ijamrsd.com JAMRSD ISSN: 2583-6404
Journal of Advanced Multidisciplinary Research Studies and Development
Peer Reviewed Referred Journal Bi-monthly Open Access

# RESPONSIBLE AI IN CENTRAL BANK COMMUNICATION ANALYSIS: EVIDENCE FROM RBI MONETARY POLICY STATEMENTS

**Ujjwal Kishor Chaudhari[1], Dr. Dhanpal Nivrutti Waghulde[2]**
*[1] Research Student, KCES's Institute of Management and Research, Jalgaon.*
*Email: ujjwalkchaudhari@gmail.com*
*[2] Assistant Professor, KCES's Institute of Management and Research, Jalgaon.*
*Email: dhanpal.tony@gmail.com*

**Abstract**
The growing application of artificial intelligence tools to central bank communications raises pressing questions about methodological responsibility. This paper investigates those questions through the lens of the Reserve Bank of India (RBI) and its Monetary Policy Committee (MPC) statements. Using a six-dimension responsible AI framework — covering transparency, accountability, fairness, safety, reliability, and human oversight — the study evaluates where AI-assisted analysis of RBI texts succeeds and where structural risks persist. A simple natural language processing pipeline and worked sentiment example ground the discussion empirically. The paper finds that while computational approaches offer real efficiency gains, risks rooted in opaque model design, language-based bias, and the absence of human review remain largely unaddressed in current practice. Five governance measures are proposed to bring that practice closer to responsible AI standards.
**Keywords:** Responsible AI, RBI Monetary Policy, NLP, Sentiment Analysis, MPC Communication, AI Governance, Central Bank Transparency.

► *Corresponding Author: Ujjwal Kishor Chaudhari*

## Introduction

For most of the twentieth century, central banks operated with deliberate opacity. Decisions were announced but rarely elaborated upon, and the reasoning behind them was considered the institution's private business. That norm shifted from the 1990s onwards & resulted in financial markets forming more accurate expectations, households and firms were able to plan more confidently, and the credibility of inflation commitments reinforced over time.

India underwent a version of this transformation in 2016. The Finance Act of that year amended the Reserve Bank of India Act to establish a six-member Monetary Policy Committee, replacing the previous arrangement under which rate decisions rested largely with the Governor. Since the MPC began operating, the RBI has published detailed policy resolutions, individual voting records, and meeting minutes — a substantial and consistently structured body of text that now spans nearly a decade of monetary policymaking.

Against this backdrop, artificial intelligence and natural language processing tools have begun to be applied to central bank communications at scale. Researchers and financial analysts use sentiment classifiers, topic models, and large language models to extract signals from MPC texts, track tonal shifts over time, and generate accessible summaries. These tools can handle large volumes of text with speed and consistency that no manual process could match. The problem is that their deployment in a domain as consequential as monetary policy has proceeded without much critical attention to whether it meets basic standards of responsible AI practice.

This paper addresses that gap. It asks how the principles of responsible AI apply to AI-assisted analysis of RBI communications, and what governance arrangements would be needed to operationalize those principles in practice. A simple NLP-based sentiment analysis, applied to illustrative MPC statement excerpts, is used to demonstrate both the potential and the pitfalls of computational approaches in this setting.

## Literature Review

### Central bank communication and monetary policy

Blinder (1998) offered an early and influential argument that communicating the reasoning behind policy decisions reduces unnecessary uncertainty and anchors expectations more effectively than rate movements alone. Indian scholarship in this space expanded considerably after the MPC framework was formalized. Studies have examined how the language of RBI resolutions signals the direction of future decisions, how dissenting votes surface in the published minutes, and how bond and equity markets respond to communication surprises. A consistent limitation of much of this work is its silence on questions of methodology: the AI tools used to analyze RBI texts are rarely subjected to the same scrutiny as the policy texts themselves.

