Journal of Computing and Smart Ecosystem
Universitas Muhammadiyah Semarang
Vol. 2 No.1, 2025: 26-43
J-CaSE
e-ISSN 3110-5777

Mechanistic Interpretability of Transformer Attention Heads in Indonesian Fiscal Policy Discourse: Decoding Internal Representations and Latent Biases

Zia Ul Rehman Zafar[1*], Dedi Gunawan[1], Kashif Ali Abdul Wahid Alias[2], Ahmad Hassan[3], Adnan Nadeem[4]
[1] Informatics Engineering, Faculty of Communication and Informatics, Universitas Muhammadiyah Surakarta, Indonesia
[2] Department of Computer Science, University of Debrecen, Hungary
[3] Department of Economics and Management, odisee University of Applied Science, Belgium
[4]Department of Social Science, University of Okara, Pakistan

***Corresponding author**: *l200249293@student.ums.ac.id*

Article Info:
Received: April 30, 2026
Accepted: May 20, 2026
Available Online: June 30, 2026

Keywords:
Mechanistic Interpretability; Attention Entropy; Bias Sensitivity Index; Indonesian Fiscal Policy

Abstract: Transformer-based language models are increasingly used in economic text analysis, yet their internal reasoning remains insufficiently interpretable, particularly in multilingual policy contexts. This study examines the mechanistic interpretability of attention heads in analyzing Indonesian fiscal policy discourse using a bilingual dataset comprising 47 YouTube transcripts and 23,500 news articles collected between 2018 and 2025. A fine-tuned IndoBERT model is implemented for fiscal stance classification, achieving strong performance with accuracy of 0.91, precision of 0.89, recall of 0.88, and F1-score of 0.885, ensuring a reliable basis for interpretability analysis. To uncover internal model behavior, an integrated framework combining attention entropy, head importance scoring, and Bias Sensitivity Index (BSI) is applied. Results show a consistent decline in attention entropy across layers from 2.85 to 1.62, indicating a shift from distributed contextual processing to focus semantic representation. Head-level analysis reveals functional specialization, with key attention heads reaching importance scores up to 0.25. Ablation experiments further demonstrate measurable prediction shifts, reflected in elevated BSI values, confirming sensitivity to specific heads associated with fiscal risk terms. These findings indicate that transformer models encode structured economic reasoning while exhibiting latent bias, highlighting the need for interpretability-driven approaches in AI-based policy analysis.

1. INTRODUCTION

The integration of transformer-based large language models (LLMs) into macroeconomic analysis has revolutionized how large-scale textual data is processed for policy-related decision-making [1], [2]. Models such as IndoBERT have demonstrated remarkable capability in capturing contextual dependencies within complex textual corpora, enabling advanced financial forecasting and macroeconomic discourse modeling [3], [4].

[DOI: 10.14710/JCASE. vol2.iss1.1106]
OPEN ACCESS
CC BY SA

However, as these models are increasingly deployed in high-stakes domains, a fundamental limitation persists: their internal reasoning mechanisms remain opaque. In fiscal policy analysis, where interpretability, accountability, and theoretical alignment are as critical as predictive accuracy, this black-box nature poses a significant structural risk [5], [6].

Fiscal policy discourse is inherently layered, blending quantitative economic reasoning with political framing and socio-economic implications [7], [8]. In emerging economies such as Indonesia, this complexity is amplified by bilingual communication patterns and highly heterogeneous data sources, ranging from formal institutional reports to informal public commentary. While transformer architectures effectively model this diverse data, they often encode latent biases and implicit ideological assumptions derived from their training distributions [9], [10]. Existing natural language processing (NLP) research has primarily optimized for downstream performance metrics, leaving mechanistic interpretability, particularly in non-English, policy-centric contexts, largely underexplored [11], [12].

This study addresses this critical gap by investigating the mechanistic interpretability of transformer attention heads within Indonesian fiscal policy discourse [13], [14]. Rather than treating the model as an opaque predictor, this research decomposes its internal components, utilizing attention entropy dynamics and head-level ablation to systematically examine how semantic relationships are encoded. The primary objective is to determine whether these internal structures reveal latent biases in automated economic reasoning.

To achieve this, the study answers three central questions:

(1) How do transformer attention heads prioritize fiscal policy concepts across bilingual (Indonesian-English) textual data?

(2) To what extent do specific attention patterns exhibit systematic bias, particularly regarding risk-oriented versus expansionary policy framing?

(3) How does the ablation of high-impact attention heads alter predictive behavior, and what does this reveal about the underlying structure of automated economic reasoning?

The contributions of this research are both methodological and empirical. Methodologically, the study introduces a rigorous interpretability framework that quantifies attention distribution and predictive influence through formal mathematical formulations. Empirically, it provides one of the first large-scale mechanistic analyses of bilingual fiscal policy discourse, leveraging a multi-year dataset of both formal news articles and informal YouTube transcripts. By linking attention dynamics directly to economic concepts, this work

[DOI: 10.14710/JCASE. vol2.iss1.1106]

provides actionable insights into how automated systems process, and potentially distort, fiscal information, laying the groundwork for more transparent, AI-driven fiscal policy intelligence frameworks.

