DISCUSSION PAPER SERIES

DP20734

REPUTATION FOR CONFIDENCE

Laura Gáti and Amy Handlan

MONETARY ECONOMICS AND FLUCTUATIONS

CEPR

ISSN 0265-8003

REPUTATION FOR CONFIDENCE

Laura Gáti and Amy Handlan

Discussion Paper DP20734
Published 13 October 2025
Submitted 15 September 2025

Centre for Economic Policy Research
187 boulevard Saint-Germain, 75007 Paris, France
2 Coldbath Square, London EC1R 5HL
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Monetary Economics and Fluctuations

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Laura Gáti and Amy Handlan

REPUTATION FOR CONFIDENCE

Abstract

We model how a central bank communicates its noisy forecasts (forward guidance) while taking into account its 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.

JEL Classification: E52, E58, C49

Keywords: N/A

Laura Gáti - laura_veronika.gati@ecb.europa.eu
European Central Bank, DG Research and CEPR

Amy Handlan - amy_handlan@brown.edu
Brown University

Reputation for Confidence*

Laura Gáti† and Amy Handlan‡

September 14, 2025

Abstract

We model how a central bank communicates its noisy forecasts (forward guidance) while taking into account its 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. Meanwhile, how the public responds to the announcement depends on whether they believe the Fed is informed about the economic situation. The tension between effective forward guidance and discretionary policy-making affects the choice of monetary policy announcement language.

This paper develops a game-theoretic model to examine how central banks manage communication to balance informational transparency with policy flexibility, which we test using text data from and about the Fed. Theoretically, we find that the bank's own uncertainty (or *confidence*) shapes how precisely it communicates. While precise announcements provide stronger forward guidance by shifting expectations, markets may overreact or underreact depending on their views of the bank's forecasting ability, which we call its *reputation for confidence*. The bank balances these tradeoffs by creating an endogenous message space that depends on both its confidence and reputation. When forecasts are more accurate – allowing the Fed to move expectations with less risk of later needing to deviate – or when reputation is lower – reducing sensitivity to announcements– the bank drafts announcements that are more precise and similar to one another. 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 from newspaper coverage of Fed uncertainty. The evidence broadly supports the model, with one exception: Fed communication is less responsive to reputation than the model predicts.

To explore the foundations for these results, we now describe the theoretical model. In the 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. The central bank only observes noisy signals of the shock, unobserved by the public, 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.

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.[1] 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). The key tension that the model emphasizes is that reputation for confidence has implications for how the market understands the Fed's announcement language, and that the Fed therefore needs to take reputation into account when drafting their statements.

Following Moscarini (2007), we conceptualize central bank communication as drafting a set of candidate announcements - what we refer to as the *message space* - from which the bank picks a single message to make public. Announcements in the model take the form of *intervals* that contain the central bank's signal of the shock, with the interval width being a measure of announcement precision. 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. We then explore what the equilibrium message space looks like, and what features can align it with Fed communication in the data. We devote particular attention to whether the central

[1] We can microfound this assumption in two ways. First, the central bank cannot credibly communicate its confidence to the public because it is always better off when the public believes it 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 the confidence as not perfectly being seen by the public.

bank ever reveals the exact value of the signal they saw, a feature we call *point revelation*.

The analysis yields two striking results. First, point revelation is possible if and only if reputation equals true confidence. What this result 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 *interpretation 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.[2] We argue this *interpretation friction* is a more accurate motive for modern central bank communication because it allows the model simulations to better reflect patterns in the data. Additionally, we find narrative-evidence from FOMC transcripts which quote policymakers choosing language while struggling with the interpretation friction.[3]

The second key result is that model predicts a time-varying message space due to the evolution of reputation, which is what we also find in the data. The intuition for this is that reputation governs the public's responsiveness to central bank announcements, because the public discounts announcements by how confident they think the bank is. As reputation varies over time through the public learning about the central bank's confidence, the public's responsiveness to announcements becomes time-varying, confronting the central bank with a slightly different communication problem every time they are preparing to make an announcement. This generates an equilibrium message space that varies over time in terms of number and similarity of candidate statements.

The model also provides predictions for when the central bank drafts many or few alternative statements, and when the alternatives contain very similar or different language.

[2]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 the the bank can secretly create inflation to generate a positive output gap through the Phillips curve.

[3]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.*"

Broadly, the model suggests that the central bank's confidence determines what range of economic scenarios to discuss, a feature we refer to as the *span* of candidate announcements. 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 small span. A less confident central bank will need to be ready to discuss scenarios that are further apart, such as inflation significantly over- or undershooting, raising the span.

Reputation instead governs the market's interpretation of the bank's announcements. When reputation is high, vague statement language can be sufficient. When reputation is low, instead, more nuanced language is necessary to make sure the market's interpretation of the announcement is the one the central bank intended. In such a case, a more vague statement such as "inflation is above target" will have to be replaced by more nuanced wording such as "inflation is at least two percentage points above target." The model thus suggests that the Fed's anticipation of the market's reaction to their words is also a key factor in Fed communication, and adjusting the number of alternative announcements is a way for the Fed to get the granularity of their wording in each candidate announcement exactly right.

To connect our model to real-world communication policy, we create novel measures of the Fed's message space (the set of candidate statements at each meeting), the Fed's confidence, and its reputation with text data for the empirical analysis. For the message space, we leverage internal Fed materials: the Bluebooks and Tealbooks. Since 2005, they include multiple candidate announcements, the *alternative statements*, from which the FOMC selects their policy announcement. We estimate the Fed's message space at each meeting using frontier text analysis methods to quantify the similarity between neighboring alternative statements. We document two novel facts about the Fed's communication choice set. First, the number of alternative statements is time-varying, ranging between two and nine. Second, the similarity between alternative statements is also time-varying, so that the most extreme alternatives sometimes convey a very similar, and at other times a starkly different message.[4] Together, the similarity and the number of alternative statements allow us to approximate

[4]Doh et al. (2020) 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.

the precision of announcements: a set of many alternatives that are all similar correspond to more detailed, precise language choices of the Fed. Conversely, having a larger span of statements (lower similarity or more diverse scenarios) and fewer statements to cover that span captures a coarser space.

For confidence, we actually measure an inversely related concept: uncertainty. We count the frequency of uncertainty words used in FOMC transcripts following Baker et al. (2016) . We invert the uncertainty index so that higher values are associated with lower uncertainty, that is, higher confidence. For reputation, we use an inversion of the Monetary Policy Uncertainty index from Baker et al. (2016). Similarly, higher values of the reputation index mean that there is less news signaling public uncertainty of the Fed, which we interpret as a higher reputation for confidence. For our new measures, we provide validation and robustness analyses to ensure they capture FOMC confidence and their time-varying message space.

To test the model predictions in the data, we consider two types of empirical specifications: a baseline OLS regression and a more dynamic, SVAR specification. They both provide the same broad conclusions, which underscores the robustness of our findings. First, we find that that increases in confidence lead to more precise announcements and a smaller span of alternative scenarios. In the SVAR, we the average precision of announcements increases after a positive shock to confidence.

Second, we find that the count of alternative statements also show that lower reputation is associated with more alternatives, after controlling for the span of messages. We see this as suggestive evidence that the Fed is drafting more alternatives when they are more concerned with how the public will respond. But, this finding is not robust across specifications. In the SVAR, we find that the precision of announcements does not respond to a reputation shock. Thus, we conclude that confidence is the main driver of the Fed's communication strategy and there is less responsive to reputation than is recommended by the model.

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 builds on several strands of literature to contribute to the macro theory of central bank communication. This 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; Sastry, 2025) and normative implications for central bank communication (Caballero and Simsek, 2025; Iovino et al., 2022; Polis et al., 2025).

The first literature we build on is the literature on cheap talk games in the spirit of Crawford and Sobel (1982), who were the first to show that misalignment in preferences leads a sender to send noisy signals to the receiver in the form of a partitioned message space. Our model borrows from the framework of Moscarini (2007), who in turn builds on Crawford and Sobel (1982) to investigate how the precision of a central bank's information - central bank confidence - affects the equilibrium message space. Our paper is thus related to others that explore various characteristics of the Crawford and Sobel (1982) cheap talk game (Frankel and Kartik, 2018; Li and Madarász, 2008), in particular its application to central bank communication (Amador and Weill, 2010; Bassetto, 2019; Cukierman and Meltzer, 1986), as well as to papers on central bank communication games in general (Gáti, 2023; Ko, 2023).

