Discovery, Articulation and Analysis of Communication Framings

Rakshitha Rao Ailneni and Sanda M. Harabagiu
Human Language Technology Research Institute, The University of Texas at Dallas
{rxa220074, sanda}@utdallas.edu

Abstract

In social media discourse, the framing of a topic dictates the way opinions are formed. This paper presents two multi-agent frameworks: one for discovering *what* issues are framed in a topic, and another articulating *interpretations* of these issues. This enables an analysis of a topic's discourse that reveals what frame issues are dominant while also providing insights about which issues have more possible interpretations. Experiments across four different topics show promising results for discovering and analyzing framings, allowing new comprehension of controversial topics. [1]

1 Introduction

In social media discourse, a variety of issues are discussed and their *framing* dictates the way these issues are understood and how opinions are formed. Therefore, it is important to discover *what* issues are framed in various ways while also discerning their *interpretations* in order to understand *why* people form certain opinions about a given topic.

The notion of *framing* originates from communication theories (Entman, 2003; Reese et al., 2001; Scheufele, 2004; Chong and Druckman, 2012; Bolsen et al., 2014), where framing is believed to enable the portrayal of an issue from one perspective to the necessary exclusion of alternative perspectives, as noted in (Boydstun and Gross, 2013). For example, when discussing the economic impact of immigration, if this issue is framed as immigrants contributing to the economy through their labor and consumption, it informs pro-immigration attitudes, excluding the perspective of stealing jobs from citizens, which in turn may lead to anti-immigration sentiment. Because nowadays, a large portion of the population forms opinions based on social media discourse, it is imperative to develop

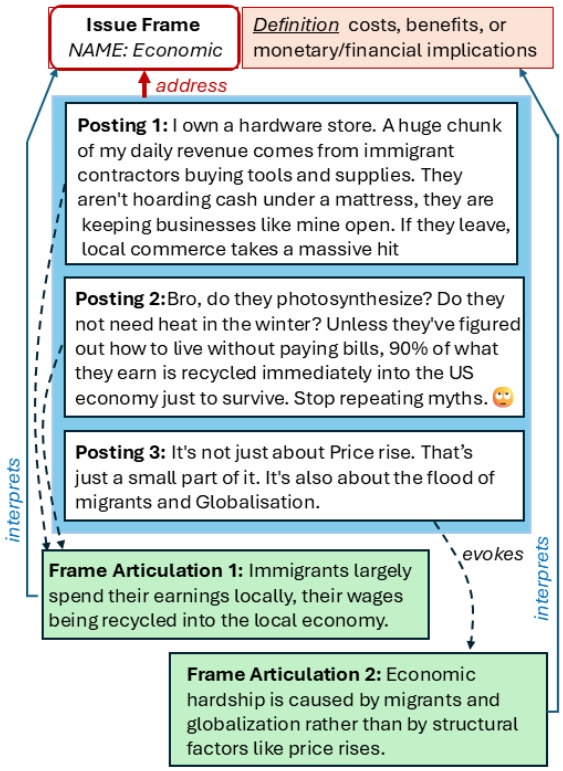


Figure 1: Example of an Issue Framing (IsF) addressed by multiple postings and two Frame Articulations (FAs) providing interpretations. The postings, IsF and FA correspond to the topic of immigration.

methods capable of discovering (a) which issues are framed on social media; and (b) what are their interpretations. Figure 1 exemplifies an issue that is framed when discussing the topic of immigration, namely the *Economic* issue. The Figure also shows three different social media postings where the issue is addressed, showcasing the fact that the same issue may have two different interpretations, which are provided by two different ways of articulating the framing.

Although there is still disagreement among communication scholars on the most appropriate defi-

[1] Code:https://github.com/rak55/Discovery-Articulation-Analysis-CFs

Findings of the Association for Computational Linguistics: EMNLP 2026, pages 47198–47242
October 24-29, 2026 ©2026 Association for Computational Linguistics

nition of framing[2], the most widely accepted definition, provided in Entman (1993), stating that *"to frame is to select some aspects of a perceived reality and make them more salient in a communicating text, in such a way as to promote problem definition, causal interpretation, moral evaluation, and/or treatment recommendation for the item described."*. Previous computational methods for framing analysis (Nguyen et al., 2015; Card et al., 2016; Naderi and Hirst, 2017; Hartmann et al., 2019; Mendelsohn et al., 2021; Otmakhova et al., 2024) focused only on the identification of the *selected aspects of the perceived reality* used in definition of framing, calling them *Issue Frames (IsFs)*.

However, Entman's definition of framing indicates that causal interpretations of the IsFs are also promoted. In this paper we generate automatically *interpretations of framings*, each interpretation being a *Framing Articulation (FA)*, which is a text that (1) addresses the Issue Frames (IsFs) identified in the posting text; and (2) contains at least one causal discourse relation. For example, FA_1 shown in Figure 1 implies a causal relation between immigrants spending their earnings locally and their spending contributing the the local economy. In FA_2 the causal relation is explicit, as the verb "cause" is used.

In this paper we claim that the comprehension of controversial topics, e.g. immigration or climate change, can be achieved through:

A. IsF Discovery: Many previous methods (Hartmann et al., 2019; Mendelsohn et al., 2021) assume that for each topic, *all* its IsFs are known. Often, these IsFs are those used the Media Frames Corpus (MFC) (Boydstun and Gross, 2013), which annotates 14 *generic* IsFs for topics involving policy issues. This allowed the problem of IsF discovery to be cast as a classification problem. However, (Mendelsohn et al., 2021) relied not only the 14 generic IsFs introduced by the MFC, but also on 13 additional IsFs which were specific to the topic of immigration (Mahtani, 2016; Hovden and Mjelde, 2019; Iyengar, 1991). This shows that even when we are aware of some of the IsFs, we still should be able to *discover* additional IsFs that help capturing all the ways a topic is framed in a given discourse. This observation allowed us to formulate the research question:

RQ1: *How could we automatically discover new, yet unknown Issue Frames (IsFs) addressed in a topic's discourse ?*

To answer RQ1, we developed **ID-IsF**, a multi-agentic LLM framework for Identifying and Discovering IsFs addressed in postings discussing a certain topic. This multi-agent framework mimics the deductive and inductive reasoning performed by social scientists when trying to find and define the IsFs. To do so, **ID-IsF** is taking advantage of recent advances in reasoning with multi-agent LLMs through: (a) an *introspective* refinement interaction between two LLM agents (Sun et al., 2025); (b) *contrastive* reasoning (Wu et al., 2024; Chia et al., 2023; Jacovi et al., 2021); and (c) multi-agent *debate* (Liang et al., 2024). In this way, **ID-IsF** is not only identifying *known* IsFs, but it also discovers *new* IsFs, and validates them. Most importantly, for each new IsF, its *name* and *definition* is generated, resembling the way IsFs have been considered in (Hartmann et al., 2019; Mendelsohn et al., 2021).