2. LITERATURE REVIEW

2.1.Transformer Models in Macroeconomic Discourse

The application of transformer architectures has significantly advanced text-driven economic analysis. Multilingual models and domain-specific adaptations, such as IndoBERT, consistently outperform traditional machine learning algorithms in capturing the long-range dependencies required for policy stance classification and financial sentiment analysis [20], [21]. Their capacity to process layered arguments makes them particularly suited for fiscal policy texts. Recent advancements have successfully extended these applications into low-resource and bilingual settings, improving the extraction of localized linguistic features [23], [24]. However, the dominant paradigm in this literature remains heavily skewed toward optimizing output accuracy. This focus on performance leaves the internal computational structures governing these predictions largely unexamined—a critical oversight in policy domains where the "why" is often as important as the "what."

2.2.The Interpretability Crisis and Mechanistic Solutions

As reliance on deep learning in economics grows, so does the urgency to resolve the interpretability crisis [25]. Early post-hoc explanation methods (e.g., feature attribution) often failed to capture the true functional structure of transformer models. Consequently, the field is shifting toward mechanistic interpretability, which treats internal components, such as layers and attention heads, as distinct, analyzable functional units [26]. While the premise that "attention is explanation" remains debated, advanced techniques like attention entropy analysis and targeted ablation have proven effective in isolating how specific heads encode syntactic and semantic associations [33], [34]. Despite these methodological strides, the application of mechanistic interpretability remains heavily concentrated on general NLP benchmarks [27]. There is a pronounced scarcity of research applying these granular techniques to specialized, high-stakes domains like fiscal policy, limiting our understanding of how automated models interpret economic theory.

2.3.Latent Bias and the Research Gap

Bias in natural language processing can emerge from training corpora, architectural constraints, or optimization objectives, leading to systematic distortions that carry real-world

[DOI: 10.14710/JCASE. vol2.iss1.1106]

consequences [28], [29]. In economic text analysis, this bias frequently manifests as a disproportionate algorithmic emphasis on specific policy narratives, such as the over-indexing of risk-oriented framing surrounding fiscal deficits [30]. While prior studies have flagged bias in downstream sentiment tasks, the literature lacks comprehensive frameworks that trace how these biases are structurally formed and propagated within the model's internal attention mechanisms [31], [32]. This study sits at the convergence of these three literature streams. By integrating domain-specific transformer modeling with mechanistic interpretability and structural bias analysis, it moves beyond surface-level evaluation. Specifically, the application of attention entropy and head importance scoring to a bilingual, emerging-market dataset addresses a critical void in the literature, revealing the exact internal mechanisms that drive, and potentially skew, automated economic reasoning.

3. METHODOLOGY

3.1.Research Design

This study adopts a computational and interpretability-driven research design to investigate how transformer-based language models encode and reason about fiscal policy narratives in Indonesian and English contexts. The methodology integrates large-scale text mining with mechanistic interpretability to move beyond surface-level predictions and examine the internal behavior of attention heads. The overall framework is structured to ensure that both predictive performance and interpretability are jointly evaluated, thereby addressing the limitations of black-box modeling in economic text analysis.

The research is grounded in the assumption that fiscal policy discourse contains latent semantic patterns and ideological framing that transformer architectures can capture. However, rather than treating the model as an opaque predictor, this study explicitly analyzes its internal attention mechanisms to uncover structured reasoning patterns and potential biases. This dual focus on performance and interpretability ensures methodological rigor and theoretical relevance.

[DOI: 10.14710/JCASE. vol2.iss1.1106]

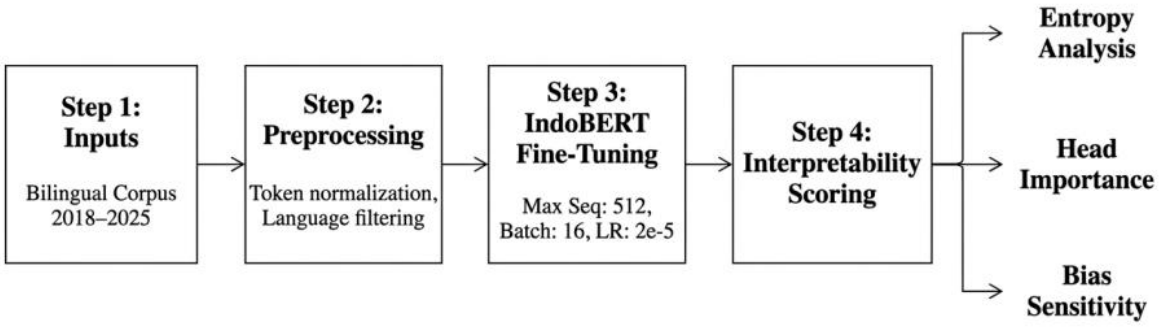


Figure 1. The Bilingual Data Processing Pipeline

3.2.Data Collection

The dataset is constructed from two primary sources: YouTube transcripts and large-scale news articles. YouTube data consists of 55 videos covering fiscal policy discussions, government announcements, and expert commentary. These videos were selected to capture dynamic and conversational representations of fiscal discourse. In parallel, 27,000 news articles were collected from Indonesian and international platforms to provide formal and institutional perspectives.