Our work also relates to papers on reputation-building about preferences in incomplete information 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), especially in conjunction with central bank communication. Our setup differs from the latter 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. Coarse communication in equilibrium does not come from preference misalignment or an inbuilt concern for reputation, but from a misinterpretation channel, related to the main ideas in the Bayesian Persuasion literature (Kamenica and Gentzkow, 2011), and to the literature on disagreement between the central bank and the public (Caballero and Simsek, 2022; Sastry, 2025). This result also echoes Iovino et al.

(2022)'s finding that when the central bank and the market are imperfectly informed, it is not always optimal to reveal the central bank's private information to the public.

The learning element of our model belongs to the Bayesian parameter learning literature, and comes from Ghofrani (2023) and Baley and Veldkamp (2022). Our work is also related to other papers where there is learning about a policymaker's type, such as King and Lu (2025), where reputation is also conceptualized as Bayesian learning about policymaker type from announcements and realized inflation. Thus our work also connects to the classical literature on the conduct of 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. Our finding that the central bank communicates coarsely due to an evolving reputation as the market learns about central bank confidence over time corroborates the evidence presented by Kostyshyna and Petersen (2023) that central banks communicate with uncertainty, that this uncertainty is relevant for markets' expectation formation (Istrefi and Piloiu, 2020), as well as the point of Sastry (2025) that models seeking to rationalize disagreements in the data between markets and the central bank need to feature varying market beliefs about the precision of the bank's information. It also lines up with the insight of Caballero and Simsek (2022) that disagreements between the central bank and the market need to be taken into account actively by the bank.

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); Aruoba and Drechsel (2025); Bholat et al. (2015); Cieslak and McMahon (2023); Doh et al. (2020); 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. (2020) and Handlan (2022) use alternative statements that the Fed considers at their policy meetings. Our innovation is to use computational text analy-

sis to discipline a structural model to better understand the strategic pressures on the Fed communication decision.

2 Background on Federal Reserve Communication

Over the past decades, the Fed has increased the mediums in which it systematically communicates its 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.

What are in the FOMC policy statements? The 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 Fed forecasts, reinforcing the view that statements signal future policy.

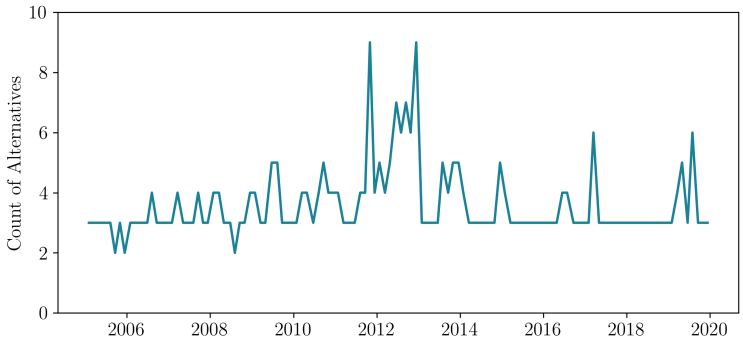
How are these statements created? Drafts are prepared by staff, revised with input from Governors and regional banks, and presented as a set of alternative statements in the the FOMC policy meeting materials (Tealbooks/Bluebooks) since 2005. At the meeting, FOMC members vote on the statement language usually by selecting one of the statement drafts from the set. Almost always "alternative B" is selected.[5]

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, each draft reflects a message regarded by staff or Committee members as substantively different from the others.

In Figure 1, we graph the number of alternative statements from 2005 through 2018. We

[5]Sometimes new phrases are added to the statement that were not in 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.[6] 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. The range of alternative statements is on spectrum from dovishness to hawkishness. Alternative A is considered more "dovish," concerned with supporting economic growth and employment, the last letter alternative, usually C or D, is considered more "hawkish," concerned with preventing high inflation. Alternative B, which is often called the "likely" alternative, lies in between the extremes with respect to 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. As investors and journalists have increasingly dedicated resources to parsing every word of the FOMC – earning the term "Fed Watchers" – the FOMC rationally has responded by choosing their words carefully to balance their own forecasts, the forward guidance they want to provide, and how they anticipate the public will receive their words. In this paper, we model these tensions in order to study the central bank's forward-guidance communication decision in the presence of uncertainty.

[6]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.

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 is a reduced-form stand-in for policy-relevant shocks[7]:

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

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

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

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. We index confidence relevant at

[7] This could be a demand shock or preference shock. For a demand shock, one could have a structural demand shock by adding a demand block to the model. But this would add expositional complexity without altering the intuition.

time t (for meeting t) with $t-1$ to reflect the idea that the central bank's confidence when going in to a policy meeting is given by their historical forecasting performance leading up to period t.

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.

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

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

The novelty of our paper 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 $\bar{H}_t \equiv \mathbb{E}_t^P[H_t]$. The key question we will ask is how the evolution of reputation affects central bank communication, and how communication affects reputation.

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

different parameters for it: $H_{t-1} \sim \Gamma(\alpha, \beta)$. Reputation then evolves as P updates beliefs on H_t 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 $k_t = \mathbb{E}_t^P[\theta_t | H_{t-1}\theta_t] = H_{t-1}\theta_t / \bar{H}_{t-1}$, and $\mu_k = 0$.

The CB has two policy tools: setting inflation, π_t, and making a public announcement. 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 to the public.

We refer to the set of candidate statements as the *message space*, and denote it by $\mathcal{M}_t$. The message space is endogenous to confidence, H, and reputation, $\bar{H}$. For a given confidence and reputation, the message space $\mathcal{M}_{H,\bar{H}}$ is either $\mathbb{R}$ itself, or a partition of $\mathbb{R}$ with non-overlapping intervals. Because the public does not see the message space, we need to place some assumptions on how they can conceive of messages for their basis for forming expectations. We assume the public knows the message space will be supported by $\mathbb{R}$ they will entertain $\mathcal{M}_{\bar{H}}$, the combined set of possible messages given reputation:

$$\mathcal{M}_{\bar{H}} \equiv \bigcup_{H} \mathcal{M}_{H,\bar{H}} \subseteq \mathbb{R}. \tag{10}$$

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 $\mathcal{M}_{\bar{H}}$.[8] 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, a particular message interval A could be part of multiple $\mathcal{M}_{H,\bar{H}}$, so the public will not be able to perfectly back out H from a single message.[9] For simplicity of notation, $\mathcal{M}_t$ is the message space at at time t as a function of H_{t-1} and $\bar{H}_{t-1}$.

Based on the observed announcement, the public takes a single action, which is to form (rational) expectations of inflation. Concretely, the CB makes an announcement $A_t(\theta_t) \in \mathcal{M}_t \subseteq \mathbb{R}$ of the signal θ_t. For instance, if the CB observes a signal of $\theta_t = 2$, then a possible example announcement is the interval $[1, 3]$. Importantly, the infinitely fine message "2" is also a possible announcement, as well as the infinitely coarse message $(-\infty, \infty)$. Note that in principle, a possible message would also be an interval that does not contain θ_t, for instance $[0, 1]$ when $\theta_t = 2$. In other words, it is possible for the CB to lie. 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 ever again. The possibility of the grim trigger results in it being optimal for the CB to always send announcement intervals that contain θ_t, that is, truthtelling messages.[10]

While most papers on communication games concentrate on whether the sender tells the truth or not, given a fixed message space, we will take a different approach. In this paper, we focus our analysis on equilibria where the message space is endogenous to reputation and confidence. The message space adjusts in order to induce the central bank to announce the truth, albeit potentially imprecisely, because it is disciplined by the grim trigger. Accordingly, the message space is allowed to vary by number and width of announcement intervals, resulting in varying precision of messages. In particular, the width of announcement intervals captures whether announcements are coarse or fine. For example, $A_t(\theta_t) = \mathbb{R}$ is the coarsest candidate message, and $A_t(\theta_t) = \theta_t$ is the finest, a situation we will refer to as *point*

[8] 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.