### Responsible AI: core dimensions

The responsible AI literature has converged on a set of overlapping principles. Transparency refers to the legibility of an AI system — how clearly its logic, training data, and outputs can be understood by those who use or are affected by it. Accountability concerns the allocation of responsibility when AI outputs are wrong or harmful. Fairness addresses whether a system treats different groups, languages, or contexts equitably. Safety covers the prevention of harm, including harm arising from erroneous outputs in high-stakes settings. Reliability refers to the consistency of a system's performance over time and across conditions. Human oversight is concerned with preserving meaningful human agency in processes that AI tools support or influence. These principles have been applied most extensively in healthcare, criminal justice, and hiring. Their application in financial communication analysis is less developed.

### AI in financial communication analysis

Dictionary-based sentiment analysis, in which policy-relevant words are assigned positive or negative scores and aggregated into an index, was among the first computational methods applied to central bank communications. More recent approaches have used transformer-based language models to capture richer semantic content. Both families of methods have shown promise, but both also carry vulnerabilities that responsible AI principles are designed to address: opaque model logic, training data that may not reflect the specific institutional context, and outputs that are often presented as more definitive than the underlying methodology warrants.

## Data and Methodology

### Institutional setting

The MPC brings together the RBI Governor, who chairs the committee and holds a casting vote, with the Deputy Governor responsible for monetary policy, one RBI executive director, and three external members appointed by the central government. Decisions are reached by simple majority. Meetings follow a bimonthly schedule — typically six rounds per year — and each round produces a policy resolution, a Governor's statement, and minutes published fourteen days after the resolution. The RBI also releases a comprehensive Monetary Policy Report twice a year. Together, these documents form the corpus that AI tools in this domain analyse.

**NLP processing pipeline**

Raw MPC text is prepared for computational analysis through a sequence of steps that are standard in the NLP literature. The text is first converted to lowercase and stripped of punctuation marks that carry no semantic content. Words that appear frequently across all documents but carry little analytical meaning, articles, prepositions, conjunctions are removed. The remaining text is split into individual tokens, and each token is reduced to its base form through lemmatization, so that inflected variants of a policy-relevant word are counted as instances of the same concept.

**Sentiment scoring model**

Policy tone is quantified using a scoring rule calibrated to monetary policy language. Words and phrases associated with a growth-supportive or accommodative orientation, including support, recovery, accommodative, liquidity, growth, and stimulus are classified as positive signals. Words associated with inflation control, monetary tightening, including inflation pressure, price stability, and inflation risk are classified as negative signals. The sentiment score S is computed as:

$$\mathbf{S = (P - N) / (P + N)}$$

where P is the count of growth-supportive keywords and N the count of inflation-focused keywords in a given document. Scores range from −1 to +1. Positive values indicate an accommodative communication stance; negative values suggest a tightening orientation; values close to zero reflect a broadly neutral tone. This method was chosen deliberately over more sophisticated alternatives because interpretability, auditability, and researcher accountability take precedence over predictive complexity in a responsible AI context. A dictionary-based approach sacrifices some analytical depth in exchange for full transparency at every step: any score can be traced back to specific words in the source text, any classification decision can be challenged on identifiable grounds, and any reader can independently verify or contest the result. That trade-off is itself a responsible AI design choice.

**Application of Responsible AI Framework**

This section evaluates how each of the six responsible AI dimensions applies specifically to the NLP-based sentiment analysis constructed in this study for RBI MPC statements. Each sub-section identifies the concrete design decision or risk that the dimension creates at the corpus, modelling, or output stage.

**Transparency**

Transparency is operationalized through full disclosure of three elements: the corpus boundaries, the complete keyword dictionary, and the scoring formula. Any reader can replicate every calculation on any MPC statement using the same inputs. This design addresses the transparency deficit in AI-based central bank communication analysis, where undisclosed proprietary models make independent verification impossible.

**Accountability**

Because keyword classifications are made by the researcher and not delegated to a black-box algorithm, a clear line of responsibility exists for every scoring decision. The choice that price stability is hawkish and that support is dovish can be contested and defended on identifiable grounds because those choices are visible. For regulated financial institutions, this traceability is the foundation on which any supervisory review of AI-assisted analysis would have to rest.