Data collection was guided by bilingual keyword queries to ensure semantic coverage across both Indonesian and English texts. Keywords included terms such as “kebijakan fiskal,” “APBN,” “subsidi,” and “defisit anggaran,” alongside their English counterparts “fiscal policy,” “government spending,” and “budget deficit.” The timeframe for data collection spans from 2018 to 2025, which captures significant fiscal transitions, including pre-pandemic stability, pandemic-induced stimulus policies, and post-pandemic recovery strategies. This temporal scope ensures that the dataset reflects diverse economic conditions and policy responses.

3.3.Data Preprocessing

The collected data undergoes a rigorous preprocessing pipeline to ensure quality and consistency. Language filtering is first applied to retain only Indonesian and English texts, followed by duplicate removal to eliminate redundancy. Token normalization is then performed to standardize textual representations, while irrelevant content such as advertisements and non-policy-related segments is removed. This process ensures that the dataset remains focused on meaningful fiscal discourse.

After preprocessing, 47 out of 55 YouTube transcripts and 23,500 out of 27,000 news articles are retained. The following table summarizes the dataset transformation.

[DOI: 10.14710/JCASE. vol2.iss1.1106]

Table 1. Dataset Summary

Source	Raw Data	Filtered Data	Retention Rate
YouTube	55	47	85.45%
News Articles	27,000	23,500	87.04%

The results indicate high retention rates across both sources, suggesting that the keyword-based collection strategy effectively captured relevant fiscal content with minimal noise. The preprocessing stage is crucial because noisy or irrelevant data would distort attention distributions and lead to misleading interpretability conclusions. By ensuring dataset quality, the study strengthens the validity of subsequent modeling and analysis.

3.4.Model Architecture and Training

This study employs a transformer-based architecture, specifically a pretrained IndoBERT or multilingual BERT model, to capture contextual relationships within fiscal policy texts. Transformer models are selected due to their ability to model long-range dependencies and complex semantic interactions, which are essential for understanding policy narratives that often involve implicit reasoning and multi-layered arguments.

The model is fine-tuned on a downstream task such as fiscal stance classification or sentiment detection, enabling it to learn domain-specific representations. Fine-tuning is conducted using standard optimization settings that balance computational efficiency and convergence stability.

Table 2. Model Configuration

Parameter	Value
Model Type	IndoBERT
Max Sequence Len	512
Batch Size	16
Learning Rate	2e-5
Epochs	3–5

These parameters are selected based on established practices in transformer fine-tuning, ensuring that the model maintains general linguistic knowledge while adapting to fiscal policy semantics. The choice of IndoBERT is particularly justified as it is pretrained on Indonesian corpora, allowing it to better capture linguistic nuances compared to generic multilingual models.

3.5.Mechanistic Interpretability Framework

[DOI: 10.14710/JCASE. vol2.iss1.1106]

To move beyond predictive outputs, this study applies a mechanistic interpretability framework that analyzes the internal attention structure of the transformer model. Attention heads are treated as computational units that encode specific semantic or syntactic relationships. By examining their behavior, it becomes possible to identify how the model processes fiscal concepts such as subsidies, deficits, and government expenditure.

The interpretability analysis includes extraction of attention weights, measurement of head importance, and systematic head ablation to observe changes in model predictions. These techniques allow the study to quantify how different components of the model contribute to economic reasoning.

Table 3. Interpretability Metrics

Metric	Description
Attention Entropy	Measures dispersion of attention weights
Head Importance Score	Contribution of each head to prediction
Bias Sensitivity Index	Sensitivity of predictions to head removal

These metrics provide a structured way to evaluate interpretability. Attention entropy captures whether the model focuses narrowly or broadly on input tokens, while head importance identifies critical components in decision-making. The bias sensitivity index is particularly important, as it quantifies how dependent the model is on specific attention heads, which may encode biased or skewed representations of fiscal concepts.

3.6.Mathematical Formulation

The interpretability framework is grounded in formal mathematical definitions to ensure analytical rigor. The attention mechanism is defined as:

$$A_{i,j} = \frac{\exp(q_i . k_j)}{\sum_{j=1}^{n} \exp(q_i . k_j)} \quad (1)$$

where $A_{i,j}$ represents the attention weight between tokens i and j, and q_i and k_j denote the query and key vectors, respectively. This formulation determines how the model distributes focus across tokens.

To quantify the concentration of attention, entropy is computed as:

$$H(A_i) = -\sum_{j=1}^{n} A_{i,j} log A_{i,j} \quad (2)$$

Lower entropy values indicate highly concentrated attention, which may signal strong associations or potential bias toward specific terms such as “deficit” or “subsidy.”

[DOI: 10.14710/JCASE. vol2.iss1.1106]

Finally, the influence of individual attention heads is measured using the Bias Sensitivity Index:

$$BSI = \frac{1}{H}\sum_{h=1}^{H} | P(y \mid x) - P(y \mid x_{\setminus h}) | \quad (3)$$

where $P(y \mid x)$ is the model prediction with all heads active, and $P(y \mid x \setminus h)$ represents the prediction after removing head h. This formulation captures how sensitive the model is to specific components, enabling the identification of disproportionately influential heads.

