To test the model predictions for the central bank's communication strategy, we also need measures of confidence and reputation. Theoretically, confidence represents the precision of the central bank's signal which is the inverse of the bank's signal variance, while reputation is the market's view of the central bank's signal variance. Accordingly, we use the common measurement strategy to count "uncertainty words" in text to create an uncertainty measure.

The underlying assumption is that more frequent discussion of uncertainty corresponds to higher uncertainty and, thus, higher variance and lower precision.

We apply this method to two different sets of text to approximate the Fed's confidence and reputation. For reputation, we start with Baker et al. (2016)'s Monetary Policy Uncertainty (MPU) index. The index reports a monthly share of articles about monetary policy that also use "uncertainty" terms relative to all articles from the top 10 US newspapers. For confidence, we use the same list of "uncertainty words" and count their usage by FOMC members in the FOMC meeting transcripts, relative to the total amount of speech from policymakers. To recap, the interpretation of these two measures is that the more the FOMC says they are uncertain, then that represents higher variance of the Fed's signals, and therefore lower confidence. And analogously for the reputation measure, the more newspapers link uncertainty and the Fed, the higher is the perceived variance of the Fed's signals, and corresponding to a lower reputation. Therefore we multiply the indices by -1 to create our confidence and reputation series, so that higher values of the measures indicate higher confidence and higher reputation.

While the Baker et al. (2016) measure of uncertainty has been validated through its use in many research papers over the past decade, our text measure for confidence needs explicit validation. In Appendix G we validate the uncertainty-word-count measure of FOMC transcripts against a question in the Summary of Economic Projections (SEP) which asks about FOMC policymaker forecast uncertainty compared to the past 20 years. The SEP solicits projections and surveys the FOMC policymakers four times a year – at every-other FOMC meeting – since mid-2007. We find that our text-based uncertainty measure from the transcripts accounts for 90 percent of the variation in the number of FOMC members who say they are more uncertain, and has a strong correlation coefficient of about 0.6. Accordingly, we see our measure of uncertainty from the transcripts as an accurate representation of FOMC uncertainty that we can then invert to create our confidence measure for the larger sample of FOMC meetings.

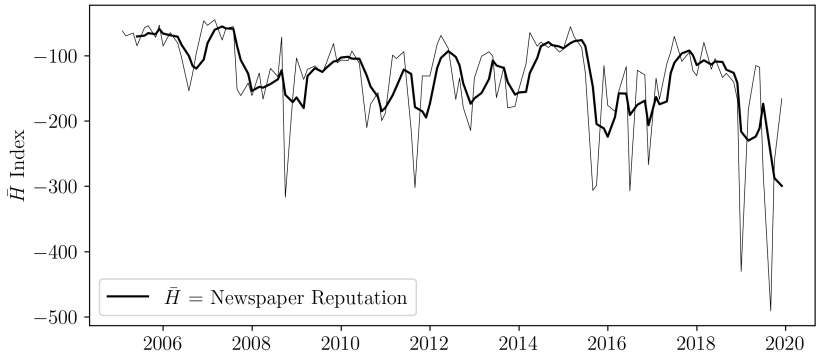
Figure 8 plots the normalized, smoothed series for confidence and reputation from 2005 through the end of 2019. Because we use the same construction strategy for both series – same dictionary and same normalization to the 1985-2010 average – we plot the ratio between

