Louder than Rates

The Systematic Nature of Central Bank Communication

Tiziana Assenza
Toulouse School of Economics, University of Toulouse Capitole 1

Fabrice Collard
Toulouse School of Economics, CNRS & CEPR

Dogukan Guney
Toulouse School of Economics, University of Toulouse Capitole 1

Philipp Wangner
University of Mannheim

October, 30, 2026

Abstract: Do central banks decide systematically how much to communicate when explaining their policy decisions? Using all U.S. Federal Open Market Committee policy statements since 1994, we measure communication effort through the change in Shannon entropy and estimate a forward-looking communication rule. We find that communication is systematic: the Federal Reserve communicates more when inflation is expected to exceed target and output is expected to fall below potential. This finding is robust across a variety of sensitivity exercises. We then develop a New Keynesian model with imperfect information showing that systematic communication acts as a second policy instrument, stabilizing expectations and complementing interest-rate policy, especially at the zero lower bound.

Keywords: Central Bank Communication, Monetary Policy, Systematic Rules.

First Version: April 2025. We would like to thank Klaus Adam, Guido Ascari, Alessandro Pavan and participants to seminar and conferences at Computing in Economics and Finance 2026 in Venezia, FEP Porto, Università Cattolica in Milano, University of Florence, 7^{th} Behavioral Macroeconomics Workshop, Heterogeneity and Expectations in Macroeconomics and Finance (University of Bamberg) for helpful comments. We also thank Joaquim Arroyos Serra and Letizia Sormani for excellent research assistance. Assenza and Collard acknowledge funding from the French National Research Agency (ANR) under the Investments for the Future (Investissements d'Avenir) program, grant ANR-17-EURE-0010 and funding from the French National Research Agency (ANR) under the project ANR-23-CE26-0011-01. Wangner thanks the UniCredit & Universities foundation for financial support. *Author contact information*: Tiziana Assenza, tiziana.assenza@tse-fr.eu, https://tizianaassenza.wixsite.com/tizianaassenza; Fabrice Collard, fabrice.collard@tse-fr.eu, http://fabcol.free.fr; Dogukan Guney, dogukan.guney@tse-fr.eu, https://sites.google.com/view/dogukanguney; Philipp Wangner, philipp.wangner@uni-mannheim.de, https://www.pwangner.com.

Introduction

Since the early 1990s, central banks, in particular the Federal Reserve, have steadily expanded the role of communication alongside their conventional instrument, the policy rate. Press conferences, forward guidance, summary of economic projections, central bankers' speeches and, at the heart of the apparatus, the post-meeting policy statement have become the principal channels through which the Federal Reserve explains and justifies its decisions. The importance of communication is by now hard to overstate. As Bernanke (2015) put it, "*When I was at the Federal Reserve, I occasionally observed that monetary policy is 98 percent talk and only two percent action.*" A large literature confirms that words move asset prices, expectations and, ultimately, macroeconomic outcomes (see Blinder, Ehrmann, Fratzscher, de Haan, and Jansen, 2008, Blinder, Ehrmann, de Haan, and Jansen, 2024, for surveys). Yet much less is known about how central banks determine the extent of their communication effort. This paper addresses this gap by examining two related questions. Do central banks decide, analogously to the policy rate (Taylor, 1993, Clarida, Galí, and Gertler, 2000), in a systematic way how much explanation to offer? And if they do, how does the systematic part of their communication shape the propagation of shocks, macroeconomic stability and the scope for stabilization policy at the zero lower bound?

We address these questions in two steps. First, we construct a novel text-based measure of central bank communication effort and use it to estimate a rule that captures its systematic component. For the Federal Reserve, systematic communication increases when expected inflation exceeds an estimated anchor and output is expected to fall below potential. Second, we integrate the estimated communication rule into an imperfect-information New Keynesian model. Importantly, the inflation and output components of the rule exert opposing forces on the cyclicality of communication: the inflation response tends to make communication procyclical, whereas the output response makes it countercyclical. Under our estimates for the Federal Reserve, the countercyclical force dominates by a substantial margin. Countercyclical communication then weakens the effect of expected future activity on current demand, relaxes the conventional Taylor principle, attenuates demand shocks, increases the output and inflation response to cost-push shocks, and can shorten or even eliminate zero lower bound episodes. Our central contribution is therefore to show that the *cyclicality of systematic central bank communication* is a key determinant of monetary-policy transmission, shock propagation, and macroeconomic stability, both away from and at the zero lower bound.

Measuring Communication Effort We measure variation in communication using the Shannon entropy (Shannon, 1948) of the word distribution in the non-technical portion of all FOMC statements issued since 1994. Entropy does not classify a statement as hawkish or dovish, does not measure sentiment, and does not indicate whether the central banks assessment is correct; rather, it is rooted in information theory and provides the standard measure of lexical diversity. Entropy is higher when word use is distributed across a more diverse vocabulary and lower when language is more repetitive. Our baseline measure of central bank communication is the quarterly change in the logarithm of entropy between consecutive FOMC statements. We interpret this measure as a proxy for changes in the

amount of linguistic elaboration in the statement and refer to it as *communication effort*, as it measures the extent to which monetary policymakers put extra effort in explaining their current decisions. More importantly, we validate this measure by showing that changes in entropy follow a stable, forward-looking rule that recovers the Federal Reserves announced inflation target from text alone - patterns that are difficult to reconcile with a purely mechanical byproduct of statement construction. Overall, the time series of communication effort displays a consistent pattern: communication rises sharply at the onset of recessions, particularly during the Global Financial Crisis, remains moderately elevated during much of the zero-lower-bound period, and fluctuates around zero in more tranquil periods.

A Communication Rule As our baseline specification, we then estimate a forward-looking communication reaction function in the tradition of Taylor (1993) and Clarida, Galí, and Gertler (2000), using iterated GMM with internal instruments over the sample 1994Q2–2022Q4. We obtain five main findings. First, the data do not reject the over-identifying restrictions, consistent with a systematic rule governing communication effort. Second, the estimated inflation anchor (1.93%) is close to the Fed's announced 2% target, even though it is identified entirely from statement text. This correspondence suggests a form of internal consistency between the Federal Reserves interest-rate and communication rules. Third, the response of communication effort to expected inflation is large and precisely estimated: a one-percentage-point upward deviation from target raises lexical diversity by approximately 7%. Fourth, the Committee communicates more when output is expected to fall short of potential: a one-percent shortfall raises communication effort by about 1.5%. Fifth, the inertia parameter is small and statistically insignificant (0.07), indicating that systematic communication responds primarily to the economic outlook rather than to its own recent history.

These baseline findings are robust to several changes in measurement and specification. Across alternative inflation measures, output-gap proxies, and forecast horizons, and when Tealbook forecasts replace rational-expectations forecasts, the inflation response remains positive and precisely estimated. The estimated anchor moves with the inflation measure, lying between roughly 1.8% and 1.9% for PCE measures and somewhat higher for CPI measures. The output response retains its sign but is less precisely estimated in some specifications. Relying on an external instrument strategy using oil-supply news from Känzig (2021) and high-frequency monetary-policy surprises from Nakamura and Steinsson (2018) yields similar estimates of the inflation anchor and a positive inflation response. These estimates also suggest that the output response varies with the source of fluctuations and becomes insignificant when identifying variation originates from oil-supply news. Finally, replacing entropy with the growth rate of statement length reduces the estimated inflation response by roughly one half and eliminates the output response, suggesting that statement length alone cannot account for the documented state dependence in communication effort. In other words, the Federal Reserve's dual mandate is carried in the intensive margin of communication effort.

A natural question is whether the documented systematic nature of communication is specific to the Federal Reserve, or whether it reflects a genuine international regularity that would emerge from the data of any central bank. To assess this question, we estimate analogous communication rules for the European Central Bank, the Bank of Japan, the Bank

of Canada, the Reserve Bank of Australia, and the Riksbank. For most of these central banks, communication rises with expected inflation and, to a lesser extent, when activity is expected to fall below potential, while the inferred inflation anchor lies near the announced target. At the same time, the estimates reveal important institutional heterogeneity. Consistent with its primary price stability mandate, communication by the ECB does not respond significantly to the output gap; the Bank of Canada responds primarily to the stance of U.S. monetary policy; and the Reserve Bank of Australia and the Riksbank communicate more when their currencies are expected to depreciate. Taken together, these findings motivate treating communication as a systematic policy margin: it is state-dependent across diverse monetary-policy environments, but the determinants of communication vary with institutional mandates and economic circumstances rather than conforming to a common, invariant rule.

A New Keynesian Model with a Communication Rule Having characterized the communication rule empirically, we next examine how systematic communication affects the propagation of shocks, macroeconomic stability and the scope for stabilization policy, both away from and at the zero lower bound. To this end, we develop a Communication Representative Agent New Keynesian (CRANK) model featuring two structural disturbances, demand shocks and cost-push shocks. Each disturbance comprises persistent and transitory components. The key departure from the benchmark full-information rational expectations RANK model (Woodford, 2003b, Galí, 2015) is informational. As in the imperfect-information New Keynesian literature (Woodford, 2003a, Lorenzoni, 2009, Blanchard, L'Huillier, and Lorenzoni, 2013) households observe the shocks but cannot perfectly infer their persistence. By contrast, the central bank observes the decomposition and provides a public signal about the persistent components. Greater communication effort increases the precision of this signal. Monetary policy is governed simultaneously by two systematic feedback rules: a Taylor rule for the nominal interest rate and a communication rule that determines signal precision. Together with the dynamic IS equation and the New Keynesian Phillips curve (NKPC), these rules characterize the models aggregate dynamics in a parsimonious system of just four equations.

In this setting, communication affects aggregate dynamics through the DIS equation via two channels. The first is a mean-distortion channel, whereby greater communication effort improves signal precision and thereby attenuates households systematic underreaction to current shocks. The second is a risk channel, whereby greater communication effort reduces posterior uncertainty, thereby lowering precautionary saving and directly stimulating aggregate demand.

We derive three main theoretical results in normal times, i.e., for economies away from the zero lower bound (ZLB). *First*, we show that communication and the real interest rate are locally substitutable policy instruments because both operate through aggregate demand: the interest rate through intertemporal substitution and communication through precautionary saving. *Second*, we establish a close link between the net cyclicality of systematic communication and equilibrium stability. The key to this characterization is that the inflation and output feedback channels operate in opposite directions. On the one hand, higher expected inflation calls for greater communication effort and, through the NKPC,

makes communication procyclical, thereby compounding the effects of expectations in the DIS equation. On the other hand, higher expected output calls for lower communication effort, making communication countercyclical and thereby attenuating the effects of expectations in the DIS equation. Because greater communication effort is expansionary, countercyclical communication helps stabilize aggregate fluctuations and thereby relaxes the Taylor principle: a weaker interest-rate response to inflation may suffice to ensure determinacy. Conversely, procyclical communication reinforces fluctuations and tightens the corresponding determinacy condition. *Third*, we show that cyclical communication shapes the propagation of DIS and cost-push shocks asymmetrically. Under a DIS shock, countercyclical communication acts as an automatic stabilizer, whereas procyclical communication acts as an automatic amplifier. By contrast, cyclical communication increases the output and inflation responses of positive cost-push shocks regardless of whether communication is procyclical or countercyclical. The reason is that systematic communication causes cost-push shocks to propagate partly like endogenous aggregate-demand shocks. When expected inflation rises, central banks increase communication effort, which stimulates aggregate demand and further raises inflation. The net cyclicality of communication – whether it compounds or discounts expectation-driven dynamics in the DIS equation – only determines then the magnitude, rather than the direction, of this additional expansionary effect.

Finally, we illustrate our theoretical findings by calibrating the model to the US economy. Under the estimated communication rule for the Federal Reserve, systematic communication is countercyclicality and gives rise to a discounted DIS equation, making current demand less than one-for-one responsive to expected future output. Relative to the extensive literature providing alternative micro-foundations for a discounted DIS equation (see, among others, Gabaix, 2020, Del Negro, Giannoni, and Patterson, 2023) our contribution is to show that such discounting can emerge from the systematic conduct of countercyclical communication policy. Importantly, this mechanism alters the conduct of monetary stabilization away from the zero lower bound. Following a contractionary demand shock, the estimated rule increases communication as the economy weakens, while the policy rate falls. Communication therefore absorbs part of the stabilization burden. In the calibrated model, the policy rate is cut by less, and the output loss is approximately 25 percent smaller at every horizon. On impact, output declines by 1.9 percent, compared with 2.6 percent in the model without state-dependent communication. Following a cost-push shock, the two instruments instead operate in opposite directions: the policy rate rises to fight inflation, while the inflation component of the communication rule increases information provision and supports demand. Communication reduces the impact output decline from 2.3 to 0.7 percent, at the cost of a modest increase in annualized inflation, from 3.8 to 3.9 percent.

The role of systematic communication becomes even more pronounced at the ZLB. In the standard model, a policy rate constrained at zero can admit self-fulfilling deflationary equilibria (Benhabib, Schmitt-Grohé, and Uribe, 2001, Eggertsson and Woodford, 2003). Sufficiently countercyclical communication, as implied by our estimated rule for the Federal Reserve, eliminates these equilibria: expectations of weaker economic activity induce greater communication effort, which supports current demand and prevents pessimistic expectations from becoming self-fulfilling. By contrast, an inflation-only communication rule, as implied by our estimated rule for the ECB, has the opposite property: lower expected

inflation calls for less communication and reinforces the contraction. In our numerical example for the US economy, economies with and without systematic communication both reach the lower bound following a large negative demand shock, but communication reduces the duration of the episode from eight quarters to four and limits the output decline at the trough by approximately 30 percent. Following a smaller shock, systematic communication keeps the policy rate above zero throughout, whereas the economy without communication remains at the bound for four quarters. State-dependent countercyclical communication thus acts as an automatic stabilizer when conventional interest-rate policy is constrained, shortening or even preventing ZLB episodes.

Overall, our results imply that the monetary stance depends not only on what a central bank communicates on average but also on how its information provision varies with economic conditions. In the model, the output component of the communication rule is what makes the estimated rule stabilizing: a rule responding only to inflation instead amplifies demand fluctuations and leaves the economy vulnerable to self-fulfilling contractions at the lower bound. This mechanism provides a rationale for communication frameworks that condition guidance on real activity as well as inflation, including outcome-based guidance such as that adopted by the Federal Reserve in September 2020. More generally, when communication effort responds systematically to the economic outlook, its systematic component should be considered alongside the policy-rate rule when assessing the monetary-policy stance.

Outline Section 1 discusses the related literature and clarifies our contribution. Section 2 constructs the communication measure and estimates the communication rule of the Federal Reserve, examines its robustness and presents the international evidence. Section 3 develops the CRANK model and studies systematic communication in normal times and at the zero lower bound. The final section concludes. Additional results and all proofs are in the online appendix.

1. Related Literature

This paper connects to four strands of the literature on central bank communication.

Measuring Central Bank Communication First, our paper contributes to the empirical literature measuring central bank communication and quantifying its effects on markets, expectations, and the macroeconomy. A large body of work isolates the surprise component of policy announcements from high-frequency movements in asset prices: Kuttner (2001) and Bernanke and Kuttner (2005) use federal funds futures around FOMC decisions, while Gürkaynak, Sack, and Swanson (2005) establishes that statements move yields above and beyond the contemporaneous rate decision. This identification underpins much of the subsequent literature on the macroeconomic effects of monetary policy (see, *e.g.* Gertler and Karadi, 2015, Nakamura and Steinsson, 2018, Bauer and Swanson, 2023). A closely related literature classifies the content of forward guidance, most prominently along the Delphic–Odyssean distinction of Campbell, Evans, Fisher, and Justiniano (2012), Campbell, Fisher, Justiniano, and Melosi (2017), separating signals about the economic outlook from commitments about the future path of rates. Our paper is in the spirit of this identification

strategy but inverts its logic. Rather than recovering the unexpected component of communication from market reactions, we recover its systematic component (the rule that governs how much the Committee chooses to say as a function of economic conditions), in direct analogy to the distinction between monetary policy shocks and the Taylor rule.

A parallel strand measures communication directly from text; our measure belongs to this one. One approach gauges readability or linguistic complexity through indices such as Flesch–Kincaid, documenting that FOMC statements have grown longer and harder to parse over time and that readability of the statements correlates with market volatility (Jansen, 2011, Bulíř, Čihák, and Jansen, 2013, Hernández-Murillo and Shell, 2014). Another extracts tone or sentiment (whether a statement reads as hawkish or dovish) through dictionary methods and supervised classifiers (Lucca and Trebbi, 2009, Picault and Renault, 2017), an effort recently extended to the paralinguistic content of communication by Gorodnichenko, Pham, and Talavera (2023), who analyze vocal emotion in FOMC press conferences. Closely related to our own exercise, Shapiro and Wilson (2022) move from measuring tone to recovering objectives, estimating the FOMC's implicit inflation target and loss function directly from the sentiment expressed in its deliberations. We depart from this literature by proposing a measure of communication effort, the change in Shannon entropy between consecutive periods, that is not systematically related to tone or sentiment. Rather than capturing directional stance, it measures a statement's informational density and novelty relative to to the previous statement. While readability indices assess clarity and sentiment measures assess hawkishness, our metric asks how much the Committee chose to elaborate, given what it had to say. This distinction is especially clear relative to Shapiro and Wilson (2022): while they infer the Fed's objective from the directional content of deliberations, we recover the inflation target from changes in lexical diversity alone through an entirely different mechanism.

NLP and Text-As-Data Tools Second, our paper relates to the literature that applies natural-language-processing (NLP) and text-as-data tools to economics (see, e.g. Gentzkow, Kelly, and Taddy, 2019). Within central banking, a first set of methods relies on topic models, most commonly Latent Dirichlet Allocation (LDA), to recover the thematic structure of policy text. For example, Hansen and McMahon (2016) and Hansen, McMahon, and Prat (2018) characterize the content of FOMC communication and its effects on markets. A second set of methods uses text classification and dictionary-based sentiment analysis to map words into discrete categories (hawkish, dovish, or neutral) either through prespecified dictionaries (Lucca and Trebbi, 2009, Apel and Blix Grimaldi, 2014, Picault and Renault, 2017, Aruoba and Drechsel, 2024) or through supervised classifiers. A more recent literature replaces sparse word counts with dense vector representations, word and document embeddings and large language models, to capture semantic content that bag-of-words methods miss (see, e.g. Handlan, 2022, Hansen and Kazinnik, 2023, Baumgärtner and Zahner, 2025)

Relative to these approaches, our approach differs methodologically. We quantify textual communication using the change in Shannon entropy, an information-theoretic measure of informational richness. Unlike topic models and sentiment classifiers, which characterize text according to its substantive content or directional tone, the change in entropy

summarizes the variation in the amount and diversity of information used by the central bank between periods independently of the message it conveys. The measure requires no labeled training data, no choice of lexicon, and no *ex ante* specification of the number of topics; it is therefore transparent and replicable, and less exposed to the overfitting and degrees-of-freedom concerns that attend supervised and embedding-based methods. The change in entropy is a parsimonious single-number summary of communication effort that complements the richer but more parameter-intensive representations developed elsewhere in this literature.

Systematic Rules Third, our paper contributes to the extensive literature estimating systematic monetary policy rules following the seminal work of Taylor (1993, 1999). Rather than focusing on the standard interest-rate instrument, we apply the logic of systematic rules to central bank communication itself. To address the endogeneity problem inherent in policy-rule estimation, we use two distinct identification strategies. First, we use the influential GMM framework with internal instruments of Clarida, Galí, and Gertler (2000), extended to identification-robust inference by Mavroeidis (2010). Second, we follow Barnichon and Mesters (2020) in using externally identified shocks as instruments, like in Romer and Romer (2004), Nakamura and Steinsson (2018), Känzig (2021).

A more recent strand of this literature incorporates text-based measures of central bank communication into policy-rule estimation (see, among others, Lucca and Trebbi, 2009, Shapiro and Wilson, 2022, Aruoba and Drechsel, 2024). Our paper departs from this strand. Rather than treating communication indicators as additional explanatory variables in an interest-rate rule, we estimate a communication rule in which the change in entropy is the dependent variable, capturing the systematic component of communication effort directly. In that sense, our approach is most closely related to Gáti and Handlan (2024), who likewise conceptualize the communication rule as a systematic mapping from the Fed's macroeconomic expectations to its language. The two exercises differ first in what they measure: Gáti and Handlan (2024) recover the content of the Fed's expectations from its language, reading its forecasts back out of the words, whereas our entropy measure captures the amount of communication effort, independently of the message conveyed. The two exercises also differ in method. Because they represent language as a high-dimensional bag of words, they estimate the inverse of this mapping by regressing the Fed's internal forecasts on the text via regularized (ridge) regressions, separately for each forecast variable. We collapse the language side to a single information-theoretic statistic and estimate the forward rule directly, with the change in entropy responding to expected inflation and the output gap. By reducing the language side to a single statistic rather than several hundred penalized regressors, our formulation sidesteps the dimensionality that ridge is designed to manage and is directly amenable to GMM estimation with both internal and external instruments, an identification route not available to a high-dimensional text regression.

New Keynesian Framework On the modeling side, we develop a four-equation Communication RANK (CRANK) model that extends the textbook framework of Woodford (2003b) and Galí (2015) by relaxing the full-information rational expectations benchmark, under which central bank communication is irrelevant. Households are assumed to have imperfect information about the underlying economic state, as in the noisy-information models

of Woodford (2003a) and Lorenzoni (2009) and the rational inattention literature initiated by Sims (2003) and Maćkowiak and Wiederholt (2009) (see Maćkowiak, Matejka, and Wiederholt, 2023, for a survey), where agents optimally choose signal precision subject to an information-processing constraint. We depart from this literature by endogenizing the precision of public signals through systematic central bank communication. This takes noisy-information and rational inattention mechanisms to the supply side of information.

Because signal precision is a second-order moment, shocks propagate partly through a time-varying precautionary saving motive; this connects our framework to a growing analytical literature on cyclical risk and aggregate demand (see, among others, Werning, 2015, Acharya and Dogra, 2020, Bilbiie, 2025, Antonova, Matvieiev, and Poilly, 2026). Our model differs from this literature in three respects. First, signal precision is endogenously determined by central bank communication effort and, as our empirical evidence suggests, is countercyclical. Second, because public signals originate from the central bank, this property turns the precautionary saving channel into a monetary policy instrument, with direct consequences for local equilibrium stability. Third, we embed a binding zero lower bound (ZLB) and show that communication effort acts as a substitute for conventional interest-rate policy, mitigating (or even preventing) the welfare losses associated with ZLB episodes relative to the canonical RANK benchmark (see, e.g. Eggertsson and Woodford, 2003, Adam and Billi, 2006, 2007, Nakov, 2008).

The ZLB analysis places our paper within the large literature on the limits of conventional monetary policy and on well-known New Keynesian puzzles, including the forward guidance puzzle and the observed macroeconomic stability of economies at the ZLB (see, among others, Eggertsson and Woodford, 2003, Campbell, Evans, Fisher, and Justiniano, 2012, Cochrane, 2018, Debortoli, Galí, and Gambetti, 2020, Del Negro, Giannoni, and Patterson, 2023). Because our resolution of these puzzles works through a discounted dynamic IS equation, the paper also connects to a growing literature on attenuation mechanisms (see, e.g., McKay, Nakamura, and Steinsson, 2017, Angeletos and Lian, 2018, Farhi and Werning, 2019, Gabaix, 2020, Meichtry, 2023, Gallegos, 2024, Bilbiie, 2025). We ground this attenuation in an empirically estimated, state-dependent central bank communication rule, giving the mechanism both a data-disciplined foundation and a clean theoretical interpretation.

