WORKING PAPERS

N° 1761

June 2026

"Louder than Rates
The Systematic Nature of Central Bank Communication"

Tiziana Assenza, Fabrice Collard, Dogukan Guney and Philipp Wangner

Toulouse School of Economics

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

June, 28, 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 Blinder, Ehrmann, de Haan, and Jansen (2024) put it, "*when central bankers talk, financial markets listen—intently*", and a large literature confirms that words move asset prices, expectations, and ultimately macroeconomic outcomes in ways that are distinct from, and sometimes larger than, the effects of rate decisions alone. Yet, despite this prominence, a basic question remains unaddressed: *do central banks decide how much explanation to offer in a systematic way?* Is the communication effort exerted by a policy statement, measured not by its hawkish or dovish tone, the sentiment it conveys, or the truthfulness of the underlying assessment, but by the increment of explanation it adds over the preceding statement, state-dependent and rule-governed in the same way that the policy rate has been since at least Taylor (1993)? This paper argues that it is. We characterize the rule empirically and embed it in a theoretical model that reveals communication to be a genuine second instrument of monetary policy.

A New Communication Measure We operationalize communication effort using Shannon entropy, computed over the distribution of words in the non-technical portion of each FOMC statement (Shannon, 1948). Shannon entropy is the canonical measure of lexical diversity: it captures both the breadth of vocabulary drawn on and the unevenness with which terms are used. A statement that draws on varied, non-repetitive language has high entropy; one that recycles familiar, predictable formulations has low entropy.

Communication effort is identified from the change in entropy between consecutive periods. A rise in entropy indicates that the Committee elaborated more, drawing on a broader and less repetitive vocabulary, while a decline indicates linguistic economizing. This interpretation is more than a labeling convention. By construction, entropy captures only the amount of distinct linguistic material and is therefore agnostic about whether a statement is, for instance, hawkish or dovish. It isolates how much the Committee elaborates from what it chooses to convey, providing a clean measure of the supply side of communication. More importantly, as we show in this paper, the measure is validated by behavior as changes in entropy follow a stable forward-looking rule and recover the Feds announced inflation target from text alone, patterns that a mechanical byproduct would not generate.

We apply this measure to all 234 FOMC statements released since 1994. Viewed through this lens, the Committee allocates communication as an informational resource, elaborating its reasoning in some circumstances while economizing on words in others. Because FOMC statements are the Committees primary vehicle for explaining its policy decisions, our objective is not to evaluate the accuracy, tone, or substantive content of the Feds assessment. Rather, it is to measure the communication effort the Committee chooses to exert when rationalizing the policy actions it has implemented. Throughout, we interpret communication effort as distinct from communication quality: greater elaboration does not necessarily imply clearer communication.

Our focus on the supply side of communication departs from most existing work, which has concentrated on the demand side of information: identifying how markets respond to central bank announcements (e.g. Gürkaynak, Sack, and Swanson, 2005, Campbell, Fisher, Justiniano, and Melosi, 2017), or extracting tone, sentiment, or topical content from central bank text to assess whether communication is hawkish or dovish, clear or obscure (e.g. Lucca and Trebbi, 2009, Hansen and McMahon, 2016a, Binder, 2017, Hansen, McMahon, and Prat, 2018). These approaches treat communication as an object to be decoded. Yet, as Blinder, Ehrmann, Fratzscher, de Haan, and Jansen (2008) and the subsequent policy literature on transparency emphasize, more communication is not necessarily better: excessive disclosure can confuse rather than clarify, and "cacophony" may impose costs on financial markets and on the public's ability to form expectations. The relevant empirical object is therefore not whether central banks talk more on average, but whether the communication effort across states is governed by a stable, systematic rule, and whether that rule can be used to extract a structural signal about monetary policy objectives from the text of policy statements alone. A communication rule also has direct implications for how markets and the public should interpret FOMC statements: if the amount of explanation embedded in a statement is itself a systematic signal about the state of the economy and the Committee's reading of its own reaction function, then the volume of communication is informative above and beyond the substantive content of the message. Knowing whether communication obeys a rule is a prerequisite for evaluating whether and when communication is doing useful work.

Visual inspection of the resulting time series reveals that the Fed's communication effort is high at the onset of recessions, reverting as the recovery is looming and displaying little variation in more tranquil times. To go beyond visual evidence and reveal the systematic component of central bank communication, we model the change in log-entropy as a forward-looking feedback rule on expected inflation and the expected output gap, in the tradition of Clarida, Galí, and Gertler (2000), and estimate it by iterated GMM. We do not impose an inflation anchor, leaving the target as a free parameter to be recovered from the communication data alone; a test of whether communication and the policy rate are governed by the same long-run price-stability objective.

A Systematic Communication Rule Estimating the baseline rule, featuring 1-step ahead inflation and output expected gaps, on 234 FOMC statements over 1994:Q2–2022:Q4 delivers four sharp results. First, the data do not reject the over-identifying restrictions, consistent with the existence of a systematic communication rule. Second, the estimated inflation anchor (1.93%) sits close to the Fed's announced 2% target, even though it is identified entirely from statement text, indicating a form of consistency between the interest rate rule and the communication rule. 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%. The inertia parameter is small and statistically insignificant (0.07), confirming that the rule is genuinely state-contingent.

These baseline findings are robust to alternative inflation measures, output-gap proxies, forecasting horizons, and treatments of expectations, including the substitution of Tealbook forecasts for rational expectations. An external instrument strategy (using Känzig (2021) oil supply shocks and Nakamura and Steinsson (2018) high-frequency monetary policy surprises) adds a causal dimension and reveals an informative asymmetry: under supply-side instruments, the response to expected inflation is strong while the output-gap response is essentially zero; under demand-side instruments, both channels are activated. The Federal Reserve elaborates on inflation when inflation is the source of the disturbance, and elaborates on both margins when a demand contraction simultaneously threatens price stability and real activity. Replacing entropy growth with length growth in the benchmark specification nearly halves the estimated inflation sensitivity and drives the output-gap coefficient to zero, confirming that it is the intensive margin of communication (the changes in lexical diversity per unit of text, not mere length) that carries the dual-mandate signal.

A natural question is whether this systematic nature of communication is specific to the Federal Reserve, or whether it reflects a broader regularity that would emerge from the data of any central bank. In the latter case, it constitutes a genuine international regularity rather than an institutional idiosyncrasy. To assess this, we re-estimate the rule for the Euro Area, Japan, Canada, Australia, and Sweden. In every case the over-identifying restrictions are satisfied and the inflation response is positive and significant.Estimated anchors, recovered from communication alone, reproduce each central bank's publicly stated target with considerable accuracy. Cross-country differences are instructive. The European Central Bank's output-gap response is statistically indistinguishable from zero, mirroring its primary price-stability mandate. The Bank of Japan's inflation response is the largest in the sample, consistent with extreme sensitivity to any deviation from near-zero inflation. The Reserve Bank of Australia and the Riksbank display a significant additional response to expected variations in the exchange rate, as expected from small open economies. The Bank of Canada's communication is governed almost entirely by the U.S. federal funds rate rather than by domestic conditions; its statements are, to a first approximation, an exercise in explaining the U.S. monetary policy environment to a Canadian audience, revealing a hierarchical dimension of monetary communication that mirrors the international monetary system itself.

A New Keynesian Model with Communication Rule Having characterized the communication rule empirically, we ask how systematic communication effort affects macroeconomic dynamics and possibly complements or substitutes for interest rate policy. The need for a theoretical framework is sharpest when set against the workhorse three-equation New Keynesian model under full-information rational expectations, which leaves essentially no role for communication beyond the announcement of the policy rule itself. If agents know the central bank's objective and information set, there is nothing to be gained from elaboration (see Woodford, 2003b, Galí, 2015). Yet this prediction sits awkwardly with the empirical fact that communication effort has expanded precisely as central banks have become more transparent and as rational expectations have become the workhorse assumption of applied macroeconomics. Reconciling the two requires embedding systematic communication in a model where it does genuine work, one that disciplines what central banks actually do when they talk and when they choose to say more rather than less.

To address this, we develop a Communication Representative Agent New Keynesian (CRANK) model in which the central bank operates two instruments simultaneously: the conventional interest rate and a communication rule. The key departure from the benchmark is informational (see Woodford, 2003a, Lorenzoni, 2009, among others); private agents cannot directly observe the current state of the economy and instead rely on a noisy public signal issued by the central bank, where the intensity of communication effort determines the precision of the signal and responds to the macroeconomic outlook. In this setting, communication affects the economy through two channels: a mean distortion channel, whereby more precise signals attenuate the systematic underreaction of agents to current shocks, and a risk channel, whereby reduced posterior uncertainty lowers precautionary savings and directly stimulates demand.

In normal times, the interaction between the two instruments produces a clear result. The dynamic IS equation features discounting, so that current demand is less than one-for-one sensitive to expected future output. This outcome obtains directly from two key features of the estimated communication rule: *(i)* the central bank pursues a dual mandate, responding to both inflation deviations and output shortfalls, and *(ii)* its communication is countercyclical, especially through the output channel.[1] When future output is expected to be high, the output-stabilization motive leads the central bank to scale back expansionary communication effort, dampening current demand and counteracting the anticipated boom. The inflation and output feedback channels pull in opposite directions; higher expected inflation calls for more elaborate communication effort while higher expected output calls for less. The net countercyclicality, driven by the output channel, is what produces discounting. This is an equilibrium property that emerges from taking the empirically estimated communication rule seriously, not something we impose on the model.[2]

The discounting generated by communication effort governs both the propagation of shocks and equilibrium determinacy. By suppressing the scope for self-fulfilling amplification, it shifts the entire determinacy corridor downward; the minimum interest rate responsiveness required to guarantee a unique equilibrium falls, relaxing the Taylor principle, while the upper bound shifts by the same amount, leaving the width of the admissible region unchanged. A central bank with a dual mandate in its communication effort can therefore achieve determinacy with a less aggressive interest rate rule than the standard benchmark requires; conversely, an inflation-only communicator shifts the corridor upward, tightening the Taylor principle. In response to a cost-push shock, communication simultaneously anchors inflation expectations and cushions the output contraction; following a contractionary demand shock, the two instruments reinforce each other, achieving a faster recovery with smaller individual movements in either. Discounting of the IS equation also resolves a prominent class of New Keynesian anomalies, most notably the forward guidance

[1] Should either the dual mandate or the countercyclicality of communication through the output channel be absent, the net effect reverses; rather than discounting, the IS equation exhibits compounding, amplifying rather than attenuating the propagation of shocks.