Figure 8: Text-Based Measures of Confidence and Reputation

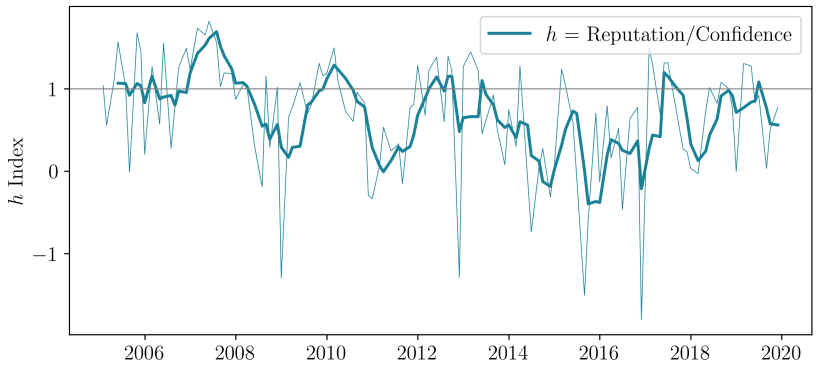
(a) Confidence



(b) Reputation



(c) Ratio of Reputation to Confidence



Note: Confidence is measured from the share of uncertainty words used by FOMC members in policy meeting transcripts. Reputation is the Baker et al. (2016) MPU index – the share of articles about monetary policy and uncertainty. Both are scaled so that the average share for 1985 to 2010 is equal to 100 (following Baker et al. (2016)), multiplied by -1 to transform uncertainty to confidence. The little h captures the ratio of reputation to confidence: constructed as the ratio of the uncertainty measures, multiplied by -1 so that higher values represent greater reputation/confidence and shifted so that $h = 1$ represents equal confidence and reputation. All graphs present the four-meeting moving average as the solid, thick line and the raw series is the thin line.

the two indices as a data proxy for our model's h variable, the reputation-to-confidence ratio. All series $H, \bar{H}, h$ are stationary.

Before going into the empirical analysis itself, we note three interesting observations in these plots. First, there is substantial volatility in confidence over time. This backs up our modeling choice to have H probabilistically change from period to period. Second, we find that changes in H precede changes in $\bar{H}$, which supports the idea that the public could be learning about the FOMC's confidence over time. Third, in the data, between 2012 to 2014, h oscillates above and below one, with the moving average very close to one. This period is also when we have a spike in the number of alternatives in Figure 1. In the model, situations where the reputation-to-confidence ratio h is close to one is the most interesting in terms of variation in count and span, and often also leads to spikes in the count. To go beyond this observational description, in the next section we statistically test for the relationships between confidence, reputation and the Fed's message space.

7.3 Specification and Results

In this section, we explore how confidence, reputation and the communication measures (count and span of alternatives) correlate. Since we are looking at relationships that have not been modeled or quantified before, we cannot be sure to have included all the relevant regressors. Therefore, on the one hand, it is important to remember not to interpret the results as causal. On the other hand, we do our best to control for confounding factors using two distinct empirical approaches. First, we employ a fixed effects model. Second, we also use a two-stage procedure where we report the correlations of residualized confidence and reputation with FOMC communication.[27]

The empirical specifications are as follows:

$$M_t = \beta_1 H_{t-1} + \beta_2 \bar{H}_{t-1} + \gamma C_t + \eta L_{t-1} + \varepsilon_t, \tag{17}$$

where t indexes the scheduled FOMC meeting. M_t is the message space variable, which is either the time-t span, the count, or the ratio between the two, $span_t/count_t$, which can be

[27]To explicitly deal with potential dynamics, we also provide a SVAR specification in Appendix H.

interpreted as the average message width (and will be denoted "Avg. Width" below). H_{t-1} and $\bar{H}_{t-1}$ (confidence and reputation) are our main regressors of interest, and C_t and L_{t-1} are vectors of contemporaneous controls and lagged controls, respectively.

We lag the main regressors such that they are measuring confidence and reputation prior to the drafting of the messages to avoid reverse causality. So H_{t-1} is measured from the transcripts from the previous FOMC meeting relative to the meeting where M_t is measured. $\bar{H}_{t-1}$ is measured from the MPU Index the month before the meeting of M_t. All series are log-transformed to account for skew. Uncertainty measures are log-transformed before multiplying by -1 to represent confidence and reputation. Additionally, we standardize, or z-score, the data so the intercept term is zero.

We include a variety of control variables. First, we have controls for mechanical features of how we measure communication. For example, we control for the count (span) of alternatives for regressions where the span (count) is the dependent variable because these terms are positively correlated; more messages are often associated with greater span of messages. We also include lags of the dependent variable and the message space control to account for persistence in the communication strategies.

We control for aggregate uncertainty in the economy with the VIX. We include the lagged level of the VIX and also lags of the change in the VIX to capture the changing aggregate uncertainty in the economy.[28] We control for other policy decisions with the change of the target federal funds rate so that our empirical specification remains focused on Fed communication policy, controlling for conventional monetary policy.