Finally, our results also connect to Eggertsson and Schüle (2024) and Ahn and Holm (2025), who resolve the forward guidance puzzle by focusing on the behavior of the central bank itself. In their view, the puzzle does not arise in practice because central banks cannot credibly commit to keeping interest-rate plans unchanged between announcement and implementation. Our argument is complementary. The puzzle dissolves partly because central banks cannot commit to a future rate path, and partly because they systematically adjust their communication effort in the interim. As our empirical results show, this adjustment is rule-based and countercyclical.

2. Empirical Analysis

This section presents empirical evidence on the systematic component of central bank communication. We construct an entropy-based measure of the communication effort exerted

by the Federal Open Market Committee (FOMC) and estimate a communication reaction function for the U.S. Federal Reserve over the period 1994–2022. We then assess the robustness of our findings to changes in specification, expectations, and identification, before exploring the existence of analogous rules at the international level.

2.1 Measuring Central Bank Communication

We develop an entropy-based approach to quantify central bank communication, drawing on three decades of FOMC statements. We begin by describing the dataset and the text pre-processing steps, and then introduce our measure of lexical diversity as a quantification of the communication effort embedded in each statement.

2.1.1 Data Description Our dataset consists of all U.S. FOMC policy statements released between January 1994 and March 2025.[1] In the early years of this period, the Committee issued statements only after meetings in which it changed the federal funds rate; a practice that continued until the end of 1999. Beginning in 2000, statements were released after every scheduled meeting, regardless of the policy decision, and occasionally after unscheduled, ad hoc meetings. We refer to the former as *regular* statements and the latter as *exceptional* statements.

Our final sample comprises 234 policy statements, of which 220 are regular and 14 are exceptional. Most statements span one to two pages, summarizing current and expected economic conditions and explaining the rationale behind the Committee's decision. They typically conclude with a short technical paragraph outlining open market operations to achieve the target federal funds rate. Following the literature (see, among others, Handlan, 2022), we focus on the non-technical portions of the statements, where the most economically meaningful communication takes place. The average FOMC statement contains 311 words, with a minimum of 36 and a maximum of 800. Exceptional statements are roughly 25 percent shorter than regular ones (243 vs. 315 words on average) but exhibit greater dispersion (standard deviation of 186 vs. 168 words). Despite their brevity, exceptional statements remain rare, accounting for only about 6 percent of all releases.

Before constructing our entropy measure, we pre-process the text of each statement using standard natural language processing (NLP) techniques. We remove punctuation, numbers, and stop-words (see Table A.2 in Appendix A) and apply stemming to reduce words to their root form; for example, expectation, expected, and expecting all become "expect".

We then construct n-grams, defined as sequences of n consecutive words that serve as basic semantic units in the text. Our analysis includes unigrams, bigrams, trigrams, and four-grams.[2] Together, the processed statements form the corpus, C, of FOMC communication, containing 1,128 unique n-grams, including 2 four-grams, 3 trigrams, and 48 bigrams.

[1] The statements are publicly available in the Federal Reserve's archives.

[2] Details on the construction of multi-word n-grams and the list of the 53 n-grams ($n > 1$) are provided in Table A.3 in Appendix A.

2.1.2 Shannon Entropy as a Measure of Lexical Diversity To measure the communication effort of the Federal Reserve, we seek to capture how much additional explanation the FOMC provides from one statement to the next. Since this effort is not directly observable, we proxy for it using changes in the lexical diversity of FOMC statements, measured by the Shannon entropy of their n-gram distributions (Shannon, 1948). Formally, let X_s denote an n-gram drawn from the non-technical text of FOMC statement s, and let $p_s(x_i)$ denote the relative frequency of n-gram x_i in that statement. Its Shannon entropy, expressed in bits, is

$$\mathcal{E}_s = -\sum_{x_i \in \mathcal{X}_s} p_s(x_i) \log_2 p_s(x_i). \tag{1}$$

This measure captures both the breadth of the vocabulary and the distribution of linguistic expressions across a text. Entropy is low when a statement relies heavily on a limited set of recurring expressions and high when its vocabulary is more varied and less predictable.[3] The entropy of a statement increases when its linguistic expressions become more varied or their frequencies become more evenly distributed. We interpret this variation as a proxy for communication effort because FOMC statements rely on a highly templated register, built around recurring formulations such as "moderate growth" and "inflation remains subdued." Under familiar conditions, the Committee can communicate using this established vocabulary. Explaining an unfamiliar or complex policy environment, by contrast, may require additional qualifications, new formulations, and more elaborate reasoning. Such elaboration entails deliberation and drafting and potentially affects how markets interpret the policy outlook. The distinction is therefore between reusing familiar language and expanding the linguistic material used to explain policy.

This interpretation makes lexical diversity a measure of communication supply rather than of the message's substantive content or its effects. Entropy does not identify the tone of the statement, nor whether a statement is hawkish or dovish, optimistic or alarmed. Nor does greater lexical diversity necessarily imply better communication: a long and varied statement may be less informative or accessible than a concise one. Our measure captures the extent to which the Committee changes the linguistic material it uses. Whether changes in lexical diversity provide a useful measure of communication effort is ultimately an empirical question.

The distinction from other commonly used indicators is conceptual. Monetary policy shocks measure unexpected policy actions, uncertainty indices measure uncertainty about economic conditions or their outlook, readability indices capture aspects of textual complexity, hawk-dove measures characterize policy stance, and statement length counts the amount of text. Changes in log-entropy instead capture variation in the diversity of the language used by the central bank. These measures may nevertheless be correlated: in particular, longer statements can contain a broader range of n-grams, and difficult economic conditions may affect both uncertainty and communication effort. Consequently,

[3] Appendix A.2 illustrates the measure with three examples. The intuition can also be understood through a box of colored balls: a box containing only one color yields no uncertainty about the color of a draw, whereas a box containing an even mix of several colors yields greater uncertainty. Similarly, a statement dominated by recurring formulations has low entropy, while one drawing on a broader range of linguistic expressions has higher entropy.

the conceptual distinction does not by itself establish statistical independence from these alternative indicators.

We compute entropy for each FOMC policy statement in our sample, obtaining a series at meeting frequency. To align the communication measure with macroeconomic indicators, we assign each calendar day the entropy of the most recently released statement and aggregate the resulting daily series to monthly and then quarterly frequency. Regular statements occur at approximately regular intervals, whereas exceptional statements can be released at different points in the meeting cycle. To treat releases symmetrically, we assume that an exceptional statement's informational footprint persists for 52 days, corresponding to the average interval between regular meetings.[4] This convention prevents exceptional statements released shortly before a scheduled meeting from receiving disproportionately little weight. Quarterly entropy is constructed as the average of monthly entropy values within each quarter, following the convention used for policy rate aggregation in the empirical monetary policy literature (see, e.g., Clarida, Galí, and Gertler, 2000).

Table 1 reports summary statistics for the natural logarithm of entropy at meeting frequency. Regular statements have higher mean log-entropy than exceptional statements, while exceptional statements exhibit somewhat greater dispersion despite their much smaller sample. The table describes entropy levels; our measure of communication effort instead focuses on their changes over time.

Table 1. Summary Statistics of the Natural Logarithm of Entropy (Meeting Frequency)

Statement Type	Mean	Std. Dev.	Min	Max	N
All Statements	1.65	0.77	-0.96	3.02	234
Regular Statements	1.70	0.74	-0.96	3.02	220
Exceptional Statements	0.87	0.85	-0.78	2.37	14

Note: Shannon entropy is computed over the non-technical portions of each FOMC statement.

We define communication effort as the change in log-entropy between consecutive quarters:

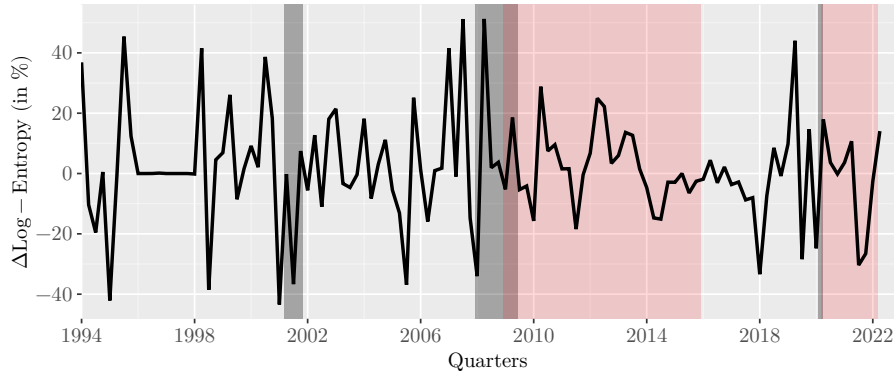
$$\Delta e_t = \log(\mathcal{E}_t) - \log(\mathcal{E}_{t-1}). \tag{2}$$

Using changes rather than levels has both a statistical and a conceptual motivation. Statistically, first-differencing helps mitigate the upward trend in entropy since 1994, which largely reflects the increasing length of FOMC statements. It does not, however, guarantee that the resulting measure is independent of statement length. Conceptually, the change captures the evolution of communication effort rather than the level of lexical diversity associated with a statement's prevailing style. A positive (resp. negative) value indicates an increase (resp. decline) in lexical diversity relative to the previous quarter. A zero value

[4]The FOMC holds eight regular meetings per year, implying an average spacing of approximately 52 days.

means that measured log-entropy is unchanged, not necessarily that the statement contains no new information. First-differencing also removes persistent differences in baseline communication style across chairpersons, although it need not eliminate all differences in their use of language. We return to the implications of these stylistic differences in Section 2.2.1.

FIGURE 1. Changes in Log-Entropy of FOMC Statements, 1994–2022



Note: Changes in log-entropy of FOMC statements at quarterly frequencies. NBER recessions and zero lower bound (ZLB) episodes are shaded in grey and light red, respectively.

Figure 1 plots changes in log-entropy from 1994 to 2022. The series displays pronounced increases around the onset of NBER recessions, with the largest movement occurring during the Global Financial Crisis in 2008. During zero lower bound (ZLB) episodes, changes are mostly positive but generally more moderate, while outside recessions and ZLB episodes the series fluctuates around zero. These patterns are consistent with the view that the FOMC increases its explanatory effort when economic conditions deteriorate or policy becomes more difficult to communicate, and adjusts that effort as conditions stabilize.

2.2 The Federal Reserve's Communication Reaction Function

We now specify and estimate a communication reaction function for the Federal Reserve, using the communication effort series constructed in the previous section.

2.2.1 A Forward-Looking Reaction Function Our specification builds on the Taylor rule literature, which links the evolution of the policy rate to deviations of expected inflation from target and to fluctuations in the output gap (Taylor, 1993, Clarida, Galí, and Gertler, 2000). We adapt this framework to model the systematic component of communication effort by specifying an analogous reaction function for changes in entropy of FOMC statements.[5] Two elements drive the dynamics of communication effort: an autoregressive component that allows for potential persistence in entropy changes over time, and a forward-looking component (for the cases $k, q > 0$) that captures how communication effort responds

[5]As illustrated in Appendix A.4, a very sizable and stable part of FOMC statements is dedicated to discussing inflation and the output gap.

to expected macroeconomic conditions. Beyond the non-stationarity already noted, specifying the rule in log-differences serves a distinct purpose. It removes any chairperson-specific shift in the average level of entropy, as, in the estimating equation, this shift plays the role of a unit-specific fixed effect that mitigates the potential confounding effects of differences in communication style across Fed chairs.[6] The rule is therefore identified off the within-chair variation of communication effort in response to economic conditions. The rule is:

$$\Delta e_t = \rho_e(L)\,\Delta e_{t-1} + (1-\overline{\rho}_e)\,\Delta e_t^* \,, \tag{3}$$

$$\Delta e_t^* = \Delta_e^* + \psi_\pi\big(\mathbb{E}\left[\pi_{t+k,t}\,|\,\mathcal{I}_t\right] - \pi^*\big) + \psi_y\,\mathbb{E}\left[\tilde{y}_{t+q,t}\,|\,\mathcal{I}_t\right] \,, \tag{4}$$

where $\rho_e(L) = \rho_{e,1} + \rho_{e,2}L + \cdots + \rho_{e,n}L^{n-1}$ and $\overline{\rho}_e \equiv \rho_e(1)$. The target Δe_t^* encodes the systematic component of the central bank's communication effort as a function of three elements: a desired steady-state change in entropy, Δ_e^*; the deviation of expected inflation from target, $\mathbb{E}\left[\pi_{t+k,t}\,|\,\mathcal{I}_t\right] - \pi^*$; and the expected output gap, $\mathbb{E}\left[\tilde{y}_{t+q,t}\,|\,\mathcal{I}_t\right]$. Here $\pi_{t+k,t}$ denotes annualized inflation between periods t and $t+k$, while $\tilde{y}_{t+q,t}$ is the average output gap between t and $t+q$, and $\mathcal{I}_t$ is the information set at t. When inflation and the output gap are both at their long-run levels, the realized change in entropy coincides with its steady-state desired level, Δ_e^*.

Equation (3) implies that the central bank adjusts its communication effort gradually, closing the gap between the current and target levels of entropy in a smoothed manner. The coefficients in $\rho_e(L)$ measure the degree of gradualism, reflecting how smoothly the Committee adapts the volume of explanation to evolving economic conditions.[7]

2.2.2 Estimation Procedure System (3)–(4) simplifies to

$$\Delta e_t = \rho_e(L)\,\Delta e_{t-1} + (1-\overline{\rho}_e)\left(\Delta_e^* + \psi_\pi\big(\mathbb{E}\left[\pi_{t+k,t}\,|\,\mathcal{I}_t\right] - \pi^*\big) + \psi_y\,\mathbb{E}\left[\tilde{y}_{t+q,t}\,|\,\mathcal{I}_t\right]\right) .$$

Adding and subtracting the realized values $\pi_{t+k,t}$ and $\tilde{y}_{t+q,t}$ yields

$$\Delta e_t = \rho_e(L)\,\Delta e_{t-1} + (1-\overline{\rho}_e)\left(\Delta_e^* + \psi_\pi\big(\pi_{t+k,t} - \pi^*\big) + \psi_y\,\tilde{y}_{t+q,t}\right) + \eta_t \,, \tag{5}$$

where

$$\eta_t \equiv -(1-\overline{\rho}_e)\psi_\pi\big(\pi_{t+k,t} - \mathbb{E}\left[\pi_{t+k,t}\,|\,\mathcal{I}_t\right]\big) - (1-\overline{\rho}_e)\psi_y\big(\tilde{y}_{t+q,t} - \mathbb{E}\left[\tilde{y}_{t+q,t}\,|\,\mathcal{I}_t\right]\big)$$

is a combination of forecast errors and is, by construction, orthogonal to any element of $\mathcal{I}_t$.

In equation (5), Δ_e^* and π^* are not separately identified. We choose to fix Δ_e^* at the sample average change in log-entropy and to estimate π^*. The advantage of this normalization is that it permits a direct test of the consistency between the announced inflation target used in the central bank's conventional rule and the target embedded in its communication

[6] Existing evidence suggests that the language of Federal Reserve communication varies across chairmen. Differences have been documented in both the complexity of FOMC statements and the sentiment of Chairs' public communications (see, e.g. De Pooter, 2021, Arismendi-Zambrano, Kypraios, and Paccagnini, 2026).

[7] For parsimony, the specification omits exogenous "communication shocks"; including them does not alter our main results.

rule. Identification and estimation of the parameter vector $\boldsymbol{\theta} \equiv (\{\rho_{e,i}\}_{i=1}^{n}, \pi^*, \psi_\pi, \psi_y)$ rely on instrumental variables. Any $\boldsymbol{z}_t \in \mathcal{I}_t$ that satisfies the exogeneity condition, $\operatorname{cov}(\boldsymbol{z}_t, \eta_t) = 0$, and the relevance condition, $\operatorname{cov}(\boldsymbol{z}_t, x_t) \neq 0$ for $x_t = (\pi_{t+k,t},\, \tilde{y}_{t+q,t})$, constitutes a valid instrument. The implied moment conditions are:

$$\mathbb{E}\left[\left[\Delta e_t - \rho_e(L)\,\Delta e_{t-1} - (1-\overline{\rho}_e)\Big(\Delta_e^* + \psi_\pi\big(\pi_{t+k,t} - \pi^*\big) + \psi_y\,\tilde{y}_{t+q,t}\Big)\right]\boldsymbol{z}_t\right] = 0\,. \tag{6}$$

We estimate $\boldsymbol{\theta}$ by iterated Generalized Method of Moments (GMM; Hansen, 1982, Hansen, Heaton, and Yaron, 1996):

$$\hat{\boldsymbol{\theta}}_s = \underset{\boldsymbol{\theta}\in\Theta}{\operatorname{argmin}}\; \overline{m}_T(\boldsymbol{\theta})^\top \hat{W}(\hat{\boldsymbol{\theta}}_{s-1})\,\overline{m}_T(\boldsymbol{\theta})\,,$$

where $\overline{m}_T(\boldsymbol{\theta})$ is the sample counterpart of the moment conditions in (6) and $\hat{W}(\hat{\boldsymbol{\theta}}_{s-1})$ is the optimal weighting matrix updated at each iteration. The estimator $\hat{\boldsymbol{\theta}}$ is the fixed point of this mapping.[8] Standard errors are computed at each step using a heteroskedasticity- and autocorrelation-consistent (HAC) estimator of the long-run variance-covariance matrix of the moment conditions, based on a vector autoregressive prewhitening of the conditions (Andrews and Monahan, 1992), which is asymptotically equivalent to the Newey and West (1987) estimator and improves on it in small samples.

2.2.3 Baseline Specification and Data We estimate the communication reaction function over the period 1994:Q2–2022:Q4. In the baseline specification, both inflation and output-gap horizons are set to one quarter, $k = q = 1$. Inflation is the annualized quarterly log-difference of the personal consumption expenditures (PCE) price index, which the FOMC has primarily monitored since 2000.[9] The PCE index accommodates substitution in the consumption basket and adapts to evolving spending patterns. The output gap is the percentage deviation of real GDP from CBO potential output.[10] Further details on data construction are provided in Appendix A.3.

We model the autoregressive dynamics of log-entropy changes with a single lag, $\rho_e(L) = \rho_{e,1} = \rho_e$, so that the baseline rule becomes

$$\Delta e_t = \rho_e\,\Delta e_{t-1} + (1-\rho_e)\left(\Delta_e^* + \psi_\pi\big(\mathbb{E}\left[\pi_{t+1,t}\,|\,\mathcal{I}_t\right] - \pi^*\big) + \psi_y\,\mathbb{E}\left[\tilde{y}_{t+1,t}\,|\,\mathcal{I}_t\right]\right)\,. \tag{7}$$

The baseline instrument set includes a constant, the first lag of the communication effort, and four lags of inflation, the output gap, the shadow federal funds rate of Xia and Wu (2016),[11] annualized quarterly commodity price inflation, and annualized quarterly M2 growth. We refer to this as the internal instrumental strategy, as it relies on lagged values

[8] The iterative procedure typically yields more efficient estimates under correct model specification than the conventional two-step GMM.

[9] See https://www.federalreserve.gov/economy-at-a-glance-inflation-pce.htm.

[10] Section 2.2.5 examines the sensitivity of the results to alternative horizons, inflation indices and definitions of the output gap.

[11] Use of the shadow rate rather than the effective federal funds rate accommodates the zero lower bound episodes in our sample. Using the effective federal funds rate instead does not materially affect our results.

of the variables entering the rule. A complementary external strategy, based on identified exogenous monetary policy and supply-side surprises, is taken up in the robustness analysis. Because, as already noted, Δ_e^* and π^* are not separately identified, we fix the steady-state change in log-entropy at its sample average and recover the implied inflation target from the estimation.

2.2.4 Baseline Estimates Table 2 reports the baseline estimates of equation (7). The Sargan–Hansen test yields a p-value of 0.28, indicating that the null hypothesis of valid over-identifying restrictions is not rejected at conventional levels. The estimates are consistent with a systematic communication rule in which the Federal Reserve adjusts its communication effort to expected macroeconomic conditions in a state-dependent manner.

TABLE 2. Communication Reaction Function — Baseline Estimates

	PARAMETERS				J-TEST
	ρ_e	π^*	ψ_π	ψ_y	p-value
Δe_t	0.07	1.93***	7.04***	-1.46**	0.28
	(0.06)	(0.17)	(1.05)	(0.68)	

Note: Baseline estimates of equation (7) with $k = q = 1$. Inflation is annualized PCE inflation; the output gap is from the Congressional Budget Office. HAC standard errors in parentheses. The rightmost column reports the p-value of the Sargan–Hansen over-identification test. For full details on the instrument set and estimation period, see Section 2.2.4. *** (**, *) denotes significance at the 1% (5%, 10%) level.

Four results emerge. First, the autoregressive coefficient ρ_e is small (0.07) and statistically insignificant: communication effort displays little intrinsic persistence once expected macroeconomic conditions are controlled for. Changes in lexical diversity are driven primarily by evolving expectations rather than by inertia in communication style. Second, the estimated inflation target π^* is 1.93 percent, highly significant and statistically indistinguishable from the Fed's announced 2 percent long-run objective. Recovered from communication data alone, this estimate provides direct evidence that the Fed's communication rule is anchored to the same nominal goal as its interest-rate rule. Third, the response coefficient on expected inflation, ψ_π, is positive (7.04) and significant at the 1 percent level: a one-percentage-point upward deviation of inflation expectations from target raises the communication effort by approximately 7 percent. The Committee's communication effort intensifies systematically when inflation expectations drift upward, consistent with an effort to clarify its policy stance and re-anchor expectations. Fourth, the coefficient on the expected output gap, ψ_y, is negative (-1.46) and significant at the 5 percent level: conditional on inflation expectations, the Committee communicates more when the economy is expected to run below potential. A one-percent shortfall in expected output prompts roughly a 1.5 percent increase in communication effort.

Re-estimating the rule for the post-2008 period still supports the existence of a systematic rule (p-value of Hansen Sargan test of 0.23), leaves the inflation parameter essentially unaffected and significant (6.80), but reinforces the output gap channel as the coefficient

rises by a factor of 4 (−6.26). This result goes together with the idea that, in the aftermath of the great financial crisis, the Fed's concern for its secondary mandate, promote growth and full employment, gained stamina.[12] These results also bear on the interpretation of the measure itself. That communication effort obeys a stable forward-looking rule, responds to expected inflation and the output gap with the signs a dual mandate predicts, and recovers an inflation anchor statistically indistinguishable from the announced 2% target, all from statement text alone, is difficult to reconcile with a purely mechanical reading of the change in entropy. The systematic, state-contingent behavior of the series is itself evidence that it captures a deliberate communication response, not an incidental property of the text.

Taken together, the baseline estimates establish that the communication dimension of U.S. monetary policy is systematic and state-contingent. Communication is a complementary instrument to the policy rate, intensifying when inflation risks are elevated and when economic slack is pronounced. These findings are consistent with viewing central bank communication effort as an endogenous, policy-active component of the monetary transmission mechanism rather than a passive by-product of interest-rate decisions.

2.2.5 Robustness Analysis We now show that the baseline findings are robust to a wide range of variations in the specifications. Before turning to specific exercises, it is useful to summarize what is invariant throughout. First, the over-identifying restrictions implied by the communication rule are never rejected at conventional significance levels, supporting the existence of a systematic communication rule. Second, the estimated inflation target π^* remains tightly centered on the Fed's 2 percent objective across all specifications. Third, the response of communication effort to expected inflation, ψ_π, is always positive and statistically significant, with economically large magnitudes. Fourth, the response to the expected output gap, ψ_y, is always negative, although its precision is somewhat sensitive to measurement choices. Intrinsic persistence is uniformly small or modestly negative, indicating that communication adjustments are driven by expected macroeconomic conditions rather than by inertia.