[2] Furthermore, this countercyclicality has a consequence that is easy to get backwards. Given our estimated rule, one would expect a central bank that communicates more when inflation threatens to communicate less when inflation falls. It actually does the exact opposite when the output channel dominates and increases its communication effort as the economy sinks into a demand-driven disinflation. The same force appears at the zero lower bound, where an inflation focused communicator would fall silent exactly when talk is most needed.

puzzle (the prediction that announcements about distant future interest rates have implausibly large effects on current output and inflation). Our resolution is conceptually distinct from existing ones, which introduce discounting through behavioral assumptions about bounded rationality: here, discounting is an equilibrium outcome of countercyclical communication effort under rational expectations, requiring no departure from the standard informational assumptions beyond the central bank's dual mandate.

The most consequential results are related to the zero lower bound (ZLB). When the interest rate is exhausted, communication effort substitutes directly for conventional monetary policy by sustaining private expectations of future demand. Determinacy at the ZLB again requires discounting, and discounting again requires a dual mandate. The output-stabilization component of the communication rule is what breaks the self-fulfilling deflationary spiral, while inflation-only communication actively undermines determinacy by amplifying rather than dampening expectational feedback. Quantitatively, for a large shock that pushes the economy to the ZLB, active communication effort delivers a faster recovery and a shorter, weaker disinflationary episode; for a shock of intermediate magnitude, communication prevents the ZLB from binding altogether, acting as an automatic stabilizer that keeps the nominal interest rate positive throughout the episode.

Paper Outline The remainder of the paper is organized as follows. Section 1 offers a review of the literature and highlights our contribution. Section 2 presents the empirical analysis, including the construction of the communication measure, the estimation and robustness analysis of the Fed's communication reaction function, and international evidence. Section 3 develops the CRANK model, characterizes local determinacy and its dual-mandate requirements, and examines the role of systematic communication effort at the zero lower bound. The final section concludes. Additional results and all proofs are collected 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 (2016b) 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.[3] 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.[4] 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.

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

[4]Details on the construction of multi-word n-grams and the list of n-grams are provided in Table A.3 in Appendix A.

2.1.2 Shannon Entropy as a Measure of Lexical Diversity To quantify the communication effort exerted by the Federal Reserve, we use a measure of lexical diversity drawn from quantitative linguistics: the Shannon entropy of the n-gram distribution of each statement (Shannon, 1948). Lexical diversity captures the breadth and unevenness of the vocabulary used in a text. A statement that draws on a wide range of words and phrases, used in roughly even proportions, has high lexical diversity; a statement that recycles a small number of formulations has low diversity.[5] We then measure the communication effort of the central bank as the change in entropy between consecutive statements. A committee that finds itself explaining a complex or unfamiliar policy environment is forced to use a wider and less stylized vocabulary, raising the lexical diversity of its statement; conversely, a Committee operating in stable, well-understood conditions can rely on familiar formulations and produce a lower-diversity statement. Moreover, the rise in lexical diversity during difficult episodes reflects more elaborate treatment of the two mandate objectives rather than a shift to new topics: inflation- and output-related language together account for roughly 80% of the sentences in every statement throughout the sample (Appendix A.4).

The interpretation of lexical diversity as communication effort rests on a specific feature of FOMC language: its baseline is highly templated. The Committee's default register is a stock of recurring formulations ("moderate growth," "inflation remains subdued," "appropriate monetary policy") reused across statements with little variation. Within this register, conveying genuinely new content is not free: a committee cannot elaborate an unfamiliar or complex policy environment while remaining inside the stylized vocabulary, and is forced to draw on a broader and less predictable set of words and combinations. Producing that material is costly it requires deliberation, drafting, and accepting greater scope for the statement to move markets so a rise in lexical diversity reflects a deliberate allocation of an informational resource rather than a mechanical artifact. Repeating a familiar phrase, by contrast, re-asserts a proposition without adding explanation and leaves entropy unchanged. It is this asymmetry novel elaboration is costly, repetition is cheap that licenses reading the change in lexical diversity as communication effort.

Two features of this measure deserve emphasis at the outset. First, by construction, lexical diversity is not related to the substantive content of the message as it does not register whether the statement is hawkish or dovish, optimistic or alarmed, or accurate or misleading. It captures the sheer amount of distinct linguistic material the Committee chose to use, conditional on whatever it intended to convey. This is precisely the property required of a supply-side measure of communication. Second, lexical diversity is monotone neither in clarity nor in policy desirability. A long, varied statement is not necessarily more informative to the listener than a short, focused one, and the broader transparency literature has long emphasized that more communication is not, by itself, better communication (Blinder, Ehrmann, Fratzscher, de Haan, and Jansen, 2008). We interpret variation in lexical diversity between two consecutive statements as a measure of the central bank's communication effort, not of communication quality; it is the object that any rule-based account of the supply of communication must take as primitive. Whether this interpretation is warranted is ultimately an empirical question. A measure that merely reflects mechanical features of

[5] Appendix A.2 provides the reader with three examples illustrating this point.

the text, or shifts in topic, would not be expected to obey a stable forward-looking rule or to recover the Committee's announced inflation objective from text alone as we show it does in Section 2.2.

Formally, for a discrete random variable X_t distributed according to $p : \mathcal{X}_t \to [0,1]$, the Shannon entropy, expressed in bits, is

$$\mathcal{E}(X_t) = -\sum_{x_i \in \mathcal{X}_t} p(x_i) \log_2 p(x_i) \,, \tag{1}$$

with $\mathcal{E}(X_t) \in \mathbb{R}_+$. In our application, X_t denotes the set of n-grams appearing in the FOMC statement released at time t, and $p(x_i)$ is the relative frequency of n-gram x_i in the corpus C. Shannon entropy of an FOMC statement is a continuous measure of its lexical diversity. Low entropy reflects a predictable, well-worn communication style consistent with stable conditions and limited need for elaboration; high entropy reflects a more varied vocabulary, typically arising when the policy environment demands additional explanation or when new instruments and strategies require justification.[6]

We compute the Shannon entropy of each FOMC policy statement in our sample, producing a time series at meeting frequency. To align the communication data with macroeconomic indicators, we aggregate the series first to monthly and then to quarterly frequency. Each calendar day is assigned the entropy of the most recently released statement. Regular statements are naturally spaced at the inter-meeting interval; the position of exceptional statements within the meeting cycle, however, is variable. To treat all releases symmetrically, we assume that the informational footprint of an exceptional statement persists for 52 days (the average spacing between regular meetings) regardless of when the next scheduled meeting occurs.[7] While being innocuous for our results, this assumption prevents exceptional statements from receiving disproportionately low weight when they happen to occur close to a regular meeting. The quarterly series is then 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).

Figure 1 plots the resulting (log-)entropy series at both meeting and quarterly frequencies. Lexical diversity is low and stable through most of the pre-2008 period, suggesting that the Committee's language was relatively predictable. The 2001 recession produces a brief uptick, but the major shift occurs during the Global Financial Crisis (GFC), when entropy rises sharply as the FOMC adapts its communication to extreme uncertainty and to the use of unconventional monetary policy tools. After 2015, entropy gradually declines, consistent with a return to standardized policy language. The COVID-19 pandemic in 2020

[6]A simple analogy clarifies the construction. Imagine a box of colored balls. If all the balls are red, drawing one tells us nothing new and the entropy of the distribution is zero. If the box contains an even mix of four colors, each draw is unpredictable and the entropy is maximized. Lexical diversity works the same way: when the FOMC repeats familiar phrases ("moderate growth", "inflation remains subdued"), the entropy of the n-gram distribution is low; when the Committee introduces fresh language or unusual combinations of words, entropy rises.

[7]There are eight regular FOMC meetings per year, implying an average spacing of approximately 52 days.

FIGURE 1. (Log-)Entropy of FOMC Statements, 1994 –2022

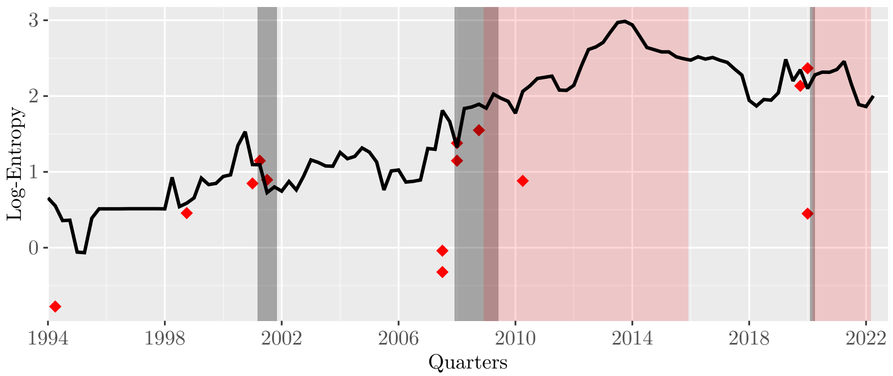

Note: Shannon entropy of FOMC statements at meeting and quarterly frequencies. Unscheduled meetings are indicated by red diamonds. NBER recessions and zero lower bound (ZLB) episodes are shaded in grey and light red, respectively.

and the onset of the Russia–Ukraine war in 2022 trigger renewed increases. Over the full sample, entropy rises and is non-stationary by standard tests.

