

# Central Bank Communication as Text Data

Measurement, public attention, strategic disclosure, and  
understandability

---

Macro-Finance · Lecture 4 · Trinity Term 2026

Fatih Kansoy

University of Oxford · Saïd Business School · Department of Economics

# The lecture follows communication from release to reception

---

1. **Measurement:** make the missing middle observable.
2. **Architecture:** ask how bundling changes attention.
3. **Unexpected content:** isolate the first public reading.
4. **Public engagement:** distinguish posting from exposure.
5. **Strategic disclosure:** compare staged records.
6. **Climate:** trace mandate translation.
7. **Understandability:** measure the language constraint.

Communication becomes data only after the economic object, source, and empirical unit are fixed.

## Part I

# Communication as measurement

---

Text and digital traces reveal the missing middle

# Standard data observe the policy input and the final response

---

$$\underbrace{i_t, \pi_t, y_t, r_t, q_t}_{\text{standard macro-finance data}} \implies \underbrace{\text{beliefs, expectations, prices}}_{\text{final outcomes}}.$$

Monetary policy is transmitted through information, but policy rates, inflation, output, yields, asset prices, and survey averages usually reveal only the endpoints.

**The empirical gap lies between the released decision and the observed outcome.**



# The hidden mechanism is a sequence of human responses

---


communication  $\longrightarrow$  attention  $\longrightarrow$  interpretation  $\longrightarrow$  belief updating.

- A message cannot update a belief if it is not noticed.
- Exposure does not guarantee a shared interpretation.
- Text and digital traces make parts of this information environment observable.

Communication data measure the mechanism that endpoint data leave hidden.

# The hidden middle is attention and interpretation

---



The broken path marks  
what standard endpoint data do not directly observe.

Editorial marker illustration generated for this lecture.

# Communication enters the information set beside macro and market data

---

$$\mathcal{I}_t = \mathcal{I}_t^{macro} \cup \mathcal{I}_t^{market} \cup \mathcal{I}_t^{communication} \cup \mathcal{I}_t^{public}.$$

$$\mathcal{I}_t^{macro}, \quad \mathcal{I}_t^{market}.$$

$$\mathcal{I}_t^{communication}, \quad \mathcal{I}_t^{public}.$$

Conventional datasets are strongest on the first pair. Text and digital traces supply partial measurements of the second pair.

**A policy message matters only after it enters an audience's information set.**

# Text data extend measurement beyond policy and prices

---

| Object                      | Standard data | Text or digital data |
|-----------------------------|---------------|----------------------|
| Policy decision             | yes           | yes                  |
| Public attention            | rarely        | yes                  |
| Public interpretation       | rarely        | yes                  |
| Disclosure choices          | rarely        | yes                  |
| Institutional admissibility | rarely        | yes                  |
| Asset-price response        | yes           | yes                  |

The added value is access to reception, selection, and admissibility—not another policy-rate series.

# The useful corpus is the one matched to the behaviour

---

| Source type              | Examples                                                   | Economic object                          |
|--------------------------|------------------------------------------------------------|------------------------------------------|
| Policy and official text | statements, minutes, speeches, transcripts, fiscal reports | priorities, guidance, disclosure         |
| Firm disclosures         | earnings calls, annual reports, filings, presentations     | expectations, risk, investment plans     |
| News and media           | newspapers, newswires, analyst reports, blogs              | information flow, sentiment, uncertainty |
| Social and public text   | Twitter/X, Reddit, survey comments, reviews                | interpretation, disagreement, beliefs    |
| Attention traces         | Google Trends, Wikipedia views, web traffic                | search and reference behaviour           |

**Availability is not validity: source choice follows the economic object.**

# A measurement rule turns a raw source into a testable variable

---

$$\text{raw source } s_i \xrightarrow{\hat{\phi}} \hat{x}_i \xrightarrow{\text{econometric design}} Y_i.$$

$s_i$ : source item ·  $\hat{\phi}$ : measurement rule ·  $\hat{x}_i$ : constructed variable ·  $Y_i$ : outcome.

The measurement rule is part of the research design. It determines what the raw material can become before the regression begins.

**Unstructured material becomes evidence only through an explicit measurement rule.**

# Different sources expose different stages of transmission

---

| Source               | Measured object                         |
|----------------------|-----------------------------------------|
| Search or page views | attention and information acquisition   |
| Official posts       | engagement with public communication    |
| Public text          | interpretation, beliefs, disagreement   |
| Institutional text   | disclosure, omission, policy priorities |

**No single source measures attention, interpretation, and disclosure at once.**

# One document can support several economic measurements

---

| Measurement | Variable                 | Question                                  |
|-------------|--------------------------|-------------------------------------------|
| Tone        | hawkish/dovish score     | how restrictive is the communication?     |
| Topic       | inflation, labour, risk  | what is the discussion about?             |
| Surprise    | residualised tone        | what was not predictable before release?  |
| Coverage    | statement–transcript gap | what was omitted from the focal release?  |
| Linkage     | mandate category         | how does an issue become policy-relevant? |

The same minutes become a different variable when the research question changes.



# The document is not the variable

---

document  $\neq$  variable.

The economic object—not the algorithm—determines whether the document is measured as tone, topic, surprise, coverage, or institutional linkage.

**Start with the quantity to be measured; only then select the text method.**

# The empirical workflow makes measurement auditable

---

source  $\longrightarrow$  cleaned corpus  $\longrightarrow$  measurement rule  
 $\longrightarrow$  structured variable  $\longrightarrow$  econometric design.

Each arrow creates a decision that must be documented: inclusion, cleaning, classification, aggregation, and identification.

**Measurement is not preprocessing; it is an explicit empirical stage.**

# Each workflow stage has a distinct failure mode

---

| Stage       | What happens                                                           |
|-------------|------------------------------------------------------------------------|
| Source      | collect posts, speeches, minutes, transcripts, searches, or views      |
| Cleaning    | remove retweets, bots, false matches, duplicates, and missing metadata |
| Measurement | classify tone, topic, interpretation, linkage, or omitted content      |
| Aggregation | move to event, meeting, institution-year, or release-window units      |
| Validation  | compare human labels, test placebos, and propagate uncertainty         |

**A credible variable records how raw material changed at every stage.**

# Three questions organise the empirical cases

---

| Research question                | Observed source                                 | Economic object                    |
|----------------------------------|-------------------------------------------------|------------------------------------|
| Who pays attention?              | searches, page views, public and official posts | attention and engagement           |
| How is the message interpreted?  | public text around announcements                | direction, disagreement, narrative |
| What becomes part of the record? | statements, minutes, transcripts, speeches      | disclosure and admissibility       |

communication → measured variable → economic mechanism.