[9] We explore this idea further in the Appendix A where we allow the public to learn about H from the message directly and through simulation find that almost always there is insufficient information for the public to want to change their prior, $\bar{H}_{t-1}$, mid period.

[10] Throughout the paper, we use "truthtelling" to mean that the announcement contains θ_t. But, truthtelling does not imply full revelation of the central bank's information.

revelation.

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

$$\bar{\theta}_t \equiv \mathbb{E}_t^P[\theta_t | A_t, \bar{H}_{t-1}], \tag{11}$$

which they use to form inflation expectations $x_t(\bar{\theta}_t(A_t, \bar{H}_{t-1}))$.[11] 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.

Given inflation expectations $x_t(\bar{\theta}_t(A_t, \bar{H}_{t-1}))$, the central bank then sets inflation, so that inflation and output are realized. This brings us to the last within-period stage of the game where reputation is updated. As foreshadowed earlier, we assume that with some probability $\gamma > 0$, the central bank's signal, θ_t, is revealed to the public. If the public catches the central bank in a lie, meaning that the announcement A_t is revealed not to have contained θ_t, the public reverts to never believing the CB again. Otherwise, $\bar{H}_t$ updates to H_{t-1}. If θ_t is not revealed, the public can still use the Phillips curve and knowledge of the equilibrium relations to back out the central bank's forecast of ω_t, $H_{t-1}\theta_t$. In this case the public uses this information to update reputation through Bayesian learning, Equation 7.

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:

$$\mathcal{L}^P(x_t, \bar{H}_{t-1}, A_t) = \mathbb{E}_t^P[L^P(x_t, \pi_t) | \bar{H}_{t-1}, A_t] = \mathbb{E}_t^P[(x_t - \pi_t)^2], \tag{12}$$

[11] In Appendix A, 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 this not a restrictive assumption.

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

$$\begin{aligned}\mathcal{L}^{CB}(A_t, \pi_t, \theta_t, H_{t-1}, \bar{H}_{t-1}) =& \mathbb{E}^{CB}[L^{CB}(y_t, \pi_t, \omega_t)|\theta_t, H_{t-1}, \bar{H}_{t-1}] \\ =& \mathbb{E}^{CB}[(y_t - b)^2 + \lambda(\pi_t - \pi^* - \omega_t)^2].\end{aligned} \tag{13}$$

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.[12]

We now describe the timing of the game. In the beginning of period t, reputation is inherited from period $t-1$. Given confidence H_t 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 also realized through the Phillips curve. Given the realizations of the observables, the public observes the CB's forecast of ω_t, $H_{t-1}\theta_t$. With probability p, 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 $1-p$, θ_t is not revealed, and the public uses the backed-out forecast $H_{t-1}\theta_t$ to update reputation.

We proceed by analyzing the communication problem in two stages. First, we 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

[12] 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 $b = 0$.

confidence, a new equilibrium message space arises.

For tractability we do not allow for the CB to anticipate how its announcements affect its reputation in future periods, beyond communicating in a way to avoid the babbling equilibrium. 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 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, the

public is solving a static problem, we can solve the stage game backwards by first computing the best-response inflation rate that minimizes the expected period loss of Equation 13

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

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[13]

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

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 lemma.

Lemma 1 *(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.

The intuition behind the lemma is the following. Our equilibrium concept implies that the message space specifies an incentive compatible announcement for every possible realization of the signal. One can thus think of the message space as the communication policy of the central bank, telling the CB what the optimal announcement would be for each possible economic outlook. In this sense, the message space is the model's notion of the set of alternative statements, each of which is an interval on the real line that is optimal for particular realizations of the signal. Once the signal is realized, the CB picks the corresponding alternative from the set of alternative statements to be released to the public, and keeps the rest of the alternative statements as internal materials.

[13]The public expectations of the bank's conditional forecast is $\mathbb{E}^P[H_{t-1}\theta_t] = \bar{H}_{t-1}\bar{\theta}_t$.

Now we provide an outline of the derivation of the equilibrium message space, suppressing the details to Appendix A. We start by computing best responses for inflation expectations x_t and inflation π_t, taking communication as given. Substituting π_t^{eqb} and x_t^{BR} into the period loss of Equation 13, and postulating that the equilibrium announcement should fulfill the incentive compatibility from the definition of equilibrium leads to an equation that characterizes the relationship between the borders of the intervals and the private sector expectations induced by the intervals. In other words, solving for the equilibrium message space becomes a problem of solving for a sequence of intervals of unknown number and width. If the solution contains K intervals, then the derivation proceeds by solving for an increasing sequence of $K-1$ cutoffs between the intervals that satisfy the incentive compatibility constraint. We do this numerically using a shooting algorithm (see Appendix B).

Examining the incentive compatibility constraint with the best responses substituted in, we arrive at the following characterization.

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.

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.[14]

The requirement that $b = 0$ is familiar from the literature on communication games which emphasize 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).

[14]The existence of point revealing equilibria does not close out the existence of other equilibria that involve coarse communication.

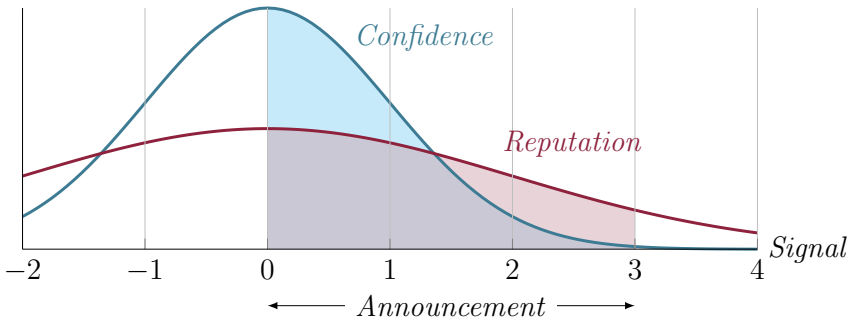
Here, instead, $b = 0$ is not sufficient for point revelation. The first of the two possible additional requirements, $\theta_t = 0$, is trivial. It simply says that if $b = 0$, so that the CB has no incentive to induce an inflation surprise, and if the CB thinks that no shock is hitting the economy, 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, namely that $\bar{H}_{t-1} = H_{t-1}$ for point revelation. The interpretation is that if the CB thinks that shocks are hitting the 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 is a contribution to the literature on central bank communication and broadens the understanding of why central banks communicate coarsely even without the 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 the same 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

Figure 2: Reputation Guides the Interpretation of Announcements



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 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. We employ a so-called shooting algorithm, outlined in Appendix B. The idea of the algorithm is to guess a sequence of interval cutoffs $\{\theta^k\}_{k=1}^{K-1}$ for a given number of alternatives K, compute the associated conditional expectations for each alternative k, $\bar{\theta}_k$, 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 $K = K_0 = 12$, and if no equilibrium is found, we decrease K by 1.[15] We keep going until $K = 2$, which is the most coarse message space. If no equilibrium message space is found, we set $K = 2$, because one can show theoretically that a two-message equilibrium always exists (see Appendix A).[16]

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 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 message space exists, coarse message spaces may also exist. Since the algorithm is designed to find coarse message spaces, in this case it will find message spaces with a finite $K \geq 2$. In the following plots, therefore, we focus on the finest coarse message space.

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.[17] Our baseline value for the output bias b is zero in order to avoid preference misalignment from driving the results. However, in certain places in the main text, we also consider a positive inflation surprise motive (Moscarini, 2007) to aid model interpretation, and we always

[15] K_0 is selected 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.

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

[17] 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).

indicate when we do so. Furthermore, robustness checks with $b = 0.02$ are in the appendix. 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.

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 start our analysis by focusing on two properties of the equilibrium message space. The first is the count of candidate messages, 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 $\theta^3 - \theta^1$.