B. Frame Articulation: Although **ID-IsF** discovers *which* IsFs are addressed in a topic's discourse, it does not account for the interpretations of the framings. This observation allowed us to formulate the research question:

RQ2: *How could we automatically articulate the interpretations of the Issue Frames (IsFs) discovered in a topic's discourse?*

To answer RQ2, we developed **Art-Frame**, a mutli-agentic LLM framework capable to (1) generate *Frame Articulations (FAs)* as interpretations of the IsFs; and (2) validate the FAs though an explainability test. **Art-Frame** is inspired by the method reported in (Weinzierl and Harabagiu, 2024), which generates FAs when it is aware of the IsFs addressed in a topic. For this reason, **Art-Frame** can operate only after **ID-IsF** has performed the IsF discovery. But, **Art-Frame**, through its multi-agent architecture, eliminates the need for human intervention in an active learning loop, as needed in (Weinzierl and Harabagiu, 2024), with the associated effort and time costs.

C. Analysis: The discovery of known and new IsFs addressed in a topic's discourse, along with all their interpretations provided by FAs enables a novel analysis of controversial topics, focused on the way topics are framed. In this paper, we showcase the analysis we performed for four different topics, allowing us to find which IsFs are most addressed in a topic's discourse as well to understanding which IsFs generate the most interpretations through FAs. The analysis also allowed

[2]Framing is termed a "fractured paradigm" Entman (1993).

us to understand the role of newly discovered IsFs in the comprehension of a controversial topic.

2 Related Work

Most of the previous methods (Card et al., 2016; Sturdza, 2018; Liu et al., 2019a; Hartmann et al., 2019; Mendelsohn et al., 2021) for computational framing analysis have conceptualized the notion of framing based on Entman (1993) definition. However, (Nguyen et al., 2015), although using Entman's definition of framings, simply equated a frame with second-level agendas. We followed the same conceptualization of Issue Frames (IsFs) as in (Mendelsohn et al., 2021) (see Appendix A), considering also both issue-generic and issue-specific IsFs.

Computationally, initial large-scale research relied on unsupervised approaches (Neuman et al., 2014; Meraz and Papacharissi, 2013; de Saint Laurent et al., 2020) which revealed only some lexical patterns, but did not discover any IsFs. In contrast, supervised NLP methods used the Media Frames Corpus (MFC) (Card et al., 2015), relying on logistic regression (Card et al., 2016), recurrent neural networks (Naderi and Hirst, 2017), lexicon induction (Field et al., 2018), and fine-tuning pre-trained language models (Khanehzar et al., 2019; Kwak et al., 2020). (Mendelsohn et al., 2021) identified IsFs by using RoBERTa (Liu et al., 2019b). However, they considered only known IsFs, ignoring the possibility to discover new IsFs, as we do in **ID-IsF**. In contrast, (Nguyen et al., 2015) ignored all known IsFs when discovering IsFs in a topic's discourse, unlike **ID-IsF**, which takes into account known IsFs when discovering new IsFs.

The only existing method capable to articulate frame interpretations was reported in (Weinzierl and Harabagiu, 2024) and later in (Weinzierl and Harabagiu, 2025; Ailneni and Harabagiu, 2025). This method requires human intervention in editing FAs in an active learning loop, unlike **Art-Frame**.

3 The Method

Framing processing is performed by two modules operating in sequential mode.

3.1 Module 1: ID-IsF

The role of this module is to (a) identify *known* IsFs and (b) discover *new* IsFs in the postings discussing a topic T. Appendix B.1 lists the known IsFs for four different topics. Clearly, classification methods (Mendelsohn et al., 2021) can identify known IsFs, but they cannot discover new IsFs. Our premise is that the discovery of *new* IsFs is made possible by (1) contrastive reasoning and (2) a verification of proposed new IsFs. The contrastive reasoning is informed by (a) *knowing* when an IsF is addressed by a posting and *knowing why*, as well as (b) *knowing* when an IsF is not addressed by a posting. To acquire this knowledge, we start with identifying *known* IsFs in the component (C1) of **ID-IsF**, illustrated in Figure 2. Component (C2) organizes this knowledge in the Issue Frame Memory; while component (C3) discovers and validates the new IsFs.

C1: Identifying known IsFs: An **Investigator Agent** is used to decide for each posting P_i which of the known IsFs are addressed in P_i. The **Investigator Agent**, using the prompt detailed in Appendix C, can decide for any IsF_x^K whether it is: (a) *applicable*, because P_i is addressing IsF_x^K; or (b) *not applicable*, because P_i is not addressing IsF_x^K; or (c) *uncertain*. In the first two cases, a rationale is also provided. When the **Investigator Agent** is uncertain about an IsF_x^K, it triggers the *Introspective Question-Answering Loop* shown in Figure 2, which resolves the decision by simulating a dialogue between two specialized sub-agents: a **Question Generator Agent** and an **Answer Generator Agent**. An introspective Question-Answering (Q/A) loop between these two agents is generated. Details of the operation of the Q/A loop and examples are presented in Appendix D. The final decision of the **Investigator Agent** on whether P_i is addressing IsF_x^K is resolved by inspecting the Q/A pairs from the Q/A loop. When the **Investigator Agent** finds sufficient evidence in the QA pairs that P_i is addressing IsF_x^K, it will change the uncertain decision and generate a *rationale*. Otherwise it decides that P_i is not addressing IsF_x^K.

