

# Do Central Banks Take Climate Change Seriously?

## Abstract

Central banks now speak openly about climate change. Whether climate enters the official record through which monetary-policy decisions are explained is another matter. This paper compares more than 37,000 speeches from 136 institutions with nearly 3,600 official meeting records from 26 central banks across 14 languages. Within the common 26-institution sample observed in both channels, benchmark comparisons show that climate is not uniquely excluded but filtered. Half of the central banks publishing meeting records never mention it in that record, and among those that do the median lag from first climate speech to first recorded mention is 18 years. When climate does appear there, it does not appear in its public form; it is framed chiefly through inflation, financial stability, and policy instruments. The paper argues that climate enters monetary-policy records through mandate translation, whereby a broad public issue becomes admissible only after it is recast in the language of the central-bank mandate.

**Keywords:** climate change, central bank communication, text as data

## Introduction

In September 2015, Mark Carney delivered what became the most consequential climate statement in the history of central banking. Speaking at Lloyd's of London as Governor of the Bank of England and Chairman of the Financial Stability Board, he warned that climate risks would crystallise beyond the horizons that govern markets, electoral cycles, and monetary policy itself, a failure he termed the “tragedy of the horizon” (Carney 2015). The speech catalysed a decade of institutional change. Research confirmed that climate-related financial risks fall within central bank mandates (Campiglio et al. 2018, Bolton et al. 2020), the Network for Greening the Financial System was founded in 2017 and now counts over 130 members, and climate risk entered the working vocabulary of financial regulators worldwide.

But speeches are not where rates are set. Central banks communicate through two distinct channels (Blinder et al. 2008). In speeches and testimony, governors signal priorities, shape expectations, and position the institution in wider public debate. After each policy meeting, committees issue an official record of their deliberation in the form of minutes, accounts, or comparable summaries. Speeches reveal what central banks want the world to hear. Meeting records reveal what institutions are willing to place on the formal record when setting rates. This distinction allows a direct test of whether the climate engagement that Carney's speech set in motion actually reached the official record through which monetary-policy decisions are explained. The paper uses that official record as the specific test of the broader question in the title, while recognising that central banks also engage climate through supervision, stress testing, collateral frameworks, and research outside the corpus studied here. The Bank of England provides the clearest answer. I obtained the Bank's complete Monetary Policy Committee transcripts for 2015, 2016, and 2017, detailed meeting records released with an eight-year delay. Across 29 meetings totalling more than 400,000 words of deliberation, climate discussion is entirely absent. Not in the month of Carney's speech. Not in the year that followed. Not in the year after that. Over the same three years, Brexit appears 449 times across 22 meetings; long-horizon structural issues were not excluded as a category. The Financial Policy Committee, which Carney also chaired, mentioned climate only once during this period, in March 2015, when it judged the risks likely to lie beyond its usual policy horizon. The topic that Carney placed at the centre of global financial-stability discourse had not entered the recorded discussion of his institution's rate-setting committee. Climate was not being debated in a different room. It had not entered any room. I call this the *Carney Puzzle*. The speech that turned climate into a central concern of public central banking came from an institution in which climate had not entered the recorded rate-setting conversation itself.

Was the Bank of England exceptional, or did the most prominent climate speech in central banking expose something more general? The question matters because central banks are not bystanders to the climate transition. Through prudential supervision, collateral policy, asset pur-

chases, and their interpretation of financial stability, they shape how climate risk is priced and how transition costs move through the financial system (Schoenmaker 2021). The climate consequences of economic governance are now subject to growing scrutiny more broadly. IMF programmes have been associated with higher deforestation, and World Bank climate finance has expanded mainly through projects where climate is bundled with other objectives rather than ring-fenced in dedicated lending (Forster et al. 2026, Kaya and Leblebicioglu 2025). Existing work on central banks has traced the rise of public climate communication and shown that it is shaped by peer effects, institutional participation, and domestic politics as much as by direct exposure to climate risk (Arseneau et al. 2022, Campiglio et al. 2025, Shears et al. 2025). Yet whether climate has crossed from public speech into the official record of monetary-policy meetings, and how it changes form when it does, has not been tested.

I address that question by comparing a multilingual corpus of more than 37,000 central-bank speeches from over 130 institutions with nearly 3,600 official meeting records from 26 central banks across 14 languages. I also benchmark climate against three other non-core topics, inequality, digitalisation, and geopolitical risk, to distinguish climate-specific patterns from general properties of committee communication. The evidence shows that climate does not move from one channel to the other unchanged. It appears less often in meeting records than in speeches, many central banks never mention it in their published records, and where it does appear it usually does so only after a long delay. More importantly, climate enters meeting records in a narrower language than the one used in public speech. It becomes legible chiefly when linked to inflation, financial stability, or concrete policy instruments.

This is the paper's central argument. Climate enters the monetary-policy record through *mandate translation*, a process through which a broad public issue becomes admissible only after being rewritten in the narrower categories that committees already recognise as legitimate. Read in this way, the Carney case is not simply a striking anecdote. It is the clearest expression of a broader tension between public climate discourse and the official record of monetary policy.

## Results

### **The Bank of England is not exceptional**

The Bank of England is not an outlier. Supplementary Table 1 shows that 13 of the 26 central banks publishing meeting records have never mentioned climate in their published monetary-policy records, despite contributing 1,570 documents. These are not generally silent committees. Twelve of the 13 mention geopolitical risk at least once in the same records. The Bank of England contributes 286 zero-climate records, the Central Bank of Brazil 244, and the Central Bank of Hungary 238.

Where climate does appear, it is concentrated rather than diffuse. The European Central Bank accounts for 21 of the 88 meeting records containing verified climate content, the People’s Bank of China for 13, and the Bank of Japan for 12. Together those three institutions account for 52.3% of all verified climate records. The Carney case is not an anecdotal contradiction between one speech and one central bank. It is the most visible instance of a wider pattern.

That wider pattern is best understood as filtering, not outright exclusion. This comparison is not cross-sample: Figure 1 and Table 1 restrict both channels to the same 26 institutions. In that common-institution sample, 5.5% of speeches mention climate compared with 2.4% of meeting records, for a speech-to-record ratio of 2.25. Table 2 reports the corresponding word-share contrast. Climate accounts for 0.565% of speech text on average but only 0.079% of meeting-record text, a factor of 7.32. Yet climate is not the topic most strongly kept out of the record. Geopolitical risk passes into committee language readily, appearing in 31.9% of meeting records and 19.5% of speeches. Digitalisation remains largely outside the rate-setting record, with a speech-to-record ratio of 11.36. Inequality, under the strict definition used in the audit, sits above climate, with a ratio of 3.77. Climate therefore occupies an intermediate position. It is not treated as unsayable. It is treated as a topic that must clear a higher threshold before appearing on the monetary-policy record. Supplementary Table 4 shows that the main frequency gap and the translated linkage profile survive leave-one-out exclusions of the ECB, the People’s Bank of China, and the Bank of Japan, including their joint removal. Inflation linkage is the most stable component of the translated profile; operational-tool linkage is present but more concentrated in a smaller number of institutions.