The lecture moves from reception to interpretation and then to institutional selection.

## Part II

# Release architecture changes attention

---

One richer event can still create fewer opportunities to notice

# The Warsh reform replaced a sequence with a bundle

---

## Before August 2015

A quarterly forecast round arrived across separate dates:

- policy decision;
- Inflation Report and press conference;
- minutes and vote.

Attention could arrive in several smaller waves.

## After August 2015

Full-release meetings bundle the decision, minutes, vote, forecast report, and press conference. Standard meetings publish the decision, minutes, and vote only. Standard meetings form the post-reform comparison group.

**The reform changed release architecture, not the existence of policy information.**

# Bundling moves attention on two opposing margins

---

## Intensive margin

- one richer event;
- decision, minutes, forecast, and conference together;
- higher salience on the bundled day.


## Extensive margin

- fewer separate occasions;
- one forecast-round opportunity disappears;
- cumulative engagement may fall.

A larger event can coexist with less attention over the whole forecast round.

# One larger release can mean fewer attention occasions

---



The focal burst is the intensive margin; the sparse occasion dots preserve the opposing extensive margin.

Editorial marker illustration generated for this lecture.



# Net quarterly attention is a gain minus a lost occasion

---

$$\Delta A_q \approx \underbrace{\text{salience gain on bundled day}}_{\text{intensive}} - \underbrace{\text{attention from removed occasion}}_{\text{extensive}}.$$

The empirical question is which margin dominates in revealed public attention once event size and event count change together.

**Communication architecture determines both the size and number of attention triggers.**

## Three traces capture three ways of noticing policy

| Trace         | Series                                           | Behaviour measured                |
|---------------|--------------------------------------------------|-----------------------------------|
| Twitter/X     | log unique users posting original BoE/MPC tweets | real-time public discussion       |
| Google Trends | UK index for BoE, interest rates, and inflation  | active information acquisition    |
| Wikipedia     | log views for BoE and monetary-policy pages      | reference and background checking |

No platform is a population-wide attention measure; agreement across traces is therefore more informative than any one series.

**The outcome is revealed attention: discussion, search, and reference behaviour.**

# Abnormal attention compares an event day with normal activity

---

$$\underbrace{A_{e,p}}_{\text{abnormal attention}} = \underbrace{Y_{t(e),p}}_{\text{event-day activity}} - \underbrace{\bar{Y}_{e,p}^{base}}_{\text{matched baseline}}.$$

$e$ : MPC event ·  $p$ : platform ·  $Y$ : transformed platform activity.

$A_{e,p} > 0$  means the event attracted more attention than the platform's matched normal level.

The design explains an event spike, not raw platform volume.

# The baseline matches weekdays and excludes nearby MPC events

| Object                          | Construction                                                           |
|---------------------------------|------------------------------------------------------------------------|
| Event activity $Y_{t(e),p}$     | log users, search index, or log page views on the event date           |
| Baseline $\bar{Y}_{e,p}^{base}$ | same-platform mean on matching non-event weekdays in the prior 30 days |
| Exclusion                       | remove days near other MPC events from the baseline                    |
| Outcome                         | event-specific abnormal attention $A_{e,p}$                            |

**Matching the attention environment separates an announcement spike from ordinary weekly rhythms.**

## Event and forecast-round analyses answer different questions

---

| Unit           | Pre-reform                             | Post-reform                          |
|----------------|----------------------------------------|--------------------------------------|
| Event          | decision day or Inflation Report day   | full-release day or standard meeting |
| Forecast round | decision day plus Inflation Report day | one full-release day                 |

---

- Event level: is the bundled day more salient?
- Forecast round: does one bundled day replace attention from two occasions?

**The unit of analysis determines whether the intensive or extensive margin is visible.**

# Cumulative attention adds event spikes within a forecast round


---

$$CumLogAtt_q = \sum_{e \in \mathcal{E}_q} A_{e,TW}.$$

The event-level measure is first constructed for each communication occasion and then aggregated to the quarterly forecast round.

Digital traces let the researcher count attention by event before aggregating it.

# Twitter attention peaks on the MPC decision day




Across  $h = -3, \dots, +3$ , event-day log unique users spike and attention decays sharply after the decision.

Kansoy (2026), source figure page 4.

# Google searches also concentrate on the event day

---




The event-day distribution shifts upward relative to adjacent days, consistent with active information acquisition around MPC announcements.

Kansoy (2026), Google Trends event-time figure.




# Wikipedia views reveal event-day reference behaviour



Reference and background-checking activity rises on the MPC event day and recedes afterward.

Kansoy (2026), Wikipedia event-time figure.


# Quarterly Twitter attention shifts left after bundling



Pre:  $N = 18$ , mean 2.966, s.d. 0.631. Post:  $N = 25$ , mean 1.742, s.d. 0.466; the post mean is below the pre minimum 1.945.

Kansoy (2026), source figure page 6.


# The reform is associated with a level shift, not a visible trend



Separate OLS trends and 95% bands show no statistically significant within-regime trend; the pattern is consistent with a break at implementation.

Kansoy (2026), source figure page 7.

# High-surprise events are more visible, not longer-lived



High shock means absolute UKMPD surprise above the sample median. Across 117 MPC events, the larger impact response still decays rapidly.

Kansoy (2026), source figure page 9.

# Bundling raises event salience but lowers round-level attention

| Object         | Finding                                              | Interpretation                      |
|----------------|------------------------------------------------------|-------------------------------------|
| Single event   | full-release days attract more discussion and search | the remaining event is more salient |
| Forecast round | cumulative Twitter attention falls about 41%         | the lost second occasion dominates  |

more attention per event  $\nRightarrow$  more attention over the round.

**Communication design changes both salience and the number of opportunities to notice.**

# Super Thursday creates an attention trade-off

---

one larger public occasion    versus    two separate attention triggers.