C2: The Issue Frame Memory (IFM): Figure 2 shows how for each known IsF_i^K, the IFM keeps track of (a) the IsF name (e.g. *Economy*); (b) its definition (e.g. *The costs, benefits, or monetary/financial implications*); (c) the list of the postings that address the IsF; and (d) the list of all postings that do not address the IsF.

C3: Discovery of new IsFs: is performed under the supervision of a central LLM **Orchestrator Agent**, which governs the life-cycle of new IsF discovery through its three steps:

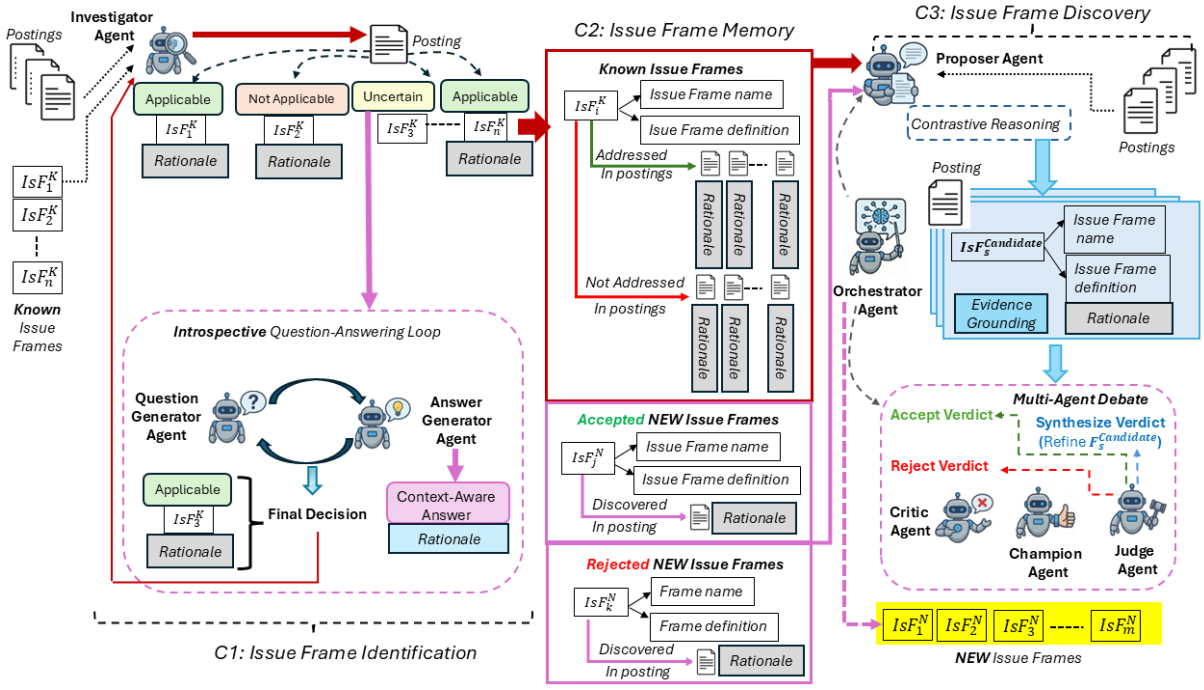


Figure 2: ID-IsF: A Multi-Agent Module for the Discovery of Issue Frames (IsFs).

Step 1: For each posting P_j, the **Orchestrator Agent** invokes a **Proposer Agent** to generate a list of candidate new IsFs $\{IsF_i^{Cand}\}$ by using *Contrastive Chain-of-Thought (CCoT)* reasoning (Chia et al., 2023). Prior work has shown that such contrastive formulations improve reasoning quality and conceptual clarity (Jacovi et al., 2021; Chia et al., 2023; Miller, 2019). Unlike simple Chain-of-Thought (CoT) reasoning (Wei et al., 2022), CCoT considers not only positive demonstrations, but also negative demonstrations. These demonstrations are retrieved by the **Proposer Agent** from the IFM, where for each IsF, examples of postings addressing it as well as examples of postings that do not address it are readily available to form demonstrations for the CCoT prompt. The CCoT prompt is detailed in Appendix C. For each IsF_i^{Cand} the **Proposer Agent**, generates: (a) the *name* of the IsF_i^{Cand}, e.g. *Victim: Family separation*; (b) the *definition* of IsF_i^{Cand}, e.g. *Migrant Children are harmed because governments use family separation as deterrents*; (c) evidence that IsF_i^{Cand} is addressed in P_j; and (d) the rationale.

Step 2: The **Orchestrator Agent** checks if any of the $\{IsF_i^{Cand}\}$ are already known. For this reason, $S_{i,j}^{Def}$, a semantic similarity between the definitions of IsF_i^{Cand} and any (known or new) $IsF_j^{K/N}$ listed in the IFM memory is computed. Appendix E details the similarity metrics we have experimented with. The **Orchestrator Agent** decides that IsF_i^{Cand} is a duplicate of an IsF already present in IFM only if $S_{i,j}^{Def} \geq S_K$, where S_K is a similarity threshold.

Step 3: The **Orchestrator Agent** invokes a multi-agent debate framework where the *discriminant validity* and *applicability* of an IsF_i^{cand} are debated. Three different agents, namely a **Champion Agent**, as **Critic Agent** and a **Judge Agent** participate in the debate. The agents roles are the following:

◇ **The Champion Agent** defends IsF_i^{cand}, arguing for (a) its *applicability* based on evidence that it is addressed in P_j; and (b) the *discriminant validity*, based on IsF_i^{cand} distinctness from all the known IsFs addressed by P_j.

◇ **The Critic Agent** acts as a strict gatekeeper for the necessity of IsF_i^{cand}. This agent operates on the principle of parsimony: new IsFs should not be created if existing ones suffice. It attempts to invalidate IsF_i^{cand} by scrutinizing the **Champion Agent**'s arguments for (a) *redundancy* (when IsF_i^{cand} is similar to other IsFs addressed in P_j; and (b) *hallucination* (when IsF_i^{cand} relies on over-reaching assumptions).