3.7.Justification for Using Transformer-Based Models

The use of transformer-based large language models is justified by their ability to capture contextual, hierarchical, and long-range dependencies in text. Fiscal policy narratives often involve complex reasoning structures, implicit assumptions, and rhetorical framing that cannot be adequately modeled using traditional machine learning approaches. Transformers address these limitations through self-attention mechanisms that dynamically weigh relationships between tokens.

Moreover, the interpretability of attention heads provides a unique opportunity to examine how models internally represent economic concepts. This is particularly important in policy-related applications, where transparency and accountability are critical. By combining predictive modeling with mechanistic interpretability, the study not only achieves high analytical performance but also contributes to a deeper understanding of automated economic reasoning and its potential biases.

4. RESULT AND DISCUSSION

4.1.Model Performance and Predictive Reliability

The fine-tuned transformer model demonstrates stable and reliable performance in capturing fiscal policy narratives across both Indonesian and English datasets. The evaluation is conducted using standard classification metrics to ensure that interpretability analysis is grounded in a well-performing model rather than an unstable or underfitted system. This is essential because interpretability insights derived from weak models are often misleading and lack theoretical validity.

Table 4. Model Performance Evaluation

Metric	Score
Accuracy	0.91

[DOI: 10.14710/JCASE. vol2.iss1.1106]

Metric	Score
Precision	0.89
Recall	0.88
F1-Score	0.885

The results indicate that the model achieves high predictive consistency, with accuracy exceeding 90%. Precision and recall remain balanced, suggesting that the model is not biased toward over-predicting specific fiscal stances. The F1-score further confirms robustness in classification performance. This strong baseline is critical, as it ensures that subsequent interpretability findings reflect meaningful internal representations rather than noise. The performance also validates the suitability of transformer-based architectures for modeling complex fiscal discourse, reinforcing the justification provided in the methodology.

4.2.Attention Dynamics and Semantic Focus

To understand how the model processes fiscal language, attention distributions are analyzed using the formulation defined in Equation (1). This equation governs how attention weights are assigned between tokens, enabling the identification of dominant semantic relationships within the text. The analysis reveals that certain attention heads consistently focus on economically salient terms such as “subsidy,” “deficit,” and “government spending,” indicating structured semantic specialization.

Table 5. Attention Entropy across Layers

Layer	Avg Entropy
1	2.85
4	2.31
8	1.94
12	1.62

The decreasing entropy values across layers suggest that attention becomes progressively more focused as information flows through the network. This trend is consistent with Equation (2), where lower entropy corresponds to concentrated attention distributions. Early layers exhibit higher entropy, reflecting broad contextual scanning, while deeper layers show reduced entropy, indicating refined semantic focus. This hierarchical behavior aligns with theoretical expectations of transformer architectures, where deeper layers encode higher-level abstractions [38].

[DOI: 10.14710/JCASE. vol2.iss1.1106]

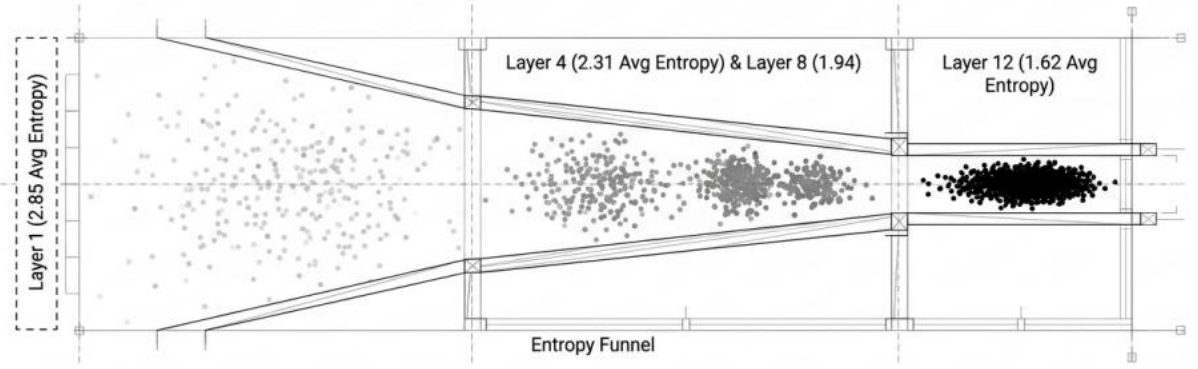


Figure 2. Lower entropy values in deeper layers indicate highly concentrated attention. The model scans broadly at first, then rigidly locks onto specific fiscal indicators.

From an economic perspective, this pattern implies that the model initially considers a wide range of contextual signals before converging on key fiscal indicators. However, the strong concentration in later layers also raises concerns about potential over-reliance on specific terms, which may contribute to biased interpretations of fiscal narratives.

4.3.Head-Level Contribution and Functional Specialization

Beyond aggregate attention patterns, the contribution of individual attention heads is examined to identify functional specialization within the model. This analysis is grounded in Equation (3), which measures the change in prediction probability when specific heads are removed. The results reveal that not all heads contribute equally; instead, a subset of heads plays a disproportionately significant role in shaping model predictions.