### **Where climate appears, it arrives late**

Even among institutions that eventually mention climate on the record, entry is slow. Figure 2 shows the institution-level lag profile. Among the 12 institutions that eventually mention climate in both channels, the mean delay from first climate speech to first climate record mention is 14.3 years. The Federal Reserve shows the longest observed lag at 29 years, followed by the ECB at 21 and the Reserve Bank of India at 20. These are not short lags associated with differences in drafting convention. They imply sustained periods in which climate is publicly discussed but has not yet entered the official record of monetary policy. Restricting the lag calculation to years in which both archives are already active lowers the mean observed lag to 10.4 years, but the qualitative conclusion remains one of slow entry rather than quick convergence (Supplementary Table 10).

Once right-censored non-entrants are incorporated, the delay looks longer still. In a survival sample of 24 institutions, the Kaplan–Meier median time from first climate speech to first climate record mention is 18.0 years (Supplementary Figure 1). Survival remains 75.3% after ten years and 56.5% after fifteen. A Cox model yields a hazard ratio of 1.30 on first speech year ( $p = 0.012$ ).

Later public-adoption cohorts cross from speech to record somewhat faster than earlier ones. The threshold for entry appears to have fallen in the post-Paris and post-NGFS period, but it has not disappeared. The Cox estimate should nonetheless be read cautiously: it rests on 24 institutions with 12 observed entrants.

The Bank of England transcripts, discussed in the introduction, confirm that the absence from published records is not merely editorial. At that institution, climate is virtually absent from the underlying discussion itself (Supplementary Table 5). The issue had not been relocated to another committee room. The Bank of England materials weaken the most direct version of the “wrong room” objection and support the use of published meeting records as a reliable indicator of what committees are willing to place on the monetary-policy record. Record forms nonetheless vary across institutions. Supplementary Table 6 classifies the 26 series by committee, record form, archive coverage, and document length, and Supplementary Table 7 shows that the speech-record gap persists in both longer and shorter record series.

### **What enters the record is translated climate**

The central result concerns not only how often climate appears but how it appears once admitted. Figure 3 shows that climate changes form when it crosses from speeches into meeting records. Inflation linkage rises from 6.3% of climate speech excerpts to 33.1% of climate record excerpts (Supplementary Table 2). Operational-tool linkage rises from 4.5% to 25.7%. Unlinked climate language, the broad framing that dominates public speech, falls from 72.2% to 33.8%. The linkage category most closely associated with public climate discourse, financial stability, also declines, from 19.9% in speeches to 8.1% in meeting records. Committees do not merely mention climate less often. They discuss it in different terms. Supplementary Table 9 gives paired examples of broad unlinked speech language and the translated meeting-record language that replaces it.

This pattern is more specific than editorial selection. Editing could remove diffuse or low-value material from published records. But editorial selection alone does not explain why the surviving climate material is so sharply concentrated in inflation-linked and tool-linked language rather than simply shorter. That content signature is difficult to explain unless climate is already being framed in those terms within the committee process itself. The published record may be selective, but what survives is not arbitrary. Climate becomes admissible chiefly when it can be expressed in the categories that structure monetary-policy deliberation.

Meeting records are also somewhat more action-oriented than speeches as a genre. Climate records show a higher prepare-and-implement share than climate speeches (30.4% versus 20.7%), and the mean action stage rises from 1.62 to 1.86 (Supplementary Figure 3). But a non-climate baseline shows that meeting records are generally more action-oriented than speeches. The action-

stage difference is therefore partly a genre effect. The linkage difference is not. The validation exercise confirms this direction. Manual coding identifies more operational-tool linkage and less unlinked content in meeting records than the automated labels do. The distinctive result is mandate translation.

### **Weak within-institution co-movement between speeches and the monetary-policy record**

If public climate speech were a leading indicator of later record content, institutions that speak more about climate should later place more climate on the monetary-policy record. Figure 4 shows the institution-year fixed-effects estimates (Supplementary Table 3). The coefficient on same-topic speech word share is 0.152 for climate ( $p = 0.299$ ),  $-0.018$  for inequality ( $p = 0.413$ ),  $-0.005$  for digitalisation ( $p = 0.480$ ), and 0.302 for geopolitical risk ( $p = 0.108$ ). None is estimated precisely at conventional levels. The results do not show mechanical within-institution co-movement between public speech and the monetary-policy record. But these regressions are thinly powered, given 26 institutions and near-zero climate word shares in most years, so they are better read as descriptive background than as decisive standalone evidence. The broader case for separate agendas rests more heavily on the Carney Puzzle, the Lagarde shift, and the benchmark comparisons.

The ECB provides the clearest within-institution illustration that the threshold for admission can shift. Of the 95 Governing Council accounts in the corpus, 39 fall in the Draghi period from 2015 to 2019 and 56 in the Lagarde period from 2020 to 2024. None of the Draghi-era accounts contains a verified climate reference. Twenty-one of the Lagarde-era accounts do; climate appears in 37.5% of published accounts under Lagarde. The same institution under the same treaty framework moved from complete absence to routine presence. This is not a clean leadership experiment, since the shift coincides with the European Green Deal, post-Paris learning, and wider institutional change as well as leadership transition. But it shows that climate's entry into the monetary-policy record is not fixed by mandate text alone.

Across institutions, the Carney Puzzle is not an isolated contradiction between one speech and one central bank. It is the visible edge of a broader process in which climate reaches the monetary-policy record selectively, belatedly, and in translated mandate language.

## **Discussion**

The question posed by the Carney Puzzle now has a clear answer. The Bank of England was not an isolated contradiction between one famous speech and one silent committee. It was the sharpest expression of a broader pattern. Across central banks, climate became publicly salient well before

it entered the monetary-policy record. The contribution of the paper is therefore not simply to document a gap between speeches and meeting records. It is to show that the gap has structure. Climate is filtered, delayed, and translated before it becomes admissible on the monetary-policy record.