**Fairness**

A major fairness problem is specific to the RBI context. A dictionary derived from Federal Reserve or ECB texts embeds a macroeconomic vocabulary in which food price dynamics, monsoon-driven supply shocks, and rural credit conditions are marginal. Since these are recurring preoccupations in RBI communications, such a dictionary would produce systematically skewed classifications. The dictionary used here was built from RBI MPC texts specifically to avoid importing those external biases.

**Safety**

This study addresses safety by presenting all scores as indicative rather than definitive, and by stating the method's scope explicitly so that users cannot mistake a narrow sentiment index for a comprehensive policy assessment.

**Reliability**

If the pipeline applied to 2018 statements uses a different keyword list than the one applied to 2026 statements, any observed shift in scores may reflect the methodological change rather than a genuine change in RBI communication. This study specifies that any future revision requires the entire time series to be recomputed from scratch.

**Human oversight**

Sentiment scores are treated as analytical inputs, not conclusions. Each score is situated against the policy decision taken at the same MPC meeting and the macroeconomic conditions prevailing at that time before any interpretive claim is made. A score of 0.33 during rising food inflation carries a different meaning than the same score in an easing phase. Human oversight converts a number into a policy-relevant observation.

## Empirical Analysis

### NLP-based sentiment analysis

To illustrate the application of the Responsible AI framework, a simple sentiment analysis is performed on excerpts from RBI MPC statements. The analysis follows a structured NLP pipeline that converts raw policy text into measurable indicators of communication tone.

### NLP processing pipeline

Figure 1 illustrates the NLP workflow used in the analysis. This process converts unstructured policy statements into structured textual data that can be analyzed quantitatively.

Figure 1: NLP pipeline — RBI MPC policy statement analysis

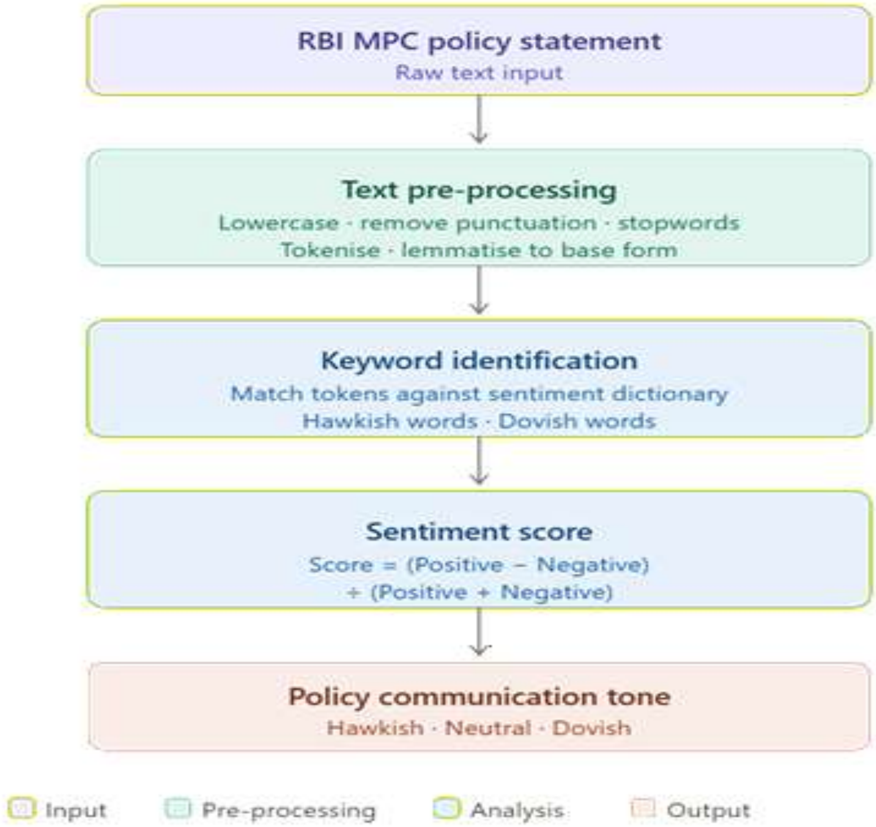


**Sentiment dictionary**

The monetary policy keyword dictionary used in this study classifies terms into two groups. Growth-supportive (dovish) terms include: support, recovery, accommodative, liquidity, growth, and stimulus. Inflation-control (hawkish) terms include: inflation pressure, tightening, price stability, and inflation risk. These selections reflect the vocabulary that has recurred across RBI MPC statements.