Measurement Table 3 reports estimates obtained for alternative measures of inflation and the output gap. Using headline PCE in the baseline, ρ_e is small and insignificant; under core PCE or core CPI instead, it turns mildly negative and significant (−0.11 and −0.12, respectively). A similar pattern arises for the unemployment-gap specification (−0.10). These mild mean-reverting tendencies suggest that, once underlying price and activity pressures are controlled for, larger-than-usual changes in lexical diversity are partially corrected at the following meeting, a pattern consistent with deliberate fine-tuning of communication effort.

The estimated inflation target π^* adjusts systematically with the choice of price index, in a direction that mirrors the well-documented CPI–PCE wedge. PCE-based measures yield estimates in the range 1.83–1.93 percent , close to the Fed's 2 percent objective; CPI-based measures yield 2.23–2.46 percent. The communication rule thus inherits the implicit yardstick of the underlying price index. The inferred target shifts with the price measure, exactly as the known CPI–PCE relationship would predict, providing a useful internal cross-validation.

[12]The results are reported in Table A.12 of the Section A.5 in the Online Appendix.

TABLE 3. Robustness to Measurement

VARIABLES	PARAMETERS				J-TEST
	ρ_e	π^*	ψ_π	ψ_y	p-value
I. Inflation					
CPI	0.02	2.46***	11.33***	-1.46	0.29
	(0.07)	(0.16)	(1.69)	(1.09)	
GDP Deflator	-0.04	1.79***	7.13***	-1.12*	0.41
	(0.04)	(0.16)	(1.33)	(0.64)	
Core PCE	-0.11***	1.83***	4.35***	-0.61	0.40
	(0.04)	(0.23)	(1.22)	(0.56)	
Core CPI	-0.12***	2.23***	5.54***	-0.62	0.75
	(0.03)	(0.16)	(1.47)	(0.55)	
II. Output Gap					
Quadratic Trend	0.03	1.82***	7.92***	-1.49***	0.40
	(0.05)	(0.14)	(1.22)	(0.34)	
Unemployment Gap	-0.10***	1.89***	3.50***	-0.57	0.46
	(0.03)	(0.27)	(0.62)	(0.51)	

Note: Specification with $k = q = 1$ and one lag of change in log-entropy throughout. Inflation is annualized PCE inflation unless stated otherwise; the output gap is the CBO measure unless stated otherwise. HAC standard errors in parentheses. The negative unemployment gap is used to preserve the expected signs of ψ_π and ψ_y. The rightmost column reports the p-value of the Sargan–Hansen over-identification test. *** (**, *) denotes significance at the 1% (5%, 10%) level.

The inflation response ψ_π is robustly positive and significant across all inflation measures, but its magnitude varies intuitively with the salience of the indicator. Headline measures elicit larger responses (CPI: 11.33; PCE: 7.04) than core measures (4.35–5.54), consistent with the view that the Committee elaborates most forcefully when inflation is most visible to households and financial markets. The output gap coefficient ψ_y retains its negative sign throughout but loses statistical significance under several core inflation measures and under the unemployment gap. Two interpretations are consistent with this pattern. First, there is a hierarchy in the rule: inflation deviations trigger a more reliable and more forceful communication response than real-activity fluctuations. Second, labour-market measures of slack introduce additional noise during periods dominated by participation shifts and sectoral reallocation, especially during the post-pandemic recovery, weakening the signal available to estimation.

Forecasting Horizon Table 4 varies the forecast horizons k and q for inflation and the output gap from zero (contemporaneous) to four quarters ahead. The inflation response ψ_π is stable across all specifications, ranging narrowly from 5.5 to 5.7 and remaining strongly significant; the Committee's communication response to deviations of inflation from target is independent of the horizon over which expectations are formed.

TABLE 4. Forecasting Horizon

SPECIFICATION	PARAMETERS				J-TEST
	ρ_e	π^*	ψ_π	ψ_y	p-value
$k=0$, $q=0$	-0.12**	1.83***	5.72***	-0.49	0.92
	(0.05)	(0.18)	(1.10)	(0.48)	
$k=4$, $q=4$	-0.04	1.64***	5.48***	-1.43*	0.25
	(0.05)	(0.26)	(1.39)	(0.78)	
$k=4$, $q=1$	-0.02	1.62***	5.60***	-1.19	0.33
	(0.05)	(0.26)	(1.45)	(0.73)	

Note: Specification with one lag of change in log-entropy; inflation is annualized PCE inflation and the output gap is the CBO measure. HAC standard errors in parentheses. The rightmost column reports the p-value of the Sargan–Hansen over-identification test. *** (**, *) denotes significance at the 1% (5%, 10%) level.

The output-gap coefficient ψ_y strengthens and becomes more precisely estimated as the horizon extends: it is small and insignificant at $k=q=0$ (-0.49), but rises to -1.43 at $k=q=4$ and is significant at the 10 percent level, closely matching the baseline. The asymmetry between the inflation and output responses across horizons points to a distinct timing in the Fed's communication strategy: it is responsive to inflation deviations at any horizon but elaborates most actively about real conditions when anticipating future output developments rather than reacting to current slack. The estimated inflation target π^* also tracks the horizon sensibly: it declines modestly from 1.83 at $k=q=0$ to around 1.62–1.64 at longer horizons, consistent with the lower long-run inflation expectations that prevailed across most of the post-2000 era.

Expectations Table 5 re-estimates the baseline replacing rational expectations with the Federal Reserve's Tealbook (formerly Greenbook) projections (the staff forecasts available to the FOMC at each meeting) under both PCE and CPI inflation. The sample shortens to 1996:Q1–2019:Q4 due to data availability,[13] and the lag length of the instrument set is reduced to two to preserve efficiency.

The communication rule is stable: both specifications pass the over-identification test and deliver π^* estimates of 1.83 and 1.93 percent, statistically indistinguishable from the announced 2 percent target. Two features differ from the baseline in instructive ways. First, the autoregressive coefficient ρ_e turns significantly negative (around -0.17 to -0.18): under internal forecasts, communication adjustments are actively corrected once the staff outlook stabilizes, a pattern consistent with deliberate fine-tuning rather than momentum. Second, the inflation response ψ_π is somewhat larger than in the baseline (8.48 and 7.74 under PCE and CPI, respectively), suggesting that when the staff projects inflation above target, the Committee responds more forcefully in its statement, communicating preemptively before

[13] For PCE inflation, Tealbook forecasts are available from 2000:Q1 only. We fill the missing observations with CPI projections for 1996:Q1–1999:Q4, since the Fed tracked CPI inflation prior to 2000, thereby limiting the loss of data points. Forecasts are only available until 2019:Q4 for CPI.

TABLE 5. Tealbook Expectations

INDICATOR	PARAMETERS				J-TEST
	ρ_e	π^*	ψ_π	ψ_y	p-value
PCE	-0.18***	1.83***	8.48**	-1.09*	0.34
	(0.06)	(0.19)	(3.80)	(0.59)	
CPI	-0.17***	1.93***	7.74**	-1.22**	0.21
	(0.06)	(0.21)	(3.90)	(0.60)	

Note: Specification with $k = q = 1$ and one lag of change in log-entropy. The estimation period is 1996:Q1–2019:Q4. HAC standard errors in parentheses. The rightmost column reports the p-value of the Sargan–Hansen over-identification test. *** (**, *) denotes significance at the 1% (5%, 10%) level.

pressures become visible in published data. The output-gap response ψ_y remains negative and significant (around -1.1 to -1.2), confirming that this channel is not an artifact of the formulation of expectations.

Instrumental Strategy The exercises above rely on lagged internal macroeconomic variables as instruments. We now assess the robustness of our findings to an alternative identification strategy based on externally identified structural shocks, following Barnichon and Mesters (2020), who use this approach to estimate the structural equations of the New Keynesian model. This alternative is informative for two reasons. First, unlike our baseline instruments, which are equilibrium outcomes reflecting the joint influence of multiple shocks, external instruments exploit variation associated with specific, independently identified structural disturbances. Second, this distinction allows us to investigate to what extent the source of macroeconomic fluctuations ought to affect the communication effort of the central banker. The identifying assumption is that the structural shocks are orthogonal to the endogenous innovation in the equation of interest, η_t, while remaining relevant for the explanatory variables. We use oil price surprises from Känzig (2021) to isolate supply-driven variation in expected inflation and activity, and high-frequency monetary policy surprises from Nakamura and Steinsson (2018)[14] to capture variation associated with monetary policy disturbances on the demand side. The internal instruments used in our baseline specification exploit equilibrium movements in inflation and activity that reflect the combined influence of multiple structural disturbances. By contrast, the external instruments isolate variation associated with oil price and monetary policy surprises, allowing us to examine whether the estimated communication rule depends on the source of the underlying disturbance. T

We consider six instrument sets (Table 6). IV1 and IV4 use four and twelve lags, respectively, of oil price surprises and their absolute values. IV2 and IV5 use the corresponding lags of monetary policy surprises and their absolute values. IV3 and IV6 pool the two types of instruments. All specifications include a constant and one lag of communication effort.

The core findings still obtain under this alternative identification strategy. The estimated inflation target, π^*, remains close to 2 percent, while the persistence parameter, ρ_e,

[14]We use the updated series of Acosta, Brennan, and Jacobson (2024).

TABLE 6. External Instrument Strategy

INSTRUMENT SET	PARAMETERS				J-TEST
	ρ_e	π^*	ψ_π	ψ_y	p-value
I. Supply-Side Instruments (Känzig 2021)					
4 lags (IV1)	-0.06	1.72***	6.73**	0.13	0.32
	(0.09)	(0.26)	(2.67)	(2.43)	
12 lags (IV4)	-0.01	1.90***	10.27***	-1.41	0.46
	(0.12)	(0.18)	(2.61)	(1.27)	
II. Demand-Side Instruments (Nakamura and Steinsson 2018)					
4 lags (IV2)	-0.27*	1.66**	7.97***	-14.41**	0.48
	(0.15)	(0.65)	(2.95)	(6.26)	
12 lags (IV5)	-0.14**	2.15***	4.86	-1.64	0.77
	(0.06)	(0.33)	(3.22)	(2.15)	
III. Combined Instruments (Supply ∪ Demand)					
IV1 ∪ IV2 (IV3)	-0.12***	1.79***	3.82***	-3.65*	0.11
	(0.04)	(0.31)	(1.28)	(1.96)	
IV4 ∪ IV5 (IV6)	-0.09**	1.96***	9.08***	-1.76***	0.10
	(0.04)	(0.10)	(0.99)	(0.66)	

Note: Specification with $k = q = 1$ and one lag of change in log-entropy. Inflation is annualized PCE inflation; the output gap is the CBO measure. HAC standard errors in parentheses. All instrument sets include a constant and one lag of entropy, plus the following external instruments. Supply-side sets use lags of Känzig 2021 oil price surprises and their absolute values; demand-side sets use lags of Nakamura and Steinsson 2018 monetary policy surprises (updated by Acosta, Brennan, and Jacobson 2024) and their absolute values. Combined sets pool both. The rightmost column reports the p-value of the Sargan–Hansen overidentification test. *** (**, *) denotes significance at the 1% (5%, 10%) level.

is small or modestly negative across specifications. Beyond this robustness result, the external instruments provide additional insight into the mechanisms underlying the communication rule by distinguishing between supply- and demand-driven sources of macroeconomic variation.

With supply-side instruments (IV1, IV4), the estimated inflation response, ψ_π, is positive, large, and statistically significant, ranging from 6.7 to 10.3. By contrast, the estimated output-gap response, ψ_y, is small and statistically indistinguishable from zero. In other words, when inflationary pressures originate on the supply side, the Committee devotes greater communication effort to explaining the implications for prices, while the associated output fluctuations elicit a more limited response. This pattern is consistent with the difficulty of offsetting the real effects of supply disturbances through monetary policy, alongside the importance of explaining their implications for inflation expectations.

The estimates obtained with demand-side instruments (IV2, IV5) reveal a less robust response to economic activity. The estimated inflation response remains positive in both

specifications, although it is statistically significant only under IV2. By contrast, the estimated output-gap response varies substantially with the instrument set: it is large and negative under IV2 ($\psi_y = -14.41$, significant at the 5% level), but much smaller and statistically insignificant under IV5 ($\psi_y = -1.64$). The results suggest that the inflation channel is consistently positive under demand-side identification, whereas the evidence for an additional response to economic activity is sensitive to the choice of instruments and lag length. The large negative estimate under IV2 is consistent with the Committee devoting additional communication effort to economic activity when fluctuations are induced by demand-related disturbances, but the much smaller estimate under IV5 indicates that this interpretation is not robust across specifications.

When supply- and demand-side instruments are pooled (IV3, IV6), the estimated coefficients combine features of the two separate specifications. The inflation response remains positive and statistically significant, while the output-gap response is negative and statistically significant in both specifications. The Sargan–Hansen test does not reject the overidentifying restrictions at the 5% level. In that case, as in the internal instrument case, higher expected inflation triggers greater the communication effort and so does an expected drop in economic activity.

The observed differences in the estimates would have important implications for the interpretation of the baseline estimates. If the communication response to a given movement in inflation or activity depends on the underlying shock, a single linear rule estimated from aggregate equilibrium variation may average over distinct responses to different disturbances. The external-instrument exercise helps expose this possibility by isolating particular sources of variation that are difficult to disentangle using internal instruments alone. At the same time, the differences across estimates do not, by themselves, establish that the communication rule is state-dependent. They may also reflect differences in the variation isolated by each instrument set, the strength of the first-stage relationships, or the precision with which the coefficients are identified. In particular, a statistically significant coefficient under one specification and an insignificant coefficient under another do not establish that the two coefficients differ statistically. The appropriate conclusion is therefore that the results suggest, but do not conclusively establish, shock-dependent communication responses. Likewise, the J-tests for the pooled specifications do not reject the overidentifying restrictions at conventional significance levels and cannot, on their own, establish that a common linear rule fails to capture the responses to both sources of variation.[15]

Taken together, the external-instrument exercise both corroborates the main findings and sharpens their economic interpretation. The estimated inflation target remains close to 2 percent, and the inflation response is positive across all six specifications, although its statistical significance varies. The output-gap response, by contrast, is more sensitive to the source of identifying variation and to the instrument specification: it is small under supply-side instruments and negative under demand-side and pooled instruments, but is statistically significant under only one of the two demand-side specifications. These findings support the view that the Committee systematically adjusts its communication effort in

[15] A formal assessment of shock-dependent responses would require a specification that explicitly allows the communication rule to vary with the source of the disturbance, together with appropriate tests of coefficient equality. We leave this extension to future work.

response to macroeconomic conditions, while suggesting that its response to economic activity may depend on the nature of the underlying disturbance. More broadly, the exercise illustrates the value of distinguishing between equilibrium movements in macroeconomic variables and the structural shocks that generate them when assessing the mechanisms underlying central bank communication.

Length vs Entropy A statement can be lexically diverse either because it is long, or because it packs unusual and varied content into a given length, or both. To control for this concern, we decompose the Shannon entropy of an FOMC statement into two conceptually distinct margins: an extensive margin and an intensive margin. Taking logarithms of $\mathcal{E}(X_t)$ yields the identity

$$\log \mathcal{E}(X_t) = \log \mathcal{L}_t + \log \left(\mathcal{E}(X_t)/\mathcal{L}_t \right) , \tag{8}$$

where $\mathcal{L}_t$ is the total number of n-gram occurrences in the statement at time t. The first term captures the extensive margin—the volume of language used—while the second is the lexical density of the statement, that is, the lexical diversity per unit of text once length is held fixed. We refer to the second term as the intensive margin of lexical diversity throughout, and we reserve the word density for this length-normalized quantity.

In first differences, changes in total entropy are strongly correlated with changes in statement length (0.91), but this relationship masks an offsetting movement in lexical density. The correlation between changes in log length and changes in log entropy per unit of text, $\Delta \log(\mathcal{E}(X_t)/\mathcal{L}_t)$, is -0.73. Thus, quarters with larger increases in statement length tend to exhibit larger declines, or smaller increases, in lexical diversity per unit of text. This negative correlation suggests that additional text tends to rely more heavily on familiar language, partly offsetting the increase in total entropy associated with greater length. These results explain why we use total entropy rather than statement length alone. Length captures the extensive margin of communication, while entropy per unit of text captures its intensive margin. Because the two can move in opposite directions, length alone may badly capture communication effort. Total entropy accounts for both margins and is therefore better suited to our analysis.

To assess how much of the rule is driven by the extensive margin alone, we re-estimate equation (7) replacing Δe_t with $\Delta \log \mathcal{L}_t$, the growth rate of statement length. The results, reported in Table 7, are instructive precisely for what they fail to find. The inflation response remains positive and significant at $\psi_\pi = 4.07$, but it is barely half the magnitude obtained under the baseline specification: the inflation sensitivity of communication effort appears muted when communication is reduced to a word count. The output-gap coefficient, $\psi_y = -1.03$, is statistically insignificant. A rule built on length alone is, in effect, blind to the second mandate. These are not minor discrepancies. Restricting attention to the extensive margin systematically understates how forcefully the Committee escalates its communication effort when price stability is at risk and misses entirely the counter-cyclical dimension of communication, the documented tendency to elaborate more when the economy weakens. Total change in entropy recovers both. By jointly encoding the volume of language and its lexical density, it captures the intensive margin (the deliberate enrichment of vocabulary beyond mere verbosity) which carries the output-stabilization signal. When the economy

runs below potential, the Committee writes differently rather than simply writing more, using a less predictable vocabulary to convey the additional nuance that slack conditions demand. The distinction reflects a genuine shift in communication effort that length, by construction, cannot capture. Ignoring the intensive margin therefore amounts to ignoring a policy-relevant dimension of central bank communication, one that is directly tied to the real economy and to the Fed's dual mandate. A change in Shannon entropy, as defined in equation (1), is a strictly more informative object than a word count, and the reaction function estimates confirm that treating it as such is empirically consequential.

TABLE 7. Communication Reaction Function — Log-Length Estimates

	PARAMETERS				J-TEST
	ρ_e	π^*	ψ_π	ψ_y	p-value
$\Delta \log(\mathcal{L}_t)$	-0.08*	2.38***	4.07***	-1.03	0.35
	(0.04)	(0.33)	(0.95)	(0.78)	

Note: Estimates of equation (7) with $k = q = 1$. Inflation is annualized PCE inflation; the output gap is from the Congressional Budget Office. HAC standard errors in parentheses. The rightmost column reports the p-value of the Sargan–Hansen over-identification test. For full details on the instrument set and estimation period, see Section 2.2.4. *** (**, *) denotes significance at the 1% (5%, 10%) level.

A back-of-the-envelope calculation makes the cost of restricting attention to the extensive margin transparent. Intrinsic persistence is small and statistically insignificant in both the entropy rule ($\rho_e = 0.07$) and the length rule ($\rho_e = -0.08$, marginally significant at 10%), and the implied inflation target is in each case not statistically different from the Fed's announced 2% objective ($\pi^* = 1.93$ and $\pi^* = 2.38$, respectively). Setting persistence to zero and the target to the announced value, the two rules reduce to a simple feedback rules in expected inflation and the expected output gap:

$$x_t \approx \Delta_x^* + \psi_\pi^x \big(\mathbb{E}\left[\pi_{t+1,t} \,|\, \mathcal{I}_t\right] - 2 \big) + \psi_y^x \, \mathbb{E}\left[\tilde{y}_{t+1,t} \,|\, \mathcal{I}_t\right] , \qquad x \in \{\Delta e, \Delta \log(\mathcal{L})\} .$$

Because the change in total log-entropy decomposes additively into the change in log-length and log-lexical density, $\Delta e_t = \Delta \log \mathcal{L}_t + \Delta \log(\mathcal{E}(X_t)/\mathcal{L}_t)$, the implied rule for the intensive margin is simply the difference between the entropy and the length rules. Its coefficients are $\psi_\pi^{\text{int}} \approx 7.04 - 4.07 = 2.97$ and $\psi_y^{\text{int}} \approx -1.46 - (-1.03) = -0.43$. The intensive margin therefore absorbs roughly 42% of the entropy rule's inflation response and 29% of its output-gap response: a sizeable share of the Fed's systematic communication reaction on both legs of its mandate is carried by how the statement is written, not by how long it is. The asymmetry between the two channels, however, is what makes the omission consequential rather than merely incomplete. In the entropy rule, both responses are statistically significant; in the length rule, the output-gap coefficient is statistically indistinguishable from zero, and the -0.43 contribution of the intensive margin is precisely what tips it into significance once it is added back. A researcher who measured communication by length alone would observe a Federal Reserve that responds systematically to inflation but appears silent about real

activity, a partial reaction function indistinguishable from that of a strict inflation targeter. The dual-mandate dimension of the FOMC's communication effort is borne almost entirely by the intensive margin and is invisible to extensive-margin measures.

2.3 International Evidence

If the rule we document for the Federal Reserve were an individual characteristic of the U.S. , for example reflecting the dual mandate, the FOMC structure, or the particular transparency culture of the Greenspan–Bernanke era, our exercise would be of limited interest beyond a single, but important, country. A broader question is whether central banks systematically adapt their communication to economic conditions. We therefore assess whether the communication rule identified for the United States extends to other advanced economies by re-estimating it for Australia, Canada, Japan, the Euro Area, and Sweden. Differences in monetary policy mandates, objectives, and economic structures allow us to examine whether state-contingent communication reflects a more general feature of central bank behavior or is specific to the US institutional and macroeconomic environment.

The baseline specification of Section 2.2.4 requires two extensions to accommodate non-U.S. economies. For small open economies (Australia and Sweden), the exchange rate may be a first-order concern: currency depreciation affects import prices and, by extension, the inflation outlook, generating a natural additional motive for communication effort that is largely absent in a closed economy of the size of the United States. For Canada, the depth of trade and financial integration with the United States raises the possibility that the Bank of Canada's communication is driven by the stance of U.S. monetary policy as well as by domestic fundamentals. The extended rule incorporates both channels:

$$\Delta e_t = \rho_e \Delta e_{t-1} + (1-\rho_e)\Big(\Delta_e^* + \psi_\pi\big(\mathbb{E}[\pi_{t+1,t}|\mathcal{I}_t] - \pi^*\big) + \psi_y\,\mathbb{E}[\tilde{y}_{t+1,t}|\mathcal{I}_t] + \psi_{\Delta s}\big(\mathbb{E}[s_{t+1,t}|\mathcal{I}_t] - \Delta s^*\big) + \psi_r\big(\mathrm{FFR}_t - r^*\big)\Big). \tag{9}$$

The exchange rate term ($\psi_{\Delta s}$) captures any change in communication effort when the domestic currency is expected to depreciate relative to its long-run benchmark; the U.S. monetary policy term (ψ_r) is included only for Canada and measures how much of the Bank of Canada's communication effort is explained by the stance of its dominant trading partner. Given the shorter samples, the instrument set is reduced to two lags each of the inflation rate, output gap, policy rate, and money growth.[16] All forecasting horizons are set to one quarter. Table 8 reports the results.

Common Features The over-identification test passes comfortably in every specification, establishing that systematic forward-looking communication rules are not specific to the U.S. Fed. Across all countries, the inflation response ψ_π is positive and statistically significant (Canada aside, discussed separately), and the output-gap response ψ_y is negative.

[16] M2 is used in all countries except Australia, where only M3 is available. For the Euro Area and Australia, the term spread between 10-year and 3-month rates is added; when exchange rate terms are included, lags of changes in the exchange rate are added; for Canada's extended specification, lags of the U.S. federal funds rate are added.