That finding changes how climate engagement by central banks should be judged. Public speeches show that governors are willing to talk about climate, signal awareness, and position their institutions in wider debate. Meeting records do something narrower and more consequential: they place on the official record the considerations used to justify decisions. The weak within-institution relation between speech salience and record salience shows that the two channels are not close substitutes. An institution can speak frequently about climate without making it part of the recorded reasoning of monetary policy. For that reason, speech counts are a poor proxy for operational integration. The harder test is whether climate enters the official record through which policy decisions are explained.

The benchmark comparison clarifies why entry is uneven. Geopolitical risk passes into meeting records readily because it maps at once onto inflation, trade disruption, commodity prices, and activity. Digitalisation largely remains elsewhere because much of the relevant institutional work sits in payments, technology, supervision, or CBDC programmes rather than in the rate-setting meeting. Climate sits between those poles. It is too prominent to remain outside central-bank speech, but not sufficiently proximate to the monetary-policy mandate to enter the record in its public form. When climate does cross the threshold, it does so chiefly through inflation linkage and operational tools. That is the substantive content of mandate translation.

How should this pattern be interpreted? One reading treats it as appropriate functional differentiation. Speeches serve signalling and coordination functions; meeting records summarise decisions. On that view, the gap reflects sensible division of labour rather than institutional reluctance. There is force to that argument, but the evidence in this paper narrows its scope. The Bank of England's treatment of Brexit shows that long-horizon structural issues are not excluded as a class. The MPC discussed Brexit 449 times over the same three years in which climate was absent. The issue is therefore not horizon length alone. What governs admission is whether a topic is judged relevant to the committee's operative categories and worth placing on the monetary-policy record.

A second reading emphasises structural constraints. Central-bank mandates rarely mention climate explicitly, legal frameworks create genuine uncertainty about operational engagement, and climate-risk analysis raises real methodological difficulties (Cochrane 2020, Cullen 2023). Those constraints plainly matter. Yet they are not sufficient to explain the observed variation. The ECB operates under treaty constraints at least as binding as those facing the Bank of England, and the People's Bank of China operates under political conditions that would not obviously favour climate discussion. Nevertheless, both institutions place climate on the monetary-policy record. The ECB's

shift under Lagarde is especially instructive. The legal framework did not change, yet the threshold for admission did. The most plausible reading is therefore not that structure is irrelevant, but that leaders and institutions interpret structure differently. Climate integration is shaped by choice within constraints, not by mandate text alone (Mishra and Reshef 2019, Shears et al. 2025).

In that sense, organisational theory remains useful, but the mechanism here is more specific than a generic account of divergence between formal commitments and internal practice (Meyer and Rowan 1977, Bromley and Powell 2012). Cho et al. (2015) describe a similar logic in sustainability reporting, where institutions respond to conflicting pressures through selective disclosure rather than straightforward consistency. Central banking exhibits an analogous process. Climate is not simply kept out, nor is it incorporated wholesale. It is admitted once it can be rewritten in the language of price stability, financial stability, and instruments.

That distinction matters for the relation to decoupling. Decoupling, as described by Meyer and Rowan (1977) and Bromley and Powell (2012), predicts a gap between public rhetoric and internal practice that typically appears as silence, symbolic compliance, or diluted implementation. The evidence here shows something more specific. Climate that survives into meeting records is not merely rarer or shorter than climate in speeches; it is restructured around inflation, financial stability, and operational tools. Silence remains part of the pattern, but the distinctive signature is transformation rather than simple absence. On that point, the data fit mandate translation more closely than a generic account of decoupling.

Two limits follow from the design. Outside the Bank of England companion archive, the analysis observes authorised meeting records rather than verbatim transcripts; wider transcript release with suitable lags would allow future work to separate silence in deliberation from silence in publication more directly. The 26-institution sample is restricted to central banks that publish usable post-meeting records, though the ECB means these institutions cover considerably more than 26 countries, and the core speech-record frequency gap is measured within that publishing sample.

The paper does not take a position on whether central banks should widen their climate role. That question remains contested (Campiglio et al. 2018, Schoenmaker 2021, Cochrane 2020, Lastra 2024). The contribution here is empirical. Whatever one thinks about the proper scope of central-bank action, the relevant question is not whether central bankers mention climate in public. It is whether climate has entered the official record of monetary policy, and on what terms.

What follows from that finding is not a mandate prescription but an information-cost argument. When admission depends on mandate translation, broad transition-risk framing, long-horizon physical-risk discussion, and cross-sectoral climate linkages are less likely to appear on the official monetary-policy record unless they can be tied to inflation, financial stability, or a concrete instrument. That filtering matters in a climate-affected economy because energy shocks, food-price

shocks, transition policies, and physical disruption increasingly feed into inflation, output, and balance-sheet risk. A record that admits climate mainly after that reframing may still contain relevant climate material, but it is liable to omit some of the wider information through which climate first becomes economically salient. Those omissions matter for environmental governance as well as for textual interpretation: they leave a thinner public trace of the information that can shape transition-risk pricing, define the scope of climate stress testing, and support wider green-finance frameworks. For that reason, public climate speech is a weak accountability metric. The official record is the harder and more consequential test because it documents the considerations institutions are willing to place behind monetary-policy decisions.

On that measure, integration remains partial. The Carney Puzzle therefore names more than a striking episode at the Bank of England. It captures a broader feature of contemporary central banking: climate became central to public discourse before it became part of the recorded language of rate-setting.

## Methods

The empirical design compares two channels of central-bank communication, public speeches and official post-meeting records, across a multilingual corpus spanning four decades. Climate content is identified in both channels, benchmarked against three other non-core topics, and coded for the stage of institutional action and the mandate category through which climate is expressed. The cross-country analysis is based on published meeting records. The Bank of England transcript archive is used separately as a companion check on the relation between the meeting itself and the public record.

### Corpora and translation

The speeches corpus combines the BIS Central Bank Speeches archive, directly collected material from official central-bank websites, and the harmonised speech collection assembled by [Campiglio et al. \(2025\)](#). It contains 37,814 speeches from 136 institutions between 1986 and 2025.

The meeting-record corpus is an original collection of 3,602 official post-meeting records from 26 central banks between 1995 and 2024. These records include minutes, accounts, and comparable summaries issued after monetary-policy meetings. Nine undated Polish files with no substantive translated text were excluded; the usable corpus contains 3,593 records. Supplementary Table 6 lists the committee names, record forms, archive coverage, and median document lengths for the 26 series.

The meeting-record corpus is multilingual. Only 44% of documents originate in English. The translated material spans Spanish (404 documents), Portuguese (244), Hungarian (238), Czech (213), Turkish (203), Hebrew (195), Polish (194), Thai (104), Chinese (61), Romanian (61), Urdu (58), Ukrainian (51), and Russian (5). This coverage matters because an English-only corpus would miss more than half of the published monetary-policy record, and the linguistic diversity guards against the Anglophone bias that pervades existing research on central bank communication.