◇ **The Judge Agent** serves as the sole decision-maker and enforces a *dynamic termination policy*. Its verdict may take one of three forms: *Accept*, *Reject*, or *Correction*. When the arguments presented by any of the other two agents are decisive,

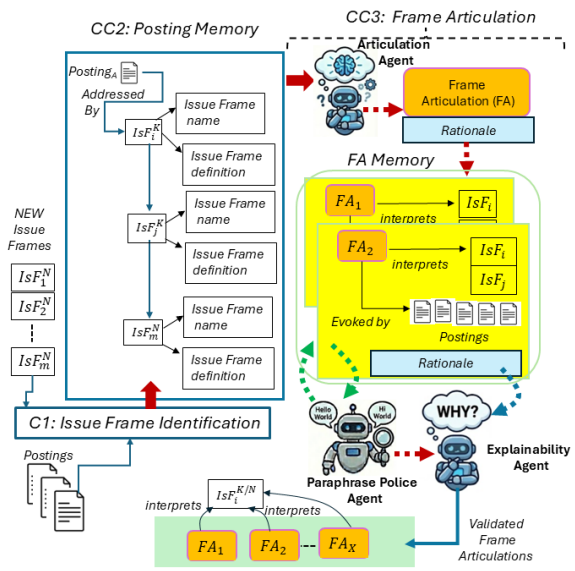


Figure 3: **Art-Frame:** A Multi-Agent Module for Frame Articulations.

the **Judge Agent** issues an early verdict to avoid unnecessary computation. Crucially, the **Judge Agent** also possesses *correction capabilities*: if the **Champion Agent** identifies that the name of IsF_i^{cand} does not correspond to its definition. In this case, the **Judge Agent** intervenes to refine the name or the definition of IsF_i^{cand}. A detailed example of the multi-agent debate where corrections are performed is shown in Appendix F. The Appendix also shows the prompts used by the **Champion, Critic** and **Judge Agents**.

The advantage of using three different agents in this debate stems from the observation that in step 1, when the **Proposer Agent** utilizes contrastive reasoning to hypothesize a IsF_i^{cand} is addressed by P_j, it uses a single generative model susceptible to Confirmation Bias (Turpin et al., 2023; Sharma et al., 2023). To mitigate this, we treat the Proposer Agent's output (IsF_i^{cand}) strictly as a Hypothesis, not a new IsF. The Multi-Agent Debate is rigorously testing the hypothesis to ensure it is not a hallucination or a subtle variant of an existing IsF (Turpin et al., 2023; Liang et al., 2024; Xiong et al., 2023). This form of debate is based on a paradigm rooted in Minsky's *Society of Mind* theory (Minsky, 1986), which posits that intelligence emerges from the interaction of conflicting sub-processes.

When the **Judge Agent** accepts IsF_i^{cand}, a new entry in the IFM memory is created. After all postings are processed, **ID-IsF** outputs the list of newly discovered IsFs.

3.2 Module 2: Art-Frame

As illustrated in Figure 3, **Art-Frame** has three components, but the first component is identical to (C1) of **ID-IsF**. The role of (C1) in **Art-Frame** is to identify all postings where the new IsFs are addressed. This allows the second component (CC2) to organize a Posting Memory listing for each posting P_i all the IsFs, known or new, that it addresses. The third component, (CC3) is used for articulating and validating the frame articulations of each IsF.

CC3: Frame Articulations: An **Articulation Agent** uses Chain-of-Thought (CoT) prompting (Wei et al., 2022) to generate all frame articulations $\{FA_i^k\}$ of each IsF_i. From the Posting Memory, the **Articulation Agent** is able to retrieve $ST_i = \{IsF_i^j\}$, all IsFs addressed by a posting P_i. ST_i along with P_i inform the prompting of the **Articulation Agent**, which generates FA_i^k and its $Rationale_i^j$ for P_i. Figure 3 shows how the *FA Memory* is created, to list all $\{FA_i^k\}$ and (1) the IsF_i^j (new or old) for which each FA_i^k provides interpretations; and (2) all the postings that evoke each FA_i^k.

The FA Memory is used by two agents: A **Paraphrase-Police Agent**, whose role is to find paraphrased FAs; and an **Explainability Agent**, whose role is to validate the articulated FAs. The **Paraphrase-Police Agent** identifies paraphrases between FAs that interpret the same IsFs, using the same few-shot CoT prompting reported in (Ailneni and Harabagiu, 2025). When two FAs are paraphrases, one of them is randomly selected and deleted from the FA Memory.

After the **Paraphrase-Police Agent** consolidates the FA Memory, the **Explainability Agent** operates on the FA Memory. We use an *Articulation Reward Model* R_{art} as a quality filter. Implemented as a discriminative encoder, R_{art} is fine-tuned using a binary cross-entropy loss on a dataset comprising both valid and invalid FAs (See Appendix 5). During inference, the model predicts a scalar validity score $s \in [0, 1]$ for each FA. A strict acceptance threshold τ is used, ensuring that only FAs with high articulation quality are considered.

In addition, we assume that a valid FA_i generated by the **Articulation Agent** should pass a *reconstruction test* performed by the **Explainability Agent**. Given rationale R_i provided for each FA_i in the FA Memory, as well as the IsFs for which FA_i provides an interpretation, the **Explainability Agent** is prompted to re-generate the FA_i^R. The

Explainability Agent has no access to the posting during reconstruction, as a sound rationale should be self-sufficient in regenerating the FA_i without requiring the original posting. The CoT prompt is detailed in Appendix C. If the **Paraphrase-Police Agent** finds that FA_i and FA_i^R are paraphrases, then the reconstruction test is passed. When the reconstruction test is not passed, FA_i is not considered valid. As shown in Figure 3, **Art-Frame** outputs a list of valid FAs, providing interpretations to IsFs.

4 The Data

We use the data that was released publicly in (Weinzierl and Harabagiu, 2025), which covers four different discourse topics:

1. Dataset for Immigration (IMM): It consists of 2,730 posts from Twitter / X annotated with 14 issue-generic IsFs introduced by the Policy Frames Codebook (Boydstun et al., 2018; Card et al., 2015), 11 issue-specific IsFs (Benson, 2013; Hovden and Mjelde, 2019), and 2 narrative-focused IsFs (Iyengar, 1990) surrounding the topic of immigration. In this study, we focus on the 14 issue-generic IsFs *(IMM-G)* and 11 issue-specific IsFs *(IMM-S)*.