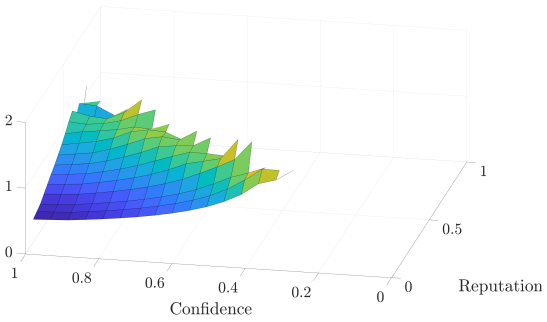
Figure 3 and Figure 4 show 3D plots comparing how the count and the span vary for different values of confidence and reputation. Higher z-axis values are indicated by lighter colors, while darker colors mean low z-axis values. While Figure 3 shows a smoothed version of the direct model output for the span, we chose to let Figure 4 convey a stylized representation of the model output for the count in order to simplify the interpretation. The smoothed version of the direct model output is shown in Appendix C, along with alternative specifications.

We start with Figure 3. Notice that the span is only defined when reputation is weakly below confidence, that is, "below" the diagonal. This is because, as can be seen on Figure 4, when reputation is above confidence, the count is generally two, so there is only a single cutoff.

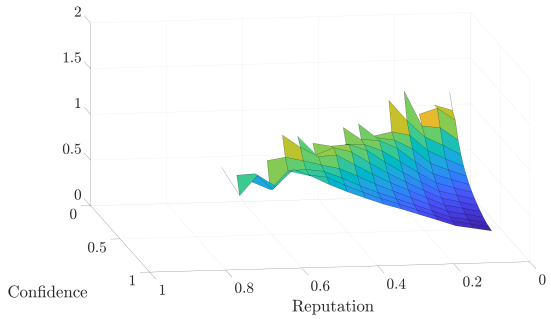
Otherwise, one observes that span is decreasing in confidence. This makes sense: as the

Figure 3: 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 versions with $b = 0.02$, see Appendix C.

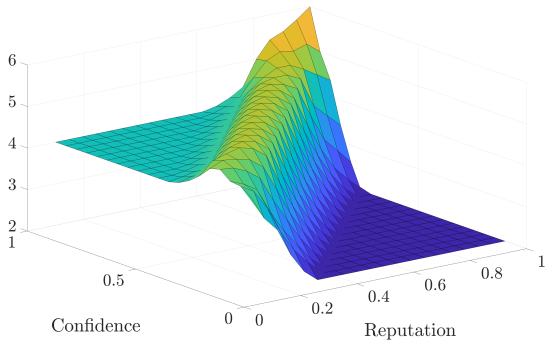
CB becomes more confident about the shocks hitting the economy, the alternative statements no longer need to span the whole space of possible shocks. If, for example, the CB is confident that inflation will be close to the target, then there is no need to draft alternatives discussing hyperinflation or deflation.

A visual illustration of this intuition is provided on the left panel of Figure 5. If the CB is highly confident that the shock was between [-1,1] (teal line), then they do not need to draft alternatives that discuss shock values of say -2 or 2. If the CB is not so confident (red line), 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.

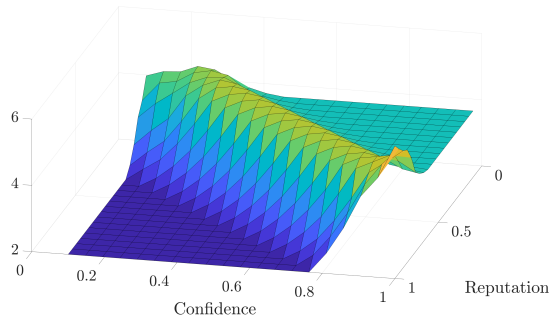
At the same time, the span is increasing in reputation. Let us try to dissect this. Leaning on the intuition conveyed by Figure 2, when reputation is below confidence, the announcements of the central bank are interpreted as less tight than what the bank intends. To get a particular message across, then, the CB needs to send a more precise message, which corresponds to a tighter interval. In such a case, the CB drafts alternatives as intervals with smaller width. As reputation increases and the market's expectations become more responsive to CB announcements, the CB can send wider intervals. For a fixed number of alternatives, a wider span of messages accomplishes this.

Figure 4: Count of Messages

(a) View 1



(b) View 2



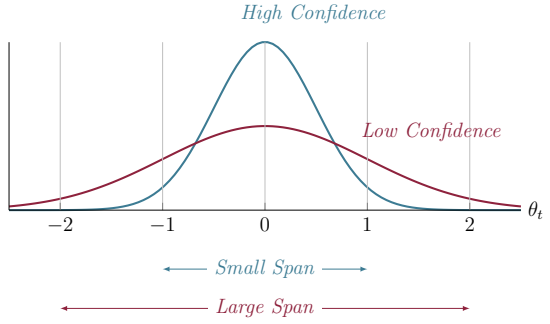
Note: The figure shows a stylized representation of the equilibrium number of alternative statements for varying levels of confidence (H_{t-1}) and reputation ($\bar{H}_{t-1}$). For the smoothed-out raw model output and versions with $b = 0.02$, see Appendix C.

Now looking at Figure 4, we see that the variation in the count is mainly attributable to whether we are below or above 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 above one, the CB always just drafts two alternatives, while when $h < 1$, the CB drafts four. In other words, when reputation exceeds confidence, fewer alternative statements are drafted than when reputation is below confidence.

The intuition is as follows. Consider a three-message equilibrium with a given span, as is shown on the right panel of Figure 5. Announcing Alternative A will generate a higher market expectation $\bar{\theta}_t$ when reputation is high (teal line) than when reputation is low (red line). If the CB has low reputation, achieving a higher $\bar{\theta}_t$ requires sending a tighter interval.

Figure 5: Intuition

(a) Confidence and Span



(b) Reputation and the Count

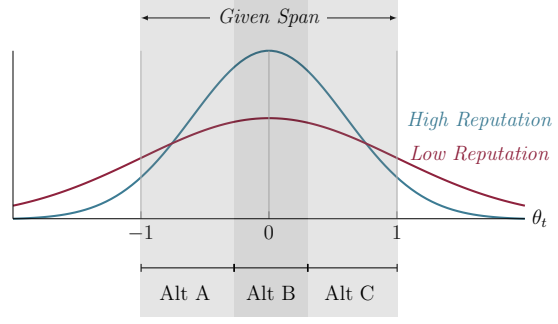
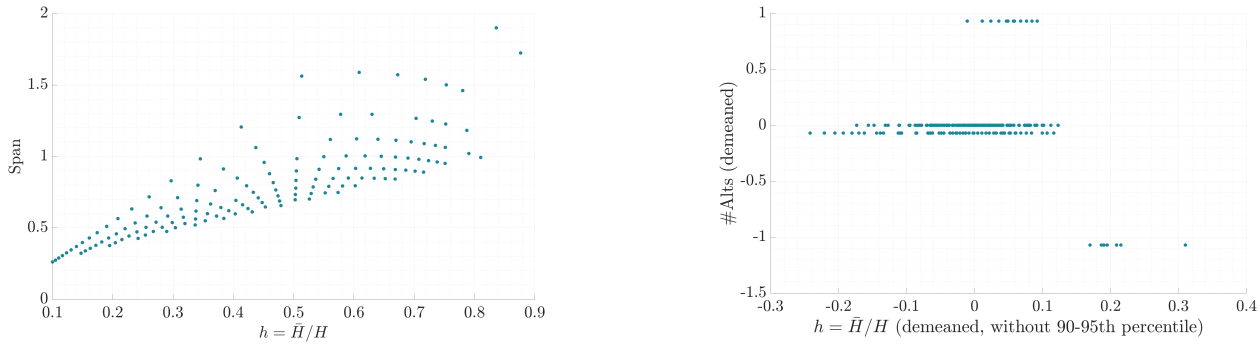


Figure 6: Span, Count and Reputation-to-Confidence Ratio (h)