## Event salience

Co-located decision, minutes, forecast, and conference make the remaining full-release day easier to notice.

## Round exposure

Removing the separate forecast occasion gives the public fewer chances to encounter the forecast round.

A bigger event-day spike does not replace repeated exposure.

# Architecture changes both topic composition and timing

---

## What attracts attention

- more forecast and outlook discussion;
- no clear shift to rate-only discussion;
- search and reference traces respond to the richer package.

## When attention arrives

- a sharp announcement-day response;
- rapid decay after the event;
- surprise raises visibility, not persistence.

communication architecture = what is bundled + how often attention is triggered.

# Attention is not belief or understanding

---

attention trace  $\neq$  beliefs  $\neq$  understanding.

---

| Limit               | Interpretation                                        |
|---------------------|-------------------------------------------------------|
| Selected users      | Twitter/X is a trace, not the whole public            |
| Normalised searches | Google Trends measures relative search intensity      |
| Bundled reform      | timing, minutes, and event frequency changed together |

---

**The credible claim is narrower: architecture changes exposure opportunities.**



## Part III

# Unexpected language changes the public reading

---

Interpretation can move without approval

# The empirical target is the first directional reading

---

BoE communication bundle → attentive-public interpretation  
→ possible wider diffusion.

| Object      | Design choice                                               |
|-------------|-------------------------------------------------------------|
| Setting     | scheduled post-2015 MPC bundles                             |
| Public text | BoE-related Twitter/X discourse around release              |
| Window      | post-release discourse relative to the pre-release baseline |
| Target      | directional interpretation, not volume or generic sentiment |

The outcome is the first public reading—not representative household beliefs.

## The classifier targets policy direction, not generic sentiment

---

$$\hat{L}_i^{direction} \in \{\text{hawkish, dovish, neither}\}.$$

Each BoE-related tweet receives a directional label. Approval is classified separately so that interpretation and evaluation are not conflated.

A positive or negative emotional tone is not automatically hawkish or dovish.

# The Public Hawkishness Index nets hawkish against dovish readings

---

$$\text{PHI}_{m,w} = \frac{N_{m,w}^{\text{hawk}} - N_{m,w}^{\text{dove}}}{N_{m,w}^{\text{all}}}.$$

$m$ : MPC meeting ·  $w$ : time window · denominator: all BoE-related tweets.

Positive PHI means that discourse reads the Bank as more hawkish. Non-directional tweets remain in the denominator.

The index measures the direction of the public reading, not its volume.

# The outcome is a post-release change from the pre-release baseline

---

$$\Delta \text{PHI}_m = \text{PHI}_{m,0-24h} - \text{PHI}_{m,pre48h} .$$

The comparison holds the meeting fixed and asks whether the distribution of directional readings shifts after the communication bundle.

**A meeting-level change is more informative than the raw post-release index.**

# Unexpected tone removes the part predictable before release

---

$$\text{Tone}_m^{BoE} = \alpha + \Pi' I_m + \theta_{qoy} + u_m, \quad \text{UnexpectedTone}_m = \hat{u}_m.$$

| Term                      | Meaning                                                                    |
|---------------------------|----------------------------------------------------------------------------|
| $\text{Tone}_m^{BoE}$     | hawk–dove tone of the minutes                                              |
| $I_m$                     | pre-release OIS expectations, rate surprise, macro data, gilt and FX moves |
| $\text{UnexpectedTone}_m$ | tone not predicted by that information set                                 |

Residualised tone is the text analogue of a policy surprise.

# The transmission regression separates language from the rate surprise

---

$$\Delta \text{PHI}_m = \beta_0 + \beta_1 \text{UnexpectedTone}_m + \beta_2 \text{RateSurprise}_m + \beta' W_m + \theta_{qoy} + \varepsilon_m.$$

$\beta_1$ : response to unexpected BoE tone ·  $\text{RateSurprise}_m$ : decision surprise ·  $W_m$ : pre-release market controls.

The coefficient is reduced-form: unexpected content may combine language with private central-bank information.

# A hawkish reading can be supportive or critical

---

$$Approval_{m,w} = \frac{N_{m,w}^{support} - N_{m,w}^{critical}}{N_{m,w}^{all}}.$$

## Interpretation

“The Bank sounds hawkish.” The direction can be hawkish, dovish, or neither.

## Evaluation


“That stance is right” or “that stance is a mistake.” Evaluation can be supportive, critical, or neither.

Persuasion requires approval to move, not only the directional reading.



# A directional reading does not determine approval

---



The same directional interpretation can lead to supportive or critical evaluation.

Editorial marker illustration generated for this lecture.

# The interpretation response is concentrated in the first six hours

$$\Delta \text{PHI}_m^{a-b} = \text{PHI}_{m,a-b} - \text{PHI}_{m,pre48h} .$$

| Window | Interpretation                                        |
|--------|-------------------------------------------------------|
| 0–6h   | immediate reading before a full news cycle            |
| 6–24h  | same-day reading after media and social circulation   |
| 24–48h | whether the interpretation survives into the next day |

The evidence is an attention burst, not a persistent shift in discussion.

## Topic shares reveal which part of the bundle travels

---

$$Share_{m,w}^k = \frac{N_{m,w}^k}{N_{m,w}^{all}}, \quad \Delta Share_m^k = Share_{m,post}^k - Share_{m,pre48h}^k.$$

The four main topics are rates and MPC decisions, inflation and prices, the economic outlook, and financial stability.

Inflation is the clearest topic response: concrete price language travels more readily than technical policy language.


# Unexpected hawkish content changes interpretation, not approval

| Margin         | Finding                                                  |
|----------------|----------------------------------------------------------|
| Interpretation | $\Delta$ PHI rises over 0–24 hours                       |
| Timing         | strongest in the first six hours; fades by 48 hours      |
| Approval       | no comparable movement in support for the Bank           |
| Topics         | inflation discussion is the clearest supporting response |

interpretive channel operates      but      persuasive channel is not detected.

The design studies public reading—not representative beliefs or trust.


# High and low unexpected-tone events diverge at release



The distance opens sharply at the release and then gradually closes.

Kansoy (2026), source figure page 1.

# Bundled communication produces the clearest same-day response




The post-2015 bundled sample shows a clearer interpretation response immediately after release, with attenuation across later windows.