2. Dataset for COVID-19 Vaccine Hesitancy (VH): consists of the test split of the Twitter / X postings released by COVAXFRAMES (Weinzierl and Harabagiu, 2022). This dataset contains 2,815 postings annotated with 7 IsFs and 113 FAs. These IsFs correspond to seven dimensions of vaccine hesitancy that influence an individual's likelihood of vaccination (Geiger et al., 2022).

3.*Dataset for Climate Change (CC) & 4. Dataset for Abortion (AB):* consist of 933 Twitter / X postings about Abortion (AB) and 564 postings about Climate Change (CC), originating from the SemEval-2016 Task 6 dataset (Mohammad et al., 2016). The known IsFs and their definitions for both topics are provided in Appendix B.1.

5 Experimental Results

Implementation Details We used GPT-5-mini both for **ID-IsF** and **Art-Frame** as the backbone LLM for all agents. All agents operate with a fixed temperature of 1.0. For **ID-IsF**, the question-answering loop is limited to (`max_loop` = 2). The multi-agent debate terminates either by the decision of the **Judge Agent** or after (`max_run` = 3). In **Art-Frame**, τ is set to 0.8, ensuring that only FAs with high articulation quality are validated.

Quantitative Results: First, we have been interested to know *how well* **ID-IsF** discovers *known* generic IsFs. We used the annotations from the **IMM** dataset, the only dataset with annotated known IsFs. The results are listed in Table 1, showing that **ID-IsF** outperforms other methods, including Zero-Shot and Few-Shot prompting approaches detailed in Appendix H.

The discovery of *new* IsFs by **ID-IsF** was evaluated against three baselines:

◇ *(Baseline 1): IsF Discovery by clustering postings* ($IsFD_{HAC}$). Assuming that all postings that address the same new IsF belong to the same cluster, we have used Hierarchical Agglomerative Clustering (HAC) (Ward, 1963) with the single-link strategy. Similarity between postings was measured with the Sentence-BERT model (Reimers and Gurevych, 2019). The name and the definition of a new IsF was generated by an LLM (GPT-5-mini) with the prompt detailed in Appendix C.

◇ *(Baseline 2) Zero-Shot Prompting (ZS-Prompt)* for IsF Discovery: A Chain-of-Thought (CoT) prompting, in a zero-shot manner, of an LLM (GPT-5-mini) was used, conditioned on known IsFs.

◇ *(Baseline 3) Few-Shot Prompting (FS-Prompt)* for IsF Discovery: An LLM (GPT-5-mini) is conditioned on five same-topic demonstrations to discover new IsFs. To construct these demonstrations, we intentionally omit all *known* IsFs addressed by postings used in demonstrations. To facilitate a fair comparison, we apply the same post-hoc LLM-based semantic de-duplication made possible by Step 2 of **ID-IsF** for all baselines.

Table 2 shows that **ID-IsF** discovers the smallest number of new IsFs, but as we show later in Table 5, the quality of its new IsFs is far superior to those produced by the other methods, indicating that "more" is not necessarily "better".

To quantify how interpretable IsFs are, we computed the number of FAs articulated by **Art-Frame** and compared them with those produced by the baselines. Table 3 reports the average number of FAs interpreting IsFs for each topic. While baseline methods often discover a larger number of IsFs, these IsFs tend to support fewer interpretations. In contrast, IsFs discovered by **ID-IsF** consistently yield a higher number of FAs per IsF, indicating that these new IsFs that are more interpretable.

Finally, we were interested to measure how well

Model	P	R	F1-score
(Mendelsohn et al., 2021)	0.73	0.72	0.71
Zero-Shot Prompt	$0.72_{\pm 0.03}$	$0.79_{\pm 0.02}$	$0.75_{\pm 0.02}$
Few-Shot Prompt	$0.74_{\pm 0.01}$	$0.80_{\pm 0.04}$	$0.77_{\pm 0.02}$
ID-IsF	$0.77_{\pm 0.02}$	$0.85_{\pm 0.03}$	$0.81_{\pm 0.01}$

Table 1: Identification of *known* Generic IsFs on the IMM-G Dataset, evaluated by Precision (P), Recall (R) and F1-score. Results are averaged over three independent runs, reported as mean ± standard deviation.

Method	IMM-G	IMM-S	CC	AB	VH
$IsFD_{HAC}$	21	19	18	19	15
ZS-Prompt	34	30	27	29	22
FS-Prompt	24	21	19	21	16
ID-IsF	**13**	**10**	**10**	**11**	**9**

Table 2: Number of new IsFs discovered for each topic.

Method	IMM-G	IMM-S	CC	AB	VH
$IsFD_{HAC}$	5.6	3.0	4.0	2.7	6.8
ZS-Prompt	4.2	3.9	3.1	2.0	5.0
FS-Prompt	6.1	4.3	4.5	2.9	7.4
ID-IsF	**9.0**	**5.7**	**6.6**	**3.8**	**11.4**

Table 3: Average number of Frame Articulations (FAs) articulated per new IsF across topics.

ID-IsF is capable to *re-discover known* IsFs. For this purpose, we experimented with randomly selecting a percentage $P = 10\%, 25\%, 50\%$ or 75% of known IsFs for each topic. Because the only topic that had generic IsFs is IMM, we considered only the issue-specific IsFs in the experiment, for consistency. We found that the value of P has no effect on the ability of **ID-IsF** to re-discovered IsF, since nearly all known IsFs were re-discovered. Appendix I provides the details.

Qualitative Results are produced for (1) the discovered new IsF; and (2) the articulated FAs.

Quality of the discovered IsFs are revealed by human evaluation using five criteria, listed in Table 4, based on several Likert scales (Robinson, 2014). Table 5 reports the qualitative evaluation results for the IMM-S data, whereas Appendix J.1 lists the results for all four topics. The results indicate that the IsFs discovered by **ID-IsF** consistently outperform the results of the baselines across all five criteria. The largest gains are observed in Salience (SAL), Posting Coverage (PC) and Semantic Distinctiveness (SD), indicating that IsFs discovered by **ID-IsF** are not only salient to the topic, but also are well addressed in the postings and they are semantically non-redundant. Improvements in Definition Clarity (DC) and Definition–Name Alignment (DNA) further suggest that **ID-IsF** produces interpretable and well-formed IsF definitions.