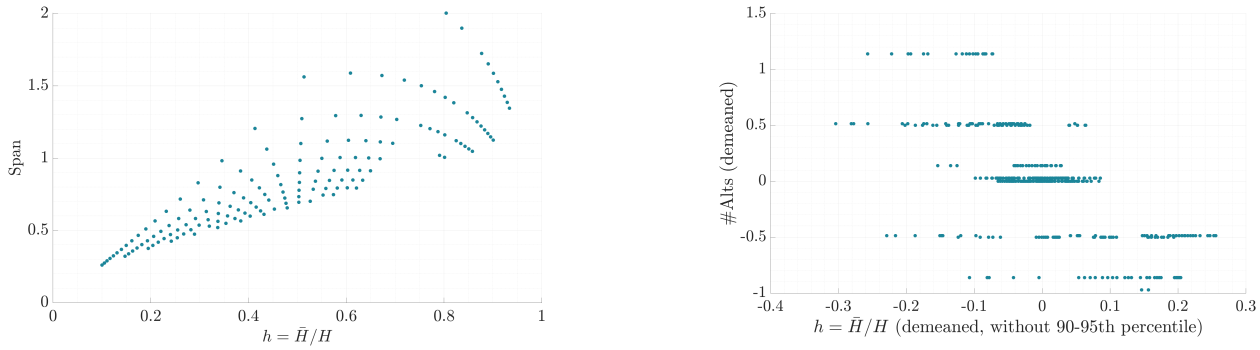
(a) Span against h, $b = 0$

(b) Count against $h, b = 0$



(c) Span against h, $b = 0.02$

(d) Count against $h, b = 0.02$



Note: The left column shows a scatter plot of the span, defined as the distance between the two outmost cutoffs ($\theta^{K-1} - \theta^1$), against the reputation-over-confidence ratio ($h \equiv \bar{H}/H$), excluding cases where $h > 1$ because those result in two-message equilibria where the span is not defined. The right column shows a scatter plot of the count of alternatives (demeaned) against the demeaned reputation-over-confidence ratio, excluding the 90-95th percentile.

For a given span, the count has to go up.

Figure 6 zooms in on the relationship between span, count and the reputation-over-confidence ratio, h. It plots scatter plots of span against h when $h < 1$ (so that the span is well-defined) on the left column, and scatter plots of demeaned count measures against h, excluding the 90-95th percentiles of the counts (the diagonal), on the right column.

Figure 6 corroborates the intuition from Figure 3 and 4: the span is positively correlated with h, while the count is either uncorrelated, or negatively correlated. This means that in a high reputation-over-confidence region, the CB drafts a low number of alternatives, while in a low reputation-over-confidence region, there will be more alternatives. Within that region, the higher confidence, the lower the span.

Overall, we see that both reputation and confidence affect the span and the count, for two reasons. First, because when either reputation or confidence changes, this affects h, determining which region we are in. Second, because both span and count affect the width of the candidate announcements, so that holding one fixed, one can adjust the other to get the desired width.

What is happening on the diagonal of Figure 4, or really close to it? Exactly on the diagonal, the conditions of Proposition 1 hold, meaning that the bank could release an equilibrium announcement with zero width, and it would be believed. While that is an equilibrium, the numerical solution only considers coarse message space equilibria. The high number of alternative statements, especially on the side of the diagonal where reputation is just below confidence reflects that on the one hand, in a point revealing equilibrium, there are infinitely many alternative statements, and on the other that arbitrarily close to point revelation, the CB drafts exceedingly precise messages to get as close to a point revelation outcome as possible. The high number of alternatives along the diagonal also raises the span somewhat to ensure that the alternatives are sufficiently distinct. However, along the diagonal, the span clearly increases when confidence drops, corroborating the previous intuition that, when more uncertain, the CB prefers to prepare draft statements for a wider variety of economic scenarios.

Thus one can summarize the predictions of the model for how confidence and reputation affect equilibrium communication with the following example. If the economy is calm, and the central bank is very confident that there are no sizable shocks to inflation, then they will only draft alternatives that consider inflation running at or close to the target. On the other hand, if the economy is volatile, the alternatives will need to cover a larger range of possible outcomes, for example discuss both a high-inflation scenario as well as a low-inflation one. How many alternatives to draft will depend on reputation. Markets will need the statement to be more precise the more uncertain they are, that is, the lower reputation is. In the volatile economy example, for a relatively high reputation, the bank may just draft three alternatives: one that concerns itself with inflation being below target, one with inflation on target, and one with inflation above target. If reputation is relatively low, the bank could draft for example five alternatives: a very low inflation alternative, a low inflation one, an

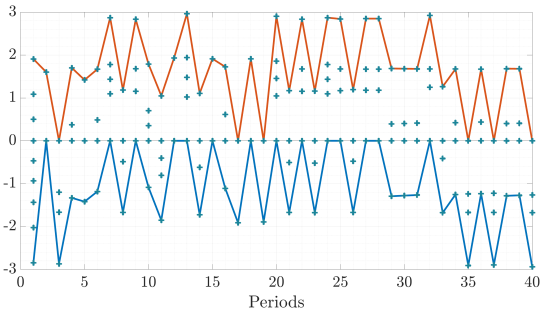
on-target one, a high inflation one, and a very high inflation one. This is just like the Harker quote from the introduction, where the idea was that with "elevated uncertainty regarding the appropriate path for the policy rate, [...] I can easily envision outcomes calling for a steeper pace of tightening. But I can also envision outcomes that would call for a prolonged pause."

6 Simulation

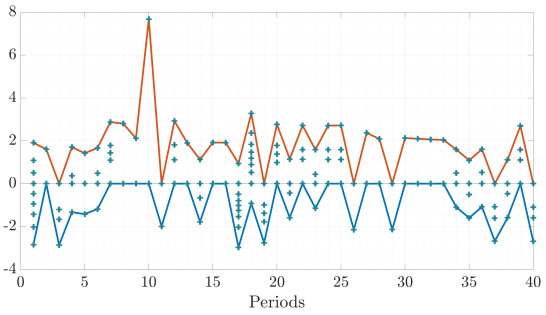
In this section we simulate the model to explore how the equilibrium message space evolves over time and to lead up to our regression analysis. We continue working with the calibration in Table 1, focusing on the case of zero output bias, and providing the positive output bias case in Appendix C, along 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 7: 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 7 shows how the message space evolves for an economy subject to the same sequence of shocks. The teal crosses represent each alternative statement, in particular,

the expectation induced by the alternative k, $\bar{\theta}_t^k$. We also use the normalization that the chosen message lies on the zero line. The red and blue lines connect the highest and lowest alternative messages in every period, respectively.

There are two main takeaways from Figure 7. 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 misinterpretation tension 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 G of the appendix, 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

This section details the text-based measurement procedure and how we use them to test the model predictions in the data. In particular, we quantitatively study how the FOMC's message space in the data varies with respect to measures for confidence and reputation. The FOMC's message space has easily measured characteristics from the alternative statements in the FOMC meeting materials: the number and similarity of alternatives. Alongside these, we create "uncertainty" measures a 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.

7.1 Message Space in Practice

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

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, which is captured by 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. Here we can see an example of the Fed using variation in word order and word choice to experiment with statements they could make, before settling on a specific one. This experimentation appears to allow the Fed to fine-tune their message to correspond exactly to what they want to say and the markets to understand.

We can 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 from words to a 768-dimension vector, called an embedding, that encodes the meaning behind the joint occurrence of words.[18]. We can then measure the differences between statements by measuring the differences between the statement vectors using Euclidean distance.[19]

This measure simplifies the comparison of alternatives down to a single dimension for approximating the span of the Fed's message space. In our model the distance between neighboring messages is proportional to the width of message intervals. In the data, we are able to measure the distance between messages and assume that this 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 the online appendix. This is evidence that the using the on dimensional distance measure captures our conceptual notion of message space span and spacing as envisioned in our model.