TABLE 8. International Evidence

RULE	PARAMETERS						J-TEST
	ρ_e	π^*	ψ_π	ψ_y	$\psi_{\Delta s}$	ψ_r	p-value
Euro Area: 2000Q1–2022Q4							
Baseline	-0.07	2.24***	3.82**	-0.27	–	–	0.23
	(0.07)	(0.54)	(1.91)	(0.59)			
Japan: 2003Q2–2022Q4							
Baseline	-0.38***	0.46*	7.16**	-3.47**	–	–	0.46
	(0.08)	(0.28)	(3.06)	(1.67)			
Australia: 1992Q1–2022Q4							
Baseline	-0.02	2.30***	4.99**	-5.28**	–	–	0.74
	(0.08)	(0.23)	(2.52)	(2.46)			
Extended	-0.18*	2.07***	2.02**	-3.97***	0.71***	–	0.14
	(0.09)	(0.53)	(0.98)	(1.33)	(0.26)		
Sweden: 2000Q1–2022Q4							
Baseline	-0.54***	0.87**	1.01***	-1.36***	–	–	0.14
	(0.04)	(0.40)	(0.20)	(0.30)			
Extended	-0.11**	1.83***	1.05***	-1.10***	0.17***	–	0.48
	(0.05)	(0.31)	(0.21)	(0.24)	(0.04)		
Canada: 1998Q1–2022Q4							
Baseline	-0.24***	2.28	0.37	1.37**	–	–	0.15
	(0.05)	(1.79)	(1.18)	(0.58)			
Extended	-0.19***	2.35***	0.76	0.83	–	0.97**	0.43
	(0.05)	(0.84)	(1.21)	(0.78)		(0.42)	
Simple	-0.19***	1.81***	–	–	–	1.28***	0.20
	(0.05)	(0.41)				(0.26)	

Note: All specifications use a forecasting horizon of one quarter both for inflation and output gap. The instrument set includes two lags each of the inflation rate, output gap, policy rate, and M2 (M3 for Australia) growth; the Euro Area and Australia additionally include the term spread; exchange rate terms are instrumented with their own lags; the Canadian extended specification adds lags of the U.S. federal funds rate. HAC standard errors in parentheses. The rightmost column reports the p-value from the Sargan–Hansen over-identification test. *** (**, *) denotes significance at the 1% (5%, 10%) level.

Central banks elaborate more when inflationary pressure builds or when output falls below potential, precisely as in the Fed baseline. Intrinsic persistence varies but is generally modest; several economies display mildly negative ρ_e, suggesting that communication effort overshoots on impact and partially reverses at the next meeting, consistent with the meeting-by-meeting structure of policy deliberation. A second shared feature is that the

estimated inflation anchors π^* are recovered from communication data alone (without using interest-rate data) and yet they align closely with each central bank's publicly stated objective: approximately 2 percent for Australia and the Euro Area, around 1.8–2 percent for Sweden once the exchange rate term is included. Japan is the notable exception and deserves separate treatment.

Japan and the Euro Area as Institutional Tests Two economies provide especially clean tests of the rule's logic, because their institutional features make sharp predictions about which coefficients should be large and which should be small.

The Bank of Japan (BoJ) is the clearest case in the table. Its estimated inflation anchor is $\pi^* = 0.46$ (barely above zero) and yet its inflation response $\psi_\pi = 7.16$ is the largest in the sample. The BoJ thus elaborates intensely when inflation moves above its near-zero implicit anchor, precisely the behavior one would expect from a central bank that spent three decades fighting deflation and for which any upward deviation from price stability is a salient event requiring explanation. The negative and significant output-gap response ($\psi_y = -3.47$) is equally consistent with an institution attentive to real activity in a low-inflation environment in which the conventional rate instrument is frequently constrained.[17]

The Euro Area offers the mirror image. The European Central Bank's (ECB) primary mandate is price stability, with output stabilization explicitly secondary, and its communication rule reflects this distinction: $\psi_\pi = 3.82$ is large and significant, while $\psi_y = -0.27$ is negligible and statistically indistinguishable from zero. The estimated anchor of 2.24% sits close to the ECB's original "below but close to 2%" definition of price stability (later revised to a symmetric 2% target in 2021). The communication rule reads each central bank's mandate back from the text of its statements.

Exchange Rates in Small Open Economies For Australia and Sweden, the extended rule yields a positive and highly significant exchange rate coefficient ($\psi_{\Delta s} = 0.71$ and 0.17, respectively): a 1% expected depreciation raises the communication effort by 0.71% in Australia and 0.17% in Sweden. Small open economy central banks face a trade-off that large, relatively closed economies do not: exchange rate movements simultaneously affect import prices and the inflation outlook, but a single interest-rate adjustment cannot cleanly signal how the central bank is weighting these competing considerations. Words therefore do part of the work, clarifying the policy reaction function before any rate change is warranted. The inclusion of the exchange rate term also disciplines the domestic coefficients: for Sweden, the implied anchor rises from 0.87% in the baseline to 1.83% in the extended rule, a more plausible estimate for an economy with a 2% inflation target. The text record corroborates these findings: the term "exchange rate" appears 137 times in Reserve Bank of Australia statements and 208 times in Riksbank statements (approximately 0.7 per statement), against 6 occurrences in ECB statements and none in Bank of Japan statements, a pattern that mirrors the estimated $\psi_{\Delta s}$ coefficients almost perfectly.

[17] The primary mandate of the BoJ is price stability, but under article 2 of the Act, price stability is pursued as a means of contributing to the sound development of the national economy.

Canada and the Hierarchy of Monetary Communication Canada is the sharpest deviation from the common pattern in the international exercise. In the baseline specification, the inflation response is positive but insignificant, and the output-gap response is positive, the opposite sign relative to every other country in the table. The baseline rule is therefore at odds with the rules we found for the other countries. Given its proximity with the U.S., we include the U.S. federal funds rate in the extended version of the rule. Then, the picture changes significantly. The domestic inflation and output-gap coefficients remain small and insignificant, while $\psi_r = 0.97$ (significant at 5%) in the extended rule and 1.28 (significant at 1%) in the simplified specification that retains only the funds rate. The Bank of Canada calibrates the communication effort of its statements primarily to the stance of U.S. monetary policy, not to domestic macroeconomic conditions. Its communication is, to a first approximation, an exercise in explaining the American monetary policy environment rather than its own domestic stance. The result is not explained by U.S.-induced domestic spillovers: the significance of ψ_r survives the inclusion of domestic inflation and output gaps in the rule, pointing to an explicit informational dependence on the Federal Reserve's policy signal. The text record corroborates this: Bank of Canada statements mention "United States" 289 times over the estimation period (approximately 1.5 per statement), against 8 in ECB statements (0.03 per statement). The Canadian case thus reveals a hierarchical dimension of monetary communication that mirrors the well-known hierarchical structure of the international monetary system itself: the dominant central bank sets the informational agenda, and smaller, highly integrated economies adapt both their policy stance and the narrative explaining it accordingly. The international evidence therefore suggests the existence of an international communication cycle, analogous to the international credit cycle.

Overall cross-country evidence confirms that the existence of a systematic communication rule is not a U.S. peculiarity: forward-looking feedback on expected macroeconomic conditions governs the communication effort of every central bank in our sample. The rule adapts sensibly to each economy's structural features: mandate (the ECB's inflation focus and negligible output-gap response), history (the BoJ's near-zero anchor and heightened sensitivity to any inflation uptick), size and openness (Australia's and Sweden's exchange rate channel), and external dependence (Canada's dominant reaction to U.S. policy). The implied inflation anchors are recovered from communication alone, without the use of interest-rate data, and yet they reproduce each central bank's publicly announced target with considerable accuracy. This cross-economy consistency provides confirmation that communication and conventional monetary policy are governed by the same long-run preferences, and that the systematic nature of central bank communication is a feature of modern monetary policy frameworks, not an artifact of any particular institutional setting.

3. Systematic Communication in a Simple RANK Model

In this section, we develop a simple conceptual framework, a Communication Representative Agent New Keynesian (CRANK) model, to study how the systematic component of central bank communication propagates through the economy. Our approach embeds a communication rule into a New Keynesian environment. This lets us characterize, in

a transparent and tractable way, the macroeconomic consequences of the central bank's efforts to manage public expectations.

3.1 The Environment

The economy is populated by an infinitely-lived representative household, a competitive final good firm, a continuum of firms in the intermediate good sectors and a central bank. The economy is subject to two aggregate shocks, a demand shock and a cost push shock. In the lines of the imperfect-information literature (see, e.g. Woodford, 2003a, Lorenzoni, 2009, Blanchard, L'Huillier, and Lorenzoni, 2013), households cannot directly disentangle the persistent from the transitory component of the disturbances that hit the economy. Instead, they observe the realized shocks contaminated by transitory noise and receive a noisy public signal about their persistent components, whose precision depends fundamentally on the communication effort of the central bank.

3.1.1 Household Time is discrete and the future is discounted at rate $\beta \in (0,1)$. There is a representative household with preferences over consumption, c_t, and hours worked, n_t, represented by the following intertemporal utility function

$$\widetilde{\mathbb{E}}_0 \left[\sum_{t=0}^{\infty} \beta^t \left(\frac{c_t^{1-\gamma}}{1-\gamma} - \nu \frac{n_t^{1+\varphi}}{1+\varphi} \right) \right] , \tag{10}$$

where $1/\gamma > 0$ is the elasticity of intertemporal substitution, $\nu > 0$ is a labor dis-utility shifter and $1/\varphi > 0$ is the Frisch labor supply elasticity. $\widetilde{\mathbb{E}}_t[\cdot]$ denotes the expectation operator that departs from the full-information benchmark.

Information Structure A large body of evidence documents that private agents have imperfect real-time information about aggregate economic conditions (see Coibion and Gorodnichenko, 2015, for a survey). One reading of this evidence is not that agents fail to observe realized fundamentals, but that they cannot tell, in real time, how persistent a given movement in those fundamentals is. We capture this idea parsimoniously by assuming that each of the two aggregate disturbances— a cost-push shock, θ_t, and a demand (DIS) shock, X_t —is the product of a persistent and a purely transitory component:[18]

$$\widehat{\theta}_t \equiv -\Psi^{-1}(\log(\theta_t) - \log(\theta)) = \tau_t^{\theta} + \omega_t^{\theta} , \qquad \widehat{x}_t \equiv \log(X_t) - \log(\bar{X}) = \tau_t^{x} + \omega_t^{x} , \tag{11}$$

where the persistent components follow independent AR(1) processes

$$\tau_t^i = \rho_i \, \tau_{t-1}^i + \epsilon_t^i , \qquad \epsilon_t^i \sim \mathcal{N}(0, \sigma_i^2) , \qquad \rho_i \in (0,1) , \qquad i \in \{\theta, x\} , \tag{12}$$

and the transitory components are innovations, $\omega_t^i \sim \mathcal{N}(0, \varsigma_i^2)$, mutually and serially independent and independent of the ϵ_t^i. It is convenient to stack the persistent components into

[18] θ and X denote the unconditional average of both shock. The cost-push disturbance is expressed in (11) directly in the units in which it enters the Phillips curve, where $\Psi > -$ is the price-adjustment cost parameter introduced below, so that the inflationary disturbance is $-\Psi^{-1}(\log\theta_t - \log\theta)$. Defining $\widehat{\theta}_t$ as in (11) is therefore a pure normalization, which leaves the informational structure untouched and delivers the familiar NKPC.

the state vector $\boldsymbol{\tau}_t \equiv (\tau_t^\theta, \tau_t^x)'$, which evolves as

$$\boldsymbol{\tau}_t = \mathbf{D}_\rho \boldsymbol{\tau}_{t-1} + \boldsymbol{\epsilon}_t \,, \qquad \boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\epsilon) \,, \tag{13}$$

with $\mathbf{D}_\rho \equiv \text{diag}(\rho_\theta, \rho_x)$ and $\boldsymbol{\Sigma}_\epsilon \equiv \text{diag}(\sigma_\theta^2, \sigma_x^2)$, and to collect the transitory components in $\boldsymbol{\omega}_t \equiv (\omega_t^\theta, \omega_t^x)' \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\omega)$, $\boldsymbol{\Sigma}_\omega \equiv \text{diag}(\varsigma_\theta^2, \varsigma_x^2)$.

The household perfectly observes the realized fundamentals, *i.e.* the sums in (11), which we gather in the vector

$$\boldsymbol{z}_t \;\equiv\; (\widehat{\theta}_t,\, \widehat{x}_t)' \;=\; \boldsymbol{\tau}_t + \boldsymbol{\omega}_t \,. \tag{14}$$

She however cannot decompose each observation into its persistent and transitory parts. This distinction is important because only the persistent component is relevant for expectations. Indeed, the transitory component is unforecastable ($\mathbb{E}_t[\omega_{t+1}^i] = 0$) and does not propagate, so a household planning consumption and savings must form a view about $\boldsymbol{\tau}_t$ from data that mix it with transitory noise. The observation (14) is therefore itself a signal about $\boldsymbol{\tau}_t$, with signal noise $\boldsymbol{\omega}_t$ of fixed variance $\boldsymbol{\Sigma}_\omega$.

Central Bank Communication On top of the fundamentals, the agent receives a public signal, $\boldsymbol{s}_t$, issued by the central bank about the persistent component $\boldsymbol{\tau}_t$

$$\boldsymbol{s}_t = \boldsymbol{\tau}_t + \boldsymbol{\eta}_t \,, \qquad \boldsymbol{\eta}_t \sim \mathcal{N}\big(\mathbf{0},\, \boldsymbol{\Sigma}_\eta(\Delta e_t)\big) \,, \tag{15}$$

where $\boldsymbol{\Sigma}_\eta(\Delta e_t)$ denotes the variance-covariance matrix of the communication noise, which, through its communication effort Δe_t, is under the control of the central bank. More precisely, we assume that $\boldsymbol{\Sigma}_\eta(\Delta e_t)$ takes the form

$$\boldsymbol{\Sigma}_\eta(\Delta e_t) = \sigma_\eta^2(\Delta e_t) \begin{pmatrix} \varkappa_\theta & 0 \\ 0 & \varkappa_x \end{pmatrix} \quad \text{with} \quad \varkappa_\theta,\, \varkappa_x > 0 \,. \tag{16}$$

Hence, communication is a noisy message about the persistent state, $\boldsymbol{\tau}_t$, that agents cannot observe directly and its informativeness is the policy object. The noises are uncorrelated and, more importantly, the precision of the signal scales with the communication effort of the central bank through the common factor $\sigma_\eta^2(\Delta e_t)$. The function $\sigma_\eta^2(\Delta e_t) : \mathbb{R}_+ \to \mathbb{R}_{++}$ is twice continuously differentiable and strictly decreasing in the communication effort Δe_t. The loadings $\varkappa_i > 0$ are the relative volatilities of communication across the two disturbances and should not be confused with the transitory variances ς_i^2 of the fundamentals, which are a structural feature of the shocks rather than a policy instrument.

The signal is public and issued by the central bank, it is therefore common to all agents. This distinguishes our framework from models of dispersed private information (Lucas, 1972, Woodford, 2003b), in which the coordination motive of Morris and Shin (2002) implies that agents discount public signals in favor of their private ones. In this model, because the signal is common, individual posterior beliefs coincide with the aggregate posterior, implying that the beauty-contest dynamics that complicate dispersed information models are absent. The welfare implications of providing public information,

whether it crowds out useful private signals, do not arise in this environment. The relationship between communication effort and aggregate outcomes is therefore clean and monotone.[19]

Belief Formation Conditional on the information inherited from the previous period, $\mathcal{I}_{t-1}$, the household holds a Gaussian prior over the current persistent state, $\boldsymbol{\tau}_t \mid \mathcal{I}_{t-1} \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_\mu)$, where the prior mean is the persistence-adjusted posterior of the previous period, $\boldsymbol{\mu}_t \equiv \mathbb{E}_{t-1}[\boldsymbol{\tau}_t] = \mathbf{D}_\rho \widetilde{\boldsymbol{\tau}}_{t-1}$, and $\boldsymbol{\Sigma}_\mu$ is the (stationary) prior variance derived below. Within the period the household confronts two pieces of information, the fundamentals $\boldsymbol{z}_t$ and the communication signal $\boldsymbol{s}_t$, and updates optimally by Bayes' rule. Because it isolates the marginal role of communication, we present the update as two sequential Gaussian steps; the order is immaterial and the resulting posterior coincides with the joint update (see Appendix B.1).

First, the household reads the fundamentals. Combining the prior with $\boldsymbol{z}_t = \boldsymbol{\tau}_t + \boldsymbol{\omega}_t$ yields the pre-communication belief $\boldsymbol{\tau}_t \mid \boldsymbol{z}_t, \mathcal{I}_{t-1} \sim \mathcal{N}(\boldsymbol{m}_t, \boldsymbol{P})$ with

$$\boldsymbol{m}_t = \boldsymbol{\mu}_t + \boldsymbol{G}(\boldsymbol{z}_t - \boldsymbol{\mu}_t)\,, \quad \boldsymbol{P} = (\boldsymbol{I} - \boldsymbol{G})\boldsymbol{\Sigma}_\mu\,, \quad \boldsymbol{G} \equiv \boldsymbol{\Sigma}_\mu(\boldsymbol{\Sigma}_\mu + \boldsymbol{\Sigma}_\omega)^{-1}\,, \tag{17}$$

where $\boldsymbol{I}$ denotes the identity matrix, $\boldsymbol{G}$ is the gain on the fundamentals and $\boldsymbol{P}$ the residual variance after reading them. Second, the household reads the central bank's message. Combining the pre-communication belief with $\boldsymbol{s}_t = \boldsymbol{\tau}_t + \boldsymbol{\eta}_t$ delivers the posterior $\boldsymbol{\tau}_t \mid \boldsymbol{z}_t, \boldsymbol{s}_t, \mathcal{I}_{t-1} \sim \mathcal{N}(\widetilde{\boldsymbol{\tau}}_t, \widetilde{\boldsymbol{\Sigma}})$,

$$\widetilde{\boldsymbol{\tau}}_t = \boldsymbol{m}_t + \boldsymbol{K}(\boldsymbol{s}_t - \boldsymbol{m}_t)\,, \quad \widetilde{\boldsymbol{\Sigma}} = (\boldsymbol{I} - \boldsymbol{K})\boldsymbol{P} = \boldsymbol{K}\,\boldsymbol{\Sigma}_\eta(\Delta e_t)\,, \quad \boldsymbol{K} \equiv \boldsymbol{P}\big(\boldsymbol{P} + \boldsymbol{\Sigma}_\eta(\Delta e_t)\big)^{-1}\,, \tag{18}$$

where $\boldsymbol{K}$ is the communication Kalman gain and the identity $\widetilde{\boldsymbol{\Sigma}} = \boldsymbol{K}\boldsymbol{\Sigma}_\eta(\Delta e_t)$ is proved in Online Appendix B.1. Since all disturbances are independent, every matrix in (17)–(18) is diagonal and the filter operates shock by shock. Writing $\sigma^2_{\eta,i}(\Delta e_t) = \varkappa_i \sigma^2_\eta(\Delta e_t)$ for the i-th communication noise variance, the component gains and variances solve

$$G_i = \frac{\sigma^2_{\mu,i}}{\sigma^2_{\mu,i} + \varsigma_i^2}\,, \quad P_i = (1 - G_i)\sigma^2_{\mu,i}\,, \quad k_i = \frac{P_i}{P_i + \sigma^2_{\eta,i}(\Delta e_t)}\,, \quad \widetilde{\sigma}_i^2 = k_i\,\sigma^2_{\eta,i}(\Delta e_t)\,, \tag{19}$$

with the stationary prior variance pinned down by the fixed point $\sigma^2_{\mu,i} = \rho_i^2 \widetilde{\sigma}_i^2 + \sigma_i^2$, equivalently $\widetilde{\sigma}_i^2 = \big(1/\sigma^2_{\mu,i} + 1/\varsigma_i^2 + 1/\sigma^2_{\eta,i}(\Delta e_t)\big)^{-1}$, which admits a unique positive solution (see Appendix B.1). As is common in the literature, we assume that households update their beliefs with the steady-state Kalman gains G_i and k_i, evaluated at the stationary prior variance and at the steady-state communication effort Δe. Communication effort therefore affects the precision of the central bank's message, and hence the accuracy of beliefs, but not the weight households place on it. When communication is highly precise ($\sigma^2_\eta \to 0$), $\boldsymbol{K} \to \boldsymbol{I}$; agents adopt the message and, since the message concerns the persistent state

[19] The extension to dispersed private information, in which a fraction of agents also receive idiosyncratic signals about $\boldsymbol{\tau}_t$, would introduce higher-order belief dynamics (Angeletos and Lian, 2018) and potentially non-monotone effects of public communication.

directly, recover the FIRE belief $\widetilde{\boldsymbol{\tau}}_t \to \boldsymbol{\tau}_t$. When communication is entirely uninformative ($\sigma_\eta^2 \to \infty$), $\boldsymbol{K} \to \mathbf{0}$ and agents fall back on the pre-communication belief $\boldsymbol{m}_t$, which still exploits the fundamentals rather than the prior alone. In the general case, the gain $\boldsymbol{K}$ summarizes the quality of central bank communication in a single sufficient statistic.

The signal extraction error, the gap between the agent's belief about the current persistent state and its realization, is

$$\boldsymbol{\delta}_t \equiv \widetilde{\boldsymbol{\tau}}_t - \boldsymbol{\tau}_t = \boldsymbol{K}\,\boldsymbol{\eta}_t - (\boldsymbol{I} - \boldsymbol{K})\,\boldsymbol{\upsilon}_t\,, \tag{20}$$

where $\boldsymbol{\upsilon}_t \equiv \boldsymbol{\tau}_t - \boldsymbol{m}_t$ is the innovation in the persistent state relative to the pre-communication belief. As such this can be interpreted as a belief wedge. The error $\boldsymbol{\delta}_t$ captures two distinct phenomena. The first term, $\boldsymbol{K}\boldsymbol{\eta}_t$, captures the effect of communication noise. Agents partially attribute signal noise to a genuine movement in the persistent state, generating expectation errors that are driven entirely by the communication rather than by fundamentals. The second term, $-(\boldsymbol{I} - \boldsymbol{K})\boldsymbol{\upsilon}_t$, implies that agents underreact to genuine innovations in the persistent state by a factor $\boldsymbol{I} - \boldsymbol{K}$: having only imperfect help in separating the persistent movement $\boldsymbol{\upsilon}_t$ from the transitory noise embedded in the fundamentals, they revise their view of the state only partially. It is instructive to expand the innovation into its underlying sources, $\boldsymbol{\upsilon}_t = (\boldsymbol{I} - \boldsymbol{G})(\boldsymbol{\tau}_t - \boldsymbol{\mu}_t) - \boldsymbol{G}\boldsymbol{\omega}_t$: the household underreacts both to true persistent surprises $\boldsymbol{\tau}_t - \boldsymbol{\mu}_t$ and, with the opposite sign, mistakes part of the transitory noise $\boldsymbol{\omega}_t$ for a persistent movement. Both sources of error are decreasing in communication quality: as $\boldsymbol{K} \to \boldsymbol{I}$ the systematic underreaction disappears, and simultaneously $\boldsymbol{\Sigma}_\eta \to \mathbf{0}$ so that the noise term also vanishes. We finally have the following result.