Non-English records were translated into English before retrieval and coding. The Microsoft Azure Translator API was used for Spanish, Portuguese, Hungarian, Czech, Turkish, Polish, Thai, Chinese, Romanian, Ukrainian, and Russian. Hebrew and Urdu files were handled through a separate Gemini-based workflow because those materials contained substantial optical-character-recognition error. This approach is consistent with evidence that contemporary neural machine translation performs well across a wide range of languages (Costa-jussà et al. 2022). Native speakers with economics training validated random samples from both pathways. Across 150 validated documents, agreement exceeded 95% for climate-specific terminology and 92% for overall meaning.

The Bank of England case also uses a companion archive of 29 MPC transcripts and 15 FPC records from 2015 to 2017. These materials are not included in the cross-country denominator. They motivate the Carney Puzzle in the introduction, and they provide the only available check on whether absence from published meeting records corresponds to absence from the underlying committee discussion.

## **Climate identification**

Climate content is identified in two stages. First, each document is scanned for 26 climate-related terms spanning direct climate references (climate change, global warming, greenhouse gas), carbon and emissions language (carbon neutral, net zero, decarbonisation), transition-risk terms (transition risk, physical risk, stranded asset, climate stress test), policy frameworks (Paris Agreement, TCFD, NGFS), and green-finance and energy-transition vocabulary (green bond, green loan, sustainable finance, renewable energy, fossil fuel). Each match triggers extraction of a 200-word context window. This step identifies 4,857 candidate speech documents and 500 candidate meeting records, with 14,700 and 623 candidate excerpts respectively.

Second, each candidate excerpt is semantically verified using OpenAI's GPT-5.2 API to determine whether the keyword appears in a genuine climate-policy context or in an unrelated discussion such as commodity pricing. This approach builds on recent advances in applying large language models to economic text analysis (Silva et al. 2025). The same vocabulary carries different meanings across channels. In speeches, references to coal, oil, or fossil fuels often concern transition risk, decarbonisation, or energy policy. In meeting records, the same terms usually refer to current

inflation, energy costs, or external price shocks. Indeed, 62% of keyword matches in meeting records derived from just two terms, coal and oil and gas, almost entirely in commodity price contexts. Without semantic verification, climate would be materially overstated on the monetary-policy record.

The design nonetheless prioritises precision over recall. Because retrieval still begins with explicit climate vocabulary, it can miss economically climate-relevant discussion that drops climate markers entirely. The inclusion of terms such as transition risk, physical risk, stranded asset, and climate stress test reduces that problem, but it does not eliminate it. Any remaining false negatives are more likely to bias against finding climate on the monetary-policy record than to inflate it.

After verification, 3,266 speech documents and 88 meeting records remain, with 12,582 and 148 verified excerpts respectively. The difference in false-positive rates makes the scale of this problem concrete. False positives account for 32.8% of keyword-matched speech documents but 82.4% of keyword-matched meeting records. Overlapping keyword windows within a document are merged before verification and trimmed to roughly 220 words where necessary. Document-level mention rates are calculated from verified-document counts. Climate word shares are calculated by merging verified excerpt intervals back onto the source documents and dividing the resulting deduplicated climate text by total document word count.

## **Benchmark topics**

The paper's argument depends on showing that climate is filtered rather than simply absent. That requires comparison. I benchmark climate against three other topics that are not traditionally considered core monetary-policy concerns, namely inequality, digitalisation (including fintech and CBDC), and geopolitical risk.

The three topics were chosen because they span the range of how readily non-core issues enter committee language. Geopolitical risk maps directly into inflation, trade disruption, and commodity prices, so it should pass into meeting records with minimal filtering. Digitalisation sits largely outside the monetary-policy record because the relevant institutional work is concentrated in payments, technology, and supervision. If climate behaves like geopolitical risk, the filtering story is wrong. If climate behaves like digitalisation, the story is trivially true. The informative case is if climate sits between them, which it does.

For each topic, candidate excerpts are retrieved from the harmonised corpora using topic-specific keyword lists and are then semantically verified using the same procedure as for climate. Inequality is coded strictly. I count discussion of inequality, poverty, income distribution, or wealth gaps as policy topics in their own right, but exclude routine passages in which inflation or interest-rate changes are merely said to affect different groups differently. This strict definition keeps

the benchmark comparable to the climate measure, which also requires substantive rather than incidental discussion. Targeted manual audits for inequality, digitalisation, and geopolitical risk are reported in Supplementary Table 11. Unlike climate, those benchmarks were not subjected to the full two-coder content validation exercise, which should be kept in mind when interpreting the comparison. The inequality audit shows higher ambiguity in meeting records than the other benchmarks, so the climate-versus-geopolitical-risk and climate-versus-digitalisation contrasts are more firmly grounded than the exact ordering between climate and inequality.

## Content coding

Verified climate excerpts are coded along two dimensions. The first is action stage, which captures how far the institution moves beyond acknowledgment. The four stages are *acknowledge* (generic recognition of climate-related risks), *assess* (analysis, monitoring, or risk evaluation), *prepare* (plans, framework development, or preparatory institutional work), and *implement* (concrete operational use of tools or policies). For the summary statistics reported in the Results, these stages are also mapped to an ordered four-point scale.

The second dimension is mandate linkage, the coding that matters most for the mandate-translation argument. Each verified excerpt is classified for explicit connection to *inflation*, *financial stability*, and *operational tools*. Inflation linkage requires an explicit connection to inflation, price stability, CPI, supply shocks, exchange-rate pass-through, or related transmission channels. Financial-stability linkage requires an explicit connection to balance-sheet risk, asset valuation, credit conditions, prudential exposure, or systemic vulnerability. Operational-tool linkage requires an explicit connection to instruments or operational frameworks such as collateral policy, asset purchases, refinancing operations, supervision, disclosure requirements, or stress testing. This category therefore spans both rate-setting instruments and broader central-bank tools. *Unlinked* is coded when none of the three linkages is explicit. Multiple linkages can be assigned to the same excerpt, so linkage shares reported in the Results do not sum to 100.

The coding is intentionally literal. Borderline cases are coded as unlinked, which makes the mandate-translation contrast conservative.

To interpret whether climate meeting records are unusually action-oriented, the analysis also codes a stratified non-climate sample of meeting-record excerpts. This baseline separates the genre effect (meeting records are generally more action-oriented than speeches) from the climate-specific finding (climate is reframed through mandate language when it enters the record).