**Worked example**

Figure 2 demonstrates how sentiment scores are calculated from an MPC statement. A positive sentiment score indicates a relatively growth-supportive communication tone, while negative values would indicate a more hawkish stance emphasizing inflation control.

**Figure 2: Example Sentiment Calculation**

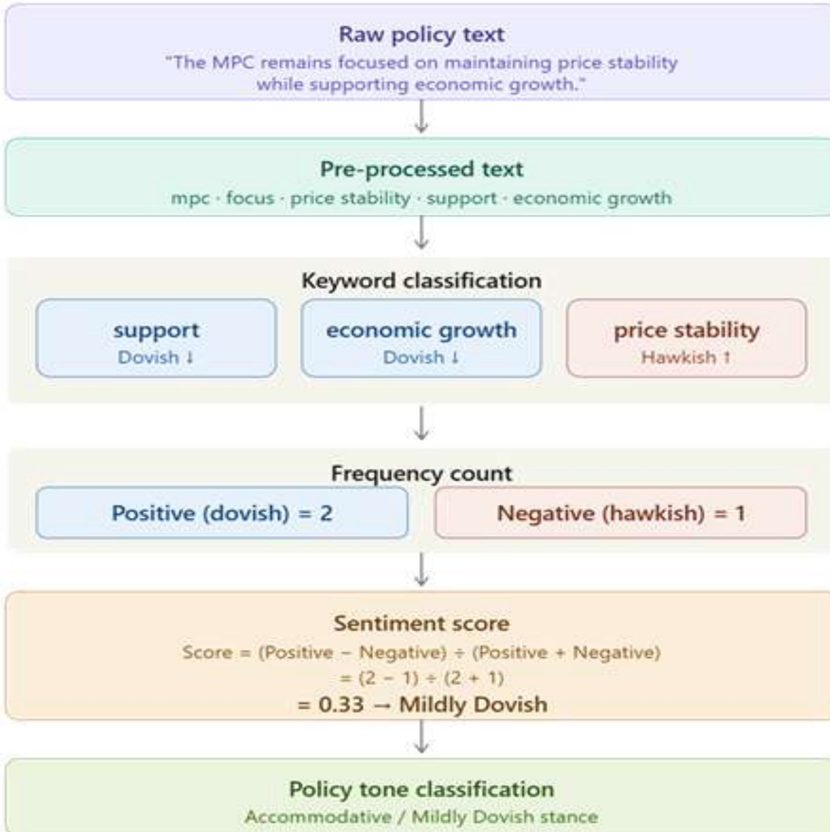


**Results**

**Framework findings**

The application of the responsible AI framework showed that none of the six dimensions remained purely theoretical when tested against the actual design decisions made in building this study's NLP pipeline.

**Transparency** was demonstrated by documenting the full pipeline, disclosing the keyword dictionary, and stating the scoring formula in a form that any reader can independently verify.

**Accountability** was demonstrated by locating responsibility for keyword classification.

**Fairness** surfaced as a corpus design question. Limiting the analysis to English-language MPC documents was a deliberate scoping decision, and its fairness implication.

**Safety** was addressed by presenting the sentiment score as an indicative measure rather than a decisive signal.

**Reliability** was built in by fixing the keyword dictionary for the entire corpus and specifying that any future revision would require the full time series to be recomputed.

**Human oversight** was operationalized by situating each sentiment score within its macroeconomic context before drawing any interpretive conclusion.

**What the worked example showed**

The sentiment calculation applied to the MPC excerpt — "The MPC remains focused on maintaining price stability while supporting economic growth" — produced a score of 0.33. This result has a limited but concrete meaning: given the keyword dictionary constructed for this study, the statement contains more growth-supportive language than inflation-control language, by a ratio of two positive tokens to one negative token. The score is consistent with the balanced framing that characterizes many RBI statements during periods when both inflation and growth concerns are simultaneously present.