Lemma 1 (Signal Extraction Error). *The signal extraction error $\boldsymbol{\delta}_t$ follows an autoregressive process of the form*

$$\boldsymbol{\delta}_t = D_\varrho \boldsymbol{\delta}_{t-1} + M_e \mathcal{E}_t\,,$$

where $D_\varrho = \mathrm{diag}(\varrho_\theta, \varrho_x)$ with $\varrho_i = \rho_i(1-\Gamma_i) < \rho_i$, where $1 - \Gamma_i \equiv (1-k_i)(1-G_i)$, $\mathcal{E}_t = (\varepsilon_t^\theta, \varepsilon_t^x, \omega_t^\theta, \omega_t^x, \eta_t^\theta, \eta_t^x)'$ and

$$M_e = \begin{pmatrix} -(1-k_\theta)(1-G_\theta) & 0 & (1-k_\theta)G_\theta & 0 & k_\theta & 0 \\ 0 & -(1-k_x)(1-G_x) & 0 & (1-k_x)G_x & 0 & k_x \end{pmatrix}.$$

Hence the error associated with each persistent component of the two structural shock is an exogenous AR(1) process which is driven by the innovations to the persistent and transient components of the structural shocks and the noise of the signal, whose volatility is under the control of the monetary authority. After a shock, the error converges back to its mean faster than the persistent components of the shocks since $\varrho_i < \rho_i$.

Household's Decision In every period, the representative household consumes and supplies hours of labor at the real wage w_t. She also receives transfers (real dividends and rebated costs) T_t from the production sector. The household can invest in a risk-free nominal bond b_{t+1} paying a gross nominal return R_t/X_t in period $t+1$, where $R_t = 1 + i_t$ is the gross nominal interest rate and $X_t = \exp(\tau_t^x + \omega_t^x)$ is an exogenous shock that acts as a DIS shock in the model,[20] whose persistent component τ_t^x follows the AR(1) process in (12). The agent therefore faces the budget constraint:

$$c_t + b_{t+1} = \frac{R_{t-1}}{X_{t-1}\pi_t} b_t + w_t n_t + T_t \,, \tag{21}$$

where $\pi_t \equiv P_t/P_{t-1}$ denotes gross price inflation. Because X_t enters the return on savings decided in t, it is observed by the household at t; what she does not observe is the split of $\log X_t$ into its persistent part τ_t^x and its transitory part ω_t^x. The household decides her consumption-saving plans and labor supply by maximizing (10) subject to (21).

3.1.2 Intermediate and Final Good Firms Production and price setting follow the standard New Keynesian setup.

Final Good Sector A competitive representative firm produces the consumption good y_t by bundling a continuum $j \in [0,1]$ of intermediate goods $y_t(j)$ using a Dixit-Stiglitz aggregator with time-varying elasticity of substitution, $\theta_t > 1$:

$$y_t = \left(\int y_t(j)^{\frac{\theta_t - 1}{\theta_t}} dj \right)^{\frac{\theta_t}{\theta_t - 1}} .$$

Profit maximization yields the demand for each intermediate variety:

$$y_t(j) = \left(\frac{p_t(j)}{p_t} \right)^{-\theta_t} y_t \,,$$

where $p_t = \left(\int_0^1 p_t(j)^{1-\theta_t} dj \right)^{\frac{1}{1-\theta_t}}$ is the aggregate price index. As discussed above, θ_t varies exogenously, with the associated cost-push wedge $\widehat{\theta}_t = -\Psi^{-1}(\log\theta_t - \log\theta) = \tau_t^\theta + \omega_t^\theta$ of (11).

Intermediate Good Sector Each intermediate good j is produced using labor $n_t(j)$, according to the production function:

$$y_t(j) = A n_t(j)^{1-\alpha} \,,$$

where $A > 0$ is a constant technology shifter and $\alpha \in [0,1)$ captures the degree of decreasing returns to scale. The corresponding real marginal costs are given by

$$mc_t(j) = \frac{w_t}{(1-\alpha)} \frac{n_t(j)}{y_t(j)} \,,$$

where the real wage is common to all firms.

[20] Here, we follow Smets and Wouters (2007) who interpret $\exp(-\hat{x}_t)R_t$ as the effective $t+1$ nominal return on bonds, resulting in a nominal bond premium of $(\exp(-\hat{x}_t) - 1)R_t$. An alternative rationale for this shock can be found in Fisher (2015); the exact representation is immaterial for the main results in this paper.

Price Setting Intermediate firms set nominal prices subject to quadratic price adjustment costs à la Rotemberg (1982). The real profit of each intermediate firm is

$$d_t(j) = \left(\frac{p_t(j)}{p_t}\right)^{-\theta_t} y_t \left[\frac{p_t(j)}{p_t} - mc_t(j)\right] - \frac{\Psi}{2}\left(\frac{p_t(j)}{p_{t-1}} - 1\right)^2 y_t \,,$$

where $\Psi \geq 0$ governs the magnitude of price adjustment costs, defined symmetrically around the zero-inflation steady state $\pi = 1$. We assume that this cost is rebated lump sum to households along with dividends.[21] Note that the adjustment is defined relative to the past aggregate price, p_{t-1}, rather than the past individual price. This assumption buys us tractability as it implies that the New Keynesian Phillips Curve (NKPC) is totally static.[22] Firms operate in the goods and labor markets and observe their own realized cost conditions; the imperfect-information friction bears on the household's expectations, not on contemporaneous price setting. Hence, in a symmetric equilibrium, the optimal price-setting condition writes

$$(\pi_t - 1)\pi_t = \frac{\theta_t}{\Psi}\left[mc_t - \frac{\theta_t - 1}{\theta_t}\right] . \tag{22}$$

3.1.3 The Monetary Authority The monetary authority is assumed to have superior information than the households, and can perfectly observe the current realization of the persistent state $\boldsymbol{\tau}_t$ (equivalently, it can decompose the fundamentals into their persistent and transitory parts) and therefore operates under FIRE. It also pursues its mandate through two systematic policy tools: *(i)* standard interest rate policy and *(ii)* communication policy.

Interest Rate Policy The first channel is conventional, i.e., the central bank sets the gross nominal interest rate R_t according to a static (contemporaneous) rule:

$$R_t = R \cdot (\pi_t/\pi)^{\phi_\pi} \cdot (y_t/y)^{\phi_y} \,. \tag{23}$$

The central bank follows a Taylor (1993) rule with positive feedback coefficients to current inflation and output, $\phi_\pi \geq 0$ and $\phi_y \geq 0$.[23] Because the central bank observes the current state, both π_t and y_t are in its information set when it sets R_t.

Communication Policy The second, and novel, channel operates through central bank communication. Specifically, variations in communication effort are governed by a communication rule inspired by, and consistent with, our empirical findings:[24]

$$\exp(\Delta e_t) = (\mathbb{E}_t[\pi_{t+1}]/\pi)^{\psi_\pi} \cdot (\mathbb{E}_t[y_{t+1}]/y)^{-\psi_y} \,, \tag{24}$$

[21] A direct implication of this assumption is that, in a symmetric equilibrium, the aggregate resource constraint of the economy is given by $c_t = y_t$.

[22] This assumption is made only for convenience as it simplifies the algebra, but does not affect qualitatively our results.

[23] In Section 3.3.2 we will also consider the possibility that the central bank hits the zero lower bound, in which case the nominal interest rate $i_t = R_t - 1$ is set at zero.

[24] Implicit in this formulation is that Δe is normalized to 0 in steady state.

where $\psi_\pi \geq 0$ and $\psi_y \geq 0$ denote the degree with which the central bank adjusts its communication to expected inflation and output respectively.[25]

Two features of this specification deserve explicit note. First, the rule does not feature any mean reversion term. This reflects our empirical finding that communication effort displays no statistically significant tendency to revert to a long-run mean, so that only the systematic response to the macroeconomic outlook drives its dynamics. Second, in line with our empirical results, the inflation target π that enters the communication rule (24) coincides with the one appearing in the Taylor rule (23), suggesting that the two instruments are anchored to a common objective.

The communication rule (24) captures two distinct motives for central bank communication, both consistent with its dual mandate. First, the central bank communicates more intensely when inflation is expected to rise ($\psi_\pi \geq 0$). By providing clearer signals about its policy intentions, it works to anchor inflation expectations and eliminate self-fulfilling inflationary spirals. Second, the central bank increases communication effort when the economy is expected to contract ($\psi_y \geq 0$): in the face of a looming recession, higher communication effort reassures the public.

3.2 The Dynamic IS Equation

This section aims at deriving the dynamic IS equation and makes explicit how communication effort will affect aggregate demand. Using the good market clearing condition $c_t = y_t$, the household's Euler equation rewrites as

$$y_t^{-\gamma} = \beta \frac{R_t}{X_t} \widetilde{\mathbb{E}}_t \left[\frac{y_{t+1}^{-\gamma}}{\pi_{t+1}} \right] , \tag{25}$$

where $\widetilde{\mathbb{E}}_t[\cdot]$ is the expectation operator under the posterior distribution (18). Denoting witg a hat $\widehat{\ }$ the log deviation of a variable from its deterministic level[26], the Euler equation rewrites

$$\exp(-\gamma \widehat{y}_t) = \exp(\widehat{\imath}_t - \widehat{x}_t) \widetilde{\mathbb{E}}_t \left[\exp(-\gamma \widehat{y}_{t+1} - \widehat{\pi}_{t+1}) \right] .$$

Taking logs and applying the moment generating function of a log-normal distribution, the Euler equation rewrites as the sum of first and second conditional moments

$$-\gamma \widehat{y}_t = \widehat{\imath}_t - \widehat{x}_t - \gamma \widetilde{\mathbb{E}}_t[\widehat{y}_{t+1}] - \widetilde{\mathbb{E}}_t[\widehat{\pi}_{t+1}] + \Omega_t , \tag{26}$$

where

$$\Omega_t \equiv \frac{\gamma^2}{2} \widetilde{\sigma}^2_{\widehat{y}_{t+1}|t} + \frac{1}{2} \widetilde{\sigma}^2_{\widehat{\pi}_{t+1}|t} + \gamma \widetilde{\sigma}_{\widehat{y}_{t+1} \widehat{\pi}_{t+1}|t}$$

can be interpreted as the risk component. Crucially, the imperfect information bears solely on the conditional moments of future outcomes, which the household forms from her posterior over the persistent state $\boldsymbol{\tau}_t$. We then obtain the following result.

[25] Appendix D analyses a variant where the communication depends on current economic outcomes rather than expected.

[26] In particular, $\widehat{x}_t = \tau_t^x + \omega_t^x$ and $\widehat{\theta}_t = \tau_t^\theta + \omega_t^\theta$, as defined in (11), are the (observed) realized shocks

PROPOSITION 1 (Risk as Communication). *Up to a first order approximation , the risk component writes*

$$\Omega_t \approx -\chi_e \widehat{\Delta} e_t \,, \tag{27}$$

where, the communication effectiveness, χ_e, takes the form

$$\chi_e \equiv -\Upsilon \frac{\partial \sigma_\eta^2(\Delta e_t)}{\partial \Delta e_t}\bigg|_{\Delta e_t = \overline{\Delta e}} > 0 \,. \tag{28}$$

As shown in the Online Appendix B.4, communication effectiveness[27] can be given a structural interpretation as products of two distinct objects. The first, Υ, is a constant term that is related to the sensitivity of future log-marginal utility, the properties of the fundamental shocks and the signals, and the Kalman gains.[28] The second and more important object is the marginal reduction in the common signal variance per unit of communication effort, $\partial \sigma_\eta^2(\Delta e_t)/\partial \Delta e_t|_{\overline{\Delta e}}$, which characterizes communication technology and is negative as more communication effort improves the precision of the signal.

Using the posterior distribution and substituting (27) into the log-linearized Euler equation yields a characterization of the log-linear version of the dynamic IS equation.

PROPOSITION 2 (Dynamic IS Equation). *To first order around the deterministic steady state, the DIS equation is given by:*

$$\widehat{y}_t = \mathbb{E}_t[\widehat{y}_{t+1}] - \frac{1}{\gamma}\left(\widehat{\imath}_t - \mathbb{E}_t[\widehat{\pi}_{t+1}]\right) + \Lambda_\delta \boldsymbol{\delta}_t + \frac{\chi_e}{\gamma}\widehat{\Delta} e_t + \frac{\widehat{x}_t}{\gamma} \,, \tag{29}$$

where $\Lambda_\delta = \frac{\gamma \boldsymbol{\Phi}_y + \boldsymbol{\Phi}_\pi}{\gamma}$, where the row vectors $\boldsymbol{\Phi}_y$ and $\boldsymbol{\Phi}_\pi$ are the one-step-ahead responses of output and inflation to the persistent state, holding the household's information fixed.

Relative to the standard New Keynesian model, the DIS equation contains two additional terms that capture the two channels through which imperfect information about the persistent component of the fundamental shocks enters aggregate demand.

The Mean Distortion Channel. Plugging the definition (Equation 20) of the signal extraction error $\boldsymbol{\delta}_t$ into the mean distortion channel, $\Lambda_\delta \boldsymbol{\delta}_t$, the latter decomposes into two components. The first one, $\Lambda_\delta \boldsymbol{K} \boldsymbol{\eta}_t$, captures a noise-driven mismatch between agents' point forecasts of future outcomes and the FIRE benchmark. Because agents partially attribute communication noise to a genuine movement in the persistent state, their expectations of future outcomes are distorted even when $\boldsymbol{\upsilon}_t = \mathbf{0}$, generating spurious demand fluctuations reminiscent of the noise-driven business cycles (see Lorenzoni, 2009, Blanchard, L'Huillier, and Lorenzoni, 2013, among others). Importantly, as the signal becomes more precise

[27] The communication effectiveness may result either from the ability of the central bank to issue better or more understandable communication, in the sense of the Flesch-Kincaid index for example, or from more attentive households, which would be reflected in the Kalman gain coefficient in an attention model.

[28] While this is essentially immaterial for our results, the interested reader is referred to Section B.4.

$(\sigma_\eta^2 \to 0)$, the noise itself vanishes and this channel closes. The two components of the mean distortion channel therefore attenuate together as communication effort increases.

The second, $-\Lambda_\delta(\boldsymbol{I}-\boldsymbol{K})\boldsymbol{\upsilon}_t$, captures a systematic mismatch between agents' point forecasts of future outcomes and the FIRE benchmark. When $\boldsymbol{\upsilon}_t > \mathbf{0}$ hits the economy, a fully informed agent—who can tell the persistent movement apart from the transitory noise in the fundamentals—would revise her expectations of future consumption and inflation upward by $\boldsymbol{\Phi}_y\boldsymbol{\upsilon}_t$ and $\boldsymbol{\Phi}_\pi\boldsymbol{\upsilon}_t$ respectively, and adjust current consumption accordingly. An imperfectly informed agent, who cannot fully disentangle the persistent movement from transitory noise even with the help of communication, revises her forecasts upward by only a fraction $\boldsymbol{K}$ of the full amount; a phenomenon Coibion and Gorodnichenko (2012) document empirically as "information rigidity" in professional forecasters. This underreaction reduces the current consumption response to persistent shocks by a factor $(\boldsymbol{I}-\boldsymbol{K})$, generating the dampening term in (29). The mechanism is isomorphic in form to the sticky-information model of Mankiw and Reis (2002). The degree of discounting, $(\boldsymbol{I}-\boldsymbol{K})$, is a policy variable rather than a fixed structural parameter: the central bank can directly calibrate the sensitivity of current demand to future expectations by adjusting the precision of its public signals about the persistent state.

The Risk Channel The term $\chi_e\widehat{\Delta}e_t/\gamma$ is the risk channel. It captures the direct stimulative effect of communication effort on aggregate demand through the compression of the risk component Ω_t. Higher communication effort reduces the posterior variance $\widetilde{\sigma}_i^2$ about the persistent state, lowering the risk component in the Euler equation and making current consumption less sensitive to precautionary motives. The effect is proportional to communication effectiveness χ_e, defined in (28). This channel is conceptually related to a risk-based demand channel, where the relevant risk is aggregate risk about the persistent component of the shock, and communication resolves it.

3.3 Equilibrium Dynamics

In the general equilibrium, the log-linear representation of the model is summarized by the following dynamic system

$$\widehat{y}_t = \mathbb{E}_t[\widehat{y}_{t+1}] - \frac{1}{\gamma}\left(\widehat{\imath}_t - \mathbb{E}_t[\widehat{\pi}_{t+1}]\right) + \Lambda_\delta\boldsymbol{\delta}_t + \frac{\chi_e}{\gamma}\widehat{\Delta}e_t + \frac{\widehat{x}_t}{\gamma}\,, \tag{30}$$

$$\widehat{\pi}_t = \kappa\widehat{y}_t + \widehat{\theta}_t\,, \tag{31}$$

$$\widehat{\imath}_t = \phi_\pi\widehat{\pi}_t + \phi_y\widehat{y}_t\,, \tag{32}$$

$$\widehat{\Delta}e_t = \psi_\pi\mathbb{E}_t[\widehat{\pi}_{t+1}] - \psi_y\mathbb{E}_t[\widehat{y}_{t+1}]\,. \tag{33}$$

Equation (30) is the standard dynamic IS equation extended for the presence of imperfect information and communication, with the observed DIS shock $\widehat{x}_t = \tau_t^x + \omega_t^x$ entering contemporaneously and the signal extraction error $\boldsymbol{\delta}_t$. Equation (31) is the New Keynesian Phillips curve obtained by log-linearizing (22), with the observed cost-push shock $\widehat{\theta}_t = \tau_t^\theta + \omega_t^\theta$ of (11) and slope $\kappa = \Psi^{-1}(\theta-1)\left[(1+\varphi)/(1-\alpha)+\gamma-1\right] > 0$. Equations (32) and (33) capture the monetary and communication policies available to the central banker.

The definition of the DIS shock, $\widehat{x}_t$, and the cost push shock, $\widehat{\theta}_t$, together with the law of motion (13) of the persistent state, completes the description of the equilibrium dynamics.

Policy Substitutability The first implication of the DIS equation (30) is a precise mapping between the two policy instruments. In the DIS equation, a unit increase in the real interest rate reduces output by $1/\gamma$, while a unit increase in communication raises output by χ_e/γ. The two instruments are therefore locally perfect substitutes, with an exchange rate equal to χ_e. One unit of interest rate tightening is exactly offset by $1/\chi_e$ units of expansionary communication. This instrument duality grounds the substitution between conventional monetary policy and communication observed during the post-2008 period (the Federal Reserve's guidance between 2008 and 2015 and the ECB's Draghi-era communication are canonical examples). The effectiveness of this substitution depends on the communication effectiveness χ_e. Institutions whose communication is impaired, due to fragmented messaging or a history of reneging on guidance, require a disproportionately greater communication effort to replicate the demand effect of a given real rate move. This substitutability extends beyond contemporaneous monetary policy actions to the future path of policy. In particular, forward guidance regarding the future trajectory of real interest rates can equivalently be implemented through announcements about the future precision of public signals. The effectiveness of interest-rate forward guidance depends on the informativeness of the accompanying public signals as announcements about future policy rates affect expectations only insofar as the precision of those signals does not deteriorate.

Discounting and Forward Guidance Plugging the communication rule into the DIS equation, and making use of the NKPC in the communication rule, the DIS equation rewrites

$$\widehat{y}_t = \Theta \mathbb{E}_t[\widehat{y}_{t+1}] - \frac{1}{\gamma}(\widehat{\imath}_t - \mathbb{E}_t[\widehat{\pi}_{t+1}]) + \frac{\chi_e \psi_\pi}{\gamma}\mathbb{E}_t[\widehat{\theta}_{t+1}] + \frac{\widehat{x}_t}{\gamma} + \Lambda_\delta \boldsymbol{\delta}_t \,, \tag{34}$$

where $\Theta \equiv 1 + \frac{\chi_e}{\gamma}(\kappa\psi_\pi - \psi_y)$ acts as a discounting/compounding factor which, as shown in the following proposition, is shaped by central bank communication policy.

Proposition 3 (DIS Equation and Discounting). *The DIS equation features discounting* $(0 < \Theta < 1)$*, if and only if* $\kappa\psi_\pi < \psi_y < \kappa\psi_\pi + \frac{\gamma}{\chi_e}$*. When* $\psi_\pi\kappa > \psi_y$*, the DIS equation features compounding* $(\Theta > 1)$*.*

Proposition 3 shows that the cyclicality of communication effort affects the effective discount factor in the DIS equation. For our standard model parameterization (see below), the upper bound for ψ_y is very large (around 15), and well beyond the estimates we obtained in our empirics (around 1.5). We therefore focus on the other side of the inequality, which determines the cyclicality of communication effort. Using the NKPC, the communication rule rewrites as

$$\widehat{\Delta} e_t = (\psi_\pi\kappa - \psi_y)\mathbb{E}_t[\widehat{y}_{t+1}] + \psi_\pi \mathbb{E}_t[\widehat{\theta}_{t+1}] \,.$$

Therefore, the sign of $\psi_\pi\kappa - \psi_y$ (or identically the sign of $\Theta - 1$) determines the cyclicality of communication. When communication is sufficiently responsive to expectations of future

output developments ($\psi_y > \kappa\psi_\pi$), communication effort becomes countercyclical and the DIS equation features discounting. Compounding obtains when communication responds relatively little to output expectations and, overall, the communication effort is procyclical.

This condition can be given a clear and intuitive interpretation. Consider an environment in which future output is expected to be high. Through the NKPC, high expected output translates into rising expected inflation. The communication rule then pulls in two directions simultaneously. The inflation-fighting motive leads the central bank to communicate more intensely, as captured by $\psi_\pi\kappa$, which stimulates demand expectations and amplifies the anticipated expansion. The output-stabilization motive pulls in the other direction, prescribing a reduction in expansionary communication, as captured by ψ_y, which dampens current demand and counteracts the boom. Discounting arises when this dampening force dominates: $\psi_y > \psi_\pi\kappa$.

Importantly, discounting requires a genuine *dual mandate*. The central bank must respond to output conditions in its communication (ψ_y large) more than it does to the endogenous inflation that high output generates through the Phillips curve. The parameter κ mediates this second effect; a steeper Phillips curve makes inflation a more powerful indirect trigger for communication, eroding the stabilizing force of the output channel. Second, discounting arises from the systematic component of communication policy rather than from households' expectational errors: for a given χ_e, Θ does not depend on the gains G_i and k_i that shape the signal extraction error. It does require imperfect information, though, since the risk channel operates only if households face residual uncertainty about the persistent state; under FIRE, $\chi_e = 0$ and $\Theta = 1$ (see (41) below). A central bank can therefore reshape the propagation mechanism of expectations and suppress the scope for self-fulfilling amplification through the design of its communication strategy alone.

By generating a discounted dynamic IS equation, our framework contributes to the broader literature on New Keynesian anomalies, including the forward guidance puzzle (Del Negro, Giannoni, and Patterson, 2023) and the muted macroeconomic response observed during prolonged episodes at the effective lower bound (Cochrane, 2018, Debortoli, Galí, and Gambetti, 2020). Existing explanations for discounted intertemporal dynamics fall into three broad categories: distributional channels, such as procyclical income inequality and income risk (Werning, 2015, Acharya and Dogra, 2020, Bilbiie, 2025); information-based and behavioral channels, including cognitive discounting (Gabaix, 2020) and sticky or dispersed information (Meichtry, 2023, Gallegos, 2024); and structural features of household balance sheets or preferences, such as utility from wealth holdings (Michaillat and Saez, 2021), life-cycle considerations (Del Negro, Giannoni, and Patterson, 2023) and household debt with endogenous default premia (Beaudry and Portier, 2018). The DIS equation (34) offers a distinct channel, countercyclical central bank communication. While the resulting reduced-form discounting is formally analogous to information-based and behavioral approaches such as Gabaix (2020) and Meichtry (2023), the economic interpretation differs in at least two fundamental ways. In our setting, discounting arises, first, from the conduct of monetary policy itself through state-dependent communication effort and, second, through the revision of beliefs. This implies that central banks can partially attenuate the effectiveness of their own forward guidance through the cyclical properties of their communication strategies. This complements recent work by Eggertsson and Schüle (2024) and

Ahn and Holm (2025) arguing that central banks did not commit to remain unresponsive to inflation and output dynamics between the announcement of future policy cuts and their implementation. In our case, the interim adjustment is operated through communication rather than shifts in the policy rate.