We also entertain two specifications targeted specifically at catching confounders. The first is a fixed effects model. We include two kinds of fixed effects. The first seeks to absorb internal institutional characteristics of the Fed that might affect Fed communication: who was Federal Reserve chair? Was there a press conference? Was the SEP released? Then we include time fixed effects to capture variation in the economic environment above and beyond that captured by the VIX that might also drive some of the communication strategy. In particular, we include fixed effects for the zero-lower-bound (ZLB) period, as that is the

[28]Specifically, we include a two-month rolling average of the VIX for the month prior to the FOMC meeting where we measure M_t as a control for aggregate uncertainty. We use a two month rolling average to capture persistent uncertainty in the inter-meeting period. Lastly, first-differences are across FOMC meetings.

most plausible candidate for a shift in economic conditions in our sample. Furthermore, we also see sizable movements in our text measures in Figure 7 starting in about 2011, and hence we seek to avoid attributing changes in communication falsely to confidence and reputation if they were actually driven by the ZLB.

As our second empirical specification, we provide a two-stage specification using the well-known residualization approach. In a first stage, we residualize the confidence and reputation variables by regressing them on each other, the VIX, their lags, and the fixed effects described above. From this first stage, we see that just under 35 percent of the variation in confidence is accounted for by these controls, while it is approximately 50 percent for reputation. The remaining residuals represent changes in confidence and reputation that are orthogonal to our controls, including the VIX, our stand-in for aggregate uncertainty. In the second stage, we run the regression in Equation 17 using the residuals from the first stage as regressors. The coefficients in this specification capture fluctuations in the message space that are correlated with plausibly exogenous changes in reputation and confidence. Ultimately, we find qualitatively similar results between both approaches.

Table 2 shows the results. We find empirical results that are consistent with the theory: lower reputation and higher confidence are associated with more precise communication. In particular, higher confidence is significantly associated with a lower span, while higher reputation with a lower count, albeit not significantly, both across the baseline and the residualized specifications. From the model's perspective, the data are consistent with being in the region where reputation is below confidence ($\bar{H} < H$): communication span decreases as confidence rises, while it increases with the bank's concern for reputation.

This provides some suggestive evidence that the intuition from the model is playing out in the data, in which on the one hand the more confident the CB is, the more narrow range of potential economic outcomes do they address in the set of alternatives drafted at a given policy meeting. Coming back to the committee disagreement interpretation of a low degree of confidence, this would suggest that the Fed staff draft alternatives that discuss more widely dispersed economic outcomes when disagreement in the FOMC is high.

On the other hand, there is also indication in the data that a higher number of alternative statements is associated with lower reputation. Although not statistically significant in our

Table 2: Message Space Measures, Confidence and Reputation

	(1) $Span_t$	(2) $Span_t$	(3) $Count_t$	(4) $Count_t$	(5) $Avg.\ Width_t$	(6) $Avg.\ Width_t$
H_{t-1}	-0.135**		-0.006		-0.149**	
	(0.059)		(0.091)		(0.071)	
$\bar{H}_{t-1}$	-0.075		-0.111		-0.065	
	(0.064)		(0.110)		(0.077)	
H^{res}_{t-1}		-0.136***		-0.030		-0.146**
		(0.051)		(0.077)		(0.062)
$\bar{H}^{\text{res}}_{t-1}$		-0.092*		-0.086		-0.088
		(0.047)		(0.080)		(0.058)
Controls:						
Message Space	$Count_t$	$Count_t$	$Span_t$	$Span_t$	-	-
VIX_{t-1}	Yes	Yes	Yes	Yes	Yes	Yes
$\Delta\,\text{FFR}_t$	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Lags	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.718	0.718	0.594	0.594	0.607	0.607
R^2 Adj.	0.674	0.674	0.531	0.531	0.554	0.554
N	118	118	118	118	118	118

Notes: HAC-robust standard errors with small sample correction in parentheses. * p<.1, ** p<.05, *** p<.01. Series are all standardized (z-scored) after log transformations. H_{t-1} is the confidence score from the previous FOMC meeting. $\bar{H}_{t-1}$ is the reputation score prior to the FOMC meeting. Regression sample covers Feb 2005 - Dec 2019. VIX controls include level and first difference. Included fixed effects for: Federal Reserve Chair Bernanke, zero-lower-bound federal funds rate, FOMC meeting with a press conference, FOMC meeting with Summary of Economic Projections release. Lagged controls include 1 period lag for the dependent variable and all regressors, except for the residualized regressors. Full regression tables and the first-stage regressions for residualized regressors are in Appendix H.