Criterion 1: **Salience (SAL)** of new IsFs: Human judges decide if an IsF represents a salient aspect of the topic, and it is not a special case on a known IsF, answering *Yes=1* or *No=0*.
Criterion 2: **Definition Clarity (DC)** of new IsFs: Human judges decide whether the IsF definition is fluent and easy to understand and rate on a 5-point Likert scale, ranging from 1 (incomprehensible) to 5 (perfectly clear).
Criterion 3: **Definition-Name Alignment (DNA)**: Judges are asked to judge the extent to which the definition explains, justifies, and faithfully reflects the meaning implied by the IsF name. Judgments use a 5-point Likert scale, with a score of 1 indicates that the definition and name are semantically mismatched; while a score of 5 indicates that the definition offers a comprehensive explanation of the IsF name.
Criterion 4: **Posting Coverage (PC)**: Judges evaluate whether postings addressing a given IsF genuinely discuss the IsF. Judges are shown the IsF name and definition along with a random sample comprising 20% of postings addressing the IsF (details are provided in Appendix K.3). For each posting, judges answer *Yes* if the posting clearly addresses the IsF and *No* otherwise. The PC score for an IsF is computed as the proportion of postings receiving a *Yes* judgment among all sampled postings for that IsF.
Criterion 5 **Semantic Distinctiveness (SD)**: measures how clearly IsFs differ in meaning. Sentence Bert (Reimers and Gurevych, 2019) is used to measure the similarity between the definitions of pairs of IsFs. After selecting the most similar pairs of IsFs, judges are asked to consider how distinctive the IsFs are, based on their definitions, using a three-point scale: *Duplicate* (0), *Overlap* (0.5), or *Distinct* (1.0).

Table 4: Evaluation Criteria for Discovered New IsFs.

Quality of the FAs generated by **Art-Frame** are revealed by the properties listed in Table 6. These properties and the metrics that evaluate them were first reported in (Weinzierl and Harabagiu, 2024).

Two computational linguists and a social scientist produced the judgments for the four properties of FAs generated across all topics. Table 7 shows the values for Soundness (Z), Clarity (A), Interpretation Support (S) and Analytical Value (V) obtained when the experts have judged the FAs generated for the issue-specific IsFs identified to the topic of Immigration (IMM-S). Appendix J.2 lists the results obtained for the other topics as well. Table 7 indicates that **Art-Frame** is judged as being better than the baselines across all measures. This demonstrates that identifying high-quality IsFs is

Model	SAL	DC	DNA	PC	SD
ID-IsF	**0.60**	**4.33**	**4.33**	**0.93**	**0.90**
$IsFD_{HAC}$	0.16	3.12	3.39	0.48	0.68
ZS-Prompt	0.10	3.81	4.07	0.55	0.63
FS-Prompt	0.19	4.06	4.20	0.66	0.75

Table 5: Qualitative results for the discovery of *new* IsFs for the topic of Immigration (IMM-S).

Property 1: **Soundness (Z)** of FAs: Human judges decide N_Z, the number of FAs deemed *sound* out of N_T, the total number FAs generated by **Art-Frame**. Soundness is evaluated as $Z = N_Z/N_T$.
Property 2: **Clarity (A)** of FAs: Human judges decide N_C the number of FAs deemed clear, computing the quality of articulation as $A = N_C/N_T$
Property 3: **Interpretation Support (S)**: Human judges are asked to decide if the FA provides a good interpretation of its corresponding IsF. N_S is the number of FAs supporting the interpretation of an IsF. Support is evaluated as $S = N_S/N_T$
Property 4: **Analytical Value (V)**: Judges evaluate if any FA_i is not a special case of another FA_j supporting the same IsF. In addition, each FA should be relevant, being evoked by at least $k = 3$ postings. N_V is the number of FAs meeting both conditions. The analytical value is measured as $V = N_V/N_T$.

Table 6: Evaluation Criteria for generated FAs.

a prerequisite for articulating logically sound FAs; because **ID-IsF** provides a robust foundation, the resulting FAs are both coherent and grounded. In contrast, the baselines struggle to match this fidelity. The $IsFD_{HAC}$ baseline performs the worst across all metrics, indicating that clustering-based approaches often produce noisy IsFs that are difficult to articulate into meaningful FAs.

The gains of **Art-Frame** are even more pronounced for Interpretation Support (S) and Analytical Value (V), where baseline methods fall substantially short.

Method	Z	A	S	V
Immigration (IMM-S)				
Art-Frame	**0.98**	**0.98**	**0.86**	**0.79**
$IsFD_{HAC}$	0.71	0.75	0.38	0.34
ZS-Prompt	0.83	0.83	0.54	0.52
FS-Prompt	0.89	0.91	0.61	0.59

Table 7: Qualitative results for the Frame Articulations (FAs) generated for the topic of Immigration (IMM-S)

Across all topics, the results of $Z \geq .97$ and $A \geq .95$ across all topics indicates high soundness and clarity for the articulated FAs. The average result of $S = 0.83$ indicates that there were some FAs that were not judged to support well the interpretation of some IsFs (see Appendix J.2 for details). The average result of $V = 0.75$ is due to many FA_i being judged as special cases of other FA_j, judges preferring the more general FA_j to be accepted, while the special case FA_i were discarded. Inter-Annotator Agreement resulted in Krippendorff's α between **0.79** and **0.84** (see Appendix K for details).

Model	SAL	DC	DNA	PC	SD
ID-IsF	**0.60**	**4.33**	**4.33**	**0.93**	**0.90**
w/o Investigator	0.50	3.82	3.76	0.81	0.81
w/o IFM	0.19	4.30	4.29	0.91	0.51
w/o Multi-Agent Debate	0.33	4.29	3.81	0.90	0.65

Table 8: Impact of various LLM agents on the *quality* of *new* IsFs from the IMM-S Dataset.

Since the **ID-IsF** system is complex, using multiple agents, we conducted ablation studies to justify the complexity of **ID-IsF** through the impact of each agent on the quality of new IsFs. Table 8 lists the results for IMM-S topic, indicating that all agents contribute significantly. Ablation studies for the other topics are detailed in Appendix L.1. Appendix L.2 details the effect of ablation experiments on **Art-Frame**.