3.3.1 Normal Times: Determinacy and Propagation We start by studying the dynamics of the economy when the central bank is solely confronted with *business as usual* fluctuations. We characterize the conditions for local determinacy when both instruments are available, showing how systematic communication effort can alter the Taylor principle, and then examine the economy's dynamic response to structural disturbances. Because the signal extraction error $\boldsymbol{\delta}_t$ enters (30) as an exogenous forcing term, it does not affect the determinacy analysis.

PROPOSITION 4 (Local Determinacy). *Using the NKPC together with both systematic policy rules, output dynamics are purely forward-looking and given by:*

$$\widehat{y}_t = \mathcal{M}\,\mathbb{E}_t[\widehat{y}_{t+1}] + \frac{1}{\gamma + \kappa\phi_\pi + \phi_y}\left[(1+\chi_e\psi_\pi)\mathbb{E}_t[\widehat{\theta}_{t+1}] - \phi_\pi\widehat{\theta}_t + \widehat{x}_t + \gamma\Lambda_\delta\boldsymbol{\delta}_t\right], \tag{35}$$

where $\mathcal{M} \equiv \frac{\gamma\Theta+\kappa}{\gamma+\kappa\phi_\pi+\phi_y}$. *Equation* (35) *admits a locally determinate solution if and only if* $|\mathcal{M}| < 1$, *that is*

$$-2\gamma - \kappa(1+\phi_\pi) - \phi_y \;<\; \chi_e\left[\kappa\psi_\pi - \psi_y\right] \;<\; \kappa(\phi_\pi - 1) + \phi_y\,. \tag{36}$$

The upper inequality in (36) is the generalization of the Taylor principle to a contemporaneous-data rule (see Bullard and Mitra, 2002). Setting $\psi_\pi = \psi_y = 0$ collapses it to the textbook condition $\kappa(\phi_\pi - 1) + \phi_y > 0$, while the lower inequality becomes $-2\gamma - \kappa(1+\phi_\pi) - \phi_y < 0$ and holds trivially. The operative requirement is therefore the upper bound, and communication modifies it by the term, $\chi_e[\kappa\psi_\pi - \psi_y]$. Hence, a dual-mandate communicator with $\psi_y > \kappa\psi_\pi$ makes the term negative and relaxes the Taylor principle, so that a policy stance that would otherwise deliver indeterminacy is rescued by communication. On the contrary, an inflation-only communicator with $\psi_y = 0$ makes it positive and tightens the principle, forcing more aggressive interest-rate policy to preserve determinacy. The latter case bears on institutions, such as the historical ECB, whose communication was disproportionately oriented toward price stability (see, e.g., Ehrmann and Fratzscher (2007) and the empirical section). Formally, the lower inequality in (36) places a bound not on the interest-rate response but on how countercyclical communication may be. To see this, it is useful to rearrange it as $\psi_y < \kappa\psi_\pi + [2\gamma + \kappa(1+\phi_\pi) + \phi_y]/\chi_e$, which for the calibrated parameters (see below) exceeds 33 and never binds in practice.

The next proposition reports the closed form solution for output dynamics.

PROPOSITION 5 (Output Dynamics). *At the locally determinate solution, the output dynamics are linear in the persistent components,* τ_t^i, *the signal extraction errors,* δ_t^i, *and the transitory components,* ω_t^i*:*

$$\widehat{y}_t \;=\; \zeta_\theta\,\tau_t^\theta + \zeta_x\,\tau_t^x \;+\; \xi_\theta\,\delta_t^\theta + \xi_x\,\delta_t^x \;-\; \frac{\phi_\pi}{\Delta}\,\omega_t^\theta + \frac{1}{\Delta}\,\omega_t^x\,, \tag{37}$$

where the cost-push and DIS loadings are given by

$$\zeta_\theta = \frac{(1+\chi_e\psi_\pi)\rho_\theta - \phi_\pi}{\Delta(1-\rho_\theta\mathcal{M})}\,, \quad \xi_\theta = \frac{\rho_\theta\big[(\gamma+\kappa)\zeta_\theta + 1\big]}{\Delta_\theta}\,, \tag{38}$$

$$\zeta_x = \frac{1}{\Delta(1-\rho_x\mathcal{M})}\,, \quad \xi_x = \frac{\rho_x(\gamma+\kappa)\zeta_x}{\Delta_x}\,, \tag{39}$$

with the auxiliary parameters

$$\Delta = \gamma + \kappa\phi_\pi + \phi_y > 0\,, \quad \textit{and} \quad \Delta_i \equiv \Delta + \chi_e\big(\psi_y - \kappa\psi_\pi\big)\rho_i(1-\Gamma_i) > 0 \quad \textit{for } i \in \{\theta, x\}\,. \tag{40}$$

Proposition 5 shows that output dynamics can be decomposed into three main forces. First, the loadings ζ_i capture the response of output to the persistent state when the household holds correct beliefs about it, that is, when the mean distortion channel is absent ($\delta_t^i = 0$).[29] For a given degree of communication effectiveness χ_e, these loadings are unaffected by communication quality, since the Kalman gains G_i and k_i do not enter their expressions. However, communication as a *systematic policy* affects them through $\mathcal{M}$ and, in addition, through $\chi_e\psi_\pi$ in the numerator of the cost-push loading. Second, the imperfect-information friction operates through the two additional state variables δ_t^i, whose loadings ξ_i depend on systematic central bank communication through ζ_i and Δ_i. Third, output dynamics also depend on the purely transitory shock components ω_t^i, whose loadings are unaffected by communication.

Given the NKPC and the systematic policy rules (31)-(33), inflation, the nominal interest rate, and communication effort can be recovered as

$$\widehat{\pi}_t = \kappa\widehat{y}_t + \widehat{\theta}_t\,, \qquad \widehat{\imath}_t = \phi_\pi\widehat{\pi}_t + \phi_y\widehat{y}_t\,, \qquad \widehat{\Delta e}_t = \big(\kappa\psi_\pi - \psi_y\big)\mathbb{E}_t[\widehat{y}_{t+1}] + \psi_\pi\rho_\theta\tau_t^\theta\,. \tag{42}$$

More generally, the impulse responses of output, $m_{y,i}(h)$, and inflation, $m_{\pi,i}(h)$, h periods after an initial shock to the persistent component of the fundamental, ϵ_0^i, are given by

$$m_{y,i}(h) \equiv \rho_i^h\left[\zeta_i - \xi_i(1-\Gamma_i)^{h+1}\right]\,, \quad m_{\pi,i}(h) \equiv \kappa m_{y,i}(h) + \rho_\theta^h\mathbb{1}_{\{i=\theta\}}\,. \tag{43}$$

[29] Because they incorporate communication as a policy rule, the loadings ζ_i should not be confused with those in an economy operating under FIRE. Under FIRE, the household observes $\boldsymbol{\tau}_t$, so that $\boldsymbol{\delta}_t \equiv \mathbf{0}$ and her posterior variance $\widetilde{\boldsymbol{\Sigma}}$ is zero, regardless of communication effort. Since communication affects the risk component Ω_t only through $\widetilde{\boldsymbol{\Sigma}}$ (see Proposition 1), the risk channel vanishes, $\chi_e = 0$, and communication has no effect on allocations, regardless of the policy rule. The output multipliers on the persistent state are then given by (see Appendix B.10):

$$\zeta_\theta^{\text{FIRE}} = \frac{\rho_\theta - \phi_\pi}{\Delta\,(1-\mathcal{M}^{\text{FIRE}}\rho_\theta)}\,, \qquad \zeta_x^{\text{FIRE}} = \frac{1}{\Delta\,(1-\mathcal{M}^{\text{FIRE}}\rho_x)}\,, \qquad \mathcal{M}^{\text{FIRE}} \equiv \frac{\gamma+\kappa}{\Delta}\,. \tag{41}$$

Under the textbook Taylor principle, $\kappa(\phi_\pi - 1) + \phi_y > 0$, which implies $\mathcal{M}^{\text{FIRE}} < 1$, we obtain $\zeta_\theta^{\text{FIRE}} < 0$ and $\zeta_x^{\text{FIRE}} > 0$. Proposition 5 also makes explicit the objects left implicit in Proposition 2:

$$\Phi_{y,i} = \rho_i\zeta_i - \rho_i(1-\Gamma_i)\xi_i\,, \qquad \Phi_{\pi,i} = \kappa\Phi_{y,i} + \rho_\theta\mathbb{1}_{\{i=\theta\}}\,, \qquad \Lambda_{\delta,i} = \frac{\Delta\,[1-\mathcal{M}\rho_i(1-\Gamma_i)]}{\gamma}\xi_i\,.$$

The output response reflects two potentially opposing forces: the response to the persistent state of the fundamental, captured by ζ_i, and the response to the signal extraction error, captured by ξ_i, which makes the mean distortion channel explicit. Under full information ($\delta_t^i = 0$), the shock propagates at a rate determined by its own persistence, ρ_i. Imperfect information introduces a distortion in this propagation, governed by the persistence of the signal extraction error, $\varrho_i = (1-\Gamma_i)\rho_i$. Since $0 < \Gamma_i < 1$ implies $\varrho_i < \rho_i$, the belief distortion dies out faster than the shock, and agents eventually learn. The mean distortion is therefore largest on impact ($h = 0$), when the filter leaves a fraction $1-\Gamma_i$ of the innovation unrevised. Relative to the impact response in the same economy without the mean distortion, namely ζ_i, the impact response of output is

$$\frac{m_{y,i}(0)}{\zeta_i} = 1 - \mathcal{R}_i\,, \quad \text{with} \quad \mathcal{R}_i \equiv \frac{\xi_i(1-\Gamma_i)}{\zeta_i}\,. \tag{44}$$

On impact, beliefs underreact by a factor of $1-\Gamma_i$, since the household updates her view of the persistent component of the shock by only a fraction Γ_i of the true innovation. For the DIS shock, (39) implies that $\mathcal{R}_x = (\gamma+\kappa)\rho_x(1-\Gamma_x)/\Delta_x$, which lies in $[0,1)$ because $\Delta_x = \Delta(1-\mathcal{M}\varrho_x) + (\gamma+\kappa)\rho_x(1-\Gamma_x)$ and $|\mathcal{M}| < 1$. Output is dampened, but by less than beliefs: at our calibration, $\mathcal{R}_x < 1-\Gamma_x$. This is because imperfect information affects only the expectational block.[30] For the cost-push shock, $\mathcal{R}_\theta$ can be negative, in which case the mean distortion amplifies the impact response, as in the numerical illustration below.

The following corollary establishes the signs of the output and inflation responses at all horizons.

Corollary 1 (Sign of Output and Inflation Responses). *Define two thresholds $\underline{\phi}_\pi$ and $\overline{\phi}_\pi$, with $\underline{\phi}_\pi < \overline{\phi}_\pi$. The following results hold:*[31]

1. *The output and inflation responses to a DIS shock are positive at all horizons, i.e., $m_{\pi,x}(h) > 0$ and $m_{y,x}(h) > 0$ $\forall h \geq 0$.*

2. *The inflation response to a cost-push shock is positive at all horizons, i.e., $m_{\pi,\theta}(h) > 0$ $\forall h \geq 0$. The sign of the output response depends on three cases: (i) if $\phi_\pi \geq \overline{\phi}_\pi$, $m_{y,\theta}(h) < 0$ $\forall h \geq 0$; (ii) if $\phi_\pi \leq \underline{\phi}_\pi$, $m_{y,\theta}(h) \geq 0$ $\forall h \geq 0$; (iii) if $\underline{\phi}_\pi < \phi_\pi < \overline{\phi}_\pi$, there exists $h^*_\theta > 0$ such that $m_{y,\theta}(h) \lessgtr 0$ for $h \lessgtr h^*_\theta$.*

The first part of Corollary 1 shows that a positive innovation to the persistent component of the DIS shock raises both output and inflation at all horizons. Although the loadings on the persistent state ($\zeta_x > 0$) and the signal extraction error ($\xi_x > 0$) work in opposite directions, the former dominates the propagation of the shock.[32]

The response to a positive innovation to the persistent component of the cost-push shock is more nuanced. Inflation rises at all horizons, but the sign of the output response

[30] The realized disturbances $\widehat{\theta}_t$ and $\widehat{x}_t$ are observed and enter the NKPC and the DIS equation at full strength.

[31] Their exact values are provided in Online Appendix XX.

[32] ζ_x and ξ_x are well defined under the generalized Taylor principle of Proposition 4. Importantly, $|\mathcal{M}| < 1$ guarantees $1 - \mathcal{M}\rho_i > 0$, so that ζ_i has the sign of its numerator, implying $\zeta_x > 0$.

depends on how aggressively the central bank raises the nominal interest rate in response to inflation. If the central bank is sufficiently hawkish ($\phi_\pi > \overline{\phi}_\pi$), output falls, preserving the standard inflation-output trade-off of the New Keynesian model under FIRE. By contrast, if the central bank is sufficiently dovish ($\phi_\pi < \underline{\phi}_\pi$), output rises. For intermediate values ($\underline{\phi}_\pi < \phi_\pi < \overline{\phi}_\pi$), output initially falls before turning positive after a few quarters. This distinction follows from the loadings on the persistent innovation and the signal extraction error in Proposition 5. The numerators of both ζ_θ and ξ_θ become positive when ϕ_π is sufficiently low. A positive output response to a cost-push shock requires systematic central bank communication that responds to inflation, together with sufficiently effective communication. Formally, the condition is $\phi_\pi < (1 + \chi_e \psi_\pi)\rho_\theta$. Without communication about inflation, this condition cannot hold under the Taylor principle $\phi_\pi > 1$. However, countercyclical communication ($\kappa\psi_\pi < \psi_y$) allows our economy to remain locally determinate even when $\phi_\pi < 1$. The latter two cases can therefore arise only in a model with systematic central bank communication.

Communication as Stabilization Proposition 5 shows that systematic central bank communication affects output dynamics, and hence inflation dynamics, through two distinct channels. First, it changes the loadings on the persistent components, ζ_i. Second, it changes the loadings on the signal extraction errors, ξ_i. Corollary 2 compares the output and inflation responses in two economies. The first features cyclical communication, $\Theta \neq 1$, and is denoted by the superscript C. The second features no (or a-cyclical) communication, $\Theta = 1$, with $\psi_y = 0, \psi_\pi = 0$, and is denoted by the superscript NC.

Corollary 2 (Stabilization vs Amplification). *The following two statements hold.*

1. *Following a DIS shock, we obtain for any horizon h:*

$$sgn(m^C_{y,x}(h) - m^{NC}_{y,x}(h)) = sgn(m^C_{\pi,x}(h) - m^{NC}_{\pi,x}(h)) = sgn(\Theta - 1)\,.$$

2. *Following a cost push shock, if $\phi_\pi \geq \rho_\theta$, we obtain for any horizon h:*

$$sgn(m^C_{y,\theta}(h) - m^{NC}_{y,\theta}(h)) = sgn(m^C_{\pi,\theta}(h) - m^{NC}_{\pi,\theta}(h)) > 0\,.$$

The first part of Corollary 2 shows that the effect of systematic central bank communication on the propagation of a DIS shock depends on its cyclicality. When communication is countercyclical ($\psi_y > \kappa\psi_\pi \Leftrightarrow \Theta < 1$), it dampens both output and inflation responses at all horizons. Conversely, when communication is procyclical ($\psi_y < \kappa\psi_\pi \Leftrightarrow \Theta > 1$), it amplifies both responses. Acyclical communication leaves them unchanged. Together with the sign results in Corollary 1, these findings indicate that the propagation of the DIS shock is driven mainly by the loading on the persistent component, ζ_x, which is strictly increasing in Θ. The loading on the signal extraction error, ξ_x, plays a smaller role.

The second part of Corollary 2 shows a different pattern for cost-push shocks. Regardless of its cyclicality, if $\phi_\pi \geq \rho_\theta$, systematic central bank communication *always* raises the output response relative to the no-communication case and *always* amplifies the inflation response at all horizons. To gain economic intuition, it is useful to consider the following two experiments. In the first, the central bank adjusts its communication effort to expected

variations in output only ($\psi_y > 0, \psi_\pi = 0$). In the second, the central bank adjusts it in response to expected variations in both output and inflation ($\psi_y > 0, \psi_\pi > 0$). In the former case, because communication is strictly countercyclical, the DIS equation always features discounting. This force mitigates the negative response of output triggered by the rise of the nominal, and hence real, interest rate that follows an inflationary cost-push shock. As a result, *ceteris paribus*, a central bank that adjusts its communication based solely on expected variations in output raises output and, through the NKPC, inflation relative to an economy without communication. This establishes the second statement for a strictly countercyclical communication. Consider then the second case where the central bank now reacts to expected variations in both output and inflation. The propagation mechanism of the cost push shock is fundamentally altered as the DIS equation now features an additional term $(\chi_e \psi_\pi/\gamma)\mathbb{E}_t\left[\hat{\theta}_{t+1}\right]$, which propagates a cost push shock, to some extent, like a demand shock. Intuitively, since the cost push shock is persistent, inflation is expected to be above target in the next period. This calls for an increase in the central bank's communication effort which lowers uncertainty and, *ceteris paribus*, raises aggregate demand. This is precisely the force that can make the output response positive in Corollary 1. Whether overall communication is procyclical or countercyclical now only matters to the extent that it amplifies or attenuates the output gain arising from this "endogenous demand shock". As a result, output and inflation will always be higher in an economy with communication relative to an economy with acyclical communication.

Communication and the Shape of Output Dynamics The following corollary shows how central bank communication affects the shape of the output response through $\mathcal{R}_i$, defined in (44) as the contribution of the mean distortion channel to the impact response, relative to the response that would prevail in its absence. This ratio therefore measures the extent to which imperfect information dampens or amplifies the initial output response relative to that benchmark.

Corollary 3 (Shape of Output Response). *Define* $\lambda_i \equiv \ln(\rho_i)/\ln(\rho_i(1-\Gamma_i)) \in (0,1)$.

1. *If* $\mathcal{R}_i \leq \lambda_i$, *the output response is monotone.*
2. *If* $\mathcal{R}_i > \lambda_i$, *the response is hump-shaped (trough-shaped if* $\zeta_i < 0$*) and has a unique interior extremum, with the same sign as* ζ_i, *at*

$$h_i^\star = \frac{\ln\left(\mathcal{R}_i/\lambda_i\right)}{\ln\left(1/(1-\Gamma_i)\right)} > 0\,. \tag{45}$$

The shape of the output response depends on the strength of the mean distortion relative to the response to the persistent component of the shock. When $\mathcal{R}_i > 0$, the mean distortion dampens the initial response, while a negative value implies amplification. If $\mathcal{R}_i$ is sufficiently small, the output response remains monotone, as in the absence of the mean distortion. When $\mathcal{R}_i$ exceeds the threshold λ_i, however, the initial dampening is sufficiently strong to generate a nonmonotone response, with an interior extremum $h_i^\star$ periods after

the shock.[33] Intuitively, the signal extraction error initially offsets part of the response to the persistent component. Since the resulting belief distortion fades faster than the shock itself, this offsetting effect gradually weakens, allowing the output response to build up before eventually declining. Finally, when $\mathcal{R}_i \geq 1$, the mean distortion fully offsets or even reverses the initial response, so that output starts at zero or moves in the opposite direction to ζ_i.

Parametrization The set of quantitative exercises reported below require parameterizing the model. Most parameters are borrowed from the New Keynesian literature and are set to conventional values (see Table 9). The quarterly discount factor is $\beta = 0.99$ and $\gamma = 1$, so that utility from consumption is logarithmic. The interest-rate rule coefficients are $\phi_\pi = 1.5$ and $\phi_y = 0.1$, the values Taylor originally proposed and close to the post-1982 estimates of Clarida, Galí, and Gertler (2000). The communication rule coefficients $\psi_\pi = 7$ and $\psi_y = 1.5$ come straight from our baseline GMM estimates in Table 2. The persistent components have persistence $\rho_\theta = 0.95$ and $\rho_x = 0.90$, and we normalize $\sigma_\theta = \sigma_x = 1$, so that a unit innovation is a one-percent shock.[34]

We set the slope of the NKPC, κ, to 0.02, below the textbook value,[35] but an order of magnitude steeper than recent US evidence supports. Hazell, Herreño, Nakamura, and Steinsson (2022), exploiting cross-state variation to separate the slope from the endogeneity of aggregate demand, estimate a coefficient on the unemployment gap of about 0.006, of the order of 0.003 in output-gap units; Del Negro, Lenza, Primiceri, and Tambalotti (2020) reach a similar verdict from aggregate data and document a marked flattening since the 1990s; and Mavroeidis, Plagborg-Møller, and Stock (2014) caution that the slope is weakly identified, the data admitting a wide range of low values. Against that background $\kappa = 0.02$ sits at the steep end of the plausible range, which we regard as the conservative choice, because our mechanism is strengthened as the curve flattens. Indeed, the discounting in the DIS curve, Θ falls with κ, and the condition for countercyclical communication ($\psi_y > \kappa\psi_\pi$) and the condition for determinacy, become easier to satisfy. Communication effectiveness χ_e has no direct empirical counterpart, since it measures the marginal effect of communication effort on private posterior uncertainty. The assumption $\chi_e > 0$ does receive support from recent empirical work, though. De Fiore, Maurin, Mijakovic, and Sandri (2024), for instance, document two links in the chain connecting central bank communication to household beliefs. Using large language models to score the sentiment of FOMC statements and

[33] The second statement of Corollary 3 treats time as continuous. At integer horizons, the extremum is attained at either $\lfloor h_i^\star \rfloor$ or $\lceil h_i^\star \rceil$, whichever yields the larger value of $m_{y,i}(h)$ when $\zeta_i > 0$, or the smaller value when $\zeta_i < 0$.

[34] The normalization is innocuous for the impulse responses, which do not depend on σ_i; it does scale ς_i^2 and $\sigma_{\eta,i}^2$ in Table 9 proportionally.

[35] For instance, in the Rotemberg formulation of Section 3.1, an adjustment cost equivalent to firms resetting prices once a year implies a slope near 0.26. Read literally, $\kappa = 0.02$ would require $\Psi = 750$ with $\theta = 6$, $\varphi = 1$ and $\alpha = 1/3$, equivalent to a Calvo non-adjustment probability of 0.93, or an average price duration of more than three years. We do not read it that way, and neither does the literature: a flat slope is better understood as the product of moderate nominal rigidity and substantial real rigidity, that is strategic complementarity in price setting, which our symmetric Rotemberg specification does not model explicitly but which κ summarizes.

press conferences, they find a tight correspondence between FOMC communication and subsequent media coverage, with Fed Chair press conference answers proving the most influential single tool for shaping newspaper coverage. They then show that this media coverage feeds into household inflation expectations: a one-standard-deviation hawkish shift in media sentiment lowers medium-term (3-year-ahead) inflation expectations by roughly 0.18 percentage points in the NY Fed's Survey of Consumer Expectations. The chain from central bank communication to media to household beliefs is precisely what $\chi_e > 0$ captures in our model. We discipline it indirectly, through the discounting it generates in the DIS equation. Gabaix (2020) sets the discounting in the DIS equation to 0.85; reluctant to attribute all of the discounting found in the data to central bank communication, we choose χ_e so that our channel delivers about sixty percent of it. This gives $\chi_e = 0.065$ and an implied $\Theta = 0.91$.