Exceptional (unscheduled) statements, marked as red diamonds in Figure 1, cluster around episodes of distress, particularly recessions and financial crises. They are shorter on average but their linguistic diversity varies more, contributing disproportionately to spikes in aggregate entropy. Table 1 summarizes the distribution of log-entropy across meetings and confirms the pattern: regular statements display higher mean entropy but lower variance than exceptional ones, indicating that routine communication follows consistent patterns whereas crisis communication becomes more diverse and less predictable.

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.

Our measure of communication effort is then given by the change in the log-entropy between periods:

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

Working with the change rather than the level is both a statistical and a conceptual choice. Statistically, first-differencing removes the non-stationarity present in the level of entropy, which has trended upward since 1994, largely because statements have grown longer. Conceptually, communication effort is a flow, not a stock: the object of interest is how much the Committee elaborates relative to the previous period, not how rich its vocabulary happens to be at a point in time. A statement that reproduces an earlier one re-asserts its content without adding explanation; its level of entropy may be high, yet the additional effort is nil, and the change correctly records zero. A positive change thus signals that the Committee drew on a broader, less repetitive vocabulary (greater effort), a negative change that it economized on language. First-differencing has the further benefit of netting out chairperson-specific differences in baseline style, a point we develop in Section 2.2.1.

FIGURE 2. Changes in Log-Entropy of FOMC Statements, 1994–2022

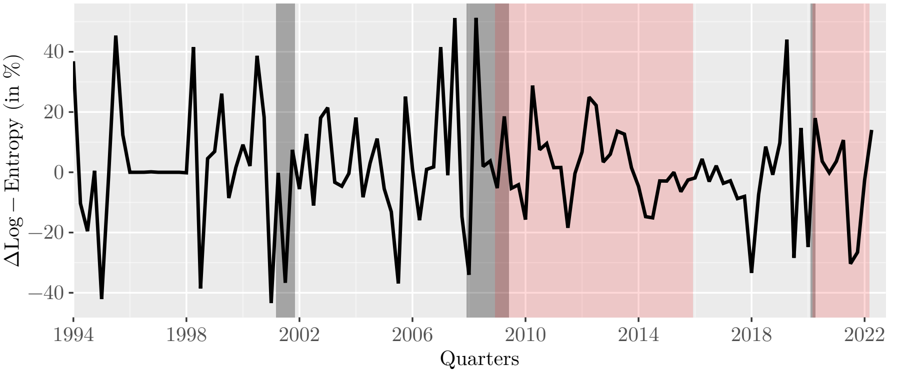

Note: Changes in log-entropy of FOMC statements at quarterly frequencies. NBER recessions and zero lower bound (ZLB) episodes are shaded in grey and light red, respectively.

Figure 2 plots the log-difference in Shannon entropy of FOMC statements from 1994 to 2022, giving a direct read on how communication effort shifted over time. The clearest pattern is that entropy spikes at the onset of each NBER recession: the Committee broadened its vocabulary and elaborated its reasoning precisely when conditions deteriorated most rapidly, which hints toward a countercyclical communication rule. The Global Financial Crisis produced the largest single move in the sample, with the communication effort surging in 2008 before partially reverting as the initial shock faded away. During ZLB episodes the series remains mostly positive, though at more moderate levels than at recession onsets, suggesting that the Committee continued to elaborate but at a steadier, less reactive pace once the initial crisis had passed and forward guidance became the primary tool. Away from recessions and ZLB periods, the change in log-entropy moves around zero; in calmer times, the Committee neither systematically expanded nor contracted its explanatory effort. The overall picture is of a Fed that treats communication as a resource to be rationed, elaborating when uncertainty is high and pulling back once policy has settled into a recognizable regime.

2.2 The Federal Reserve's Communication Reaction Function

We now specify and estimate a communication reaction function for the Federal Reserve, using the quarterly aggregated entropy 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 extend this framework to model the systematic component of communication effort by specifying an analogous reaction function for changes in entropy of FOMC statements.[8] 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 that captures how communication effort responds to expected macroeconomic conditions. We model the change in log-entropy as a function of expected economic outcomes. 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, so that the estimated coefficients are not confounded by differences in communication style across Fed chairs. The rule is therefore identified off the within-chair variation of communication effort in response to economic conditions. The rule is:

$$\Delta e_t = \rho_e(L)\,\Delta e_{t-1} + (1-\overline{\rho}_e)\,\Delta e_t^* \,, \tag{2}$$

$$\Delta e_t^* = \Delta_e^* + \psi_\pi\left(\mathbb{E}\left[\pi_{t+k,t}\,|\,\mathcal{I}_t\right] - \pi^*\right) + \psi_y\,\mathbb{E}\left[\tilde{y}_{t+q,t}\,|\,\mathcal{I}_t\right] \,, \tag{3}$$

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 (2) 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.[9]

2.2.2 Estimation Procedure System (2)–(3) simplifies to

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

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

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

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

$$\Delta e_t = \rho_e(L)\,\Delta e_{t-1} + (1-\overline{\rho}_e)\left(\Delta_e^* + \psi_\pi\left(\pi_{t+k,t} - \pi^*\right) + \psi_y\,\tilde{y}_{t+q,t}\right) + \eta_t\,, \tag{4}$$

where

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

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

In equation (4), 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 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, $\text{cov}(\boldsymbol{z}_t, \eta_t) = 0$, and the relevance condition, $\text{cov}(\boldsymbol{z}_t, x_t) \neq 0$ for $x_t = (\pi_{t+k,t}, \tilde{y}_{t+q,t})$, constitutes a valid instrument. The implied moment conditions are:

$$\mathbb{E}\left[\left[\Delta e_t - \rho_e(L)\,\Delta e_{t-1} - (1-\overline{\rho}_e)\Big(\Delta_e^* + \psi_\pi\left(\pi_{t+k,t} - \pi^*\right) + \psi_y\,\tilde{y}_{t+q,t}\Big)\right]\boldsymbol{z}_t\right] = 0\,. \tag{5}$$

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

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

where $\overline{m}_T(\boldsymbol{\theta})$ is the sample counterpart of the moment conditions in (5) 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.[10] 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.[11] 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.[12]

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

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

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

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:[13]

$$\Delta e_t = \rho_e \, \Delta e_{t-1} + (1-\rho_e)\left(\Delta_e^* + \psi_\pi\left(\mathbb{E}\left[\pi_{t+1,t} \,|\, \mathcal{I}_t\right] - \pi^*\right) + \psi_y \, \mathbb{E}\left[\tilde{y}_{t+1,t} \,|\, \mathcal{I}_t\right]\right) \,. \tag{6}$$

The baseline instrument set includes a constant, the first lag of the change in log-entropy, and four lags of inflation, the output gap, the shadow federal funds rate of Xia and Wu (2016),[14] annualized quarterly commodity price inflation, and annualized quarterly M2 growth.[15] We refer to this as the internal instrumental strategy, as it relies on lagged values 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. As noted above, Δ_e^* and π^* are not separately identified; we therefore 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 (6). The Sargan–Hansen test yields a p-value of 0.28, indicating that the null 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 (6) 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

[13] Specifications featuring more lags are not supported by the data, as changes in entropy display little persistence.

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

[15] Further details on data construction are provided in Appendix A.3.

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.[16] These results also bear on the interpretation of the measure itself. That the change in lexical diversity 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 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 individual 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. Against this stable backdrop, each exercise below adds to this picture.

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

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.

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 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.

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,[17] and the lag length of the instrument set is reduced to two to preserve efficiency.

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.

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 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 expectations proxy.

Instrumental Strategy The exercises above all rely on lagged internal macroeconomic variables as instruments. We now assess whether the rule survives an alternative strategy that uses externally identified, exogenous shocks as instruments, following Barnichon and Mesters (2020), who develop this approach to estimate the structural equations of the New Keynesian model. The key argument is that identified structural shocks are, by construction, orthogonal to the endogenous innovation in the equation of interest (here, η_t) and shift the right-hand-side variables, satisfying both the exogeneity and relevance conditions for valid instruments. We exploit oil price surprises from Känzig (2021) to isolate supply-driven

[17]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.

variation in expected inflation and activity, and high-frequency monetary policy surprises from Nakamura and Steinsson (2018)[18] to isolate demand-driven variation. Six instrument sets are considered (Table 6): IV1 and IV4 use four and twelve lags of oil price surprises and their absolute values; IV2 and IV5 use four and twelve lags of monetary policy surprises and their absolute values; IV3 and IV6 pool both types of shocks. All sets include a constant and one lag of change in log-entropy.

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 value (volatility); demand-side sets use lags of Nakamura and Steinsson 2018 monetary policy surprises (updated by Acosta, Brennan, and Jacobson 2024) and their absolute value. Combined sets pool both. The rightmost column reports the p-value of the Sargan–Hansen over-identification test. *** (**, *) denotes significance at the 1% (5%, 10%) level.

The core findings are preserved: π^* clusters tightly around 2 percent and ρ_e remains small or modestly negative across all six specifications. The exercise also reveals an informative heterogeneity in the Committee's communication response across the source of the macroeconomic shock.

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

Under supply-side instruments (IV1, IV4), the inflation response ψ_π is positive, large and significant (6.7 to 10.3), while the output-gap response is small and statistically indistinguishable from zero. When inflationary pressure originates from supply shocks, the Committee concentrates its additional communication effort on explaining the price movement, without elaborating extensively on the associated output fluctuations. This is consistent with the view that output losses following supply shocks reflect technological constraints on which communication has limited traction, while the inflation implications require active clarification to prevent expectations from becoming unanchored.