6 Framing Analysis

Framing analysis is performed by considering (a) both newly discovered and known IsFs; and (b) FAs that were judged acceptable. When analyzing the salient IsFs as well as the accepted FAs produced for each topic, we were able to discover which IsFs dominate the framing of each topic, by the number of postings that address them. Moreover, we were able to find (a) how many accepted interpretations each IsF has; and (b) to have access to these interpretations through the corresponding FAs. Furthermore, the analysis allowed us to find which IsF interpretations are dominant for each topic.

We note that for all topics **ID-IsF** discovered new issue-specific IsFs, as detailed in Appendix M. However, the only topic for which generic IsFs are applicable is the topic of IMM (Immigration). This prompted us to produce for IMM a separate analysis of the issue-specific framings, which is illustrated in Figure 4. The definitions of known immigration-specific IsFs are provided in Appendix B.1, while the definitions of the newly discovered immigration-specific IsFs are provided in Apendix B.2.

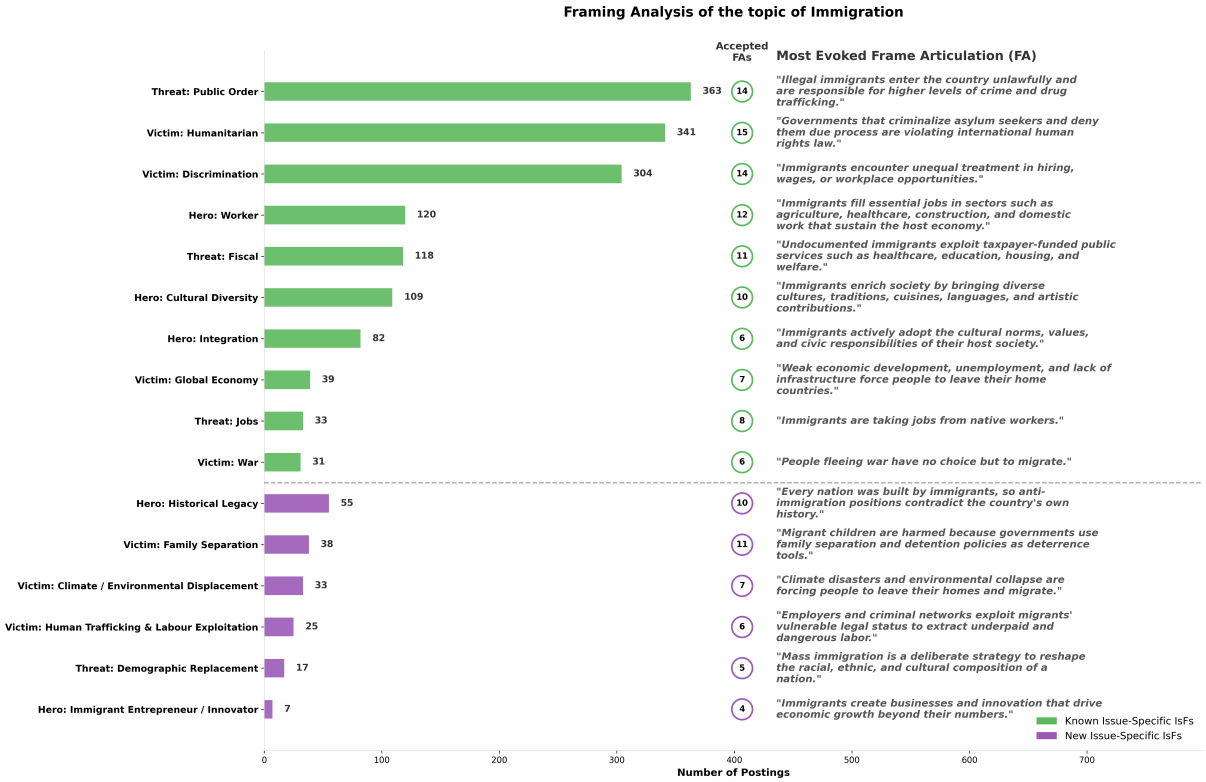


Figure 4: Framing Analysis of the topic of Immigration (IMM-S).

Figure 4 is showing that the dominant immigration-specific IsF is *Threat: Public Order*, which has 14 accepted FAs interpreting it, detailed in Appendix N. Figure 4 also shows that some of the newly discovered IsFs (e.g. *Hero: Historical Legacy*) are addressed in more posts than some of the known IsFs (e.g. *Victim: Global economy*). Therefore, **ID-IsF** discovered some new IsFs that are addressed in many social media postings. Moreover, Figure 4 shows that **ID-IsF** discovered salient IsFs clearly producing many accepted interpretations generated by **Art-Frame**, revealing unknown framings of the topic.

Figure 4 lists for each IsF the number of postings where it is addressed. The figure also lists not only the number of accepted interpretations of each IsF (new or known), available through the accepted FAs generated by **Art-Frame**, but it also showcases the FA that was evoked the most in the topic's corpus when interpreting each IsF. This allows us to quickly grasp the dominant interpretation of each IsF when analyzing a topic.

Figure 4 shows that for the Immigration topic, six new issue-specific IsFs were discovered. Among the newly discovered issue-specific IsFs, *Hero: Historical Legacy* dominates the new framings. For this IsF, Figure 4 indicates that the dominant interpretation is the FA *"Every nation was built by immigrants, so anti-immigration positions contradict the country's own history"*.

In Appendix M we show that salient IsFs with accepted FAs have been produced for the other three topics we have analyzed. Interestingly, in Appendix M we also show that **ID-IsF** has also discovered a new generic IsF, namely *Misinformation and Conspiracy* for the topic of IMM, which is surprising, given that the generic known IsF, originating from (Card et al., 2015), have been considered to provide full coverage for this topic.

7 Conclusion

This paper is describing **ID-IsF**, a new multi-agent framework capable to discover how both new and known topic issues are framed. It also details **Art-Frame**, a second multi-agent framework capable to articulate the interpretations of each framed issue. The operation in tandem of these two LLM architectures provides a novel way of understanding and analyzing the discourse of a controversial topic through the lens of framing.