TABLE 9. Parametrization

A. Preferences, pricing and policy			*B. Implied equilibrium coefficients*		
β	0.99	discount factor	Θ	0.91	DIS discounting factor
γ	1.00	log utility	$\mathcal{M}$	0.82	multiplier on $\mathbb{E}_t\widehat{y}_{t+1}$
κ	0.02	Phillips curve slope	$\Theta + \kappa/\gamma$	0.93	the same, at the ZLB
$\phi_\pi,\ \phi_y$	1.50, 0.10	Taylor rule			
$\psi_\pi,\ \psi_y$	7.00, 1.50	communication rule			
χ_e	0.065	communication efficacy			
C. Shocks and information			*D. Implied by C*		
$\rho_\theta,\ \rho_x$	0.95, 0.90	persistent persistence	$\Gamma_\theta,\ \Gamma_x$	0.48, 0.44	total belief revision
$\sigma_\theta,\ \sigma_x$	1.00, 1.00	normalization	$\varsigma^2_\theta,\ \varsigma^2_x$	3.50, 4.27	transitory variances
$G_\theta,\ G_x$	0.35, 0.30	fundamentals gain	$\sigma^2_{\eta,\theta},\ \sigma^2_{\eta,x}$	4.90, 5.12	communication noise
$k_\theta,\ k_x$	0.20, 0.20	communication gain	$\varrho_\theta,\ \varrho_x$	0.49, 0.50	belief-wedge persistence

Panels A and C list the parameters that are set; panel B the implied equilibrium coefficients and panel D the structural variances implied by the stationary filter.

Finally, the informational parameters are disciplined as follows. Each disturbance carries four informational parameters: the persistence ρ_i and volatility σ_i of the persistent component, the volatility ς_i of the iid component, and the precision $\sigma^2_{\eta,i}$ of the central bank's message. Rather than set the last two directly, we calibrate the two objects the data actually speak to, namely the gains of the two-stage update (17)–(19), G_i, the weight the household places on the realized fundamental when she revises her prior, and k_i, the weight she then places on the message. Together they drive the total revision of beliefs toward the persistent state within the quarter, $\Gamma_i = 1-(1-G_i)(1-k_i)$, which governs belief dynamics. The two noise variances then follow in closed form.[36]

Since G_i is the share of an observed surprise that the household attributes to the persistent component, calibrating it amounts to taking a stand on the persistent–transient decomposition of the process driving the fundamental shocks. Stock and Watson (2007) find the transient component of inflation to be large, increasingly so after the mid-1980s.

[36] Explicitly, $\sigma^2_{\mu,i} = \sigma^2_i/[1-\rho^2_i(1-\Gamma_i)]$, then $\varsigma^2_i = \sigma^2_{\mu,i}(1-G_i)/G_i$, $P_i = (1-G_i)\sigma^2_{\mu,i}$ and $\sigma^2_{\eta,i} = P_i(1-k_i)/k_i$.

Benhima and Poilly (2021) identify the fundamental and noise components of supply and demand signals separately for the United States and find demand signals to be the noisier of the two, which is why G_x sits below G_θ. Specifically, we set $G_\theta = 0.35$ and $G_x = 0.30$ as a baseline but will assess the sensitivity of our results to variations in these parameters in the next section.

Coibion, Gorodnichenko, and Weber (2022) run an information-provision experiment in which they hand US household subjects a public signal, the Federal Reserve's own inflation forecast, and measure how far their posterior moves toward it, which is k_i by construction. Across such studies posteriors typically move twenty to fifty percent of the way to the signal. We take the bottom of that range and set $k_\theta = k_x = 0.2$, because experimental treatments are far more salient than the routine communication k_i stands for, and because our signal concerns the persistent component of a disturbance rather than the level of inflation. The implied total revisions are $\Gamma_\theta = 0.48$ and $\Gamma_x = 0.44$. In other words, households close a little under half the gap within the quarter, and communication is responsible for roughly 45% of that revision.[37] As for G_i, we investigate the sensitivity of our results to k_i, $i \in \{\theta, x\}$, in the next section.

Quantitatively, with our estimates $\chi_e[\kappa\psi_\pi - \psi_y] \approx -0.088$, while the upper and lower bounds in (36) are respectively 0.11 and -2.15. The communication term sits comfortably inside the corridor $(-2.15, 0.11)$, with a discounting of 0.912 and the constant $\mathcal{M}$ is $0.82 < 1$. The economy is determinate.

Numerical Illustration We now illustrate the predictions of Corollary 1 - 3 within our parameterized benchmark economy. For this purpose, Figure 2 reports impulse response functions of output, inflation, the nominal interest rate, and communication effort to a 1% innovation in the persistent component[38] of the DIS shock (Panel (a)) and of the cost-push shock (Panel (b)), under two policy regimes. In the first one(dashed line), the central bank relies exclusively on the Taylor rule and as such implements a purely acyclical communication. In the second, both the Taylor rule and the communication rule are simultaneously active (solid line). The dash-dotted line corresponds to the FIRE benchmark (41) ($\Gamma_i = 1$). Since communication is inert under FIRE, this benchmark is common to the two policy regimes. The two imperfect-information economies are evaluated at the baseline values of the fundamentals and communication gains, G_i and k_i. We discuss at the end of this section how these gains shape the responses.

DIS Shock A 1% contractionary innovation in the persistent component of the DIS shock depresses output and, through the NKPC, inflation, and, as prescribed by the monetary

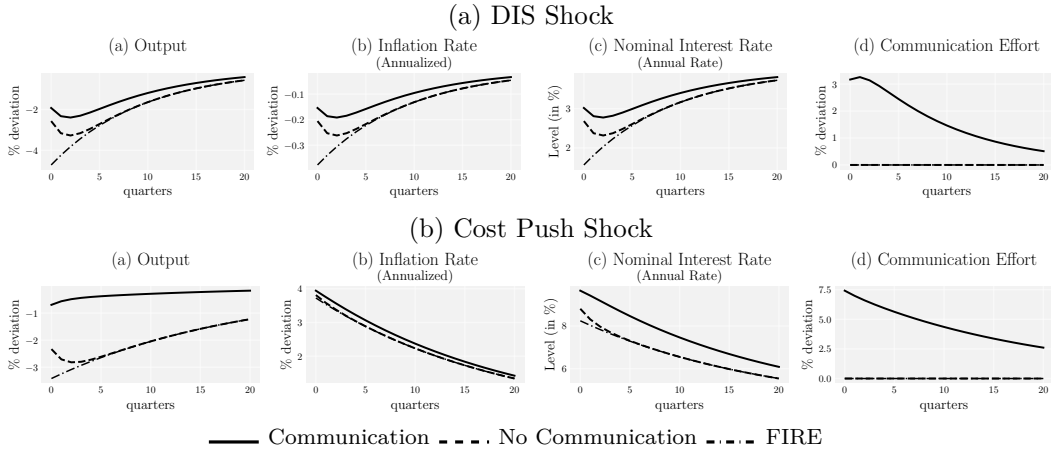
[37] One may be tempted to pin Γ_i down instead from the Coibion and Gorodnichenko (2015) regression of forecast errors on forecast revisions, the standard measure of information rigidity. That would however not be valid here. Because our signal is public and the household is a Bayesian, her forecast error is orthogonal to her own information set, and hence to her own revision, so that coefficient is exactly zero in this model at every horizon. Positive estimates in survey data reflect dispersed private information, from which we abstract. The gains are instead identified by projecting forecast revisions on contemporaneous surprises: the coefficient on the surprise in the observed fundamental is G_i, and adding the surprise in the message recovers k_i.

[38] This is the empirically relevant case, since only persistent movements propagate and thus call for a forward-looking response.

policy rule, the central bank responds by cutting the policy rate. The household observes the fall in $\widehat{x}_t$ but cannot tell how much of it is persistent. On impact, she revises her view of the persistent component by only a fraction of the true innovation, and therefore underestimates how long the shock will last. Compared with a fully informed household, she expects a milder recession and a smaller fall in inflation, and hence a lower interest rate adjustment. Both expectations lead her to cut current spending by less. The mean distortion channel thus dampens the impact response. In the communication economy, output falls by 1.9% on impact, a little more than half of the 3.4% contraction the same economy would experience absent the mean distortion ($1-\mathcal{R}_x = 0.56$, not plotted). The proportion is similar without communication, 2.6% against 4.7% under FIRE (dash-dotted line), which is then the relevant comparison since $\zeta_x = \zeta_x^{\text{FIRE}}$ when communication is inert ($1-\mathcal{R}_x = 0.55$). Measured against FIRE, $\widehat{y}_0/\widehat{y}_0^{\text{FIRE}} = 0.41$ in the communication economy, the product of the policy-rule factor $\zeta_x/\zeta_x^{\text{FIRE}} = 0.73$ and the mean-distortion factor $1-\mathcal{R}_x = 0.56$. The dampening is temporary as the persistence of the signal extraction error, ϱ_x, is much lower than that of the shock. After about five quarters, the response without communication has essentially merged with its FIRE counterpart. This delayed learning is what makes the responses in Panel (a) hump-shaped (see Corollary 3). In both economies $\mathcal{R}_x$ (0.44 with communication, 0.45 without) exceeds $\lambda_x^d = 0.20$, and output reaches its trough two quarters after the shock.

Without communication effort (dashed line), the central bank cuts the policy rate from 4% to 2.7% (annual rate) on impact, but the instrument alone cannot prevent a sharp and persistent contraction (output falls by 2.6% on impact and by 3.3% at the trough). This obtains because, absent communication, the effective multiplier $\mathcal{M}$ of Proposition 4, which governs the pass-through of expectations, increases. Expected future output passes almost one for one into current output, and the high persistence of the persistent component generates a protracted recession.

FIGURE 2. Impulse Response Functions (Normal Times)



With active communication effort, two forces pull in opposite directions which maps onto the discounting condition of Proposition 3. The inflation channel works against stabilization. The deflationary pressure induced by the negative DIS shock prescribes a reduction in communication effort, which leaves agents with less precise signals about the economic outlook, induces greater caution and curtails expenditures. The output channel works in the opposite direction. Falling output calls for more communication effort, reassures agents about the future path of demand and provides a stabilizing stimulus. Under our calibration, communication is countercyclical ($\psi_y = 1.5 > \kappa\psi_\pi = 0.14$) and the output channel dominates. The central bank raises its communication effort by 3.2% on impact and 3.3% one quarter later.[39] The additional effort supports demand through the risk channel. Output falls by 1.9% on impact and by 2.4% at the trough, and the output loss is about a quarter smaller than without communication at every horizon.

Both monetary policy instruments thus move in the same direction; the policy rate falls and communication effort rises. The policy rate falls by less than in the no-communication economy (to 3.0% against 2.7% on impact), because communication absorbs part of the stabilization burden. This illustrates the instrument duality established above, the 3.2% rise in communication effort is worth a cut of about 0.8 percentage point in the annualized policy rate. An economy with an effective communication channel therefore achieves a milder recession with a less extreme use of the interest rate.

Cost-Push Shock A 1% positive innovation in the persistent component of the cost-push shock generates the standard output-inflation trade-off of the New Keynesian framework. Inflation rises by close to 4 percentage points (annualized) on impact while output contracts, confronting the central bank with conflicting stabilization objectives. Since the interest-rate rule responds to realized inflation, the central bank tightens on impact. In the absence of communication (dashed line), the Taylor rule alone manages this trade-off, producing a persistent rise in inflation and a protracted output contraction. The policy rate jumps to 8.8%, and output falls by 2.3% on impact and by 2.8% at the trough, two quarters later.

As with the DIS shock, households revise their view of the persistent component by only a fraction of the innovation. How this underreaction affects output, however, depends on the policy regime. Underestimating the persistence of the shock, households expect a milder future recession, which supports current spending, but also lower future inflation, which raises the real interest rate they perceive and depresses current spending. Unlike after a DIS shock, the two expectations pull in opposite directions.[40] Without communication, the

[39] Although the economy is disinflating in the aftermath of the negative demand shock, the central bank communicates more, not less. Taken alone, the inflation channel would prescribe the opposite, since a below-target inflation outlook calls for less communication. That communication intensifies as disinflation sets in is therefore not a property of its inflation response but of its countercyclicality since the output channel dominates ($\psi_y > \kappa\psi_\pi$), *i.e.* the DIS equation exhibits discounting. Had the central bank operated in the compounding regime ($\psi_y < \kappa\psi_\pi$), communication would instead fall on net, further depressing demand and amplifying the recession.

[40] Formally, the loading $\Lambda_{\delta,\theta} = \Phi_{y,\theta} + \Phi_{\pi,\theta}/\gamma$ of Proposition 2 adds a negative term, the response of expected output, and a positive term, the response of expected inflation. By (42) iterated forward and (39), $\Lambda_{\delta,\theta}$ has the sign of ξ_θ, that is of $(\gamma+\kappa)\zeta_\theta + 1$. With $\zeta_\theta < 0$, the mean distortion therefore amplifies the impact response ($\mathcal{R}_\theta < 0$) if and only if $\zeta_\theta > -1/(\gamma+\kappa) = -0.98$, which holds with communication ($\zeta_\theta = -0.48$) but not without ($\zeta_\theta = \zeta_\theta^{\text{FIRE}} = -3.42$).

expected recession is deep and the first effect dominates. The underreaction then dampens the impact response, as it does for the DIS shock. In this case, output falls by 2.3% against 3.4% under FIRE (dash-dotted line), and delayed learning produces the trough two quarters after the shock. With communication, the expected recession is milder and the second effect dominates. The mean distortion channel then amplifies the output contraction. More precisely, output falls by 0.70% on impact against 0.48% in the same economy absent the mean distortion ($\mathcal{R}_\theta = -0.45$, not plotted), and against 3.4% under FIRE. Output then recovers monotonically as the wedge closes, as Corollary 3 implies for $\mathcal{R}_\theta < 0$.

Active communication improves markedly on the no-communication outcome. Communication effort jumps by 7.4% on impact and remains elevated as long as inflation is expected to stay above target. About 90% of the impact response is the direct reaction of the rule to the expected persistence of the shock (the term $\psi_\pi \rho_\theta \tau_t^\theta$ in Equation 42). The remainder is the net countercyclical response to the expected fall in output. The increase in the precision of the signal reduces the uncertainty households face about the persistent component of the shock and supports demand. Discounting further limits propagation. Output falls by 0.7% on impact, against 2.3% without communication. This comes with no material sacrifice in inflation stabilization because of the flatness of the NKPC.

Unlike after a DIS shock, the two policy instruments move in opposite directions. Because output falls less and inflation is marginally higher, the central bank raises the policy rate by more than in the no-communication economy (to 9.7% against 8.8% on impact), while the rise in communication effort is expansionary. By the instrument duality, the 7.4% rise in effort is worth a cut of about 1.9 percentage points in the annualized policy rate. The interest rate leans against inflation and communication cushions output. Active communication thus complements interest rate policy, and addresses the output consequences of the shock through the risk channel and discounting.

The Role of the Gains G_i and k_i The responses in Figure 2 are computed at the baseline gains, G_i and k_i, $i \in \{\theta, x\}$. Figures C.1 and C.2 in the Online Appendix C show how the responses in the economy with central bank communication change with the gains k_i and G_i, holding the other gain, the policy rules and communication effectiveness χ_e at their baseline values, so that the gains operate only through the signal extraction error.[41] The expression of the IRFs (Equation 43) makes clear how the gains will distort the dynamics. The loadings ζ_i do not depend on the gains, which affect the responses to the persistent innovation only through the total belief revision $\Gamma_i = 1 - (1 - G_i)(1 - k_i)$. This term sets both the signal extraction error on impact, $-(1 - \Gamma_i)$, and its persistence, $\varrho_i = (1 - \Gamma_i)\rho_i$. Raising the communication gain or the fundamentals gain therefore has the same qualitative effect, namely to shrink the mean distortion and bring the response closer to $\zeta_i \rho_i^h$. With χ_e held fixed, this is not the FIRE response, under which the risk channel is also shut down (see (41)). Since $\partial \Gamma_i / \partial k_i = 1 - G_i$, the central bank's message moves beliefs most when the fundamentals are least informative about the persistent state (G_i small).[42]

[41] In the structural expression (B.6) in the Online Appendix, χ_e depends on the gains, notably through the factor k_i^2 in Υ. Holding χ_e fixed amounts to adjusting the communication technology $\partial \sigma_\eta^2 / \partial \Delta e_t$ accordingly.

[42] The split of Γ_i between G_i and k_i matters only for the responses to the transitory components $\boldsymbol{\omega}_t$ and to the communication noise $\boldsymbol{\eta}_t$, which load on the signal extraction error with weights $(1 - k_i)G_i$ and k_i (Lemma 1). These responses are not shown.

In the case of the DIS shock, a higher k_x implies that the household places more weight on the signal issued by the Central Bank, and therefore learns more about the persistent state. She therefore reacts more on impact. The trough deepens as well, but the hump flattens and comes earlier, because she learns at a faster pace. In the limit, as $k_x \to 1$, the household relies fully on the signal, learns on impact and the persistence of the shock fully drives the dynamics. Accordingly, the hump totally disappears. On the contrary, when the households puts a low weight on the message of the central bank, she learns slower about the shock, the impact effect is milder and the trough response occurs later. The gain G_i plays a similar role, and its effects are similar. The key difference though comes from the fact that G_i controls the weight the agent places on the observed fundamental in her inference regarding its persistent component. When $G_x \to 1$, the agent learns much about the persistent shock by the mere observation of the fundamental and the learning is fast. On the contrary a low G_x indicates that the mere observation of the shock is barely informative about the persistent component. The household is a slow learner and the dynamics exhibits a hump-shaped behavior.

For the cost-push shock, the gains matter for output only, and only over the first few quarters. They also work in the opposite direction. Since the mean distortion channel amplifies the output contraction in the communication economy, a higher k_θ (similarly G_θ) lowers the output loss on impact and the response remains monotone for all gains. Inflation, the policy rate and communication effort are essentially unaffected. Inflation is driven by the observed shock, and the flat Phillips curve transmits little of the output differences. Communication effort is dominated by the direct response to the expected persistence of the shock, $\psi_\pi \rho_\theta \tau_t^\theta$, which the central bank computes under full information, independently of households' beliefs.

These experiments separate the informational role of communication from its role as a policy rule. A more informative message, that is, a higher k_i, makes demand more responsive to DIS shocks and lowers the output cost of cost-push shocks. The systematic response of communication effort, by contrast, stabilizes output after both shocks (Figure 2).

3.3.2 Communication at the Zero Lower Bound We now extend the analysis to the empirically relevant case in which a sufficiently large negative DIS shock drives the economy to the zero lower bound (ZLB), rendering the interest rate instrument inoperative. To consider this possibility, the nominal interest rate rule rewrites

$$R_t = \max\Big(R \cdot (\pi_t/\pi)^{\phi_\pi} \cdot (y_t/y)^{\phi_y}\,,\, 1\Big)\ .$$

In the standard model without communication, the ZLB gives rise to a well-known pathology: pessimistic expectations of future output can be self-fulfilling through a deflationary spiral, and the equilibrium is locally indeterminate. With communication, the Central Bank regains some control over the economy which mitigates the forces that generate local indeterminacy. Consider an economy in which the central bank follows a strict interest rate peg, or is permanently at the zero lower bound, such that, from a technical point of view $\phi_\pi = \phi_y = 0$. It follows from Proposition 4 imply that the equilibrium is locally determinate if and only if:

$$-2\gamma - \kappa < \chi_e\left[\kappa\psi_\pi - \psi_y\right] < -\kappa\,. \tag{46}$$

This establishes that, even when the interest rate instrument is not available, the central bank retains a potent stabilization tool through systematic communication effort. By actively managing private expectations, it can steer beliefs and influence the dynamics of output and inflation, preserving equilibrium determinacy in an environment where conventional wisdom would predict instability.

Condition (46) has a transparent interpretation that connects directly to the literature on ZLB pathologies. Note that $\chi_e[\kappa\psi_\pi - \psi_y]/\gamma = \Theta - 1$. Dividing (46) through by γ and rearranging yields the equivalent condition:

$$-1 < \Theta + \frac{\kappa}{\gamma} < 1. \tag{47}$$

The quantity $\Theta + \kappa/\gamma$ is the effective multiplier on expected future output in the reduced-form DIS equation at the ZLB, aggregating both the direct effect of expectations (captured by Θ) and the indirect inflation feedback through the Fisher equation (captured by κ/γ). Condition (47) is therefore simple and intuitive: the economy is locally determinate if and only if this effective multiplier lies strictly between -1 and 1.

To appreciate this result, it is useful to compare it with the standard model at the ZLB, where $\Theta = 1$ and the effective multiplier equals $1 + \kappa/\gamma > 1$. In this case, the expectational feedback is self-reinforcing. A pessimistic expectation of future output lowers expected inflation, raises the expected real interest rate, and depresses current demand, partially validating the initial pessimism. This is the classical ZLB indeterminacy result (see, *e.g.* Benhabib, Schmitt-Grohé, and Uribe, 2001, Eggertsson and Woodford, 2003) associated with an effective multiplier exceeding one. Countercyclicality of communication reduces Θ below unity, bringing the effective multiplier inside the determinacy corridor. The mechanism is clear. When future output is expected to fall, the dual-mandate communication rule prescribes an increase in communication effort. This reassures agents about the future path of demand and partially offsets the contractionary dynamic. Rather than allowing the deflationary spiral to feed back into lower current output with a multiplier exceeding one, the communication response attenuates this feedback and reduces the effective multiplier below the critical threshold.

Inflation-only communication is, by contrast, destabilizing at the ZLB. Setting $\psi_y = 0$ and $\psi_\pi > 0$ yields $\chi_e[\kappa\psi_\pi - \psi_y] = \chi_e\kappa\psi_\pi > 0$, which violates the upper bound in (46). Equivalently, from (47), inflation-only communication implies $\Theta + \kappa/\gamma > 1 + \kappa/\gamma > 1$. Intuitively, by amplifying the inflation-response channel while providing no offsetting output stabilization, such a policy reinforces rather than dampens the self-fulfilling dynamics that make the ZLB problematic. An institution like the historical ECB, whose communication framework was anchored almost exclusively on price stability, therefore foregoes the stabilization benefits of communication at the ZLB and undermines equilibrium determinacy.

A complementary failure mode imposes an upper bound on the output-sensitivity of communication effort. When ψ_y is sufficiently large, $\chi_e[\kappa\psi_\pi - \psi_y]$ becomes increasingly negative, pushing the effective multiplier $\Theta + \kappa/\gamma$ below -1 and violating the lower bound in (46). Determinacy is then lost from the opposite direction. Intuitively, excessive output-sensitivity leads the central bank to sharply scale back expansionary communication when

the private sector anticipates a recovery, more than offsetting the recovery signal and depressing current demand. Anticipating this overreaction, agents revise down their expectations, triggering further adjustments in communication and an ever-widening feedback loop. Rearranging the lower bound in (46) yields the upper bound[43]

$$\psi_y < \kappa\psi_\pi + \frac{\kappa + 2\gamma}{\chi_e}. \tag{48}$$

The dual mandate is therefore necessary but not sufficient for ZLB stabilization. Output-sensitivity must be calibrated to keep the communication rule within the interior of the determinacy region.

All in all, the dual mandate emerges as a requirement rather than a preference. At the ZLB, inflation-centered communication is destabilizing, and determinacy requires a sufficiently strong response to output gap developments. Rearranging the upper bound in (46) gives the minimum threshold, $\psi_y > \kappa(1+\chi_e\psi_\pi)/\chi_e$. The output-stabilization coefficient must exceed this threshold to offset the destabilizing effects of the inflation-response channel as they propagate through the Phillips curve. This result provides a formal rationale for arguments that central banks operating at the ZLB should extend forward guidance beyond inflation targets to include explicit output or employment commitments, as the Federal Reserve did with its outcome-based guidance in September 2020.