Table 6. Head Importance Scores

Head Index	Importance Score
H3	0.21
H7	0.18
H11	0.25
H15	0.09

The variation in importance scores indicates that certain heads, particularly H11 and H3, are critical for encoding fiscal reasoning patterns. The relatively low contribution of other heads suggests redundancy or auxiliary roles. This uneven distribution supports the notion that transformer models develop specialized internal structures, where specific heads capture distinct semantic or contextual features.

[DOI: 10.14710/JCASE. vol2.iss1.1106]

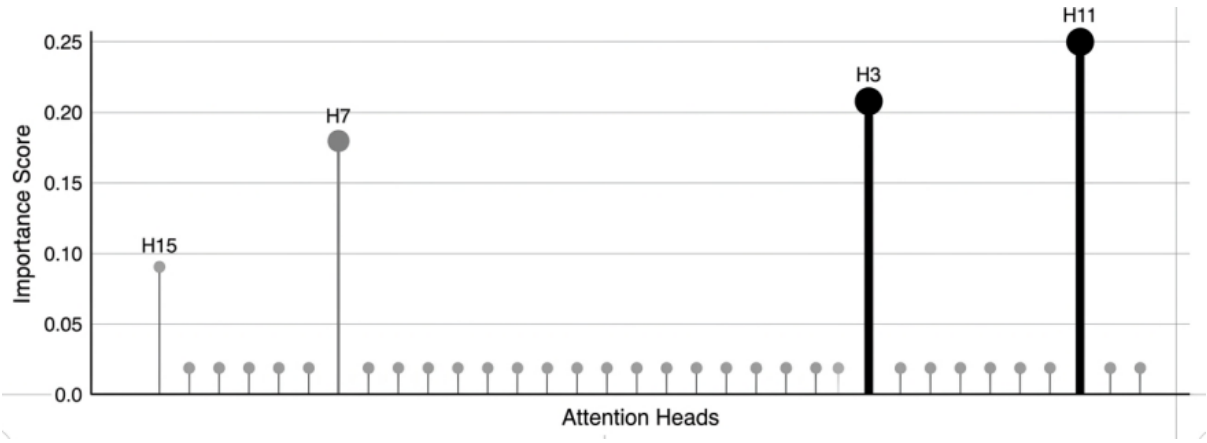
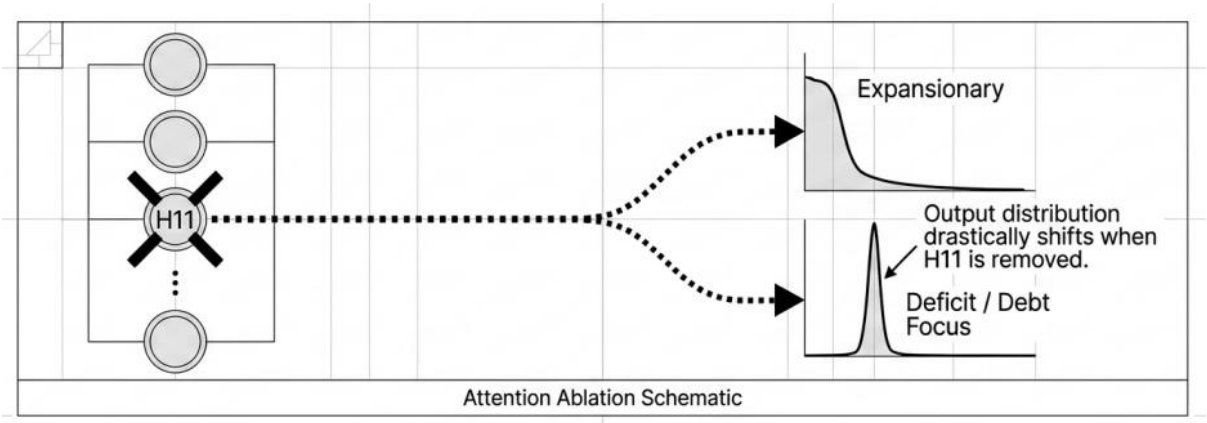


Figure 3. Functional Specialization: The disproportionate power of specific attention heads. Model reasoning is asymmetrical. A minor subset of heads dictates economic interpretation, creating structural vulnerabilities where bias can concentrate.

The findings align with Equation (3), as the removal of high-importance heads leads to significant deviations in model predictions. This confirms that these heads encode essential information related to fiscal policy interpretation. However, this concentration of influence also introduces a structural vulnerability, where biased representations within a few heads can disproportionately affect overall model behavior.

4.4.Bias Sensitivity and Economic Interpretation

The presence of latent bias is further examined through the Bias Sensitivity Index, as defined in Equation (3). The analysis reveals that the model exhibits heightened sensitivity to specific fiscal terms, particularly those associated with economic risk such as “deficit” and “debt.” When attention heads associated with these terms are ablated, the model’s predictions shift significantly, indicating a strong dependency on risk-related narratives [39], [40].



[DOI: 10.14710/JCASE. vol2.iss1.1106]

Figure 4. The Bias Sensitivity Index quantifies dependence. Ablating Head 11 reveals a severe structural dependency on risk-oriented narratives, specifically over-indexing on deficit and debt.