Kansoy (2026), source figure page 2.

# Inflation is the clearest topic carried into public discussion

---



The public does not absorb every component equally. Price and cost-of-living language is the most direct bridge to everyday experience.

Kansoy (2026), source figure page 3.

## Part IV

# Publication is not exposure


---

Official posts compete for voluntary attention



# Publication does not guarantee exposure

---



The official message  
must first be selected from a crowded information stream.

Editorial marker illustration generated for this lecture.

# The public first has to encounter the official message

---

conditional on the Bank posting, which messages attract interaction?

| Object     | Measurement                                                  |
|------------|--------------------------------------------------------------|
| Source     | complete official @bankofengland record, July 2011–July 2022 |
| Unit       | individual official tweet                                    |
| Outcome    | likes, retweets, replies, and quote tweets                   |
| Covariates | format, timing, reply/original mode, readability, length     |

**Exposure is a prior margin: a message must win attention before it can inform or persuade.**

# The model compares counts within similar attention environments

$$\mathbb{E}[Y_i \mid X_i] = \exp(X_i' \beta), \quad Y_i = Likes_i + Retweets_i + Replies_i + Quotes_i.$$

| Comparison           | Purpose                                          |
|----------------------|--------------------------------------------------|
| All tweets           | raw feature–engagement association               |
| Same broad period    | absorb follower growth and slow platform changes |
| Same week, day, hour | compare more similar attention environments      |

These are conditional associations. The Bank chooses format and timing; neither is randomly assigned.

# Large design margins predict more visible engagement

---

| Margin        | Finding                                                | Reading                                            |
|---------------|--------------------------------------------------------|----------------------------------------------------|
| Visual format | photo and video posts attract much more interaction    | presentation helps the message compete in the feed |
| Policy timing | MPC-day posts attract more engagement                  | policy is already salient                          |
| Macro state   | the MPC-day premium rises in the high-inflation period | public attention is state-dependent                |

**The largest engagement margins come from format, timing, and context—not tiny wording changes.**

# Reach, dialogue, and readability are different objectives

---

## Reply mode

Replies receive far less visible engagement than original posts. Dialogue and broadcast reach therefore measure different communication objectives.


## Readability

Easier text is positively associated with interaction, but the result is less robust than the format and timing margins.

Engagement is not comprehension, trust, approval, or expectation updating.

Visible interaction reveals a gate into communication—not the downstream success of the message.

# Official posting volume rises across policy and platform regimes




Monthly volume trends upward, with visible changes around governorship, the pandemic, and platform-resolution milestones.

Kansoy and Mundy (2025/2026), source figure page 1.

# Thursday carries the highest official posting volume


---



The weekday distribution is uneven: Thursday has the largest count, while weekend posting is sparse.

Kansoy and Mundy (2025/2026), source figure page 2.

# Original posts dominate, while reply activity grows over time




Original posts remain the principal mode; replies become a more visible part of the account's output in later years.

Kansoy and Mundy (2025/2026), source figure page 3.



# Visible engagement varies sharply by posting hour



The hour profile is highly uneven, consistent with a changing attention environment across the day.

Kansoy and Mundy (2025/2026), source figure page 4.

# Monthly engagement rates are volatile across interaction types




Likes, retweets, replies, and quotes do not move in lockstep; aggregate engagement can conceal which interaction margin changed.

Kansoy and Mundy (2025/2026), source figure page 5.

# Hashtag use is concentrated in a small vocabulary

---



#BoE is dominant, followed by a much smaller group of recurring campaign, institutional, and topical tags.

Kansoy and Mundy (2025/2026), source figure page 7.

## Public-facing language centres on the Bank, rates, and inflation



Bank, England, rate, UK, and inflation anchor the vocabulary, alongside people, work, money, policy, and economy.

Kansoy and Mundy (2025/2026), source figure page 8.

## Part V


# Strategic communication is staged disclosure

---

A statement is focal; minutes reveal what the focal text compressed

# Minutes restore part of what the focal statement compressed

---



Only selected excerpts travel into the later public record; much of the underlying discussion remains submerged.

Editorial marker illustration generated for this lecture.

# FOMC minutes are the second public layer of a meeting



$$\mathcal{I}_t^{public} = \{\text{statement}_t, \text{minutes}_t, \dots\}.$$

The statement announces the decision. The minutes later disclose the discussion, alternatives, risks, and views of voting and non-voting participants.

Later does not mean redundant: the information set can still expand.

# A meeting generates more than one public document




Read the release sequence as a disclosure architecture—not as repeated copies of one message.

ML4Z source asset min3.png.




# Minutes contain many more words than the focal statement



Length creates room for nuance, but also raises the cost of finding the policy-relevant sentence.

ML4Z source asset fomcminuteswords.png.


# The release lag changed after 2004



| Regime      | Statement         | Minutes                                                |
|-------------|-------------------|--------------------------------------------------------|
| Before 2005 | at meeting $X$    | 6–8 weeks later, near meeting $X + 1$                  |
| After 2004  | decision and vote | exactly three weeks later; voting and non-voting views |

## The April 2016 minutes changed the market's June reading

---



The May release produced a large market response: the dollar reached a nine-week high as traders brought forward the expected timing of a rate rise.

A delayed document can revise the expected path of future policy.

000000business/economy/federal-reserve-interest-rates-meeting-minutes.html



Janet L. Yelver, the Fed chairwoman. Officials have repeatedly said that they want to get out of the business of selling markets when rates will rise, but they also appear unable at times to resist the temptation. *Saul Loblitz/Agence France-Press — Getty Images*

By Benjamin Appelbaum

May 18, 2016




**WASHINGTON** — The Federal Reserve sent a sharp, simple message to financial markets on Wednesday: Pay attention. The Fed is thinking seriously about raising its benchmark interest rate at its next meeting, in June.

The unusually frank bulletin was delivered in the [official account](#) of the Fed's April meeting, which said explicitly that most officials thought "it likely would be appropriate" to raise rates in June if the economy shows clear signs of a rebound from a weak winter.


That message was sharply at odds with the expectations of

## 66 of 121

# Different outlets converged on the same June signal



The Guardian  
A June rise was possible.



Reuters  
A June move was firmly on the table.

Media compression amplified a common forward-policy reading.