Under demand-side instruments (IV2, IV5), the inflation response remains positive and significant, but is now accompanied by a large and significant negative output-gap response (around -14 under IV2, declining to more moderate values as lags are added). Demand-driven contractions appear to mobilize both channels simultaneously: the Committee both addresses its inflation stance and signals active concern for the real economy, a pattern that matches the dual nature of demand-side stabilization, in which recessions and disinflationary pressures co-exist and require joint communication.

When supply and demand instruments are pooled (IV3, IV6), the estimates settle between the two extremes, with positive inflation responses and negative output-gap responses of moderate magnitude. The Sargan–Hansen p-values for the combined specifications (0.11 and 0.10) are close to conventional rejection thresholds, suggesting some tension in the moment conditions when the two shock types are pooled. The cleanest identification, and the best-performing over-identification tests, are obtained when supply and demand instruments are used separately, where the source of variation is well defined.

The divergence in estimates across instrument sets deserves a careful interpretation. The fact that our supply- and demand-side estimates differ (particularly for ψ_y) is informative. Supply shocks shift expected inflation and output along one dimension; monetary policy surprises shift them along another. If the true rule responds differently to these two types of disturbances, as the estimates suggest, then each instrument set captures a distinct behavioral response of the central banker, and the two diverge. The divergence is evidence of state-dependence in the communication rule, not of instrument invalidity. The linear specification is an approximation to a richer function in which the sensitivity of communication effort to real activity depends on the nature of the underlying shock. The pattern of J-tests for the pooled specifications (IV3, IV6) is consistent with this reading: constraining a single linear coefficient to satisfy moment conditions from two different local responses creates the near-rejection observed.[19]

Taken together, the external-instrument exercise corroborates the baseline and adds a causal dimension: the Committee's communication effort responds systematically to exogenous shifts in macroeconomic conditions, with the inflation channel active in both supply- and demand-driven episodes, and the output-gap channel specifically mobilized when the disturbance originates on the demand side. This asymmetry clarifies the dual role of FOMC communication as both an inflation-management and an expectations-stabilization instrument.

[19] A complete treatment of this state-dependence would require a non-linear specification of the rule, which we leave for future work.

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

$$\log \mathcal{E}(X_t) = \log \mathcal{L}_t + \log\left(\mathcal{E}(X_t)/\mathcal{L}_t\right) \,, \tag{7}$$

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.

Two empirical regularities motivate our use of joint entropy rather than either margin in isolation. First, the correlation between $\log \mathcal{E}(X_t)$ and $\log \mathcal{L}_t$ is 0.97 in log-levels and 0.91 in first differences, confirming that the extensive margin dominates low-frequency entropy dynamics and that meeting-to-meeting changes in lexical diversity are largely accounted for by changes in communication volume. Second, lexical density is not redundant: its correlation structure differs sharply across frequencies. In log-levels, episodes of longer statements are also episodes of genuinely higher per-unit diversity; the Committee adds distinct words, not merely more words. In first differences, the relationship reverses, with a correlation of -0.73 between $\Delta \log \mathcal{L}_t$ and $\Delta \log(\mathcal{E}(X_t)/\mathcal{L}_t)$: meeting-to-meeting, additional length dilutes rather than enriches lexical diversity, as marginal words are drawn from familiar vocabulary.

These regularities are mutually consistent. At low frequencies, crises produce statements that are simultaneously longer and more lexically varied. At higher frequencies, additional length is achieved by elaborating existing phrasing rather than introducing novel concepts, and the intensive margin partially offsets the extensive one. A measure that conflates the two (such as raw statement length) will capture the broad sweep of FOMC communication effort but will systematically mismeasure the direction of change at cyclical frequencies, precisely the dimension of interest for identifying the macroeconomic and financial consequences of policy communication. Total Shannon entropy retains both margins and is a strictly richer object than either alone.

To assess how much of the rule is driven by the extensive margin alone, we re-estimate equation (6) 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 (6) 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 static 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 a U.S. idiosyncrasy, 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 more ambitious claim is that systematic, state-contingent communication is a general regularity of inflation-targeting central banks: a behavioral pattern that adapts its form to the structural features of each economy but shares a common architecture. We test this claim by re-estimating the rule for five additional industrialized economies (Australia, Canada, Japan, the Euro Area, and Sweden) spanning a range of mandates, objectives, sizes and degrees of external openness.

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:

$$\Delta e_t = \rho_e \Delta e_{t-1} + (1-\rho_e)\Big(\Delta_e^* + \psi_\pi \big(\mathbb{E}[\pi_{t+1,t}|\mathcal{I}_t] - \pi^*\big) + \psi_y\, \mathbb{E}[\tilde{y}_{t+1,t}|\mathcal{I}_t] + \psi_{\Delta s}\big(\mathbb{E}[s_{t+1,t}|\mathcal{I}_t] - \Delta s^*\big) + \psi_r\big(\mathrm{FFR}_t - r^*\big)\Big)\,. \quad (8)$$

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.[20] All forecasting horizons are set to one quarter. Table 8 reports the results.

[20] 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 exchange rate changes 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.

Common Features The over-identification test passes comfortably in every specification, establishing that systematic forward-looking communication rules are not a U.S. institutional idiosyncrasy. 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. 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.

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.

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 simply does not fit the Bank of Canada's communication when restricted to domestic conditions. Once the U.S. federal funds rate is added, the picture changes dramatically. 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. Communication in the 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 an otherwise standard 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, a government 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 do not directly observe the realization of aggregate shocks. Instead, they receive noisy signals about the underlying state of the economy whose precision depends fundamentally on the communication effort of the monetary authority.

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 defined over consumption, c_t, and hours worked, n_t. Her preferences are represented by the following intertemporal utility function

$$\widetilde{\mathbb{E}}_0 \left[\sum_{t=0}^{\infty} \beta^t \left(\frac{c_t^{1-\gamma}}{1-\gamma} - \nu \frac{n_t^{1+\varphi}}{1+\varphi} \right) \right] , \tag{9}$$

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 from full-information rational expectations (FIRE) 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): professional forecasters disagree substantially about inflation, current output, and the output gap; households' inflation expectations are dispersed and systematically biased; and high-frequency revisions to preliminary data releases suggest that even sophisticated agents cannot fully reconstruct the current state from observable prices and quantities alone. We capture this evidence parsimoniously by assuming that agents do not directly observe the aggregate state of the economy, $\boldsymbol{\Xi}_t$, which, as will become clear later, consists of a cost-push shock, θ_t and a DIS shock, x_t. Both shocks are modeled as AR(1) processes with persistence ρ_i and innovation volatility σ_i^2, $i \in \{\theta, x\}$. Instead, the representative agent observes a public signal, $\boldsymbol{s}_t$, issued by the central bank

$$\boldsymbol{s}_t = \boldsymbol{\Xi}_t + \boldsymbol{\eta}_t , \qquad \boldsymbol{\eta}_t \sim \mathcal{N}(0, \boldsymbol{\Sigma}_\eta(\Delta e_t)) , \tag{10}$$

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

$$\boldsymbol{\Sigma}_\eta(\Delta e_t) = \sigma_\eta^2(\Delta e_t) \begin{pmatrix} \varsigma_\theta & 0 \\ 0 & \varsigma_x \end{pmatrix} \quad \text{with} \quad \varsigma_\theta, \varsigma_x > 0 . \tag{11}$$

This formulation implicitly assumes that the noises are uncorrelated and, more importantly, that communication effort affects the precision of the signal at once, leaving the relative precision of the various signals unaffected. 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. Conditional on $\mathcal{F}_{t-1}$, the information set at $t-1$, the prior over $\boldsymbol{\Xi}_t$ is Gaussian: $\boldsymbol{\Xi}_t \mid \mathcal{F}_{t-1} \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_\mu)$, where $\boldsymbol{\mu}_t \equiv \mathbb{E}_{t-1}[\boldsymbol{\Xi}_t]$ and $\mathbb{E}_{t-1}[\cdot]$ denotes the FIRE conditional expectation based on $\mathcal{F}_{t-1}$.

Two features deserve attention. First, the signal is public: it is issued by the central bank and 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. Here, because the signal is common, individual posterior beliefs coincide with the aggregate posterior, and 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.[21] Second, the signal precision $\sigma_\eta^{-1}(\Delta e_t)$ is strictly increasing with communication effort, capturing the empirical observation that increases in the informativeness of central bank statements are associated with larger reductions in forecast disagreement (Ehrmann and Fratzscher, 2007) and macroeconomic uncertainty (Husted, Rogers, and Sun, 2020).

Upon observing $\boldsymbol{s}_t$, the household updates optimally via Bayes' rule. The posterior is:

$$\boldsymbol{\Xi}_t \mid \boldsymbol{s}_t, \mathcal{F}_{t-1} \sim \mathcal{N}\left(\widetilde{\boldsymbol{\Xi}}_t, \widetilde{\boldsymbol{\Sigma}}_t\right) , \tag{12}$$

with posterior moments:

$$\widetilde{\boldsymbol{\Xi}}_t = \boldsymbol{\mu}_t + \boldsymbol{K}_t(\boldsymbol{s}_t - \boldsymbol{\mu}_t) , \quad \text{and} \quad \widetilde{\boldsymbol{\Sigma}}_t = \boldsymbol{K}_t \boldsymbol{\Sigma}_\eta(\Delta e_t) , \tag{13}$$

where $\boldsymbol{K}_t \equiv (\boldsymbol{\Sigma}_\mu + \boldsymbol{\Sigma}_\eta(\Delta e_t))^{-1}\boldsymbol{\Sigma}_\mu$ denotes the Kalman gain. When communication is highly precise ($\sigma_\eta^2 \to 0$), $\boldsymbol{K}_t$ reduces to the identity matrix; agents discard the prior and adopt the signal, recovering the period-t FIRE expectation. When communication is entirely uninformative ($\sigma_\eta^2 \to \infty$), $\boldsymbol{K}_t \to \boldsymbol{0}$ and agents ignore the signal and rely entirely on the prior $\boldsymbol{\mu}_t = \mathbb{E}_{t-1}[\boldsymbol{\Xi}_t]$. In the general case, the Kalman gain summarizes the quality of central bank communication in a single sufficient statistic, and, as we show below, it is $\boldsymbol{K}_t$ that governs the mean distortion channel in the Dynamic IS equation.