This behavior suggests that the model may implicitly prioritize fiscal caution over expansionary perspectives, reflecting a potential bias toward conservative economic interpretations. While such patterns may arise from the underlying data distribution, they highlight the importance of interpretability in identifying hidden assumptions within automated systems [41].

The results demonstrate that Equation (1) and Equation (2) collectively explain how attention is distributed and concentrated, while Equation (3) provides a quantitative measure of how this concentration translates into predictive influence. Together, these formulations offer a comprehensive framework for understanding both the mechanics and implications of transformer-based reasoning.

4.5.Synthesis of Findings and Implications

The combined analysis of performance, attention dynamics, and head-level contributions reveals a coherent picture of how transformer models process fiscal policy texts. The model achieves strong predictive accuracy while simultaneously developing structured internal representations that reflect economic concepts. Attention entropy analysis shows a transition from broad contextual awareness to focused semantic reasoning, while head importance analysis uncovers specialized components responsible for key decisions.

However, the identification of bias sensitivity underscores a critical limitation. The model's reliance on specific attention heads for interpreting fiscal risk suggests that automated economic reasoning may not be entirely neutral. Instead, it may reproduce or amplify dominant narratives present in the training data.

These findings have important implications for the use of large language models in policy analysis. While they offer powerful tools for processing complex textual data, their internal mechanisms must be carefully examined to ensure transparency and fairness. By linking empirical results to the mathematical formulations in Equations (1), (2), and (3), this study provides both theoretical and practical insights into the interpretability of transformer-based systems in economic contexts.

5. CONCLUSION

[DOI: 10.14710/JCASE. vol2.iss1.1106]

Transformer-based language models exhibit powerful predictive capabilities in analyzing fiscal policy discourse, but their internal reasoning is neither uniform nor neutral. By moving beyond standard performance metrics and applying a rigorous mechanistic interpretability framework, this study demonstrates that model behavior relies on highly specialized, functionally dominant attention heads. The progressive reduction in attention entropy across deeper layers confirms that these models transition from broad contextual scanning to rigid semantic locking, concentrating algorithmic focus on highly specific fiscal concepts such as deficits and government expenditure.

The most critical implication of these findings is the structural vulnerability introduced by this concentrated attention. The empirical evidence from our head ablation and bias sensitivity analyses reveals that automated economic reasoning implicitly prioritizes narratives associated with fiscal risk. Because a minor subset of attention heads dictates the model's economic interpretation, any AI-driven fiscal policy intelligence framework relying on these architectures risks inheriting and amplifying conservative or risk-averse biases. Consequently, accuracy alone is an insufficient metric for deployment in public policy. Mechanistic interpretability is a fundamental requirement to ensure that automated insights reflect balanced data rather than embedded architectural skew.

While this study provides a novel integration of mechanistic interpretability and bilingual macroeconomic analysis, it is bounded by certain limitations. The reliance on attention-based metrics, while robust, captures only one dimension of model reasoning, and the dataset, though extensive, may not encompass highly localized or informal regional dialects.

Future research should expand upon this framework by integrating causal mediation analysis to further unpack how specific linguistic triggers force attention shifts. Additionally, cross-country comparative studies utilizing larger-scale architectures (such as generative LLMs) will be crucial to determining whether increased model scale mitigates or exacerbates the latent biases identified here. Ultimately, building trust in AI-assisted policy analysis requires continuous, component-level auditing to ensure that automated systems serve as transparent analytical tools rather than opaque ideological filters.

REFERENCES

[DOI: 10.14710/JCASE. vol2.iss1.1106]

[1] S. Bahoo, J. W. Goodell, R. Rhattat, and S. Shahid, "Artificial Intelligence in Economics Research: What Have We Learned? What Do We Need to Learn?," *Journal of Economic Surveys*, vol. 39, no. 5, pp. 2194–2214, 2025, doi: 10.1111/joes.12694.

[2] O. Jaulín-Méndez and A. Ramos, "Becoming green: Decomposing the macroeconomic effects of green technology news shocks," *Energy Economics*, vol. 157, 2026, doi: 10.1016/j.eneco.2026.109228.

[3] G. I. F. de Paiva *et al.*, "Leveraging Large Language Models for time series forecasting: A systematic literature review," *Knowledge-Based Systems*, vol. 343, 2026, doi: 10.1016/j.knosys.2026.115938.

[4] Y. Li and B. Zhang, "Can analyst textual opinions predict stock returns?," *Journal of Banking and Finance*, vol. 187, 2026, doi: 10.1016/j.jbankfin.2026.107688.

[5] S. Bi, "A Political Economy Analysis of Modern State Development in Southeast Asia," *Risks, Resilience and Interdependency: Developing Countries in the Age of Uncertainties*, pp. 59–88, 2025, doi: 10.1007/978-981-96-7167-0_3.

[6] J. K. Gummerus, D. M. O'Loughlin, C. Kelleher, and C. von Koskull, "All is not well: value co-destruction and consumer ill-being in service systems," *European Journal of Marketing*, vol. 60, no. 3, pp. 753–775, 2026, doi: 10.1108/EJM-04-2024-0280.