# The June decision did not validate the press forecast

---

minutes signal  $\longrightarrow$  earlier expected tightening  $\nrightarrow$  June rate increase.


On 15 June, the Fed did not raise the target rate and reaffirmed the prevailing range.

---

| What the episode shows                                                    | What it does not show                                              |
|---------------------------------------------------------------------------|--------------------------------------------------------------------|
| lagged minutes can contain path-relevant information and change positions | that the May 2016 reaction was typical across all minutes releases |

---

# Measure volatility in ten-minute intervals around release



$$SD(r) = \sqrt{\frac{\sum_{t=1}^T (r_t - \bar{r})^2}{T - 1}}.$$

— Minutes day — Non-announcement day

Source panels compare the S&P 500, EUR/USD, and the ten-year Treasury yield.

The event-window question is whether volatility rises at the release time.

## Baseline volatility: non-announcement afternoons

| Asset         | Obs.   | Mean | SD  | Min   | Max  | Skew. | Kurt. |
|---------------|--------|------|-----|-------|------|-------|-------|
| EUR/USD (%)   | 63,708 | 0    | .04 | -.79  | .58  | .04   | .04   |
| S&P 500 (%)   | 98,371 | 0    | .14 | -2.42 | 3.66 | .14   | .14   |
| 10-year (bps) | 89,019 | 0    | .04 | -5.6  | 5.5  | .04   | .04   |
| 3-month (bps) | 41,974 | 0    | .71 | -35   | 30   | .71   | .71   |

Ten-minute observations between 13:00 and 16:00 on non-announcement days.

This is the counterfactual intraday variability.



## Minutes days show a modestly different volatility profile

| Asset         | Obs.  | Mean | SD  | Min   | Max  | Skew. | Kurt. |
|---------------|-------|------|-----|-------|------|-------|-------|
| EUR/USD (%)   | 1,919 | 0    | .06 | -.56  | .40  | .06   | .06   |
| S&P 500 (%)   | 1,938 | 0    | .15 | -1.29 | 1.22 | .13   | .13   |
| 10-year (bps) | 1,930 | 0    | .05 | -3.9  | 3.8  | .05   | .05   |
| 3-month (bps) | 1,917 | .01  | .48 | -3    | 6    | .48   | .48   |

The minutes sample covers announcement afternoons over the same 13:00–16:00 window.

Minutes create an event, but the unconditional SD comparison is not identification.

## Statement days carry the largest unconditional volatility

| Asset         | Obs.  | Mean | SD   | Min   | Max  | Skew. | Kurt. |
|---------------|-------|------|------|-------|------|-------|-------|
| EUR/USD (%)   | 2,256 | 0    | .11  | -1.08 | .76  | .11   | .11   |
| S&P 500 (%)   | 3,529 | 0    | .23  | -3.76 | 2.75 | .23   | .23   |
| 10-year (bps) | 3,278 | 0    | .10  | -9.8  | 22.0 | .10   | .10   |
| 3-month (bps) | 1,613 | .06  | 1.44 | -10   | 27   | 1.44  | 1.44  |

Statements are immediate and highly focal; minutes are delayed and longer.

**The lower minutes-day SD does not imply that minutes contain no path news.**

# Regress the ten-minute return on target and path surprises

---

$$\Delta r_{x,d,\tau} = \alpha + \beta_1 TS_{d,\tau} + \beta_2 PS_{d,\tau} + \varepsilon_{x,d,\tau}.$$

| Object                   | Meaning                                                                                       |
|--------------------------|-----------------------------------------------------------------------------------------------|
| $\Delta r_{x,d,\tau}$    | return of asset $x$ in the ten-minute interval at release time $\tau$ on announcement day $d$ |
| $TS_{d,\tau}$            | target surprise                                                                               |
| $PS_{d,\tau}$            | path surprise                                                                                 |
| $\varepsilon_{x,d,\tau}$ | residual return                                                                               |

## Minutes surprises load most consistently on the path factor

| Asset         | Target (SE)   | Path (SE)       | Constant (SE) | <i>N</i> | <i>R</i> <sup>2</sup> |
|---------------|---------------|-----------------|---------------|----------|-----------------------|
| S&P 500 (%)   | −1.77** (.81) | −.94*** (.22)   | −.01 (.03)    | 103      | .16                   |
| EUR (%)       | 1.22 (.88)    | 1.18*** (.36)   | −.04** (.02)  | 103      | .13                   |
| 10-year (bps) | 16.55 (20.74) | 17.79*** (4.23) | .84 (.54)     | 103      | .16                   |
| 3-month (bps) | 9.81 (8.39)   | 5.62 (3.26)     | .29 (.27)     | 103      | .06                   |

\*\*\* $p < .01$ , \*\* $p < .05$ , \* $p < .10$ . Sample: 2004–2017.

The repeated significant coefficient is the path surprise.


## Statements produce stronger immediate reactions

| Asset         | Target (SE)     | Path (SE)        | Constant (SE) | $N$ | $R^2$ |
|---------------|-----------------|------------------|---------------|-----|-------|
| S&P 500 (%)   | -3.18*** (.66)  | -2.82*** (.71)   | .05 (.04)     | 103 | .28   |
| EUR (%)       | .24 (.77)       | 3.63*** (.58)    | -.00 (.02)    | 103 | .35   |
| 10-year (bps) | 18.77 (31.30)   | 59.81*** (18.51) | .10 (.79)     | 103 | .15   |
| 3-month (bps) | 18.83*** (4.78) | 3.18 (3.72)      | -.05 (.37)    | 103 | .14   |

\*\*\* $p < .01$ , \*\* $p < .05$ , \* $p < .10$ . Sample: 2004–2017.


**Focality strengthens the first release; it does not eliminate later news.**

# At the ZLB, ten-year yield volatility spikes more on minutes days



The source comparison describes roughly  $2\times$  versus  $5\times$  release-day volatility.

# The ZLB made lagged dollar-path news unusually salient



QE  $\rightarrow$  larger Fed balance sheet  
 $\rightarrow$  greater USD-news sensitivity.

During the ZLB, the source describes a roughly fourfold balance-sheet expansion and a depreciating, globally influential dollar.

Context can amplify the value of information even when it is three weeks old.