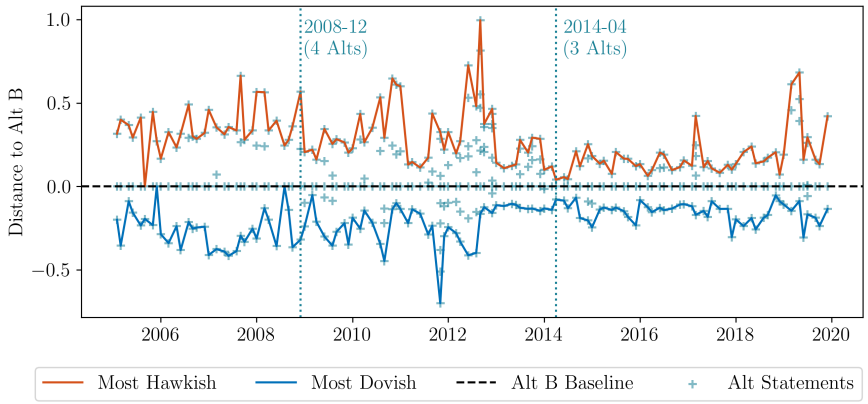
Figure 8 plots the alternative messages in relation to each other, normalized around Alt. B. The orange line shows the relative position of the most hawkish alternative (Alt. C, or the last available letter), and the blue line shows the relative position of the most dovish alternative (Alt. A). This ordering, from hawkish to dovish, is taken directly from the Tealbooks and is assigned by the 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, it represents that the Fed is considering a wider range of messages at that meeting. Later, we will use the distance between the most hawkish and most dovish alternatives to approximate the span of the Fed's message space. The span comes from adding the relative distances to alternative B.[20]

[18] Appendix E includes more background information on the alternative policy statements and the full text of alternative statements from the December 2008 FOMC meeting. Appendix F includes a detailed discussion of text processing and analysis.

[19] We also provide 1-cosine similarity versions of the graphs in Appendix F.

[20] In Appendix F, we provide robustness analysis to this measure where we simply directly measure the

Figure 8: 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. Crosses, +, represent the other alternative statements for that meeting.

In Section 2 we plot the number of alterative statements at each meeting Figure 1. In Figure 8, the number of alternatives is captured as the points, +, plotted between the hawkish and dovish contours. Here we already can already see interesting patterns of Fed communication: the number and distances of alternative statements seem weakly correlated, but exhibit distinct variation. Accordingly, we include both in our analysis below.

These two features of the set of alternatives statements - the number of alternatives and the distance between the alternatives - capture different variations which are useful to study the monetary policy communication strategy. When alternative statements are farther apart, then that implies the Fed is choosing from more opaque statements. Conversely, when the alternative statements are closer together, it reflects that policymakers are trying to make a more nuanced point. When the Fed staff provide many alternatives it could reflect either uncertainty over the economic situation if the alternatives are very different or uncertainty over how the public will respond to their words if the alternatives are closer together. We think of the set of alternative statements for the Fed as a function of those two dimensions, as in the model, and test them later in Section 7.

dissimilarity between the extreme alternatives (the first A and the last C), and we find the same qualitative results.

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 and reputation represent the precision of the central bank's signal which is the inverse of the 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 measure of their 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 from the FOMC meeting transcripts relative to the total amount of speech from policymakers. Again, the interpretation of the measure is that the more the FOMC says they are uncertain or the more newspapers link uncertainty and the Fed, then that represents higher variance (or perceived variance) of the Fed's signals, and therefore lower confidence (or reputation). Therefore, we multiply the indices by -1 to create our confidence and reputation series.

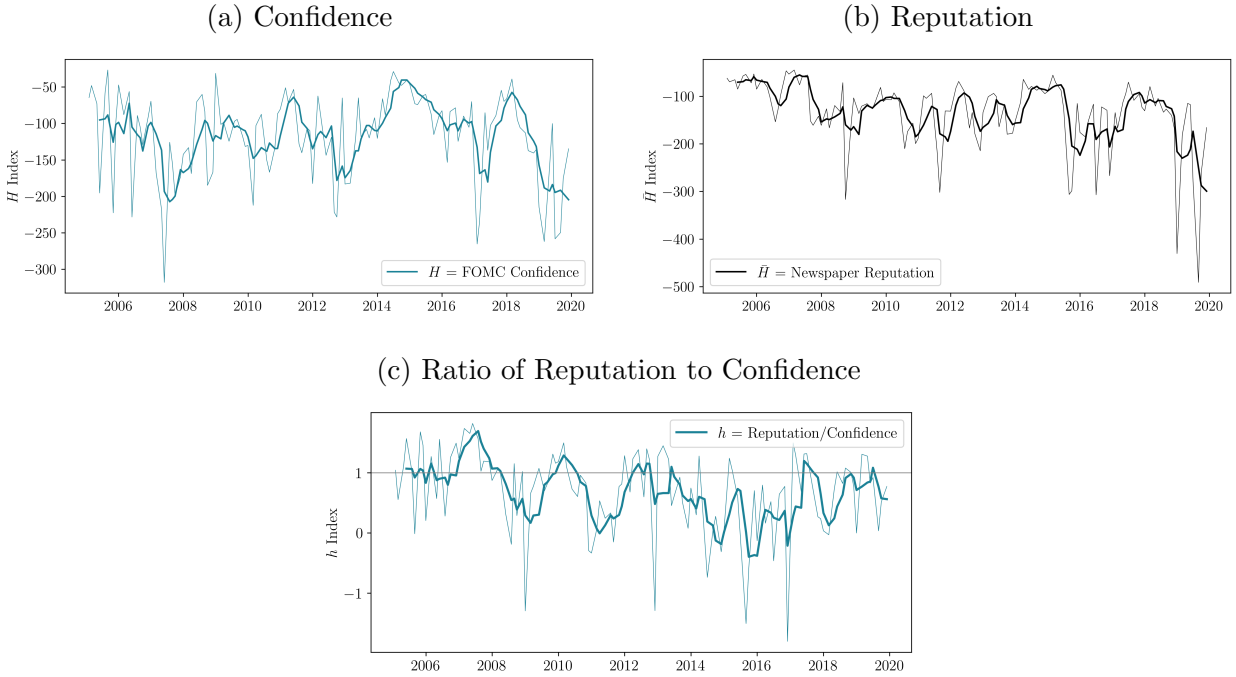
While the Baker et al. (2016) measure of uncertainty has been validated through its use in many research papers over the past decade, our text measure for confidence needs explicit validation. In Appendix F we validate the uncertainty-word-count measure of FOMC transcripts against a question in the *Summary 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, that is, every-other FOMC meeting, since mid-2007. We find that our text-based transcript measure accounts for 90 percent of the variation in the number of FOMC members who say they are more uncertain, and directly 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 9 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 the two indices as a data proxy for our model's h variable, reputation-to-confidence. All series $H, \bar{H}, h$ are stationary.

In these plots we have two interesting findings: first, there is greater volatility in confidence than reputation meeting to meeting. This backs up our modeling choice to have H probabilistically change from period to period. Second, we find that changes in H lead changes in $\bar{H}$, which supports the idea that the public is learning about the FOMC's confidence over time. Third, we find that the h fluctuates between positive and negative values

Figure 9: Text-Based Measures of Confidence and Reputation



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, and is also multiplied by -1 so that higher values represent greater reputation/confidence and shifted up two 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.

throughout our sample. In our model, when h is close to one this leads to spikes in the number of alternatives. 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. To go beyond this observational connection, in the next section we statistically test for the relationships between confidence, reputation and the Fed's message space.

7.3 Baseline Specification: OLS Approach

We measure the effect of changes in confidence and reputation on the communication strategy, as approximated by the message space measures, with two empirical approaches. First, we consider a simplified specification to understand the contemporaneous effects (Equation 17). In the next section we estimate an SVAR to allow for more dynamic relationships.

The simplified specifications are as follows:

$$M_t = \alpha + \beta_1 h_{t-1} + \beta_2 C_t + \beta_3 \Delta vix_{t-1} + \varepsilon_t, \tag{17}$$

$$M_t = \alpha + \beta_1 H_{t-1} + \beta_2 \bar{H}_{t-1} + \beta_3 H_{t-1} \times \bar{H}_{t-1} + \beta_4 C_t + \beta_5 \Delta vix_{t-1} + \varepsilon_t, \tag{18}$$

where t indexes the scheduled FOMC meeting and $M_t \in \{count_t, span_t, span_t/count_t\}$. We lag the 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. Additionally, we standardize, or z-score, the data so the intercept term α is zero.