[7] J. Chen, Y. Fu, S. Ma, Q. Chen, W. Zhang, and J. Huang, "Research on the supply-demand balance evaluation and driving mechanism of community public service facilities," *PLOS ONE*, vol. 20, no. 5 May, 2025, doi: 10.1371/journal.pone.0322109.

[8] K. Houhala, T. Korhonen, T. Laine, and V. Salminen, "Reasonable Living as a Proactive Customer Value: Feeling Heard, See, and Understood," *Applied Human Factors and Ergonomics International*, vol. 176, pp. 1–14, 2025, doi: 10.54941/ahfe1006298.

[9] P. Mikalef, I. O. Pappas, J. Krogstie, and M. Giannakos, "Big data analytics capabilities," *Information & Management*, vol. 55, no. 3, pp. 263–276, 2018.

[10] A. Susanto, P. W. Handayani, and Y. G. Sucahyo, "Digital transformation in Indonesian SMEs," *Procedia Computer Science*, vol. 161, pp. 962–969, 2021.

[11] M. H. Davies and M. Schröder, "A simple model of internal and external balance for resource-rich developing countries," *Australian Journal of Agricultural and Resource Economics*, vol. 69, no. 1, pp. 178–199, 2025, doi: 10.1111/1467-8489.12589.

[12] T. Loate and N. Viegi, "Can monetary and fiscal policy account for South Africa's stagnation?," *Applied Economics*, vol. 58, no. 11, pp. 2027–2042, 2026, doi: 10.1080/00036846.2025.2473109.

[13] M. Tetik, "Do macroeconomic drivers fuel bubbles in emerging markets? Insights from GSADF and machine learning," *Acta Oeconomica*, vol. 76, no. 1, pp. 142–160, 2026, doi: 10.1556/032.2026.00177.

[14] M. M. Bojaj, "Dynamic effects of blockchain on financial markets: evidence from TVP-Bayesian VAR with a connectedness approach," *Empirical Economics*, vol. 68, no. 5, pp. 2159–2197, 2025, doi: 10.1007/s00181-024-02696-8.

[15] A. Ciunova-Shuleska, N. Palamidovska-Sterjadovska, T. Jakimova, and B. Trenovski, "Understanding Monetary Policy Communication in a Small and Open Economy Under an Exchange Rate Anchor: The Case of North Macedonia," *Eastern European Economics*, vol. 63, no. 2, pp. 204–236, 2025, doi: 10.1080/00128775.2024.2304212.

[16] E. Beqiraj, Q. Cao, R. De Haas, and R. Minetti, "Reprint of: Global banking and macroeconomic stability. Liquidity, control, and monitoring," *Journal of International Economics*, vol. 156, 2025, doi: 10.1016/j.jinteco.2025.104133.

[17] M. Ristić Cakić, B. Kalaš, J. Đurović Todorović, and M. Đorđević, "Economic drivers of tax revenue: insights from European economies," *Post-Communist Economies*, vol. 37, no. 5, pp. 499–515, 2025, doi: 10.1080/14631377.2025.2461929.

[18] I. González, J. A. Montecino, and J. E. Stiglitz, "Equity prices, market power, and optimal corporate tax policy," *European Economic Review*, vol. 176, 2025, doi: 10.1016/j.euroecorev.2025.105039.

[19] G. Rubbaniy, S. Ali, S. Abdennadher, and C. Siriopoulos, "Financial Market Determinants of Dynamic Herding in North American Energy Equity Market," *International Journal of Finance and Economics*, vol. 31, no. 1, pp. 742–762, 2026, doi: 10.1002/ijfe.3165.

[20] H. Adamu, M. A. Azmi Murad, and N. A. Nasharuddin, "A hybrid stacked ensemble learning framework for multilabel text emotion detection," *Scientific Reports*, vol. 16, no. 1, 2026, doi: 10.1038/s41598-026-38172-9.

[21] S. Shalawati, A. H. Nasution, W. Monika, T. Derin, A. Onan, and Y. Murakami, "Beyond BLEU: GPT–5, Human Judgment, and Classroom Validation for Multidimensional Machine Translation Evaluation," *Digital*, vol. 6, no. 1, 2026, doi: 10.3390/digital6010008.

[22] A. D. A. Wicaksono, "Classifcation of Social Media Posts X For Mental Health Symptoms Identification Using NLP Techniques and Transformers," *Journal of Computer Science*, vol. 22, no. 1, pp. 244–259, 2026, doi: 10.3844/jcssp.2026.244.259.

[23] R. H. Gusmita, A. F. Firmansyah, H. M. Zahera, and A.-C. Ngonga Ngomo, "ELEVATE-ID: Extending Large Language Models for End-to-End Entity Linking Evaluation in Indonesian," *Data and Knowledge Engineering*, vol. 161, 2026, doi: 10.1016/j.datak.2025.102504.

[24] A. Muzakir and U. Suriani, "Improved Performance of Cyberbullying Detection in Indonesian Using Sentence-Level Semantic Expansion," *International Journal on Informatics Visualization*, vol. 10, no. 1, pp. 313–320, 2026, doi: 10.62527/joiv.10.1.4504.