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

$$\boldsymbol{\delta}_t \equiv \widetilde{\boldsymbol{\Xi}}_t - \boldsymbol{\Xi}_t = \boldsymbol{K}_t \boldsymbol{\eta}_t - (\boldsymbol{I} - \boldsymbol{K}_t)\widehat{\boldsymbol{\Xi}}_t , \tag{14}$$

where $\boldsymbol{I}$ denotes the identity matrix and $\widehat{\boldsymbol{\Xi}}_t \equiv \boldsymbol{\Xi}_t - \boldsymbol{\mu}_t$ is the unpredictable component of $\boldsymbol{\Xi}_t$. The error $\boldsymbol{\delta}_t$ captures two distinct phenomena. The first term, $\boldsymbol{K}_t\boldsymbol{\eta}_t$, captures the effect of noise. Agents partially attribute signal noise to a genuine movement in the state, generating expectation errors that are driven entirely by the communication rather than by fundamentals. The second term, $-(\boldsymbol{I}-\boldsymbol{K}_t)\widehat{\boldsymbol{\Xi}}_t$, implies that agents underreact to the current shock by a factor of $\boldsymbol{I} - \boldsymbol{K}_t$, as they can only partially distinguish a genuine innovation in $\boldsymbol{\Xi}_t$ from the level they expected. Both sources of error are decreasing in communication

[21] The extension to dispersed private information, in which a fraction of agents also receive idiosyncratic signals about θ_t, would introduce higher-order belief dynamics (Angeletos and Lian, 2018) and potentially non-monotone effects of public communication; we leave this extension for future work.

quality: as $\boldsymbol{K}_t \to \boldsymbol{I}$ the systematic underreaction disappears, and simultaneously $\boldsymbol{\Sigma}_\eta \to \mathbf{0}$ so that the noise term also vanishes. This monotone relationship between communication effort and expectation accuracy is the foundation of the mean distortion channel.

Household's Decision In every period, the representative household consumes and supplies hours of labor at the real wage w_t. She also receives real dividends d_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 $X_{t-1} = \exp(\hat{x}_{t-1})$ is an exogenous shock that will act as a DIS shock in the model,[22] where $\hat{x}_t$ follows an AR(1) process with persistence $\rho_x \in [0,1)$ and Gaussian innovation $\epsilon_{xt} \sim \mathcal{N}(0, \sigma_x^2)$. The agent therefore faces the budget constraint:

$$c_t + b_{t+1} = \frac{R_{t-1}}{X_{t-1}\pi_t} b_t + w_t n_t + d_t \,, \tag{15}$$

where $\pi_t \equiv P_t/P_{t-1}$ denotes gross price inflation. The household decides her consumption-saving plans and labor supply by maximizing (9) subject to (15).

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 will vary exogenously following an AR(1) process, as a way to capture a cost-push shock.

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

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

[22] 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.

where the real wage is common to all firms.

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.[23] 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.[24] Hence, in a symmetric equilibrium, the optimal price-setting condition writes

$$(\pi_t - 1)\pi_t = \frac{\theta_t}{\Psi}\left[mc_t - \frac{\theta_t - 1}{\theta_t}\right] \ . \tag{16}$$

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 aggregate state of the economy and therefore operates under FIRE. It also pursues its mandate through two channels: *(i)* standard monetary policy and *(ii)* communication.

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

$$R_t = R \cdot (\mathbb{E}_t[\pi_{t+1}]/\pi)^{\phi_\pi} \cdot (\mathbb{E}_t[y_{t+1}]/y)^{\phi_y} \ . \tag{17}$$

The central bank follows a Taylor (1993) rule with positive feedback coefficients to expected inflation and output, $\phi_\pi \geq 0$ and $\phi_y \geq 0$.[25]

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:[26]

$$\exp(\Delta e_t) = (\mathbb{E}_t[\pi_{t+1}]/\pi)^{\psi_\pi} \cdot (\mathbb{E}_t[y_{t+1}]/y)^{-\psi_y} \ , \tag{18}$$

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.[27]

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

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

[25] 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.

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

[27] Note that, implicit in Equation (18), communication is not subject to a bound akin to the ZLB that affects the nominal interest rate, implying that the central bank always communicates in the model, and extension of the model would be to allow for a communication cost that may limit the ability of the central banker to talk during certain events.

Two features of this specification deserve explicit note, as both are grounded in the empirical evidence we presented above. First, the rule is static and 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 (18) coincides with the one appearing in the Taylor rule (17), suggesting that the two instruments are anchored to a common objective.

The communication rule (18) 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 forestall self-fulfilling inflationary spirals. Second, the central bank ramps up 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, helps coordinate private expectations around a recovery path and provides direct stimulus to aggregate demand through the expectation channel.

3.2 The Dynamic IS Equation

This section aims at deriving the dynamic IS equation of the model under imperfect information, and makes explicit how communication effort will affect aggregate demand in an equilibrium. Using the good market clearing condition $c_t = y_t$, the household's first-order condition yields the Euler equation:

$$y_t^{-\gamma} = \beta \frac{R_t}{X_t} \widetilde{\mathbb{E}}_t \left[\frac{y_{t+1}^{-\gamma}}{\pi_{t+1}} \right] , \tag{19}$$

where $\widetilde{\mathbb{E}}_t[\cdot]$ is the expectation under the posterior distribution (12). Denoting $\widehat{v}_t$ the log deviation of variable v_t from its deterministic level, 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

$$-\gamma \widehat{y}_t = \widehat{\imath}_t - \widehat{x}_t - \gamma \widetilde{\mathbb{E}}_t[\widehat{y}_{t+1}] - \widetilde{\mathbb{E}}_t[\widehat{\pi}_{t+1}] + \Omega_t , \tag{20}$$

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 premium component. We then obtain the following result.

Proposition 1 (Risk as Communication). *Up to a first order, the risk premium writes*

$$\Omega_t \approx -\chi_e \, \widehat{\Delta} e_t , \tag{21}$$

where, the communication effectiveness, χ_e, takes the form

$$\chi_e \equiv -\Upsilon \frac{\partial \sigma_\eta^2(\Delta e_t)}{\partial \Delta e_t}\bigg|_{\Delta e_t = \overline{\Delta} e} > 0\,. \tag{22}$$

Communication effectiveness[28] can be given a structural interpretation as products of two distinct objects. The first, Υ, is actually related to the sensitivity of future log-marginal utility $(-\gamma\widehat{y}_{t+1} - \widehat{\pi}_{t+1})$ to the current state $\boldsymbol{\Xi_t}$. It depends on the properties of the shocks, their persistence and volatility, on the volatility of the signals and the Kalman gains.[29] It can be identified directly from the model's impulse response functions.What governs the precautionary premium, underlying the communication channel, is therefore how strongly the *entire intertemporal marginal cost of consumption* is exposed to risk the household cannot resolve about current conditions. The second object is the marginal reduction in posterior 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 (21) 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:*

$$\widehat{y}_t = \mathbb{E}_t[\widehat{y}_{t+1}] - \frac{1}{\gamma}\left(\widehat{\imath}_t - \mathbb{E}_t[\widehat{\pi}_{t+1}]\right) + \Lambda_\delta \Big(\underbrace{\boldsymbol{K}\boldsymbol{\eta}_t - (\boldsymbol{I}-\boldsymbol{K})\widehat{\boldsymbol{\Xi}}_t}_{\boldsymbol{\delta}_t}\Big) + \frac{\chi_e}{\gamma}\widehat{\Delta} e_t + \frac{\widehat{x}_t}{\gamma}\,, \tag{23}$$

where $\Lambda_\delta = \frac{\gamma\Phi_y + \Phi_\pi}{\gamma}$, $\boldsymbol{K} \equiv (\boldsymbol{\Sigma}_\mu^2 + \boldsymbol{\Sigma}_\eta^2(\overline{\Delta} e))^{-1}\boldsymbol{\Sigma}_\mu^2$ is the steady-state Kalman gain, $\widehat{\boldsymbol{\Xi}}_t \equiv \boldsymbol{\Xi}_t - \boldsymbol{\mu}_t$ is the surprise component of the current aggregate state, and $\boldsymbol{\eta}_t$ is the signal noise in (10). Φ_y and Φ_π are the one step ahead response of output and inflation to the shocks.

Relative to the standard New Keynesian DIS equation, it contains two additional terms that capture the two channels through which imperfect information about current conditions enters aggregate demand. We discuss each in turn.

The Mean Distortion Channel The first new term, $+(\gamma\Phi_y + \Phi_\pi)\boldsymbol{K}\boldsymbol{\eta}_t/\gamma$, captures a noise-driven mismatch between agents' point forecasts of future outcomes and the FIRE benchmark. Because agents partially attribute signal noise to a genuine shock, their expectations of future outcomes are distorted even when $\widehat{\boldsymbol{\Xi}}_t = 0$, generating spurious demand fluctuations reminiscent of the noise-driven business cycles (see Lorenzoni, 2009, Blanchard,

[28]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.

[29]While this is essentially immaterial for our results, the interested reader is referred to Section B.2.