We include select controls in these parsimonious specifications. We control for the count (span) of alternatives when regressing the span (count), indicated as C_t. These terms are positively correlated; more messages are often associated with greater span of messages. Finally, we include the monthly change in a rolling average of the VIX for the month prior

Table 2: Message Space Measures, Confidence and Reputation

	(1) $count_t$	(2) $count_t$	(3) $count_t$	(4) $span_t$	(5) $span_t$	(6) $span_t$	(7) $\frac{span_t}{count_t}$	(8) $\frac{span_t}{count_t}$	(9) $\frac{span_t}{count_t}$
h_{t-1}	-0.192**			0.310***			0.362***		
	(0.084)			(0.076)			(0.088)		
H_{t-1}		0.046	-0.619		-0.313***	-0.499		-0.360***	-0.543
		(0.080)	(0.743)		(0.083)	(0.682)		(0.095)	(0.783)
$\bar{H}_{t-1}$		-0.314***	-0.935		0.236**	0.061		0.284**	0.112
		(0.103)	(0.766)		(0.098)	(0.659)		(0.114)	(0.756)
$H_{t-1} \times \bar{H}_{t-1}$			-1.070			-0.300			-0.296
			(1.227)			(1.053)			(1.214)
$span_t$	0.573***	0.518***	0.512***						
	(0.121)	(0.121)	(0.120)						
$count_t$				0.529***	0.516***	0.514***			
				(0.059)	(0.055)	(0.057)			
Δvix_{t-1}	-0.134**	-0.183***	-0.179***	0.095	0.085	0.086	0.112	0.105	0.106
	(0.053)	(0.055)	(0.056)	(0.088)	(0.087)	(0.087)	(0.101)	(0.100)	(0.100)
R^2	0.312	0.364	0.368	0.365	0.366	0.366	0.134	0.135	0.136
R^2 Adj.	0.295	0.342	0.340	0.348	0.344	0.339	0.119	0.113	0.106
BIC	306.468	301.992	305.968	297.047	301.578	306.294	329.132	333.745	338.479
F-stat	7.742	6.791	5.565	37.360	30.174	23.891	8.458	5.725	4.441
F p-value	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001	0.001	0.002
N	119	119	119	119	119	119	119	119	119

Notes: HAC-robust standard errors with small sample correction in parentheses. * p<.1, ** p<.05, *** p<.01. Series are all log-transformed and standardized (z-scored). $\bar{H}$ reputation is -1(Monetary Policy Uncertainty (MPU) Index) from Baker et al. (2016). Using the top 10 U.S. newspapers, the MPU index is the frequency of articles containing terms related to "uncertainty" and "monetary policy" words, scaled so the 1985-2010 average is 100. The H is the Fed confidence measured as -1(FOMC Uncertainty Index). This index is the number of uncertainty words in the FOMC meeting transcripts spoken by FOMC members divided by total words and is scaled so the 1985-2010 average is 100. Then h_{t-1} is 2-1(MPU Index/FOMC Uncertainty Index). $t-1$ for monthly variables is the month prior to the FOMC meeting t, where as H_{t-1} is measured from the previous FOMC meeting. Regression sample covers 2005-2019.

to the FOMC meeting where we measure M_t as a control for aggregate uncertainty.[21] This control accounts for variation in confidence and reputation that are both driven by aggregate uncertainty.

We find empirical results that are consistent with the theory: lower reputation and higher confidence are associated with more precise messages. Before we detail the results, at a high level these patterns correspond to what our model predicts. Our findings suggest that the central bank's reliance on coarse communication reflects not only the broad picture implied by the model but also its individual predictions, as summarized in the previous sections and in Appendix G. In particular, the data are consistent with being in the region where $H > \bar{H}$: communication span decreases as confidence rises, while it increases with the bank's concern for reputation. Moreover, when $h < 1$, we observe Tealbooks with more alternative

[21] We use a two month rolling average to capture persistent uncertainty changes in the inter-meeting period.

statements than when $h > 1$, and the span of communication is positively correlated with h. Finally, the count of statements is negatively correlated with h, further reinforcing the interpretation that coarse communication is driven by strategic considerations about how messages will be interpreted.

Now we discuss the empirical results more in more detail. First, the count of alternative statements is negatively correlated with our measure of reputation. Because we control for the span of messages, we interpret the increased count as representing slightly more precise messages with a fall in reputation. This also shows up directly in specification (8) where reputation is positively correlated with $span_t/count_t$. Recall from the model that more precise messages are desirable for a low reputation environment because this is when public expectations are less sensitive to announcements overall. This causal argument relies heavily on the timing assumptions, which we relax in the next section when we estimate the SVAR.

To interpret scale, the coefficient in the column (2) implies that we would need a large increase in reputation, about three standard deviations, in order to decrease the number of alternatives by one from the sample average of three alternatives. A jump of this magnitude is rare in the data, but we do see a few instances in 2019. If we consider the period leading up to the end of explicit zero-lower-bound (ZLB) forward guidance (November 2011 - December 2012), there is an average of six alternatives per meeting. Here, it would take a one standard deviation increase in reputation to predict a decrease of one alternative statement. Some examples of this level of change in reputation would be like the difference in monetary policy uncertainty in news at the end of 2008 and 2011 where, respectively, the Fed went to the ZLB for the first time and used more explicit forward guidance (stating we would stay at the ZLB until mid-2013).

Second, the span of scenarios drafted decrease when the FOMC's confidence increases. That is, we find that the FOMC considers more precise announcements when they are less uncertain because we control for the number of alternatives in the regressions. This is consistent with our theoretical predictions. A one-standard-deviation increase in confidence decreases the span of alternatives by 31 percent. A one-standard deviation increase in reputation, on the other hand, would be associated with an 24 percent increase in the span, but this is offset if we allow for reputation to interact with confidence as in column (6). How

do we conceptualize the change in span? Referring back to the examples from Section 7.1, the span of messages for December 2008 is about 200 percent larger than the span in April 2014.

What is really novel about these results is showing that variation in the public's opinion of the Fed affects the precision of communication. In the model this is driven by anticipating the sensitivity of the public (again where higher reputation means greater sensitivity of expectations to central bank announcements). In the data, we find that the reputation-to-confidence ratio has a correlation that is consistent with the model: higher reputation relative to confidence is associated with fewer alternative announcements that cover a larger span. Although not all estimates are statistically significant, the point estimates point in the right direction. Interestingly, the dynamics in the data seem most consistent with model dynamics where reputation is below confidence in level terms. This is consistent with the concept that financial markets are opinionated and think less of the the Fed's ability to forecast compared to their own (Caballero and Simsek, 2022).

When we separately look at reputation and confidence, we still find evidence that a coarser message space from the Fed is caused by a mix of the Fed's uncertainty and their reputation. In particular, it looks like the number of alternatives more closely varies with reputation and aggregate volatility while the span of alternatives has a stronger relationship with confidence. Comparing the specifications with and without the interaction between reputation and confidence, we see the sign of the reputation coefficient and interaction coefficient are opposite: inducing an attenuated effect for reputation when confidence is higher. While the single-equation regressions is useful for interpreting contemporaneous effects, it imposes assumptions that do not allow for the endogenous relationships between confidence, reputation, and the communication strategy over time which we relax in the next section.

7.4 Dynamic Specification: SVAR Approach

In this section we allow all variables to be endogenous and depend on lags. In particular, let the vector of endogenous variables be

$$Y_t = [\Delta vix_{t-1},\ H_{t-1},\ \bar{H}_{t-1},\ span_t/count_t,\ tffr_t], \tag{19}$$

and the dynamic behavior of Y is captured by

$$Y_t = B(L)\ Y_{t-1} + \varepsilon_t, \tag{20}$$

where $B(L) = B_1 L + B_2 L^2 + \cdots + B_p L^p$ is a lag polynomial of order p, and ε_t denotes the structural shocks. The reduced-form representation is below, where u_t are the reduced-form residuals with covariance matrix $\Sigma_u = E[u_t u_t']$ and $A(L)$ is another lag operator:

$$Y_t = A(L)Y_{t-1} + u_t. \tag{21}$$

The ordering of variables matches that described in Equation 19 and was chosen to be consistent when the variables are observed. We assume there is an aggregate level of uncertainty captured by the change in the rolling average of the VIX which may affect other variables in the system.[22] We impose the assumption that H may affect $\bar{H}$ on impact – motivated by the timing of the two measures and to capture the public potentially learning about the Fed's confidence through announcements on either policy or its uncertainty – but we do not allow reputation to affect confidence on impact. This means we are ruling out situations where the public, newspapers, or other actors could pressure the Fed into changing their signal precision. However, we do allow the Fed to respond to the reputation effects in their messages and policies. Next is the message space variable: the span divided by the count of messages. This represents the average precision of messages for a given meeting.[23] Finally, the target policy rate, $tffr_t$, is chosen – this timing is consistent with the fact that the alternatives are drafted the week before the policy rate is voted on in the meeting.