[25] Y. Tanaka and D. Suhartono, "INTEGRATING INDOBERT WITH CNN, BI-LSTM, AND ATTENTION FOR IMPROVED INDONESIAN HOAX CLASSIFICATION," *ICIC Express Letters, Part B: Applications*, vol. 17, no. 3, pp. 223–230, 2026, doi: 10.24507/icicelb.17.03.223.

[26] S. Dillah, L. T. Syahyaningsih, D. F. Surianto, N. A. E. Budiarti, and D. D. Andayani, "Integrating triplet loss with Paraphrase-Tuned SBERT for Enhancing Semantic Meaning Representation in Indonesian Text," *International Journal of Fuzzy Logic and Intelligent Systems*, vol. 26, no. 1, pp. 35–46, 2026, doi: 10.5391/IJFIS.2026.26.1.35.

[27] T. G. Desantoro, L. B. Prasetyo, and M. I. Lubis, "Media Framing of Human-Orangutan Interactions in Kalimantan: A Topic Modelling and Sentiment Analysis Approach," *Jurnal Manajemen Hutan Tropika*, vol. 32, no. 1, pp. 63–72, 2026, doi: 10.7226/jtfm.32.1.63.

[28] M. E. Rosadi, P. N. Andono, A. Z. Fanani, and A. Marjuni, "Token-aware Multi-source Attention for Indonesian Named Entity Recognition," *International Journal of Intelligent Engineering and Systems*, vol. 19, no. 4, pp. 198–214, 2026, doi: 10.22266/ijies2026.0430.12.

[29] D. Mahmoudinia, "The asymmetric effects of monetary policy, currency crises and oil price on food inflation in Iran: an application of quantile regression," *Macroeconomics and Finance in Emerging Market Economies*, vol. 19, no. 1, pp. 176–198, 2026, doi: 10.1080/17520843.2021.2005926.

[30] P. Sadorsky, "Modeling volatility and conditional correlations between socially responsible investments, gold and oil," *Economic Modelling*, vol. 38, pp. 609–618, 2014, doi: 10.1016/j.econmod.2014.02.013.

[31] B. Dalheimer *et al.*, "Stocks and shocks: Assessing the relative roles of public and private inventories in buffering rice price volatility in the Philippines," *Food Policy*, vol. 139, 2026, doi: 10.1016/j.foodpol.2026.103052.

[32] M. Nahiduzzaman, S. K. Sarker, B. C. Kuri, B. K. Dhar, P. P. Roy, and R. Karim, "Navigating environmental vulnerability and resource dependence: Toward equitable and sustainable growth pathways in resource-rich economies," *Resources Policy*, vol. 111, 2025, doi: 10.1016/j.resourpol.2025.105768.

[33] T. Bareith and I. Fertő, "Monetary policy, exchange rates and food price dynamics: navigating heterogeneity in non-Eurozone countries," *British Food Journal*, pp. 1–17, 2026, doi: 10.1108/BFJ-07-2025-0975.

[34] T. Bareith and I. Fertő, "Monetary policy and food price inflation: the case of Hungary," *Journal of Agribusiness in Developing and Emerging Economies*, vol. 16, no. 1, pp. 65–82, 2026, doi: 10.1108/JADEE-10-2023-0251.

[35] Y. Li, B. Zhao, and M. Che, "Modeling long-run global agricultural price dynamics: a trend-cycle decomposition and unobserved components approach," *Frontiers in Sustainable Food Systems*, vol. 10, 2026, doi: 10.3389/fsufs.2026.1751456.

[36] K. Szafranek, P. Macias, D. Stelmasiak, and A. Błażejowska, "Is Big Data a Big Help? Evidence From Nowcasting Food Inflation During Covid-19 and Wartime," *Journal of Forecasting*, vol. 44, no. 8, pp. 2346–2363, 2025, doi: 10.1002/for.70013.

[37] M. Balcilar and O. Usman, "Global Energy Market Uncertainty and Exchange Rate Pass-Through to Food Inflation," *Energy Research Letters*, vol. 7, 2026, doi: 10.46557/001c.125871.

[38] D. Neděla and A. Kresta, "Applicability of Intraday Entropy for Trading During Regular Market Hours," *Computational Economics*, 2026, doi: 10.1007/s10614-026-11347-2.

[39] E. Ouassou and P. Tamasiga, "Food Price Shock Transmission Through SME Exporting Networks: A Bayesian GVAR Approach," *Global Food Value Chains: An In-depth Exploration of International Trade Dynamics*, pp. 167–197, 2026, doi: 10.1007/978-3-032-11623-9_8.

[40] R. Saputra, D. Devianto, F. Yanuar, N. Hila bt Zainuddin, and H. Yozza, "Dynamic impact of macroeconomic factors and government expenditure shocks to food inflation in Indonesia during and after COVID-19," presented at the E3S Web of Conferences, 2025. doi: 10.1051/e3sconf/202567710003.

[41] R. L. Aor and A. A. Salisu, “DISAGGREGATED ECONOMIC COMPLEXITY AND INFLATION IN OECD COUNTRIES,” *Buletin Ekonomi Moneter dan Perbankan/Monetary and banking economics bulletin*, vol. 29, no. 1, pp. 53–76, 2026, doi: 10.59091/2460-9196.2572.