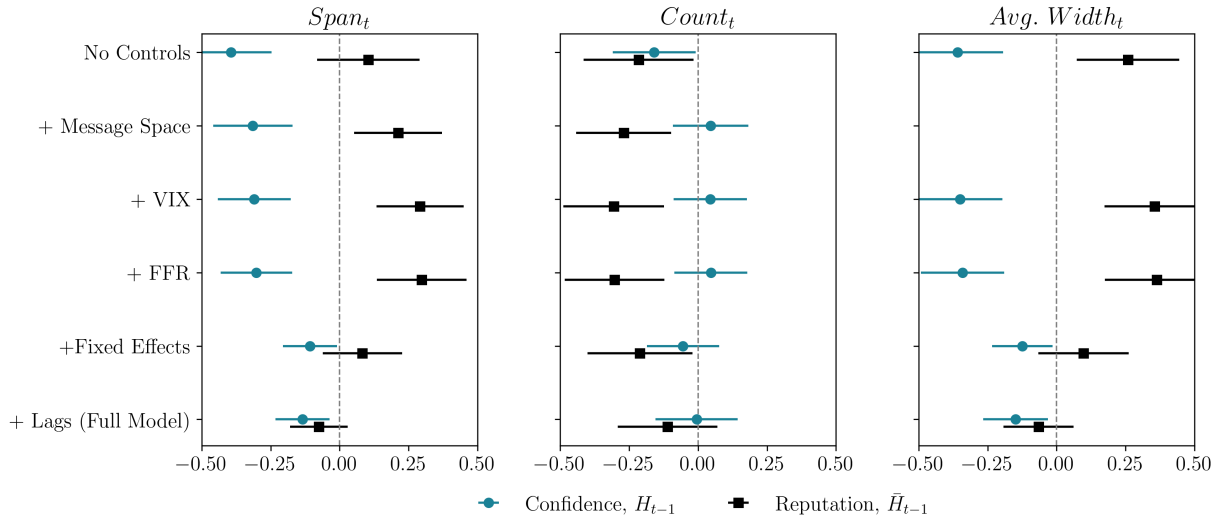
full specification, as we vary our controls (Figure 9) we can see that reputation is tied more closely to the count of messages. In the model, this happens because it takes more precise communication to move market expectations when the market thinks that the Fed is not very confident. Writing more alternatives allows the Fed to tailor each alternative more specifically, thus communicating more precisely. There is some suggestion in the data that this strategic concern may be present to some degree in actual Fed communication, although the lacking significance indicates that it is certainly not as strong as predicted by the model.

Having a look at the coefficients for average message width, we see a similar implication. Recalling that message width has an interpretation as message precision in the model, with narrower messages being more precise, and that we tested and obtained support for this in the data, we can now interpret average message width as average message precision. We can

see from columns 5 and 6 that average message width is significantly lower when confidence is higher, but is not correlated in a significant way with reputation. While this supports the model prediction that higher confidence is associated with more precise communication, it also corroborates the finding that the reputation effect is weaker in the data than predicted by the model.

The important takeaway from these empirical results is that in the data, confidence seems a primary driver of span and reputation is more of a driver for the count. Thus we find empirical support for the model's suggestion that the Fed drafts alternatives that focus on a more narrow set of economic outcomes the more confident they are about the outlook, and some support for the idea that they draft more alternatives, and thus more precise statements, when reputation is low and they need to work harder to move market expectations.

Figure 9: Communication, Confidence, and Reputation Sensitivity to Controls



Note: This plots the coefficients on confidence (teal circles) and reputation (black squares) for specifications with varying controls. Dots are the coefficients, and bars represent 90 percent confidence intervals. The "no controls" specification regresses the message space measure (span, count, or average width) on the lag of confidence and reputation. The "full model" specification matches the coefficients in Table 2 for col. 1, 3, and 5. Regression tables for these plots are in Appendix H.

What is really novel about these results is showing that variation in the public's opinion of the Fed affects the precision of communication. However, in the data we find a smaller role

for reputation in the message space compared to confidence. Thus we conclude that the Fed is reacting less to reputation in the data than our model would recommend. Interestingly, and possibly related with this point, the dynamics in the data seem most consistent with model dynamics where reputation is below confidence in level terms. This lends support to the idea that financial markets are opinionated and think less of the Fed's ability to forecast compared to their own (Caballero and Simsek, 2022). The finding that the Fed responds less to reputation than implied by our model further supports the story in Caballero and Simsek (2022) that the Fed does not bow to the market's opinion.