L'Huillier, and Lorenzoni, 2013, among others). Importantly, as the signal becomes more precise ($\sigma_\eta^2 \to 0$), 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 new term, $-(\gamma\Phi_y + \Phi_\pi)(\boldsymbol{I} - \boldsymbol{K})\widehat{\boldsymbol{\Xi}}_t/\gamma$, captures a systematic mismatch between agents' point forecasts of future outcomes and the FIRE benchmark. When a shock $\widehat{\boldsymbol{\Xi}}_t > 0$ hits the economy, a fully informed agent would revise her expectations of future consumption and inflation upward by $\Phi_y\widehat{\boldsymbol{\Xi}}_t$ and $\Phi_\pi\widehat{\boldsymbol{\Xi}}_t$ respectively, and adjust current consumption accordingly. An imperfectly informed agent, who cannot fully disentangle the shock from noise in the signal, 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 the shock by a factor $(\boldsymbol{I} - \boldsymbol{K})$, generating the dampening term in (23). 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.

The Risk Channel The term $\chi_e\widehat{\Delta}e_t/\gamma$ is the risk channel identified in the model. It captures the direct stimulative effect of communication effort on aggregate demand through the compression of the precautionary savings premium Ω_t. Higher communication effort reduces the posterior variance $\widetilde{\sigma}_{it}^2$, lowering the risk premium in the Euler equation and making current consumption less sensitive to precautionary motives. The effect is proportional to χ_e, defined in (22), which governs how many units of future marginal utility variance are eliminated per unit of communication effort. This channel is conceptually related to the risk-based demand channel of Acharya and Dogra (2020), though it arises here in a representative-agent framework. In our setting, the relevant risk is aggregate, rather than idiosyncratic, risk about the current state and communication resolves it.

3.3 Equilibrium Dynamics

In the general equilibrium, the log-linear representation of the model is summarized by the following dynamics:

$$\widehat{y}_t = \mathbb{E}_t[\widehat{y}_{t+1}] - \frac{1}{\gamma}\left(\widehat{\imath}_t - \mathbb{E}_t[\widehat{\pi}_{t+1}]\right) + \Lambda_\delta \boldsymbol{\delta}_t + \frac{\chi_e}{\gamma}\widehat{\Delta}e_t + \frac{\widehat{x}_t}{\gamma}\,, \tag{24}$$

$$\widehat{\pi}_t = \kappa\widehat{y}_t + \widehat{\theta}_t\,, \tag{25}$$

$$\widehat{\imath}_t = \phi_\pi\mathbb{E}_t[\widehat{\pi}_{t+1}] + \phi_y\mathbb{E}_t[\widehat{y}_{t+1}]\,, \tag{26}$$

$$\widehat{\Delta}e_t = \psi_\pi\mathbb{E}_t[\widehat{\pi}_{t+1}] - \psi_y\mathbb{E}_t[\widehat{y}_{t+1}]\,. \tag{27}$$

Equation (24) is the standard dynamic IS equation extended for the presence of imperfect information and communication. Equation (25) is the New Keynesian Phillips curve, Equation (26) and (27) capture the monetary policy apparatus available to the central banker: the nominal interest rate rule and the communication rule. The definition of the DIS shock, $\widehat{x}_t$, and the cost push shock, $\widehat{\theta}_t$ completes the description of the equilibrium dynamics.

Policy Substitutability The first implication of the DIS equation (24) is a precise mapping between the two policy instruments. In the DIS equation, a unit increase in the nominal 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 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 nominal 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: 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

$$\widehat{y}_t = \Theta\mathbb{E}_t[\widehat{y}_{t+1}] - \frac{1}{\gamma}(\widehat{\imath}_t - \mathbb{E}_t[\pi_{t+1}]) + \frac{\chi_e\psi_\pi}{\gamma}\mathbb{E}_t[\widehat{\theta}_{t+1}] + \frac{\widehat{x}_t}{\gamma} + \Lambda_\delta\boldsymbol{\delta}_t\,, \tag{28}$$

where $\Theta \equiv 1 + \frac{\chi_e}{\gamma}(\kappa\psi_\pi - \psi_y)$ denotes the 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. This is fully determined by the sign of $\psi_\pi\kappa - \psi_y$, where the first term captures, through the NKPC, the effect of a variation of output on the inflation channel of communication, controlled by ψ_π, and the second term is the direct output channel of communication, controlled by ψ_y. 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.

The economic content of this condition deserves examination. Consider an environment in which future output is expected to be high. Through the Phillips curve, 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, independent of departing from full-information rational expectations per se. 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 (28) offers a distinct channel: countercyclical central bank communication. While the resulting reduced-form discounting is mathematically analogous to information-based and behavioral approaches such as Gabaix (2020) and Meichtry (2023), the economic interpretation differs: discounting arises from the systematic conduct of monetary policy itself through state-dependent communication effort. The implication is direct: 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.

Output Dynamics The model admits a compact reduced-form representation in terms of output dynamics that makes the macroeconomic role of communication effort transparent. Substituting further the NKPC, the nominal interest rule and the communication rule into the transformed DIS equation (28), we obtain the following proposition.

Proposition 4 (Output Dynamics). *The aggregate output dynamics are described by:*

$$\widehat{y}_t = \Big[1 + \Lambda_{\Delta e} - \Lambda_r\Big]\mathbb{E}_t[\widehat{y}_{t+1}] + \Lambda_\theta \mathbb{E}_t[\widehat{\theta}_{t+1}] + \frac{\widehat{x}_t}{\gamma} + \Lambda_\delta \boldsymbol{\delta}_t\,, \tag{29}$$

where $\Lambda_{\Delta e} \equiv \frac{\chi_e}{\gamma}(\psi_\pi \kappa - \psi_y)$, $\Lambda_r \equiv \frac{\kappa(\phi_\pi - 1) + \phi_y}{\gamma}$, *and* $\Lambda_\theta \equiv \frac{\chi_e}{\gamma}\psi_\pi - \frac{\phi_\pi - 1}{\gamma}$.

The scalar Λ_r captures the effects of the real interest rate in determining output dynamics through adjustments in the nominal interest rate and the Fisher effect. Together with Λ_r, $\Lambda_{\Delta e}$ is a key object to determine the stability properties of the economy. It encodes, in a single expression, how the systematic component of communication policy reshapes the propagation of macroeconomic disturbances. Specifically, this term aggregates four structural ingredients: the communication effectiveness (χ_e), the responsiveness of communication to the inflation and output outlooks (ψ_π, ψ_y), the slope of the Phillips curve (κ) and risk aversion, γ. In a standard New Keynesian model without communication, the coefficient in front of expected output in the DIS equation is unity. Communication shifts this coefficient to $\Theta \equiv 1 + \Lambda_{\Delta e}$, a modification with far-reaching consequences.

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. The quarterly discount factor $\beta = 0.99$ implies an annual steady-state real interest rate of roughly 4%. The coefficient of relative risk aversion is $\gamma = 1$ (log utility), and the slope of the Phillips curve is $\kappa = 0.05$, consistent with the ballpark of estimates for the U.S. with moderate price stickiness. The shock persistence parameters are $\rho_\theta = 0.95$ for the cost-push shock and $\rho_x = 0.90$ for the demand shock. The interest rate rule coefficients $\phi_\pi = 2.0$ and $\phi_y = 0.15$ satisfy the Taylor principle and are close to the estimates of Clarida, Galí, and Gertler (2000) for the post-1982 period. The communication rule coefficients $\psi_\pi = 7$ and $\psi_y = 1.5$ are taken directly from the baseline GMM estimates reported in Table 2.

The communication effectiveness parameter χ_e is harder to pin down from standard macro data, since it reflects the marginal impact of communication effort on private posterior uncertainty, a quantity not directly observable. 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 press conferences, they find a tight correspondence between FOMC communication and subsequent media coverage (correlation of 0.83; regression coefficient of 0.65), 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.

As will become clear, for the estimated communication rule, the model generates discounting in the IS equation. Gabaix (2020) sets the degree of discounting to 0.85. Because we do not want to attribute all the discounting to our central bank communication channel, we adopt a more conservative value, halfway between no discounting and 0.85, and set it to 0.925. This implies $\chi_e = 0.065$.

Finally, the Kalman gains on supply (cost-push) and demand shocks are set based on estimates of the fundamental and noise components of supply and demand signals by Benhima and Poilly (2021). Their estimates imply a Kalman gain of 0.43 for the supply shock, 0.37 for the demand shock.

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.

Proposition 5 (Local Determinacy). *The model admits a locally stable solution if the following condition is satisfied:*

$$2\gamma + \chi_e \left[\kappa\psi_\pi - \psi_y\right] > \kappa(\phi_\pi - 1) + \phi_y > \chi_e \left[\kappa\psi_\pi - \psi_y\right]. \tag{30}$$

Condition (30) is a direct generalization of the standard Taylor principle for forward-looking rules (see Bullard and Mitra, 2002). For instance, setting $\psi_\pi = \psi_y = 0$ collapses it to the benchmark condition:

$$2\gamma \; > \; \kappa(\phi_\pi - 1) + \phi_y \; > \; 0\,.$$

In the presence of systematic communication, both bounds on $\kappa(\phi_\pi - 1) + \phi_y$ shift by the common amount $\chi_e[\kappa\psi_\pi - \psi_y]$. Three aspects of this shift deserve attention. First, the direction is determined by the sign of $\kappa\psi_\pi - \psi_y$: a dual-mandate communicator with $\psi_y > \kappa\psi_\pi$ shifts both bounds downward, while an inflation-only communicator shifts them upward. Second, the magnitude of the shift scales with χ_e; a more effective communicator exerts a larger influence on the determinacy region. Third, the width of the admissible region remains 2γ regardless of communication policy. In other words, communication shifts the determinacy corridor without changing its width.