## Verification, validation, and reproducibility

Manual verification of roughly 15,000 candidate climate excerpts across 14 languages was not feasible without computational assistance, so semantic verification and content coding were conducted using OpenAI's GPT-5.2 with structured prompts in separate reproducible passes (Silva et al. 2025). One pass distinguished genuine climate discussion from incidental keyword occurrence. A second pass assigned action-stage and mandate-linkage labels. The exact prompts, model parameters, and implementation details are reported in the replication materials. An independent audit re-verified 100 climate candidate excerpts drawn equally from accepted and rejected keyword matches and agreed with the original accept-reject decision in 97% of cases; Supplementary Table 8 reports the corresponding confusion summary and channel-specific agreement rates.

The coding scheme was validated on a stratified sample of 100 verified climate excerpts, split evenly across speeches and meeting records, and labelled independently by two coders with economics training. For action stage, Cohen's kappa between human coders is 0.826, with 87 of 100 labels matching exactly. Agreement is higher in meeting records (92%) than in speeches (82%), and disagreements cluster at adjacent categories rather than distant ones. Model-human agreement averaged 0.85, within the range established as acceptable in the methodological literature.

For mandate linkage, kappas are 0.929 for inflation linkage, 0.962 for financial-stability linkage, 0.818 for operational-tool linkage, and 0.807 for unlinked status. Mean agreement across the four binary linkage indicators is 95%. Human coding assigns slightly more operational-tool linkage and slightly less unlinked content in meeting records than the automated procedure. This directional bias means that the full-corpus mandate-translation contrast is conservative. The true shift from unlinked to linked language when climate enters meeting records is, if anything, larger than what the automated labels show.

## Statistical approach

The analysis reports document-level mention rates, document-level word shares, institution-level paired comparisons across topics, timing of first entry, and within-institution speech-record associations over time. Mention rates are defined as the share of documents containing at least one verified excerpt. Word shares are defined as the share of total document text contained in verified excerpts. The paired  $p$ -values reported in the benchmark table come from two-sided paired  $t$ -tests on institution-level comparisons within the common sample of 26 institutions observed in both channels.

For the timing analysis, first speech and first meeting-record years are measured from the verified corpora. First observed mention is not necessarily true first mention, so these lag estimates

are potentially left-censored. The observed entrant-lag summary uses the 12 institutions that eventually mention climate in both channels. Supplementary Table 10 reports a sensitivity in which the lag clock starts only once both archives are active. The survival analysis expands this to a 24-institution sample, treating institutions that never mention climate in meeting records by the end of the sample as right-censored. One institution with verified climate in meeting records has no verified climate speech in the harmonised speech corpus, and one institution's meeting-record sample ends before its first verified climate speech, so those two are excluded from the speech-to-record survival sample. The Cox model uses first speech year as a cohort covariate and should be interpreted cautiously given the small number of observed entrants.

For the institution-year panel, the unit of observation is institution-year for the 26 institutions observed in both channels. The regressions use meeting-record word share as the dependent variable and same-topic speech word share as the regressor of interest, with institution and year fixed effects and governor-turnover controls from the KOF central-bank-governor database. These regressions are descriptive rather than causal. Their purpose is to test whether public communication and the monetary-policy record move together within institutions over time, not to provide a high-power estimate of that relation. With only 26 institutions and many near-zero institution-years, they should be read as a complement to the qualitative and comparative evidence rather than as a decisive standalone test.