## Outside the ZLB, minutes path news still moves markets

| Asset   | Target (SE)   | Path (SE)      | Constant (SE) | $N$ | $R^2$ |
|---------|---------------|----------------|---------------|-----|-------|
| S&P 500 | -1.51 (.90)   | -.95*** (.27)  | -.04 (.04)    | 46  | .22   |
| EUR     | .75 (.80)     | .78** (.34)    | -.05* (.02)   | 46  | .11   |
| 10-year | 28.46 (21.48) | 12.87** (4.89) | .10 (.81)     | 46  | .18   |
| 3-month | 13.05 (12.64) | 8.09 (4.85)    | .39 (.57)     | 46  | .10   |

Non-ZLB: 2004–2008 and 2016–2017. Returns are percent for equities and FX; basis points for yields.

**Path coefficients matter even outside the constrained-rate regime.**



## During the ZLB, the path response shifts toward EUR and ten-year yields


| Asset   | Target (SE)    | Path (SE)       | Constant (SE) | <i>N</i> | <i>R</i> <sup>2</sup> |
|---------|----------------|-----------------|---------------|----------|-----------------------|
| S&P 500 | −8.04** (3.88) | −.61 (.41)      | .02 (.04)     | 57       | .14                   |
| EUR     | 1.12 (3.39)    | 2.53*** (.80)   | −.02 (.03)    | 57       | .16                   |
| 10-year | 25.82 (79.93)  | 31.33*** (8.52) | 1.76** (.73)  | 57       | .23                   |
| 3-month | 29.62 (18.82)  | −1.16 (2.01)    | .11 (.17)     | 57       | .05                   |

ZLB period: 2008–2015.

A constrained current rate can make information about the future path more valuable.

# At the ZLB: less surprise, more reaction


---



The size of the measured surprise and the size of the market response need not move together.

ML4Z source asset `ps_surprises.png`.

# Minutes and statements differ in positive sentence share



Document tone is another dimension along which the delayed and focal releases can diverge.

Blue: minutes; green: statements; moving averages and recession shading are embedded in the source chart.

# The same meeting leaves three disclosure records

---

same meeting

{ statement: immediate, short, highly focal,  
minutes: delayed, longer, less focal,  
transcript: delayed archival benchmark.

| Layer      | Research use            | Communication margin                   |
|------------|-------------------------|----------------------------------------|
| Statement  | focal public signal     | selection into the short message       |
| Minutes    | later public disclosure | partial restoration of omitted content |
| Transcript | archival benchmark      | meeting-level proposition set          |

## Measure what the statement leaves out

---



$$g_t^k = 1 - \frac{N_t^{k,\text{covered}}}{N_t^{k,\text{all}}}.$$

$N_t^{k,\text{covered}}$ : transcript items in dimension  $k$  covered by the statement ·  $N_t^{k,\text{all}}$ : all transcript items in dimension  $k$ .

$g_t^k = 0$ : full statement coverage;  $g_t^k = 1$ : complete omission.

# Omission is measured separately by policy dimension

---

| Dimension $k$ | What the transcript propositions describe                 |
|---------------|-----------------------------------------------------------|
| OUTLOOK       | current conditions and the expected macroeconomic outlook |
| PATH          | expected future policy path and policy inclination        |
| RISK          | risks, uncertainty, and contingencies around the outlook  |
| DISAGREEMENT  | conflicting views, reservations, and alternatives         |

**A single document-level word count cannot reveal selective compression.**

## Statements omit disagreement most heavily

| Dimension    | Mean gap | Reading                |
|--------------|----------|------------------------|
| OUTLOOK      | .658     | most covered           |
| PATH         | .713     | selectively compressed |
| RISK         | .782     | frequently omitted     |
| DISAGREEMENT | .852     | most omitted           |

$$g^{DISAGREEMENT} > g^{RISK} > g^{PATH} > g^{OUTLOOK}.$$

## Later channels close more than half the residual gap

---

$$\bar{g}^{statement} = .745 \quad \longrightarrow \quad \bar{g}^{statement+minutes} = .334.$$

Minutes restore propositions that were absent from the focal statement; hearings and later channels can close the record further.

Selective compression is not permanent secrecy. It is staged disclosure across channels with different timing and focality.



# Omitted outlook and path content predicts minutes repricing

minutes repricing  $\leftarrow$  omitted outlook/path  $\times$  intermeeting activity news.

| Outcome                  | Source result                     |
|--------------------------|-----------------------------------|
| Treasury and real yields | respond to the interaction        |
| Equities and breakevens  | no corresponding robust response  |
| Content dimensions       | outlook and path survive controls |

The evidence is reduced-form: omitted content may correlate with unobserved meeting information.

# The pricing test interacts the outlook gap with later activity news


---

$$Y_t = \alpha + \beta(g_t^{OUT} \times z_t^y) + \gamma_1 g_t^{OUT} + \gamma_2 z_t^y + \varepsilon_t.$$

$Y_t$ : minutes-window asset-price response ·  $g_t^{OUT}$ : statement's omitted-outlook gap ·  $z_t^y$ : intermeeting activity surprise.

$\beta > 0$  means that activity news receives a larger market response when more meeting outlook content was left out of the statement.

# The communication record closes in stages



$P$  is the focal public statement; adding channels  $C$ ,  $H$ , and  $M$  progressively reduces the residual binary gap.

126 meetings, 2005–2020; source Figure 3.

# Policy tools differ in the language wrapped around them

---




Counting a tool mention misses whether it is framed as active, contingent, constrained, or merely discussed.

ML4Z source topic1.png.

# Tool frequency is the first layer, not the full message

---



Frequency locates attention; context determines the policy meaning of that attention.

ML4Z source topic2.png.

## Part VI


# Climate concern becomes policy only through the mandate

---

Public salience and rate-setting deliberation are different communication layers

# Climate risk reaches central banks through several channels

## Two channels, but many potential impacts.




Physical and transition shocks can affect prices, output, financial institutions, and the transmission mechanism.

ML4Z source asset risk1.