The policy implications are sharp. When the central bank places sufficient emphasis on output in its communication ($\psi_y > \kappa\psi_\pi$) both bounds shift downward. The lower bound falls below zero, meaning that a Taylor rule that would ordinarily generate indeterminacy can be rescued by active communication policy. A dual-mandate communicator effectively disciplines expectations through the communication channel; the interest rate instrument can therefore be used less aggressively without sacrificing equilibrium stability. Quantitatively, using our empirical estimates, $\chi_e[\kappa\psi_\pi - \psi_y] \approx 0.065 \times (-1.15) \approx -0.075$. The admissible interval for $\kappa(\phi_\pi - 1) + \phi_y$ shifts from the standard $(0,\, 2)$ to approximately $(-0.075,\, 1.925)$: the effective Taylor principle is materially relaxed.

Conversely, when communication is anchored exclusively on inflation ($\psi_y = 0,\; \psi_\pi > 0$) both bounds shift upward, tightening the requirements on the Taylor rule. An institution with a narrowly-defined price-stability mandate in its communication must compensate with more aggressive interest rate policy to achieve the same equilibrium stability. The result has practical force as several authors have documented that the ECB's historical communication was disproportionately oriented toward price stability (see, e.g., Ehrmann and Fratzscher, 2007). In the model, such a communication posture raises the minimum interest rate responsiveness required for a determinate equilibrium.

Figure 3 reports impulse response functions of output, inflation, the nominal interest rate, and communication effort to two structural shocks, the cost-push and DIS shocks, under two policy regimes: an economy relying exclusively on the Taylor rule (dashed line) and one in which the Taylor rule and the communication rule are simultaneously active (solid line).

Figure 3. Impulse Response Functions (Normal Times)

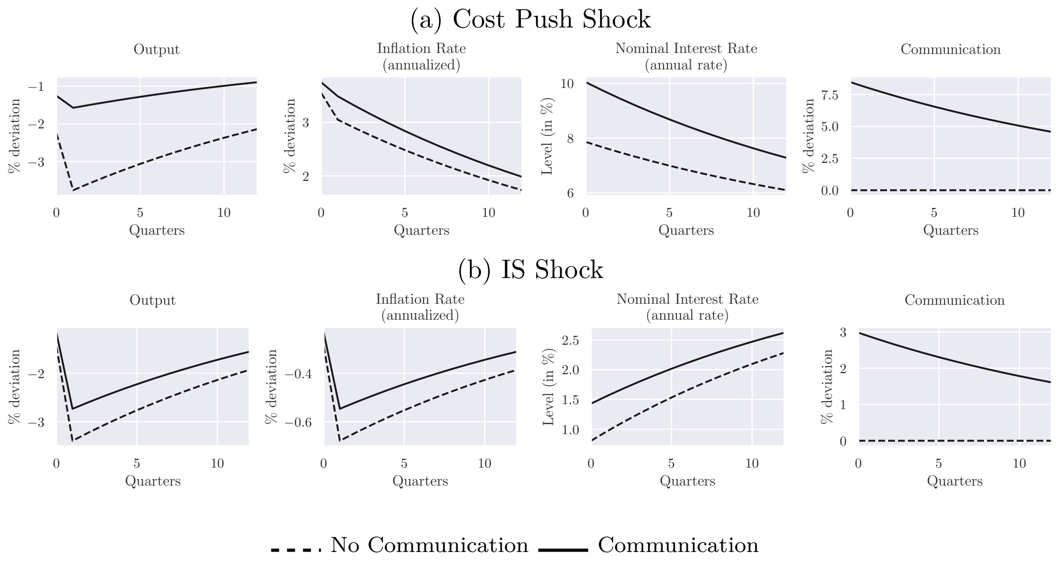

Cost-Push Shock Compared to the response obtained in the canonical NK model, the impact effect of a cost push shock is somewhat dampened by the mean distortion channel that makes agents more cautious. A 1% positive cost-push shock generates the canonical output-inflation trade-off of the New Keynesian framework: inflation rises while output contracts, confronting the central bank with conflicting stabilization objectives. In the absence of communication (dashed line), the Taylor rule alone manages this trade-off, producing a persistent inflation hike and a protracted output contraction. Active communication improves on this outcome through two simultaneous channels. As inflation rises, the communication rule intensifies central bank communication, increasing the precision of the signals the central bank sends to agents and anchoring inflation expectations more rapidly. This tames the output response. Simultaneously, the output-stabilization component of the communication rule provides a direct demand stimulus through the expectation channel that cushions the recessionary impact of the shock. The net effect, visible in Panel (a) of

Figure 3, is a smaller and shorter-lived output contraction with no material sacrifice in inflation stabilization. Active communication effort thus complements interest rate policy and addresses both the inflation and output consequences of the shock through the mean distortion channel and discounting.

Systematic communication effort during inflation episodes can therefore cushion the recessionary consequences of a supply shock even as the central bank simultaneously tightens its interest rate instrument. The two actions are not contradictory: the Taylor rule disciplines inflation by raising rates, while communication effort disciplines expectations by conveying a recovery commitment.

DIS Shock As with the cost-push shock, the mean distortion channel mitigates the impact effect of the DIS shock. Following a 1% contractionary DIS shock, output falls and so does inflation. Without communication effort (dashed line), the central bank cuts the interest rate below steady state (4%), but the instrument alone cannot prevent a sharp and persistent contraction. The economy's response is governed by the standard DIS multiplier of unity: expected future output passes through one-for-one into current output, and the high shock persistence ($\rho_x = 0.9$) generates a protracted recession.

With active communication effort, two forces operate simultaneously in opposing directions, a tension that maps onto the discounting condition of Proposition 3. The inflation channel works against stabilization as the deflationary signal induced by the IS shock prescribes a reduction in communication effort, leaving agents with less precise signals about the economic outlook, inducing greater caution and curtailing expenditures. The output channel works in the opposite direction. Falling output calls for more communication effort, reassuring agents about the future path of demand and providing a genuine stabilizing stimulus. The net effect is determined by the relative magnitudes of these forces. The inflation channel generates $\psi_\pi \kappa$ units of communication drag per unit of output contraction since the Phillips curve translates each unit fall in output into only κ units of deflation, which the communication rule then amplifies by ψ_π. The output channel generates ψ_y units of stimulus per unit of contraction. The output channel dominates if and only if $\psi_y > \psi_\pi \kappa$; precisely the discounting condition $\Theta < 1$ of Proposition 3. This equivalence is not a coincidence: the same structural condition that produces discounting in the DIS equation also governs whether the communication response to demand shocks is net stabilizing or net destabilizing. With the calibrated parameters, $\psi_\pi \kappa = 7 \times 0.05 = 0.35 < \psi_y = 1.5$, so communication rises on balance in Panel (b) of Figure 3, boosting aggregate demand and partially offsetting the contractionary impulse.[30]

Both instruments thus move in the same expansionary direction: the interest rate falls and communication rises, reinforcing rather than substituting for each other in direction

[30]The direction of the response is worth stating plainly. Although the economy is disinflating in the aftermaths of the negative demand shock, the central bank elaborates more, not less. Taken on its own, the inflation channel would prescribe the opposite since a below-target inflation outlook calls for less communication. That the Committee instead leans harder into communication precisely as disinflation sets in is therefore not a property of its inflation response but of its countercyclicality as the output channel dominates ($\psi_y > \kappa\psi_\pi$), *i.e.* when the DIS equation exhibits discounting. Note that, had the central bank operated in the compounding regime ($\psi_y < \psi_\pi \kappa$), communication would instead fall on net, further depressing demand and amplifying the recession.

while interacting through the instrument duality established above. The nominal interest rate in the communication economy falls by less than in the no-communication economy, because communication absorbs a share of the stabilization burden. The result is a materially faster recovery of both output and inflation, achieved with smaller swings in each instrument individually. The potency gain from instrument duality: an economy with an effective communication channel achieves the same stabilization outcome as a pure interest-rate economy, with a less extreme use of either instrument.

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

$$R_t = \max(R \cdot (\mathbb{E}_t[\pi_{t+1}]/\pi)^{\phi_\pi} \cdot (\mathbb{E}_t[y_{t+1}]/y)^{\phi_y}, 1)\,.$$

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. We ask whether, and under what conditions, communication policy can restore determinacy and serve as a genuine substitute for conventional monetary policy. We start by considering an economy which lies in the ZLB regime, as captured by setting $\phi_\pi = \phi_y = 0$; hence shutting down the interest rate channel altogether.

Proposition 6 (Local Determinacy (ZLB)). *At the zero lower bound, the model with central bank communication admits a locally stable solution if and only if:*

$$-2\gamma - \kappa < \chi_e\left[\kappa\psi_\pi - \psi_y\right] < -\kappa\,. \tag{31}$$

The proposition establishes that, even when the interest rate instrument is exhausted, 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.

Interpretation and the Dual Mandate Condition (31) 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 (31) through by γ and rearranging yields the equivalent condition:

$$-1 < \Theta + \frac{\kappa}{\gamma} < 1. \tag{32}$$

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 (32) 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 $1 + \kappa/\gamma > 1$. There, 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. Communication effort with a sufficiently strong output channel 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 (31). Equivalently, from (32), inflation-only communication implies $\Theta + \kappa/\gamma > 1 + \kappa/\gamma > 1$: 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.

An Upper Bound on Communication Aggressiveness There is, however, a complementary failure mode that imposes an upper bound on the output-sensitivity of communication effort. When ψ_y is very large, $\chi_e[\kappa\psi_\pi - \psi_y]$ becomes very negative, and the effective multiplier $\Theta + \kappa/\gamma$ falls below -1, violating the lower bound in (31). The economy then loses determinacy from the opposite direction. Intuitively, a communication rule that overreacts to output developments causes the central bank to aggressively scale back expansionary communication whenever the private sector anticipates a recovery; more than offsetting the recovery signal and depressing current demand. Agents anticipate this overreaction, revise down their expectations, triggering yet more communication expansion in the opposite direction, generating an ever-widening feedback loop. The model therefore identifies a precise upper bound on the output-sensitivity of communication by rearranging the lower bound in (31):

$$\psi_y \; < \; \kappa\psi_\pi + \frac{\kappa + 2\gamma}{\chi_e} \, . \tag{33}$$

This constraint is generally not binding (for the calibrated parameters, the right-hand side exceeds 32), but tightens in economies with high communication efficiency (χ_e large) or a flat Phillips curve (κ small). The dual mandate is thus a necessary but not sufficient condition for ZLB stabilization: the output-sensitivity of communication must be calibrated so that the communication strategy sits in the interior of the determinacy corridor, not at its boundary.