# Figures

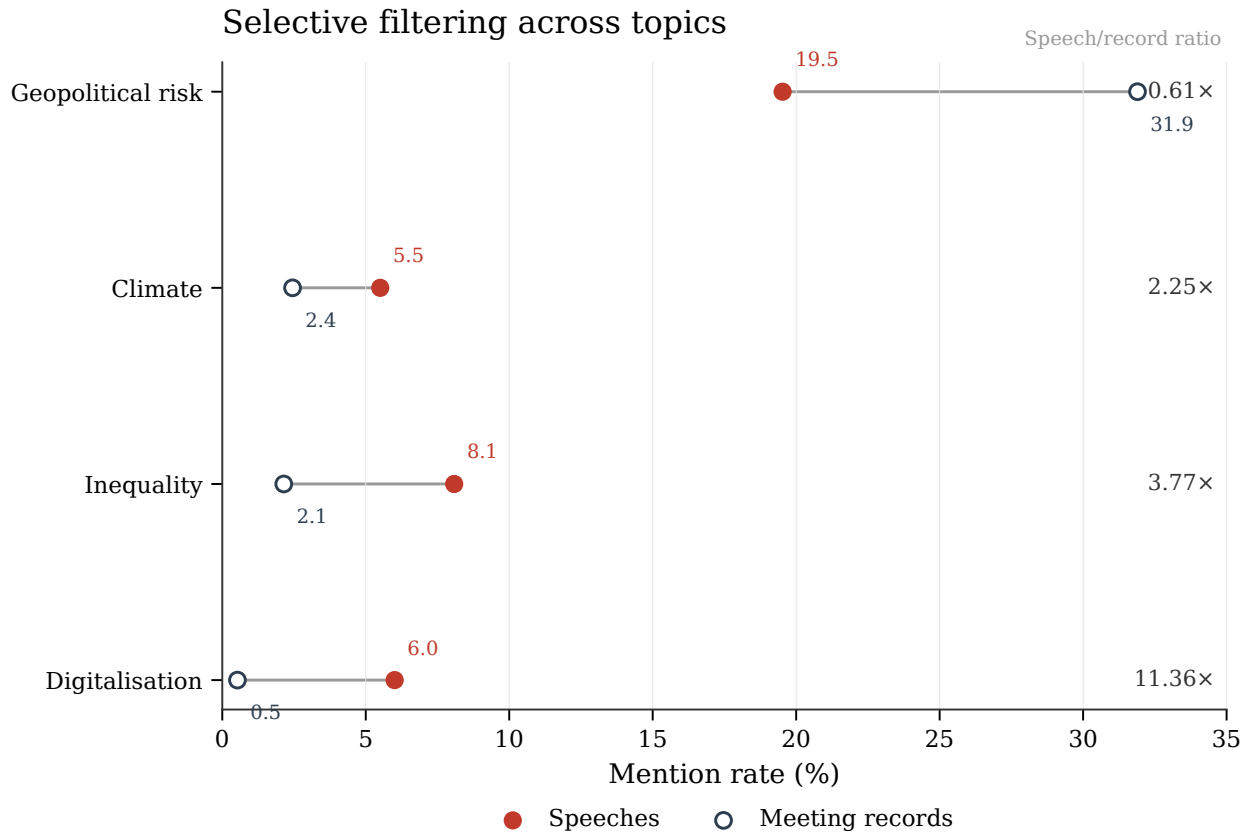
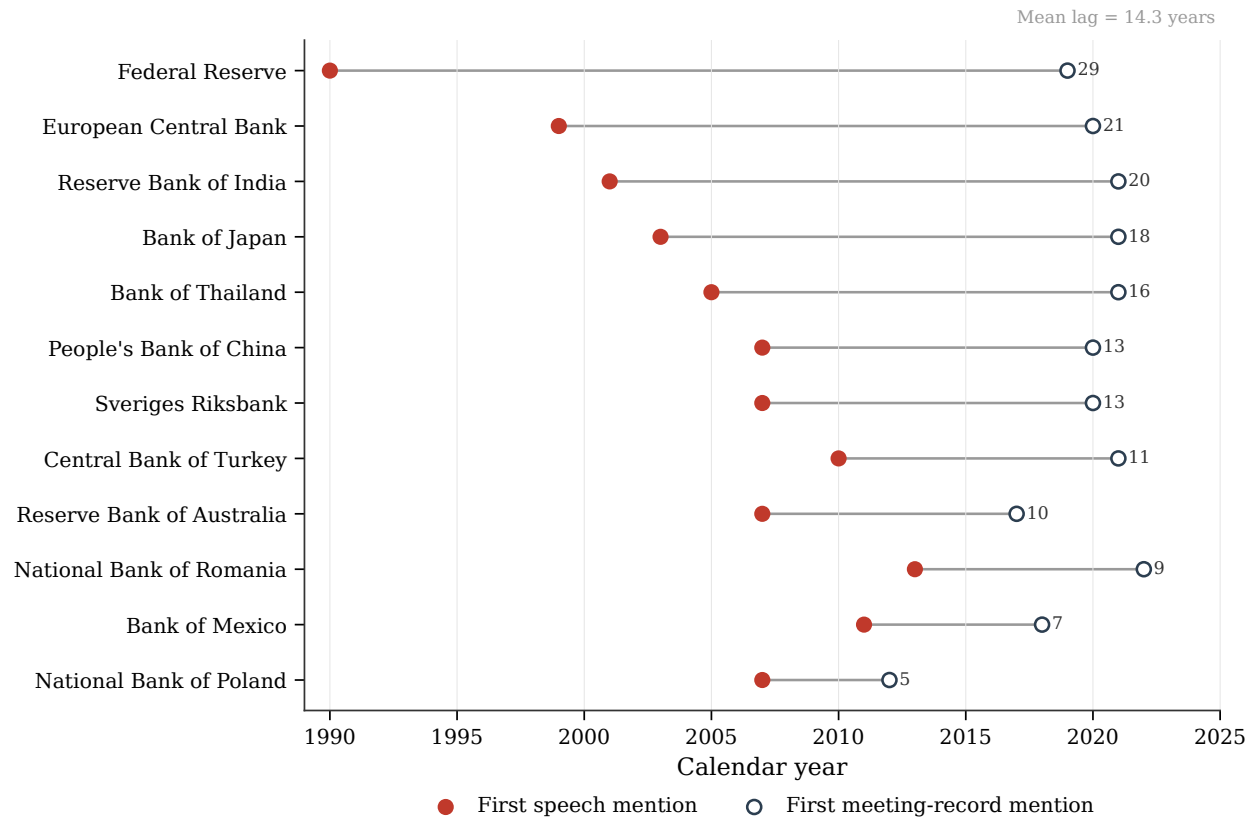
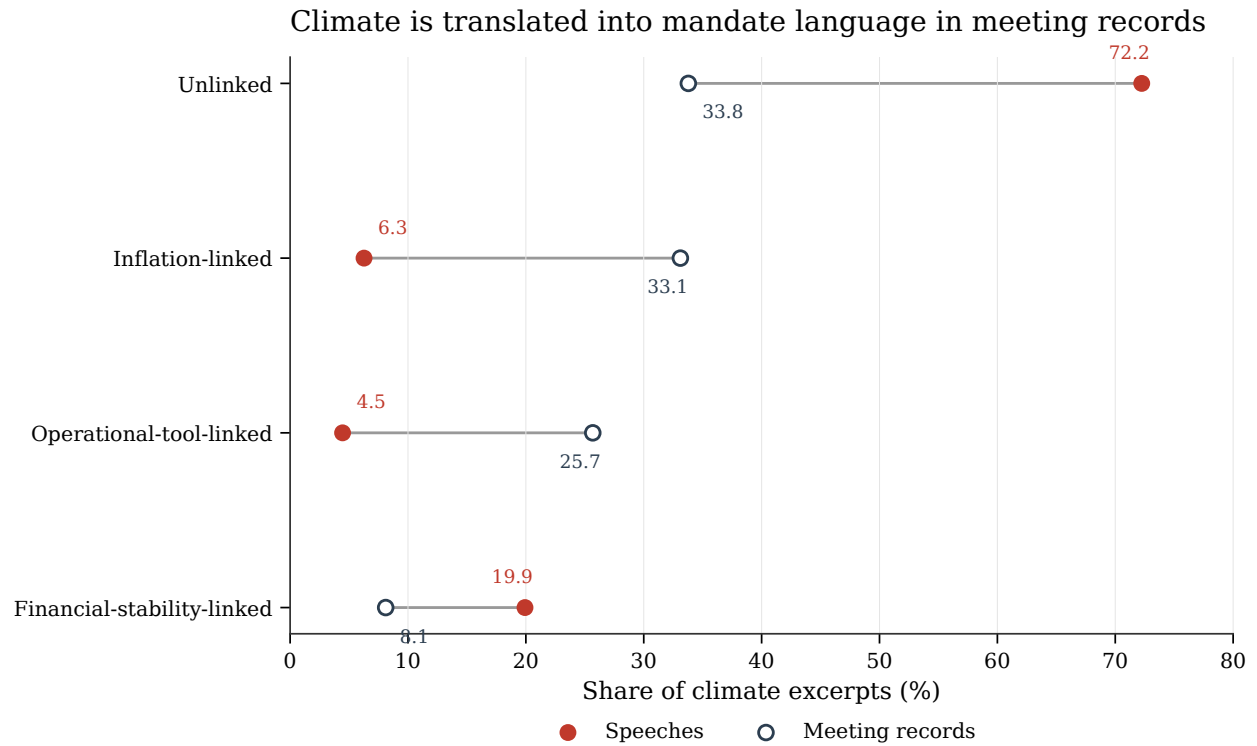