8 Conclusion

When a central bank makes announcements, not only do markets pay attention to the words, but the public's belief at how good the central bank is at forecasting - the reputation for confidence - matters for how markets respond to announcements and guidance. In this paper, we model this interaction, and explore how it affects central bank communication and monetary policy.

We find that even absent a concern for reputation - so that the central bank ignores that their actions may affect the evolution of reputation - the model generates a central bank communication policy that is in line with how the Fed communicates, with coarse transparency. This involves drafting a set of candidate messages and picking from this set a single announcement that is released to the public. Importantly, we find that the evolution of reputation drives the central bank to consider a varying number of alternative statements that are varying in their similarity to each other.

The reason this happens is that the central bank creates their announcements by anticipating how the public will receive that information and balancing that against their incentives to adjust public expectations. We show that this strategic tradeoff measurably affects the message space at the FOMC meetings, and thus ultimately the announcements they choose. However, our empirical evidence shows the Fed responds less to reputation than is recommended by the model. With the growing literature studying information effects, it is important to take these strategic language motives into account.

References

Ahrens, Maximilian, Deniz Erdemlioglu, Michael McMahon, Christopher J. Neely, and Xiye Yang (2025) "Mind Your Language: Market Responses to Central Bank Speeches," *Journal of Econometrics*, 249, 105921.

Alexopoulos, Michelle, Xinfen Han, Oleksiy Kryvtsov, and Xu Zhang (2024) "More Than Words: Fed Chairs' Communication During Congressional Testimonies," *Journal of Monetary Economics*, 142.

Amador, Manuel and Christopher Phelan (2021) "Reputation and Sovereign Default," *Econometrica*, 89 (4), 1979–2010.

Amador, Manuel and Pierre-Olivier Weill (2010) "Learning from Prices: Public Communication and Welfare," *Journal of Political Economy*, 118 (5), 866–907.

Aruoba, S. Borağan and Thomas Drechsel (2026) "Identifying Monetary Policy Shocks: A Natural Language Approach," *American Economic Journal: Macroeconomics (forthcoming)*.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis (2016) "Measuring Economic Policy Uncertainty," *Quarterly Journal of Economics*, 131 (4), 1593–1636.

Baley, Isaac and Laura Veldkamp (2022) "Bayesian Learning," in *Handbook of Economic Expectations*: Elsevier.

Barro, Robert J and David B Gordon (1983) "Rules, Discretion and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics*, 12 (1), 101–121.

Bassetto, Marco (2019) "Forward Guidance: Communication, Commitment, or Both?" *Journal of Monetary Economics*, 108, 69–86.

Bholat, David, Stephen Hansen, Pedro Santos, and Cheryl Schonhardt-Bailey (2015) "Text Mining for Central Banks: Handbook," *Centre for Central Banking Studies Handbook* (33), 1–19, Publisher: Bank of England.

Caballero, Ricardo J. and Alp Simsek (2022) "Monetary Policy with Opinionated Markets," *American Economic Review*, 112 (7), 2353–92.

——— (2025) "FCI-plot: Central Bank Communication Through Financial Conditions," *Working Paper*.

Camous, Antoine and Dmitru Matveev (2025) "Central Bank Strategic Communication and the Dynamics of Reputation," *Working Paper*.

Chari, Varadarajan, Patrick Kehoe, and Ellen McGrattan (2000) "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?" *Econometrica*, 68 (5), 1151–1179.

Cheng, Haw and Alice Hsiaw (2022) "Distrust in Experts and the Origins of Disagreement," *Journal of Economic Theory*, 200.

Cieslak, Anna and Michael McMahon (2023) "Tough Talk: The Fed and the Risk Premium," *Working Paper*.

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

Crawford, Vincent P. and Joel Sobel (1982) "Strategic Information Transmission," *Econometrica*, 50 (6), 1431–1451.

Cukierman, Alex and Allan H. Meltzer (1986) "A Theory of Ambiguity, Credibility, and Inflation Under Discretion and Asymmetric Information," *Econometrica*, 54 (5), 1099–1128.