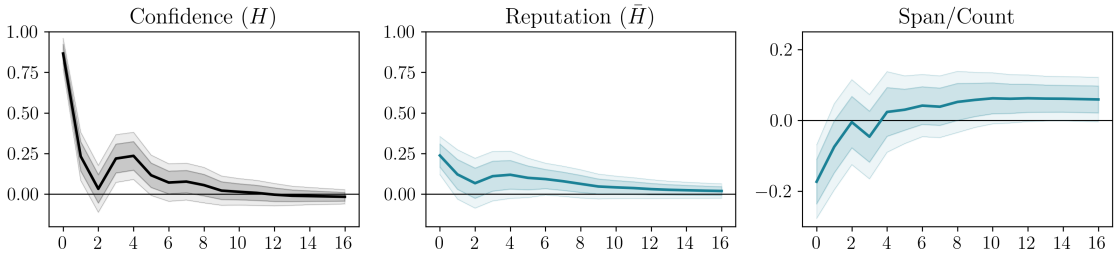
We estimate the reduced-form VAR using OLS with four lags (approximately six months). Structural parameters are recovered through the Cholesky decomposition of the residual covariance matrix Σ_u, where $\Sigma_u = PP'$ and P is the lower-triangular Cholesky factor. We compute orthogonalized impulse response functions are computed for up to 16 FOMC meetings ahead (approximately eight quarters), with standard errors obtained analytically

[22] Because we use a two month rolling average, there is information from prior to the last FOMC meeting through the month prior to the current FOMC meeting in vix_{t-1}.

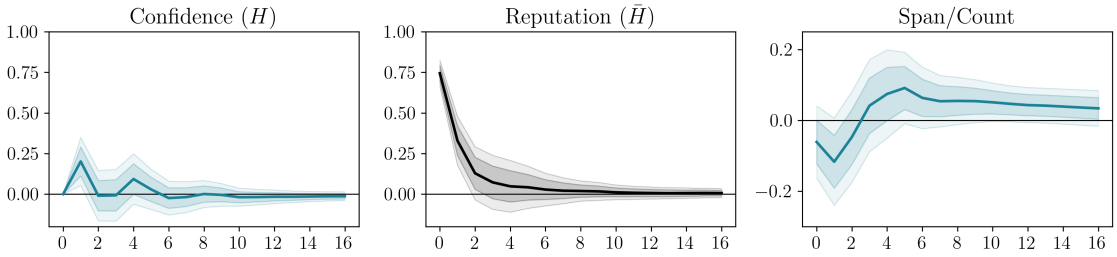
[23] In Appendix G we provide specifications with count and span as separate variables.

Figure 10: Impulse Responses to Confidence and Reputation

(a) Responses to Confidence Shock



(b) Responses to Reputation Shock



Note: Figures report impulse responses to a one standard deviation impulse to confidence or reputation. Confidence intervals at 90 percent (lighter area) and 68 percent (darker area) levels are shown. Periods represent FOMC meeting frequency. There are eight FOMC meetings per year, thus 16 periods corresponds to two years. For simplicity we suppress the plots for the VIX and TFFR here, but they are available in Appendix G

from the VAR coefficient estimates. When searching for the optimal number of lags, we found between one to two lags were optimal depending on the criterion. However, we double the lag size to four periods as our baseline specification to absorb dynamics of the system (Olea et al., 2025).[24]

How do the message space metrics respond to shocks to confidence and reputation in the data? In short, exactly as our model implies. Figure 10 summarizes the impulse responses where we allow all variables to respond dynamically. The dynamics of confidence and reputation show that shocks die out fairly quickly, but there is a positive correlation indicative of the public's learning of confidence. An unexpected, one-standard deviation increase in confidence increases reputation by 0.25 standard deviations. The pass-through is statistically significant on impact and it is far less than one-to-one. Again supporting that

[24]We provide robustness to different lags in Appendix G.

confidence – or uncertainty – is not precisely conveyed to the public. In plot (b), we also see a small positive increase in confidence after an anticipated change in reputation for one meeting. The comovement of reputation and confidence poses an interesting implication for the message space, because theoretically they create opposing effects for the message space.

Ultimately, we find that confidence is the main driver of patterns in the message space. The $span_t/count_t$ falls on impact after an unanticipated increase to confidence that quickly returns to zero. The small bump in reputation is not enough to offset the effect from the impulse to confidence. On the other hand, following a reputation shock, there is not a statistically significant change in the $span_t/count_t$. When looking at the point estimates, the pattern is also less clear. There is an initial drop in $span_t/count_t$ that corresponds to how confidence increases after the reputation shock. After that, $span_t/count_t$ rebounds to a peak at about five periods. Theory would suggest that reputation increases should lead to wider messages with larger $span_t/count_t$. But we do not see that in these impulse response plots. Compared to the model, in the data we see an increase in reputation – that is a decrease in the public's monetary policy uncertainty – causes attenuated reactions in the message space features compared to the model.

We also document a few ancillary findings on the effects of shocks to aggregate uncertainty (the VIX) and the policy decision, which corroborate findings from the literature. A shock to the VIX decreases confidence, reputation, and the target interest rate. Because we define confidence and reputation to be the inverse of an uncertainty word-count index, then those series should theoretically be negatively correlated with aggregate uncertainty. And indeed, that is what we find. The negative response of the policy rate to an increase in uncertainty is likely picking up dynamics from the Great Financial Crisis where the policy rate goes effectively to zero following the pikes in uncertainty of the crisis. More broadly, this negative relationship between uncertainty and policy rates is commonly found in the literature. Thus, we see this as a result that matches the benchmark established from the literature.

Finally, following shocks to confidence or reputation there is an increase in the target policy rate. This is an interesting link to the literature on central bank credibility which ultimately equates credibility with preferences for high interest rates. Here, our measure for confidence is our notion of a credible central bank (similar to Moscarini (2007)). Recall that

a high-confidence bank is not incentivized to create inflation surprises because it can better minimize its loss through forward guidance and hitting its policy instrument target. Therefore, a high-confidence bank is more likely to have higher interest rates (lower inflation) than a low-confidence bank. However, this dynamic is not reciprocal. Following a contractionary policy rate shock, there is no change in reputation but there is a sizable drop in confidence. Although out of the scope for this paper, we think this asymmetry garners future research.

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.

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 (2025) "Identifying Monetary Policy Shocks: A Natural Language Approach," *Working Paper*.

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 (2020) "Deciphering Federal Reserve Communication via Text Analysis of Alternative FOMC Statements," *The Federal Reserve Bank of Kansas City Research Working Papers*.

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 Uncertainty," *Working Paper*.

Granziera, Eleonora, Vegard H Larsen, Greta Meggiorini, and Leonardo Melosi (2025) "Speaking of 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.

Istrefi, Klodiana and Anamaria Piloiu (2020) "Public Opinion on Central Banks when Economic Policy is Uncertain," *Working Paper*.

Kamenica, Emir and Matthew Gentzkow (2011) "Bayesian Persuasion," *American Economic Review*, 101 (6), 2590–2615.

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.

Kostyshyna, Olena and Luba Petersen (2023) "The Effects of Communicating Inflation Uncertainty on Household Expectations," *Bank of Canada Working Paper*.

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.

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.

Olea, José Luis Montiel, Mikkel Plagborg-Møller, Eric Qian, and Christian Wolf (2025) "Local Projections or VARs? A Primer for Macroeconomists," *NBER Working Paper Series* (w33871).

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 (2025) "Disagreement About Monetary Policy," *Working Paper*, 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.