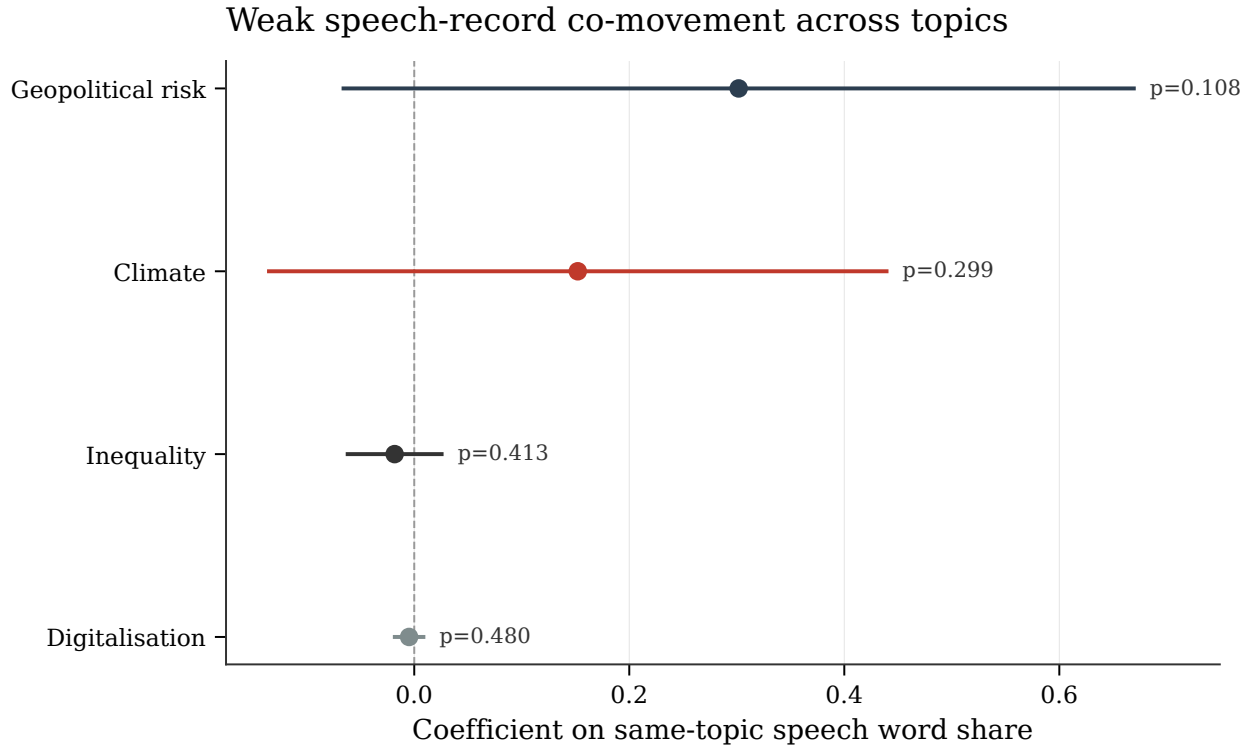
FIGURE 1. **Selective filtering across benchmark topics.** Common-institution speech and meeting-record mention rates are shown for geopolitical risk, climate, inequality under the strict definition, and digitalisation. The value at the right of each row is the speech-to-record ratio.



**FIGURE 2. Long delays separate first climate speech from first climate mention on the monetary-policy record.** Each row shows the first verified climate mention in speeches and the first verified climate mention in meeting records for one of the 12 institutions that eventually mention climate in both channels. The dashed vertical line marks the mean observed lag of 14.3 years.



**FIGURE 3. Climate enters the monetary-policy record through mandate translation.** Speech and meeting-record shares are shown for each linkage category. Inflation and operational-tool linkage rise sharply when climate enters meeting records, while unlinked climate language falls.



**FIGURE 4. Weak within-institution co-movement between speeches and the monetary-policy record.** Points and horizontal lines show coefficients and 95% confidence intervals from institution-year fixed-effects regressions of meeting-record word share on same-topic speech word share. The estimated slopes are imprecise for all four topics; public speech salience does not mechanically map into the published monetary-policy record.

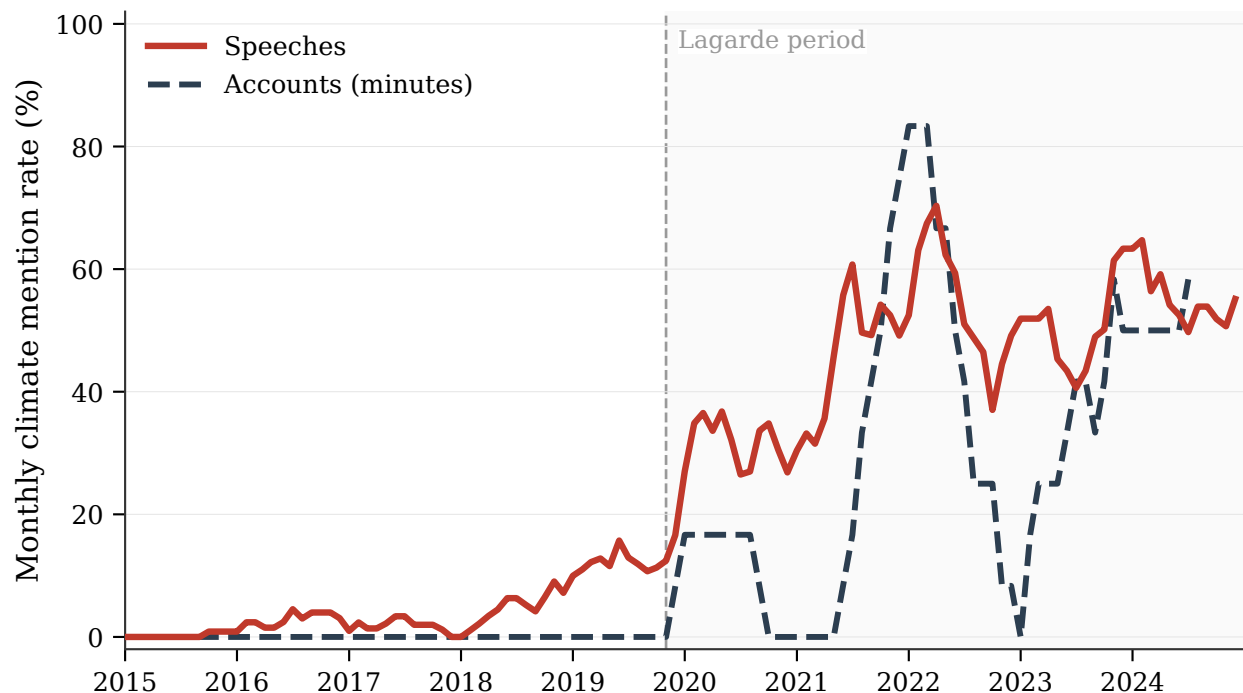


FIGURE 5. **Climate enters ECB Governing Council accounts under Lagarde.** Monthly climate mention rates in ECB speeches and Governing Council accounts (minutes) from 2015 to 2024 are shown on a common 0–100% scale. Both lines plot six-month rolling averages of monthly mention rates. The vertical marker denotes 1 November 2019, the beginning of the Lagarde period. Governing Council accounts contain no verified climate references under Draghi; under Lagarde, 37.5% do.