These results establish the dual mandate as a requirement rather than a preference. At the ZLB, an inflation-centered communication strategy is destabilizing; the only path to a determinate equilibrium runs through a communication rule that responds sufficiently aggressively to output gap developments. The minimum threshold, obtained by rearranging the upper bound in (31), is $\psi_y > \kappa(1 + \chi_e\psi_\pi)/\chi_e$: the output-stabilization coefficient must exceed a level that compensates for the destabilizing force of the inflation-response channel as it propagates through the Phillips curve. This provides a formal micro-foundation for recent arguments that central banks operating at the ZLB should expand the scope of their forward guidance beyond inflation targets to encompass explicit output or employment commitments; precisely the approach the Federal Reserve adopted with its outcome-based guidance in August 2020.

Figure 4 illustrates the quantitative implications of these results.[31] It reports impulse response functions of output, inflation, the nominal interest rate, and communication following a contractionary DIS shock, comparing the no-communication economy (dashed lines) to the active-communication economy (solid lines).[32] Two shock magnitudes are examined. Panel (a) considers a large shock of -2.0%, severe enough to push both economies to the ZLB. Panel (b) considers a milder shock of -1.53%, which binds the constraint only in the economy without communication.

Large shock (-2.0%) When the shock is large enough to push both economies to the ZLB, the dynamics diverge sharply. In the no-communication economy, the interest rate hits zero immediately and remains there for an extended episode, 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, stimulating private expectations of future demand. By sustaining a favorable expectation of future 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 χ_e/γ units of communication expansion. The outcome is a faster recovery of output and inflation, a shorter ZLB episode, and a less severe disinflationary trough. 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 effort rise as the economy sinks toward the bound. The dual mandate is thus not incidental

[31] The calibrated parameters satisfy the ZLB determinacy condition (31): $\chi_e[\kappa\psi_\pi - \psi_y] \approx -0.075$ lies in the interval $(-2.05, -0.05)$, so that the effective multiplier $\Theta + \kappa/\gamma \approx 0.975 < 1$. The equilibrium is determinate and stable, in contrast to the indeterminate standard ZLB equilibrium where the corresponding multiplier is $1.05 > 1$. The quantitative difference between the two multipliers is modest, but it is the sign of the departure from unity, not its magnitude, that governs whether sunspot fluctuations can be supported in equilibrium.

[32] The model is solved by a shooting method, starting the economy from the *business as usual* steady state and perturbing the economy by a negative IS shock.

FIGURE 4. Impulse Response Functions (Zero Lower Bound)

(a) Large Negative IS shock (-2.00%)

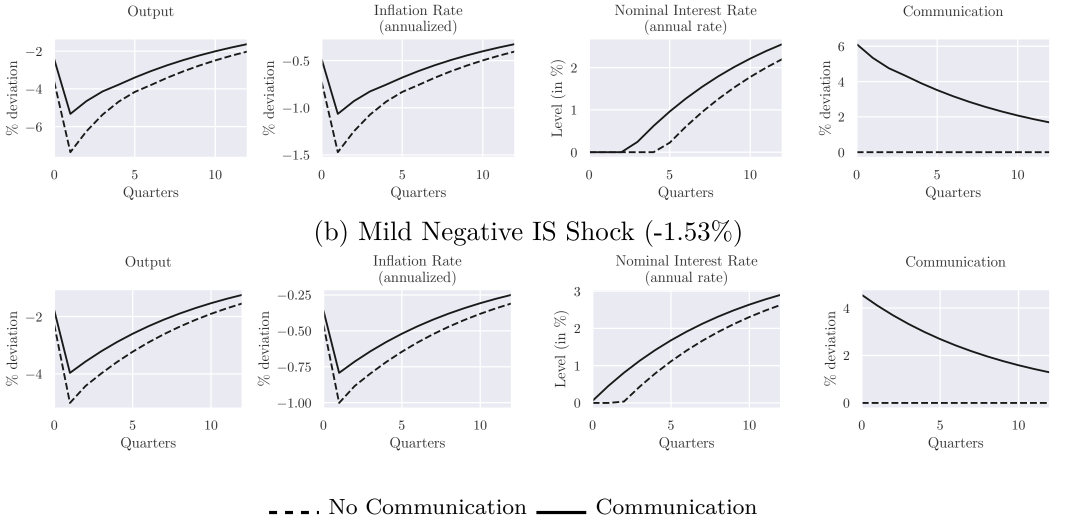

to the zero-lower-bound result, it is what keeps the central bank talking when a purely inflation-focused rule would call for silence.

Mild shock (−1.53%): The Option Value of Communication Panel (b) reveals the clearest implication of the model. Under the mild −1.53% shock, the no-communication economy still reaches the ZLB: 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, therefore acting as an automatic stabilizer. The ZLB is never reached, and the contraction is milder.[33]

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 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. 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.

Lessons for Policy Taken together, these results yield three lessons.

First, the efficacy of systematic communication is potentially highest during a ZLB episode, as greater communication effort helps reduce the time spent at the ZLB by anchoring expectations better.

[33]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.

Second, the content of the communication strategy matters as much as its intensity. An institution that communicates actively but exclusively about inflation ($\psi_y = 0$) fails to provide additional stabilization relative to the Taylor-rule-only benchmark and undermines determinacy at the ZLB. The dual mandate must be an explicit and operational component of the communication framework, with a positive and sufficiently large ψ_y.

Third, communication effectiveness is key. The technology χ_e governs both the impact multiplier of communication effort and its rate of substitution with conventional monetary policy. Institutions with high χ_e (those whose statements demonstrably shift private beliefs, as evidenced by large market reactions to central bank announcements documented in the empirical literature) can achieve a given stabilization objective with smaller movements in each instrument. A central bank that has allowed its communication framework to deteriorate, or that has communicated inconsistently across the business cycle, faces a double disadvantage at the ZLB: it loses both the demand effect of communication and the expectation-anchoring benefit that shortens ZLB episodes. The case for systematic central bank communication rests ultimately on its capacity to forestall ZLB episodes, not on its performance once the constraint has already bound.

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, 35, 37]

Acosta, M., C. M. Brennan, and M. M. Jacobson (2024): "Constructing High-Frequency Monetary Policy Surprises from SOFR Futures," *Economics Letters*, 242, 111873. [21]

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, 37]

Andrews, D. W. K., and J. C. Monahan (1992): "An Improved Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimator," *Econometrica*, 60(4), 953–966. [15]

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]

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, 20]

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. [37]

Benhabib, J., S. Schmitt-Grohé, and M. Uribe (2001): "Monetary Policy and Multiple Equilibria," *The American Economic Review*, 91(1), 167–186. [43]

Benhima, K., and C. Poilly (2021): "Does Demand Noise Matter? Identification and Implications," *Journal of Monetary Economics*, 117, 278–295. [39]

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, 37]

Binder, C. C. (2017): "Measuring Uncertainty Based on Rounding: New Method and Application to Inflation Expectations," *Journal of Monetary Economics*, 90, 1–12. [2]

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. [29, 34]

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. [2, 10]

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. [2, 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. [2, 7, 11, 14, 38, 46, – 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, 37]

Coibion, O., and Y. Gorodnichenko (2012): "What Can Survey Forecasts Tell Us about Information Rigidities?," *Journal of Political Economy*, 120(1), 116–159. [35]

——— (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *The American Economic Review*, 105(8), 2644–2678. [29]

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. [38]

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, 37]

Del Negro, M., M. P. Giannoni, and C. Patterson (2023): "The Forward Guidance Puzzle," *Journal of Political Economy Macroeconomics*, 1(1), 43–79. [8, 37]

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, 37]

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. [8, 43]

Ehrmann, M., and M. Fratzscher (2007): "The Timing of Central Bank Communication," *European Journal of Political Economy*, 23(1), 124–145. [30, 40]

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. [31]

Gabaix, X. (2020): "A Behavioral New Keynesian Model," *The American Economic Review*, 110(8), 2271–2327. [8, 37, 38]

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, 37]

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. [2, 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. [15]

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. [15]

Hansen, S., and M. McMahon (2016a): "First Impressions Matter: Signalling as a Source of Policy Dynamics," *The Review of Economic Studies*, 83(4 (297)), 1645–1672. [2]

——— (2016b): "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. [2, 6]

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]

Husted, L., J. Rogers, and B. Sun (2020): "Monetary Policy Uncertainty," *Journal of Monetary Economics*, 115, 20–36. [30]

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. [3, 7, 20, 21]

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. [4, 8, 29, 34]

Lucas, R. E. (1972): "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4(2), 103–124. [30]

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. [2, 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. [35]

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

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, 37]

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. [37]

Morris, S., and H. S. Shin (2002): "Social Value of Public Information," *The American Economic Review*, 92(5), 1521–1534. [30]

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. [3, 5, 7, 21]

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. [15]

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. [32]

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. [31]

Taylor, J. B. (1993): "Discretion Versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214. [1, 7, 14, 32, 46]

——— (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, 37]

Woodford, M. (2003a): *Imperfect Common Knowledge and the Effects of Monetary Policy* pp. 25–58. Princeton University Press. [4, 8, 29]

——— (2003b): *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton University Press, 1 edn. [3, 7, 30]

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. [16, – xiv –]