# Risk differs by horizon, probability, and severity

|                    | Climate Risks                                        | Economic Risks                                | Financial Risks                                                      | Financial Stability Risks                    |
|--------------------|------------------------------------------------------|-----------------------------------------------|----------------------------------------------------------------------|----------------------------------------------|
| <b>Real Estate</b> | Rising sea levels, frequency of storm surges         | Increased inundation of coastal parcels       | Decreased value of coastal real estate                               | Abrupt repricing of mortgage lending markets |
| <b>Insurance</b>   | More frequent and severe hurricanes, wildfires, etc. | Greater disruption to local economic activity | Pressure for higher rates, lower supply of insurance and reinsurance | Greater uninsured losses, spillover effects  |




A long horizon does not remove a risk from the central-bank problem; it changes how the risk is monitored and translated.

ML4Z source asset risk2.



# Climate shocks can become financial-stability shocks



Asset repricing, collateral losses, credit exposures, and insurance gaps can move the issue inside a conventional stability mandate.

ML4Z source asset risk3.

# Climate-related instruments mix support and penalty

---

## **Brown-penalising**

Raise the relative financing or portfolio cost of carbon-intensive activity.

Example: exclusion of coal firms from the Norwegian sovereign wealth fund's investment universe.

## **Green-supporting**

Improve funding, quotas, collateral treatment, or asset demand for lower-carbon activity.

Examples: Bank of Japan climate loans, Bangladesh Bank green credit policies, and ECB climate tilting.

**Instrument choice exposes the boundary between risk management and allocation.**

# The institutional puzzle is one mandate with many demands

## Too many jobs to do...

- Economic Growth / Full Employment
- Financial Stability
- Banking Supervision and Regulation
- Dealing with Pandemic
- Micro & Macroprudential Policies
- Inequality
- ....
- ....
- and
- Climate crisis



## Is it a central banker's job?

The answer cannot be read from the issue's importance alone. It depends on the price-stability, financial-stability, supervisory, and operational mandates of the institution.

Public concern is not yet mandate linkage.

# Carney made climate risk a central-bank signal in 2015

Bank of England Topics ▾ About ▾ News and publications ▾ Museum Contact Search the Bank of England 🔍

Published on 29 September 2015

Download speech >

Breaking the tragedy of the horizon – climate change and financial stability - speech by ... Copy link



Mark Carney  
Former Governor, Bank of England from 2013 to 2020 >

Convert this page to PDF

## Breaking the tragedy of the horizon

The speech connected climate change to financial stability and to the mismatch between long-horizon risks and shorter institutional decision horizons.

**The speech is a public signal; it is not itself a policy record.**

# Public salience can precede internal policy reasoning

2015 public speech  $\nrightarrow$  immediate climate language in the MPC record.

| Evidence                   | Count | Reading                                 |
|----------------------------|-------|-----------------------------------------|
| MPC transcripts, 2015–2017 | 29    | more than 400,000 words                 |
| Climate references         | 0     | absent from the sampled internal record |
| Brexit references          | 449   | present in 22 meetings                  |

An issue may be salient before it enters monetary-policy deliberation.

# The corpus spans public speech and official deliberation

| Dimension    | Speeches              | Policy records |
|--------------|-----------------------|----------------|
| Collected    | 37,814                | 3,602          |
| Usable       | 37,814                | 3,593          |
| Institutions | 136                   | 26             |
| Period       | 1986–2025             | 1995–2024      |
| Languages    | predominantly English | 14             |

Nine undated Polish translation stubs are excluded from the usable policy-record corpus.

The two channels differ in scale, purpose, and institutional coverage.

# Keyword candidates dwarf verified climate evidence

---

| Measure                  | Speeches | Policy records |
|--------------------------|----------|----------------|
| Climate candidates       | 4,857    | 500            |
| Verified documents       | 3,266    | 88             |
| Verified excerpts        | 12,582   | 148            |
| False-positive rate      | 32.8%    | 82.4%          |
| Common-sample mention    | 5.5%     | 2.4%           |
| Common-sample word share | .565%    | .079%          |

The common sample contains 26 institutions.

**Verification matters most in the policy record.**

# Define policy seriousness as movement across three layers

---

public issue  $\longrightarrow$  monetary-policy record  $\longrightarrow$  policy-relevant language.

## Speeches

37,814 documents from 136 institutions: a broad public agenda.

## Official records

3,593 records from 26 institutions: closer to deliberation and rate-setting policy.

Mention is an entry condition; mandate linkage is the substantive test.



# Climate keyword searches create severe false positives

---


$$\{\text{oil, gas, coal, energy}\} \neq \{\text{climate context}\}.$$

In policy records, these words often describe commodity prices, energy inflation, or near-term supply conditions rather than climate change.

82.4% of policy-record candidates are false positives after contextual verification.

A dictionary can find candidates; it cannot determine the policy meaning.

# Only context-verified mentions cross the mandate gate




One inspected excerpt proceeds; unverified keyword candidates remain outside the policy-evidence set.

Editorial marker illustration generated for this lecture.

## Verification turns a mention into coded mandate evidence

---



---

| Code       | What the excerpt links climate to                  |
|------------|----------------------------------------------------|
| Inflation  | prices, supply, or the policy outlook              |
| Stability  | financial risks and resilience                     |
| Tools      | collateral, assets, lending, or operational design |
| No linkage | climate is mentioned without a mandate bridge      |

---

Verified policy-record evidence: 88 documents and 148 excerpts.

# Climate enters the official record slowly and selectively

---

| Finding                         | Evidence                                                                                             |
|---------------------------------|------------------------------------------------------------------------------------------------------|
| Lower official-record incidence | 5.5% of speeches versus 2.4% of records                                                              |
| Institutional non-entry         | half of the 26 common institutions never mention climate in the policy record                        |
| Long translation delay          | median speech-to-record lag: 18 years                                                                |
| Benchmark contrast              | geopolitics maps quickly; climate slowly and selectively; digitalisation mostly outside rate setting |

**Public discussion and official policy reasoning are not interchangeable corpora.**

# The language narrows as climate moves toward the mandate

---

## Public speech

Climate risk, transition, sustainable finance, and long-horizon responsibility.

## Policy record

Inflation, supply shocks, financial stability, collateral, asset purchases, and supervision.

broad issue language → mandate-specific vocabulary.

**Translation, not repetition, is the evidence of policy entry.**


# Mandate linkage rises sharply inside official records

| Climate excerpt code | Speeches | Policy records |
|----------------------|----------|----------------|
| No mandate linkage   | 72.2%    | 33.8%          |
| Inflation linkage    | 6.3%     | 33.1%          |
| Operational linkage  | 4.5%     | 25.7%          |

policy seriousness  $\approx$  appearance in mandate-relevant language.