## Tables

TABLE 1. **Benchmark comparison across topics.**

Topic	Speech rate	Record rate	Speech/record ratio	Paired $p$ -value
Geopolitical risk	19.5%	31.9%	0.61	< 0.001
Climate	5.5%	2.4%	2.25	0.001
Inequality	8.1%	2.1%	3.77	< 0.001
Digitalisation	6.0%	0.5%	11.36	< 0.001

*Note:* Rates are computed in the common-institution sample. Paired  $p$ -values come from two-sided paired  $t$ -tests across institutions. Inequality uses the strict verification definition described in the Methods.

**TABLE 2. Sample construction and verification.**

	Speeches	Meeting records
Total collected records	37,814	3,602
Usable records	37,814	3,593
Institutions	136	26
Period	1986–2025	1995–2024
Languages	Predominantly English	14
Climate candidate documents	4,857	500
Climate verified documents	3,266	88
Climate verified excerpts	12,582	148
False positive rate	32.8%	82.4%
Common-sample climate mention rate	5.5%	2.4%
Common-sample climate word share	0.565%	0.079%

*Note:* The common sample comprises the 26 institutions observed in both channels. Nine undated Polish translation stubs are excluded from usable meeting records.

## References

- David M. Arseneau, Alejandro Drexler, and Mitsuhiro Osada. Central bank communication about climate change. Finance and Economics Discussion Series 2022-031, Board of Governors of the Federal Reserve System, May 2022. URL <https://www.federalreserve.gov/econres/feds/central-bank-communication-about-climate-change.htm>.
- Alan S. Blinder, Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen. Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature*, 46(4):910–945, December 2008. ISSN 0022-0515. doi: 10.1257/jel.46.4.910.
- Patrick Bolton, Morgan Després, Luiz Awazu Pereira da Silva, Frédéric Samama, and Romain Svartzman. *The Green Swan: Central Banking and Financial Stability in the Age of Climate Change*. Bank for International Settlements, Basel, January 2020. ISBN 978-92-9259-326-1. BIS Books No. 31.
- Patricia Bromley and Walter W. Powell. From smoke and mirrors to walking the talk: Decoupling in the contemporary world. *Academy of Management Annals*, 6(1):483–530, 2012.
- Emanuele Campiglio, Yannis Dafermos, Pierre Monnin, Josh Ryan-Collins, Guido Schotten, and Misa Tanaka. Climate change challenges for central banks and financial regulators. *Nature Climate Change*, 8(6): 462–468, 2018.
- Emanuele Campiglio, Jérôme Deyris, Davide Romelli, and Ginevra Scalisi. Warning words in a warming world: Central bank communication and climate change. *European Economic Review*, 178, 2025.
- Mark Carney. Breaking the tragedy of the horizon – climate change and financial stability. Speech at Lloyd’s of London, September 2015. URL <https://www.bankofengland.co.uk/speech/2015/breaking-the-tragedy-of-the-horizon-climate-change-and-financial-stability>. Governor of the Bank of England and Chairman of the Financial Stability Board.
- Charles H. Cho, Matias Laine, Robin W. Roberts, and Marine Rodrigue. Organized hypocrisy, organizational façades, and sustainability reporting. *Accounting, Organizations and Society*, 40:78–94, January 2015. doi: 10.1016/j.aos.2014.12.003.
- John H. Cochrane. Central banks and climate: A case of mission creep. *Defining Ideas*, November 2020.
- Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Baez, Gabriel Baines, Geza Barabber, Mikel Bober-Irizar, et al. No language left behind: Scaling human-centered machine translation. *Nature*, 2022. doi: 10.1038/s41586-022-05189-z. NLLB Team.
- Jay Cullen. Central banks and climate change: Mission impossible? *Journal of Financial Regulation*, 9(2): 174–209, October 2023. ISSN 2053-4833. doi: 10.1093/jfr/fjad003.
- Timon Forster, Rishikesh Ram Bhandary, and Kevin P Gallagher. Cutting trees to balance budgets: International monetary fund programs are associated with increased deforestation. *One Earth*, 2026.
- Ayse Kaya and Aslı Leblebicioglu. Understanding the world bank’s role in climate finance. *Climatic Change*, 178(9):160, 2025.
- Rosa María Lastra. Mission creep, democratic accountability and central bank independence. *SUERF Policy Note*, 2024.
- John W. Meyer and Brian Rowan. Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology*, 83(2):340–363, 1977.
- Prachi Mishra and Ariell Reshef. How do central bank governors matter? regulation and the financial sector. *Journal of Money, Credit and Banking*, 51(2-3):369–402, 2019. doi: 10.1111/jmcb.12577.
- Dirk Schoenmaker. Greening monetary policy. *Climate Policy*, 21(4):581–592, 2021. doi: 10.1080/14693062.2020.1868392.
- Esther Shears, Jonas Meckling, and Jared J. Finnegan. How central banks manage climate and energy transition risks. *Nature Energy*, 10(4):470–478, 2025. doi: 10.1038/s41560-025-01724-w.
- Thiago Christiano Silva, Kei Moriya, and Romain Michel Veyrune. From text to quantified insights: A large-scale LLM analysis of central bank communication. IMF Working Paper WP/25/109, International Monetary Fund, 2025.

## **Data and Code Availability**

Replication materials accompany this submission as a compressed archive and include the verified climate corpus, topic dictionaries, prompt documentation, classification code, generated tables and figures, the institution-year panel used in the regressions, and the companion Bank of England archive used for the motivating case.

## **AI Use Disclosure**

Generative AI tools were used only in the analytic workflow for semantic verification of candidate excerpts and for action-stage and mandate-linkage coding. Prompts, model versions, and parameters are archived in the replication materials, and AI-generated classifications were audited against human coding before submission. No generative AI was used to draft manuscript text or to create or alter figures. All numerical claims, tables, figures, and manuscript text were checked against the project data and scripts by the author before submission.

## **Author Contributions**

Fatih Kansoy designed the study, assembled the data, conducted the analysis, and wrote the manuscript.

## **Competing Interests**

The author declares no competing interests.