Figure 3 illustrates the quantitative implications. It reports impulse response functions for output, inflation, the nominal interest rate, and communication following a contractionary DIS shock, comparing the no-communication economy (dashed lines) with the active-communication economy (solid lines).[44] Two shock magnitudes are considered. Panel (a) shows a large shock of -3.75%, while Panel (b) shows a milder shock of -2.70%.[45]

Large shock: The shock is large enough (-3.75%) to push both economies to the ZLB. Then the dynamics diverge sharply. In the no-communication economy, the interest rate hits zero immediately and remains there for an extended episode of eight quarters, while output and inflation follow a deep, prolonged contractionary path. Without any expectation-management tool, private beliefs can only be anchored by the eventual return of the interest rate to its unconstrained path, which happens slowly given the high persistence of the shock.

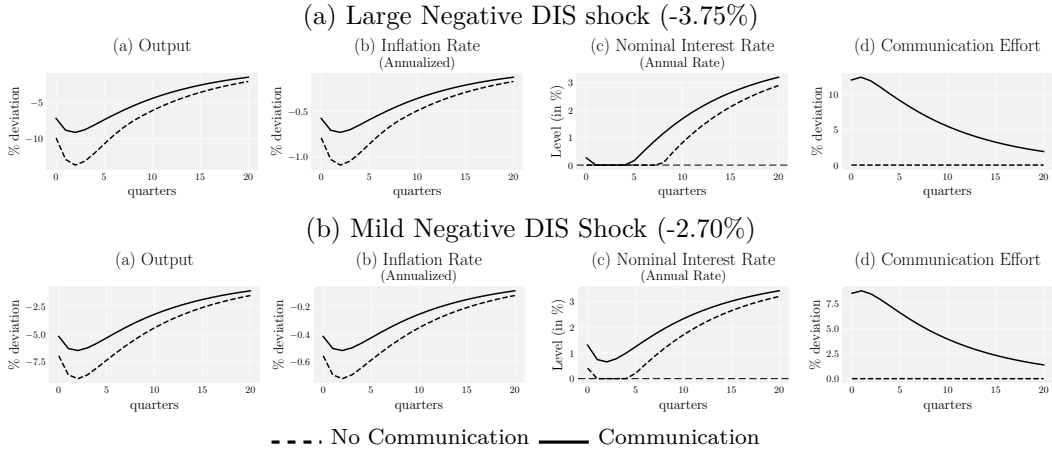
With communication active, the output channel leads the central bank to increase its communication effort aggressively as the outlook deteriorates (by 12% on impact), stimulating private expectations of future demand. By sustaining a favorable expectation of future

[43]This constraint is generally not binding. Under the calibrated parameters, its right-hand side is approximately 31, but it tightens when communication is more efficient (χ_e is large) or the Phillips curve is flatter (κ is small).

[44]The calibrated parameters satisfy the ZLB determinacy condition (46): $\chi_e[\kappa\psi_\pi - \psi_y] \approx -0.088$ lies in the interval $(-2.02, -0.02)$, ensuring a determinate and stable equilibrium.

[45]The model is solved using a shooting method, starting from the *business as usual* steady state and introducing a negative DIS shock. Households forecast at their posterior mean, as in Lemma B.2: their expectation of next-period output is the equilibrium output, ZLB included, of an economy whose persistent DIS component equals their current belief $\widetilde{\tau}_t^x$. When the ZLB does not bind, this reproduces the linear solution of Proposition 5.

FIGURE 3. Impulse Response Functions (Zero Lower Bound)



output and inflation, it reduces the effective real interest rate through the Fisher equation even with the nominal rate fixed at zero. This is precisely the instrument substitutability result operating at the ZLB. Each forgone unit of interest rate accommodation is replaced by $1/\chi_e$ units of communication expansion. The outcome is a faster recovery of output and inflation, a shorter ZLB episode (four quarters instead of eight), and a less severe disinflationary trough (annualized inflation bottoms out at -0.7% instead of -1.1%, and output at -9.1% instead of -13.6%). Interestingly, the same mechanism as in normal times is at work (Section 3.3.1). Lower inflation outlook, routed through the inflation channel alone, would call for lower communication effort, and it is only the countercyclical output channel that overturns this when $\psi_y > \kappa\psi_\pi$ and makes communication effort rise as the economy tanks toward the bound. The dual mandate is thus not incidental to the zero-lower-bound result, it is what keeps the central bank talking when a purely inflation-focused rule would call for silence.

Mild shock: Panel (b) reveals the clearest implication of the model. Under the mild shock (-2.70%), the no-communication economy still reaches the ZLB, for four quarters, as the recessionary impulse is sufficient to exhaust the interest rate instrument. The communication economy does not. By raising communication effort as the outlook deteriorates, the central bank shores up private expectations sufficiently to sustain positive nominal rates throughout the episode (the policy rate bottoms out at 0.7%), therefore acting as an automatic stabilizer. The ZLB is never reached, and the contraction is milder (output bottoms out at -6.5% against -9.0%).[46]

This prevention result is qualitatively distinct from mere mitigation. It identifies communication as an automatic stabilizer with a nonlinear payoff: for shocks below a critical

[46]This stabilizing property of communication hinges on the countercyclicality property and only unfolds when the central bank pursues a dual mandate with a strong enough response to output; in particular, communication would lead to lower government spending multipliers at the ZLB relative to a no-communication economy.

threshold, active communication eliminates the ZLB constraint entirely; above that threshold, communication reduces the depth and duration of the episode without preventing it from binding. At our calibration, the ZLB binds for any persistent DIS shock larger than 2.36% without communication, but only beyond 3.23% with communication. The option value of systematic communication is therefore asymmetric and highest for intermediate shocks, where prevention is feasible and the benefit of avoidance is largest.

4. Conclusion

This paper asks whether the Fed's communication effort is systematic: is the change in the amount of explanation embedded in a policy statement state-contingent and rule-governed, in the same way the policy rate has been since Taylor (1993)? Our answer is yes.

We measure communication effort via changes in the Shannon entropy of FOMC policy statements, a gauge of lexical diversity that is not related to tone or sentiment. Estimating a forward-looking feedback rule on expected inflation and the output gap à la Clarida, Galí, and Gertler (2000) by iterated GMM, we find strong and robust evidence of a systematic communication rule. The over-identifying restrictions are never rejected; the implied inflation target is not distinguishable from the Fed's 2% target; the inflation response and the output-gap response are both large and significant, reflecting a genuine dual communication mandate. This dual-mandate signal is borne almost entirely by the intensive margin of communication, lexical diversity per unit of text, and is invisible to extensive-margin measures such as statement length. These findings are robust to alternative inflation indices, output-gap measures, forecasting horizons, Tealbook expectations, and an external-instrument strategy using oil supply and high-frequency monetary policy surprises. Analogous rules govern every central bank in our international sample, covering the Euro Area, Japan, Canada, Australia, and Sweden, with cross-country differences that faithfully mirror each institution's mandate, history, and degree of external openness.

To characterize how these two instruments interact and what their joint operation implies for policy design, we embed the empirical rule in a Communication Representative Agent New Keynesian (CRANK) model in which the central bank operates a conventional Taylor rule and a communication rule simultaneously. Private agents do not directly observe the state of the economy; they rely instead on a noisy public signal issued by the central bank, whose precision rises with communication effort. Greater precision reduces agents' systematic underreaction to current shocks and compresses precautionary saving. Because both instruments operate through aggregate demand, they are locally perfect substitutes. This grounds the observed post-2008 substitution between interest rates and communication and implies that a credible, consistent communication framework is a precondition for ZLB resilience rather than a byproduct. A central result of the model is that a dual-mandate with countercyclical communication endogenously generates discounting in the dynamic IS equation, provided communication effort responds sufficiently to output developments, a condition our estimates satisfy. This discounting, an equilibrium consequence of communication policy rather than of behavioral assumptions or household heterogeneity, relaxes the Taylor principle and resolves the forward guidance puzzle. At the zero lower bound, it is also what preserves equilibrium determinacy: a central bank that communicates exclusively about inflation amplifies rather than dampens self-fulfilling deflationary

expectations, so the dual mandate must be reflected in the communication framework itself; talking about inflation alone is destabilizing. Quantitatively, active communication shortens ZLB episodes and, for shocks of intermediate magnitude, prevents the bound from binding altogether. Because the mechanism works through forward-looking expectations, it is an automatic stabilizer: communication effort is most valuable when used preemptively, before conditions deteriorate.

In light of our empirical and theoretical results, two avenues seem worth pursuing. First, it would be interesting to determine the optimal simple communication rule for a closed economy, in particular at the zero lower bound, and compare it with the actually implemented one. Second, the Canadian evidence suggests that Federal Reserve communication propagates internationally through an informational channel, shaping the signal environment in which other central banks operate, independently of conventional interest rate spillovers. Characterizing this channel in a two-country CRANK model, and deriving the optimal communication strategy for a small open economy that faces a dominant central bank's public signals, would clarify the dynamics of international communication cycles. This goes beyond the scope of this paper and is left for future research.

References

Acharya, S., and K. Dogra (2020): "Understanding HANK: Insights From a PRANK," *Econometrica*, 88(3), 1113–1158. [8, 38]

Acosta, M., C. M. Brennan, and M. M. Jacobson (2024): "Constructing High-Frequency Monetary Policy Surprises from SOFR Futures," *Economics Letters*, 242, 111873. [19, 20]

Adam, K., and R. M. Billi (2006): "Optimal Monetary Policy under Commitment with a Zero Bound on Nominal Interest Rates," *Journal of Money, Credit and Banking*, 38(7), 1877–1905. [8]

——— (2007): "Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates," *Journal of Monetary Economics*, 54(3), 728–752. [8]

Ahn, S., and M. Holm (2025): "There is no Forward Guidance Puzzle," Discussion paper, University of Oslo. [8, 39]

Andrews, D. W. K., and J. C. Monahan (1992): "An Improved Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimator," *Econometrica*, 60(4), 953–966. [14]

Angeletos, G.-M., and C. Lian (2018): "Forward Guidance without Common Knowledge," *The American Economic Review*, 108(9), 2477–2512. [8, 30]

Antonova, A., M. Matvieiev, and C. Poilly (2026): "Supply Shocks in the Fog: The Role of Endogenous Uncertainty," Discussion Papers 21197, CEPR. [8]

Apel, M., and M. Blix Grimaldi (2014): "How Informative Are Central Bank Minutes?," *Review of Economics*, 65(1), 53–76. [6]

Arismendi-Zambrano, J., E. Kypraios, and A. Paccagnini (2026): "Informativeness of the Federal Reserve Chair Communication's Sentiment on the Monetary Policy Uncertainty," *Annals of Operations Research*, 357(1), 633–677. [13]

Aruoba, S. B., and T. Drechsel (2024): "Identifying Monetary Policy Shocks: A Natural Language Approach," Discussion paper, CEPR Discussion Paper. [6, 7]

Barnichon, R., and G. Mesters (2020): "Identifying Modern Macro Equations with Old Shocks*," *The Quarterly Journal of Economics*, 135(4), 2255–2298. [7, 19]

Bauer, M. D., and E. T. Swanson (2023): "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," *NBER Macroeconomics Annual*, 37, 87–155. [5]

Baumgärtner, M., and J. Zahner (2025): "Whatever It Takes to Understand a Central Banker—Embedding Their Words Using Neural Networks," *Journal of International Economics*, 157, 104034. [6]

Beaudry, P., and F. Portier (2018): "Real Keynesian Models and Sticky Prices," NBER Working Papers 24223, National Bureau of Economic Research, Inc. [38]

Benhabib, J., S. Schmitt-Grohé, and M. Uribe (2001): "Monetary Policy and Multiple Equilibria," *The American Economic Review*, 91(1), 167–186. [4, 51]

Benhima, K., and C. Poilly (2021): "Does Demand Noise Matter? Identification and Implications," *Journal of Monetary Economics*, 117, 278–295. [46]

Bernanke, B. S. (2015): "Inaugurating a New Blog," Ben Bernanke's Blog, Brookings Institution. [1]

Bernanke, B. S., and K. N. Kuttner (2005): "What Explains the Stock Market's Reaction to Federal Reserve Policy?," *The Journal of Finance*, 60(3), 1221–1257. [5]

Bilbiie, F. O. (2025): "Monetary Policy and Heterogeneity: An Analytical Framework," *The Review of Economic Studies*, 92(4), 2398–2436. [8, 38]

Blanchard, O. J., J.-P. L'Huillier, and G. Lorenzoni (2013): "News, Noise, and Fluctuations: An Empirical Exploration," *The American Economic Review*, 103(7), 3045–3070. [3, 28, 35]

Blinder, A. S., M. Ehrmann, J. de Haan, and D.-J. Jansen (2024): "Central Bank Communication with the General Public: Promise or False Hope?," *Journal of Economic Literature*, 62(2), 42557. [1]

Blinder, A. S., M. Ehrmann, M. Fratzscher, J. de Haan, and D.-J. Jansen (2008): "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence," *Journal of Economic Literature*, 46(4), 910–945. [1]

Bulíř, A., M. Čihák, and D.-J. Jansen (2013): "What Drives Clarity of Central Bank Communication About Inflation?," *Open Economies Review*, 24(1), 125–145. [6]

Bullard, J., and K. Mitra (2002): "Learning About Monetary Policy Rules," *Journal of Monetary Economics*, 49(6), 1105–1129. [39]

Campbell, J. R., C. L. Evans, J. D. M. Fisher, and A. Justiniano (2012): "Macroeconomic Effects of Federal Reserve Forward Guidance," Discussion paper, Brookings Papers on Economic Activity, Spring, 1–80. [5, 8]

Campbell, J. R., J. D. M. Fisher, A. Justiniano, and L. Melosi (2017): "Forward Guidance and Macroeconomic Outcomes since the Financial Crisis," *NBER Macroeconomics Annual*, 31, 283–357. [5]

Clarida, R., J. Galí, and M. Gertler (2000): "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *The Quarterly Journal of Economics*, 115(1), 147–180. [1, 2, 7, 11, 12, 44, 54, – xiv –]

Cochrane, J. H. (2018): "Michelson-Morley, Fisher, and Occam: The Radical Implications of Stable Quiet Inflation at the Zero Bound," *NBER Macroeconomics Annual*, 32, 113–226. [8, 38]

Coibion, O., and Y. Gorodnichenko (2012): "What Can Survey Forecasts Tell Us about Information Rigidities?," *Journal of Political Economy*, 120(1), 116–159. [36]

——— (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *The American Economic Review*, 105(8), 2644–2678. [28, 46]

Coibion, O., Y. Gorodnichenko, and M. Weber (2022): "Monetary Policy Communications and Their Effects on Household Inflation Expectations," *Journal of Political Economy*, 130(6), 1537–1584. [46]

De Fiore, F., A. Maurin, A. Mijakovic, and D. Sandri (2024): "Monetary Policy in the News: Communication Pass-Through and Inflation Expectations," BIS Working Papers 1231, Bank for International Settlements. [44]

De Pooter, M. (2021): "Questions and Answers: The Information Content of the Post-FOMC Meeting Press Conference," FEDS Notes 2021-10-12, Board of Governors of the Federal Reserve System. [13]

Debortoli, D., J. Galí, and L. Gambetti (2020): "On the Empirical (Ir)Relevance of the Zero Lower Bound Constraint," *NBER Macroeconomics Annual*, 34, 141–170. [8, 38]

Del Negro, M., M. P. Giannoni, and C. Patterson (2023): "The Forward Guidance Puzzle," *Journal of Political Economy Macroeconomics*, 1(1), 43–79. [4, 8, 38]

Del Negro, M., M. Lenza, G. E. Primiceri, and A. Tambalotti (2020): "What's Up with the Phillips Curve?," *Brookings Papers on Economic Activity*, pp. 301–357. [44]

Eggertsson, G. B., and F. D. Schüle (2024): "The Forward Guidance Puzzle is not a Puzzle," Working Paper 33180, National Bureau of Economic Research. [8, 38]

Eggertsson, G. B., and M. Woodford (2003): "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 34(1), 139–235. [4, 8, 51]

Ehrmann, M., and M. Fratzscher (2007): "The Timing of Central Bank Communication," *European Journal of Political Economy*, 23(1), 124–145. [39]

Farhi, E., and I. Werning (2019): "Monetary Policy, Bounded Rationality, and Incomplete Markets," *The American Economic Review*, 109(11), 3887–3928. [8]

Fisher, J. D. M. (2015): "On the Structural Interpretation of the Smets-Wouters Risk Premium Shock," *Journal of Money, Credit and Banking*, 47(2-3), 511–516. [32]

Gabaix, X. (2020): "A Behavioral New Keynesian Model," *The American Economic Review*, 110(8), 2271–2327. [4, 8, 38, 45]

Galí, J. (2015): *Monetary Policy, Inflation, and the Business Cycle - An Introduction to the New Keynesian Framework*. Princeton University Press, 2 edn. [3, 7]

Gallegos, J.-E. (2024): "HANK beyond FIRE: Amplification, Forward Guidance, and Belief Shocks," Working Papers 2418, Banco de España. [8, 38]

Gáti, L., and A. Handlan (2024): "Monetary Communication Rules," Discussion paper, CEPR DP18730. [7]

Gentzkow, M., B. Kelly, and M. Taddy (2019): "Text as Data," *Journal of Economic Literature*, 57(3), 535–574. [6]

Gertler, M., and P. Karadi (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7(1), 44–76. [5]

Gorodnichenko, Y., T. Pham, and O. Talavera (2023): "The Voice of Monetary Policy," *The American Economic Review*, 113(2), 548–584. [6]

Gürkaynak, R. S., B. Sack, and E. T. Swanson (2005): "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 1(1), 55–93. [5]

Handlan, A. (2022): "Text Shocks and Monetary Surprises: Text Analysis of FOMC Statements with Machine Learning," mimeo, Brown University. [6, 9]

Hansen, A. L., and S. Kazinnik (2023): "Can ChatGPT Decipher Fedspeak?," Working paper, Federal Reserve Bank of Richmond, Available at SSRN 4399406. [6]

Hansen, L. P. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50(4), 1029–1054. [14]

Hansen, L. P., J. Heaton, and A. Yaron (1996): "Finite-Sample Properties of Some Alternative GMM Estimators," *Journal of Business & Economic Statistics*, 14(3), 262–280. [14]

Hansen, S., and M. McMahon (2016): "Shocking Language: Understanding the Macroeconomic Effects of Central Bank Communication," *Journal of International Economics*, 99, 114–133, 38th Annual NBER International Seminar on Macroeconomics. [6]

Hansen, S., M. McMahon, and A. Prat (2018): "Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach*," *The Quarterly Journal of Economics*, 133(2), 801–870. [6]

Hazell, J., J. Herreño, E. Nakamura, and J. Steinsson (2022): "The Slope of the Phillips Curve: Evidence from U.S. States," *The Quarterly Journal of Economics*, 137(3), 1299–1344. [44]

Hernández-Murillo, R., and H. G. Shell (2014): "The Rising Complexity of the FOMC Statement," Economic Synopses 23, Federal Reserve Bank of St. Louis. [6]

Jansen, D.-J. (2011): "Does the Clarity of Central Bank Communication Affect Volatility in Financial Markets? Evidence from Humphrey-Hawkins Testimonies," *Contemporary Economic Policy*, 29(4), 494–509. [6]

Kawamoto, T., T. Ozaki, N. Kato, and K. Maehashi (2017): "Methodology for Estimating Output Gap and Potential Growth Rate: An Update," *BoJ Reports & Research Papers*. [– xvi –]

Kuttner, K. N. (2001): "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market," *Journal of Monetary Economics*, 47(3), 523–544. [5]

Känzig, D. R. (2021): "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *The American Economic Review*, 111(4), 1092–1125. [2, 7, 19, 20]

Lee, K., N. Olekalns, and K. Shields (2013): "Meta Taylor Rules for the UK and Australia; Accommodating Regime Uncertainty in Monetary Policy Analysis Using Model Averaging Methods," *The Manchester School*. [– xv –, – xvii –]

Lorenzoni, G. (2009): "A Theory of Demand Shocks," *The American Economic Review*, 99(5), 2050–2084. [3, 8, 28, 35]

Lucas, R. E. (1972): "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4(2), 103–124. [29]

Lucca, D. O., and F. Trebbi (2009): "Measuring Central Bank Communication: An Automated Approach with Application to FOMC Statements," Working Paper 15367, National Bureau of Economic Research. [6, 7]

Maćkowiak, B., F. Matejka, and M. Wiederholt (2023): "Rational Inattention: A Review," *Journal of Economic Literature*, 61(1), 226–273. [8]

Maćkowiak, B., and M. Wiederholt (2009): "Optimal Sticky Prices under Rational Inattention," *The American Economic Review*, 99(3), 769–803. [8]

Mankiw, N. G., and R. Reis (2002): "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 117(4), 1295–1328. [36]

Mavroeidis, S. (2010): "Monetary Policy Rules and Macroeconomic Stability: Some New Evidence," *The American Economic Review*, 100(1), 491–503. [7]

Mavroeidis, S., M. Plagborg-Møller, and J. H. Stock (2014): "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature*, 52(1), 124–188. [44]

McKay, A., E. Nakamura, and J. Steinsson (2017): "The Discounted Euler Equation: A Note," *Economica*, 84(336), 820–831. [8]

Meichtry, P. (2023): "Sticky Information, Heterogeneity, and Aggregate Demand," mimeo, Université de Lausanne. [8, 38]

Michaillat, P., and E. Saez (2021): "Resolving New Keynesian Anomalies with Wealth in the Utility Function," *The Review of Economics and Statistics*, 103(2), 197–215. [38]

Morris, S., and H. S. Shin (2002): "Social Value of Public Information," *The American Economic Review*, 92(5), 1521–1534. [29]

Nakamura, E., and J. Steinsson (2018): "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect," *The Quarterly Journal of Economics*, 133(3), 1283–1330. [2, 5, 7, 19, 20]

Nakov, A. (2008): "Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate," *International Journal of Central Banking*, 4(2), 73–127. [8]

Newey, W. K., and K. D. West (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55(3), 703–708. [14]

Picault, M., and T. Renault (2017): "Words Are Not All Created Equal: A New Measure of ECB Communication," *Journal of International Money and Finance*, 79, 136–156. [6]

Romer, C. D., and D. H. Romer (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *The American Economic Review*, 94(4), 1055–1084. [7]

Rotemberg, J. J. (1982): "Sticky Prices in the United States," *Journal of Political Economy*, 90(6), 1187–1211. [33]

Shannon, C. E. (1948): "A Mathematical Theory of Communication," *The Bell System Technical Journal*, 27(3), 379–423. [1, 10]

Shapiro, A. H., and D. J. Wilson (2022): "Taking the Fed at Its Word: A New Approach to Estimating Central Bank Objectives Using Text Analysis," *The Review of Economic Studies*, 89(5), 2768–2805. [6, 7]

Sims, C. A. (2003): "Implications of Rational Inattention," *Journal of Monetary Economics*, 50(3), 665–690. [8]

Smets, F., and R. Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *The American Economic Review*, 97(3), 586–606. [32]

Stock, J. H., and M. W. Watson (2007): "Why Has U.S. Inflation Become Harder to Forecast?," *Journal of Money, Credit and Banking*, 39, 3–33. [45]

Taylor, J. B. (1993): "Discretion Versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214. [1, 2, 7, 12, 33, 54]

——— (1999): *A Historical Analysis of Monetary Policy Rules* pp. 319–348. University of Chicago Press. [7]

Werning, I. (2015): "Incomplete Markets and Aggregate Demand," Working Paper 21448, National Bureau of Economic Research. [8, 38]

Woodford, M. (2003a): *Imperfect Common Knowledge and the Effects of Monetary Policy* pp. 25–58. Princeton University Press. [3, 8, 28]

——— (2003b): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, 1 edn. [3, 7, 29]

Xia, F. D., and J. C. Wu (2016): "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit and Banking*, 48(2-3), 253–291. [14, – xiv –]