Doh, Taeyoung, Dongho Song, and Shu-Kuei Yang (2026) "Deciphering Federal Reserve Communication via Text Analysis of Alternative FOMC Statements," *American Economic Journal: Macroeconomics (forthcoming)*.

Frankel, Alex and Navin Kartik (2018) "What Kind of Central Bank Competence?" *Theoretical Economics*, 13 (2), 697–727.

Gáti, Laura (2023) "Talking Over Time – Dynamic Central Bank Communication," *Journal of Money, Credit and Banking*, 55 (5), 1147–1176.

Gáti, Laura and Amy Handlan (2025) "Monetary Communication Rules," *Working Paper*.

Ghofrani, Erfan (2023) "Learning-Induced Uncertainty and Investment Dynamics: Evidence from Bayesian Updating in Firms," *Working Paper*.

Granziera, Eleonora, Vegard H Larsen, Greta Meggiorini, and Leonardo Melosi (2025) "Speaking of i=Inflation: the Influence of Fed Speeches on Expectations," *Working Paper*.

Guembel, Alexander and Silvia Rossetto (2009) "Reputational Cheap Talk with Misunderstanding," *Games and Economic Behavior*, 67 (2), 736–744.

Haldane, Andrew and Michael McMahon (2018) "Central Bank Communications and the General Public," *AEA Papers and Proceedings*, 108, 578–583.

Handlan, Amy (2020) "FedSpeak Matters: FOMC Statements and Monetary Policy Expectations," *Working Paper*.

——— (2022) "Text Shocks and Monetary Surprises: Text Analysis of FOMC Statements with Machine Learning," *Working Paper*.

Hansen, Stephen, Michael McMahon, and Andrea Prat (2018) "Transparency and Deliberation Within the FOMC: a Computational Linguistics Approach," *Quarterly Journal of Economics*, 133 (2), 801–870.

Iovino, Luigi, Jennifer La'O, and Rui Mascarenhas (2022) "Optimal Monetary Policy and Disclosure with an Informationally-Constrained Central Banker," *Journal of Monetary Economics*, 125, 151–172.

King, Robert G and Yang K Lu (2025) "Evolving Reputation for Commitment," *Working Paper*.

Ko, Eunmi (2023) "Optimal Central Bank Forward Guidance," *International Journal of Central Banking*, 19 (4), 401–447.

Kreps, David M and Robert Wilson (1982) "Reputation and Imperfect Information," *Journal of Economic Theory*, 27 (2), 253–279.

Li, Ming and Kristóf Madarász (2008) "When Mandatory Disclosure Hurts: Expert Advice and Conflicting Interests," *Journal of Economic Theory*, 139 (1), 47–74.

Lorenzoni, Guido (2010) "Optimal Monetary Policy with Uncertain Fundamentals and Dispersed Information*," *Review of Economic Studies*, 77 (1), 305–338.

Mailath, George and Larry Samuelson (2001) "Who Wants a Good Reputation?" *Review of Economic Studies*, 68 (2), 415–441.

Moscarini, Giuseppe (2007) "Competence Implies Credibility," *American Economic Review*, 97 (1), 37–63.

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

Polis, Andrea De, Leonardo Melosi, and Ivan Petrella (2025) "The Taming of the Skew: Asymmetric Inflation Risk and Monetary Policy," *ECB Working Paper*.

Rogoff, Kenneth (1985) "The Optimal Degree of Commitment to an Intermediate Monetary Target," *Quarterly Journal of Economics*, 100 (4), 1169–1189.

Rotemberg, Julio and Michael Woodford (1997) "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual*, 12, 297–346.

Sastry, Karthik (2026) "Disagreement About Monetary Policy," *American Economic Journal: Macroeconomics (forthcoming)*, Conditionally accepted American Economic Journal: Macroeconomics.

Shapiro, Adam Hale and Daniel J Wilson (2022) "Taking the Fed at Its Word: A New Approach to Estimating Central Bank Objectives Using Text Analysis," *Review of Economic Studies*, 89 (5), 2768–2805.

Sobel, Joel (1985) "A Theory of Credibility," *Review of Economic Studies*, 52 (4), 557–573.

Woodford, Michael (2003) *Interest and Prices: Foundations of a Theory of Monetary Policy*: Princeton University Press.