This operational definition measures translation into the mandate; it does not prove that policy is optimal or sufficient.


# Climate attention rises at different times across channels



Public speech typically moves first; official policy records follow with a slower and smaller response.

ML4Z climate temporal-trends figure.

# Leadership can change public climate communication




The source figure isolates a visible change around Lagarde's ECB presidency; leadership is one possible accelerator of public signalling.

ML4Z Lagarde-effect figure.




# Commitment language evolves over time



The content question is not only whether climate appears, but how strongly the institution commits itself when it appears.

ML4Z temporal-commitment figure.


# Commitment is slightly stronger in the official record



Mean commitment: speeches 1.48, minutes 1.53; Mann-Whitney  $p = .029$ .

ML4Z grouped commitment figure.

# The first official mention often arrives decades after the first speech



The mean speech-to-record lag is 14.3 years; the source chart shows especially long delays for institutions such as the Fed and ECB.

ML4Z first-mention-lag figure.


## Part VII

# Understandability is a communication constraint

---

Flesch–Kincaid measures sentence and word difficulty—not comprehension

# FOMC speech readability varies over time



The source chart tracks Flesch–Kincaid grade level across FOMC speeches; a higher score indicates more years of schooling implied by the text’s sentence and syllable structure.

ML4Z source asset fk1.png.

# The grade-level score is a mechanical text statistic

---

$$FK = 0.39 \left( \frac{\text{words}}{\text{sentences}} \right) + 11.8 \left( \frac{\text{syllables}}{\text{words}} \right) - 15.59.$$

## Higher score

Longer sentences and more syllables per word imply a higher school-grade level.


## What is absent

Prior knowledge, policy concepts, document structure, charts, trust, and whether the audience inferred the intended message.

Readable syntax helps, but readability is not understanding.


# Mechanical readability is only one tool for understanding

---



The folded-map route is a neutral policy-logic cue—not evidence that an audience reached the intended interpretation. Editorial marker illustration generated for this lecture.

# Document type changes the measured reading burden




A statement, minutes, transcript, and speech perform different institutional jobs; comparing their grade levels should preserve that functional difference.

ML4Z source asset fk2.png.



# The communicator and the channel both matter



Chair communication and press-mediated communication are selected and edited under different constraints, so a grade-level gap is descriptive rather than causal.

ML4Z source asset fk3.png.

# Use readability as a diagnostic, not a verdict

---


| Question   | FK can help                                     | FK cannot answer                                                 |
|------------|-------------------------------------------------|------------------------------------------------------------------|
| Syntax     | Are sentences and words mechanically demanding? | Is the policy logic conceptually clear?                          |
| Comparison | Do channels or speakers differ descriptively?   | Did the channel cause understanding?                             |
| Design     | Where should an editor inspect complexity?      | Did households, markets, and specialists infer the same message? |

lower  $FK \not\Rightarrow$  greater understanding.

The communication objective is comprehension, not merely a lower score.

# Strategic communication is disclosure across time

---



---

| Lesson                | Evidence in this lecture                               |
|-----------------------|--------------------------------------------------------|
| Later can be new      | path surprises in minutes move asset prices            |
| Omission is selective | disagreement and risk have the largest statement gaps  |
| Disclosure is staged  | the residual gap falls from .745 to .334 after minutes |

---

**A short focal statement and a complete public record are different objectives.**

# Climate communication becomes policy through translation

---

public salience → official record → inflation, stability, or operational linkage.

## Necessary checks

Verify context; separate speech from deliberation; code the mandate bridge; measure the lag.

## Limits

A mention is not commitment. A mandate link is not proof of an optimal policy response.

Policy seriousness appears in mandate-relevant language, not keyword volume alone.

# Three constraints shape what audiences finally receive

| Constraint  | Empirical object | Correct interpretation               |
|-------------|------------------|--------------------------------------|
| Selection   | statement gap    | what the focal text omits            |
| Translation | mandate linkage  | how an issue becomes policy relevant |
| Readability | FK grade level   | mechanical syntactic burden          |

attention  $\neq$  understanding  $\neq$  persuasion.

Market regressions are reduced-form; text scores are descriptive; official mentions must be context-verified.

## Sources: FOMC disclosure and market responses

---



Board of Governors of the Federal Reserve System.  
Minutes of the Federal Open Market Committee: FAQ.



Board of Governors of the Federal Reserve System.  
FOMC calendars and information.



Kansoy, F. (2022).  
FOMC Minutes and Financial Markets, Warwick Economics Research Paper Series 1436.



Federal Reserve.  
Minutes of the 26–27 April 2016 meeting.

## Sources: contemporary press and ZLB context

---



The New York Times (19 May 2016).

Federal Reserve minutes and the June-rate interpretation.



The Guardian (18 May 2016).

US Fed says a June rise is possible.



Reuters (18 May 2016).

June rate increase described as on the table.



Rajan, R. (2014).

Competitive monetary easing—is it yesterday once more?

# Sources: climate communication and instruments

---



Carney, M. (2015).

Breaking the tragedy of the horizon.



Kansoy, F. (2026).

“Do Central Banks Take Climate Change Seriously?”

Working-paper listing: [fatih.ai/research](https://fatih.ai/research).



Bank of Japan; European Central Bank.

Climate-response financing; climate tilting of corporate bond holdings.



# Asset credits and provenance

---

| Asset group              | Provenance and treatment                                                                                               |
|--------------------------|------------------------------------------------------------------------------------------------------------------------|
| Strategic raster figures | copied from active ML4Z includegraphics targets; normalised losslessly to PNG; the exact min2 duplicate is stored once |
| Press screenshots        | all three active D4-F060 placements retained; NYT, Guardian, and Reuters article addresses recorded                    |
| Climate raster figures   | copied from active ML4Z targets and normalised to PNG without cropping or recolouring                                  |
| Six PDF figures          | strategic gap closure plus five climate results copied losslessly from active source targets                           |
| Readability figures      | fk1–fk3 copied from active ML4Z targets and normalised to PNG                                                          |

Flesch–Kincaid reference: Kincaid et al. (1975), Derivation of New Readability Formulas.