What this example does not show — and what this paper does not claim — is how the score would change under a different dictionary, how the RBI's overall tone has shifted since 2016, or whether the dictionary-based approach outperforms alternative methods. Those are empirical questions that require a more extensive data exercise than the scope of this paper permits.

**Governance Recommendations**

Five targeted measures would bring AI-assisted analysis of RBI communications closer to responsible AI standards.

**Methodological disclosure** – Any published or commercially distributed AI analysis of RBI texts should document the corpus it drew on, the model architecture it used, and the limitations it acknowledges. Academic journals and regulated financial institutions disseminating it should treat this as a non-negotiable condition.

**India-specific evaluation benchmark** – A shared benchmark for AI tools applied to RBI communications — developed jointly by researchers, the RBI, and the financial industry — would enable users to compare tools on a like-for-like basis and would surface systematic failures before they reach market participants.

**Mandatory human review** – Regulated financial institutions that use AI-derived readings of MPC statements in investment or risk decisions should be required to document that a qualified analyst reviewed those outputs before they were acted upon.

**Multi-lingual data infrastructure** – A curated, annotated corpus of RBI communications in both English and Hindi, maintained by the RBI in collaboration with academic institutions, would address the language-based fairness gap that individual researchers cannot close working independently.

**AI error reporting** – When AI-derived analysis of RBI communications produces a material financial error, that incident should be reported and examined under a formal process. Such mechanisms already exist for other categories of financial failure and would establish the accountability chain this domain currently lacks.

**Limitations**

The analysis is conceptual and normative in its approach; it does not include a head-to-head empirical comparison of specific AI tools applied to the RBI corpus. Such a comparison would require access to proprietary NLP platforms and their model specifications and represents the most important direction for future empirical work in this area. The governance measures proposed would require engagement from the RBI, SEBI, and the Ministry of Finance; their practical viability depends on regulatory priorities and institutional appetite that cannot be assessed from the outside. The study focuses on the use of AI to read and interpret central bank communications, and does not examine the separate question of whether AI tools might appropriately be used by the RBI itself in drafting or optimizing its public statements — a scenario that merits dedicated research. Despite these limitations, the framework developed here is not specific to India in its logic: the same responsible AI concerns arise wherever computational tools are applied to the communications of a major central bank operating in a diverse, multilingual economy.

## Conclusion

Central bank communications are among the most consequential documents that AI tools are now routinely applied to. The RBI's MPC statements are a particularly demanding case: layered with institutional context, attentive to India-specific economic dynamics, and capable of moving financial markets when read in unexpected ways. This paper has argued that deploying AI in this setting without structured attention to responsible AI principles is not a minor methodological oversight but a governance problem with real financial stakes.

The worked sentiment example demonstrated that even a straightforward dictionary-based pipeline can meet responsible AI standards — on transparency, explainability, accountability, and human oversight — when those principles are built into the design from the start rather than added as an afterthought. The five governance measures extend that logic to the broader ecosystem of researchers, regulated institutions, and policymakers who shape how AI engages with central bank communications.


## Acknowledgments

The authors would like to thank KCES's Institute of Management and Research, Jalgaon and HDFC Bank for providing the institutional support necessary to conduct this research.


## References


1. Blinder, A. S. (1998). Central banking in theory and practice. MIT Press.
2. Reserve Bank of India. (2016). Monetary Policy Committee: Constitution, mandate and procedures. Reserve Bank of India. https://www.rbi.org.in
3. Reserve Bank of India. (2020). Monetary Policy Report, October 2020. Reserve Bank of India. https://www.rbi.org.in
4. Reserve Bank of India. (2021). Report on currency and finance: Reviewing the monetary policy framework. Reserve Bank of India. https://www.rbi.org.in
5. Reserve Bank of India. (2023). Annual Report 2022–23. Reserve Bank of India. https://www.rbi.org.in
6. Reserve Bank of India. (2024). Monetary Policy Report, April 2024. Reserve Bank of India. https://www.rbi.org.in