8 Ethical Statement

We respected the privacy and honored the confidentiality of the authors who wrote the social media posts in the datasets from (Weinzierl and Harabagiu, 2022) and (Mendelsohn et al., 2021). We received approval from the Institutional Review Board at the University of Texas at Dallas for working with these Twitter/X social media datasets, and no attempt was made to identify or profile individual users.

All experimental settings, configurations, and procedures are clearly documented in this paper, the supplemental material, and the linked GitHub repository to support transparency and reproducibility. Human evaluations were conducted to assess the quality of the discovered Issue Frames (IsFs) and the Frame Articulations (FAs), and to ensure the responsible and reliable use of human judgment, we measured and reported inter-annotator agreement (IAA).

We also recognize that research on discovering and analyzing communication frames in politically sensitive or controversial discussions carries potential downstream risks. Although the methods are intended to support understanding of public discourse, similar techniques could be misused for profiling populations, identifying persuasive strategies, targeted messaging, or strategic communication aimed at shaping opinions at scale. Such risks are not unique to our approach but are relevant to computational methods that analyze framing and discourse patterns more broadly.

Our work does not identify individual users, infer personal characteristics, or develop systems aimed at targeted influence. Instead, it focuses on understanding broader patterns in how issues are discussed and framed across online discourse. We encourage future applications of these methods to be developed with appropriate ethical oversight and careful consideration of potential misuse, particularly in politically sensitive settings. Recognizing these limitations and risks is an important part of responsible research practice.

9 Limitations and Future Research

This work has several limitations that point to directions for future research. First, our experiments are conducted exclusively on social media posts from a single platform, Twitter/X, and rely solely on textual content. In practice, many social media posts also incorporate images, videos, hyperlinks, and other forms of multimedia. In future work, we plan to extend our method to support multimodal data, which would enable its application to other social media platforms such as Instagram and Facebook, where discourse is often expressed through a combination of text, images, and videos.

Second, the proposed framework assumes the availability of a predefined set of known Issue Frames (IsFs) for each topic, which serves as the starting point for **ID-IsF**. As described in Section 3.1, the Proposer Agent relies on some known IsFs to perform Contrastive Chain-of-Thought (CCoT) reasoning. While **ID-IsF** is designed to discover new IsFs, it cannot function without several known IsFs. However, our experiments with the re-discovery of known IsFs indicate that **ID-IsF** is capable of discovering IsFs that it was not aware of when provided with a small set of known IsFs. We believe that **ID-IsF** is a viable alternative to (1) methods that only discover known IsFs, e.g. (Mendelsohn et al., 2021) and (2) methods that discover IsFs without any other knowledge, e.g. (Nguyen et al., 2015), but the discovered IsFs are not readily interpretable. However, for topics where no prior knowledge of IsF is available, methods such as (Nguyen et al., 2015) can provide a starting point, as names and definitions of some IsFs discovered by this method can be discerned with human help, and provide input to **ID-IsF**.

Although the multi-agent architecture is designed to mitigate common LLM failure modes such as hallucination and redundancy, the framework remains dependent on the reasoning capabilities of large language models. As a result, errors in interpretation or subtle model biases may propagate through the discovery and articulation pipeline.

The proposed framework has two distinct scalability considerations that warrant explicit discussion. First, the multi-agent architecture of **ID-IsF** and **Art-Frame** involves multiple sequential LLM calls per posting, including introspective Q/A loops and multi-agent debate rounds, which increases computational cost and inference latency compared to single-pass methods. While this cost is manageable for the dataset sizes considered in this paper, it may limit direct applicability to web-scale corpora without modification. A practical mitigation strategy is to run **ID-IsF** on a stratified random sample of postings stratified by time period, platform, or user ideology until no new IsFs are discovered across successive batches of postings. Second, a

natural downstream application of the discovered IsFs is to map them to large-scale corpora beyond the datasets considered in this paper. Running the full **ID-IsF** identification pipeline at web scale for this purpose would be computationally prohibitive. A practical solution is to fine-tune a lightweight supervised classifier on the discovered IsFs, using the postings identified as addressing them during the discovery phase as training data, which can then be applied efficiently to large-scale corpora. We leave the operationalization of this approach as an important direction for future work.

References

Greg D. Adams. 1997. Abortion: Evidence of an issue evolution. *American Journal of Political Science*, 41(3):718–737.

Rakshitha Rao Ailneni and Sanda M. Harabagiu. 2025. Automatically discovering how misogyny is framed on social media. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 12189–12208, Albuquerque, New Mexico. Association for Computational Linguistics.

Mohammad Ali and Naeemul Hassan. 2022. A survey of computational framing analysis approaches. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9335–9348, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Arnav Arora, Srishti Yadav, Maria Antoniak, Serge Belongie, and Isabelle Augenstein. 2025. Multi-modal framing analysis of news. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 31531–31553, Suzhou, China. Association for Computational Linguistics.

Steven E. Barkan. 2014. Gender and abortion attitudes: Religiosity as a suppressor variable. *Public Opinion Quarterly*, 78(4):940–950.

Rodney Benson. 2013. *Shaping Immigration News: A French-American Comparison*. Communication, Society and Politics. Cambridge University Press.

Toby Bolsen, James N. Druckman, and Fay Lomax Cook. 2014. How Frames Can Undermine Support for Scientific Adaptations: Politicization and the Status-Quo Bias. *Public Opinion Quarterly*, 78(1):1–26.

Amber E. Boydstun, Dallas Card, Justin Gross, Paul Resnick, and Noah A. Smith. 2018. Tracking the development of media frames within and across policy issues. *Carnegie Mellon University*.

Amber E. Boydstun and Justin H. Gross. 2013. Identifying media frames and frame dynamics within and across policy issues.

Dallas Card, Amber E. Boydstun, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. The media frames corpus: Annotations of frames across issues. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444, Beijing, China. Association for Computational Linguistics.

Dallas Card, Justin Gross, Amber Boydstun, and Noah A. Smith. 2016. Analyzing framing through the casts of characters in the news. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1410–1420, Austin, Texas. Association for Computational Linguistics.

Yew Ken Chia, Guizhen Chen, Anh Tuan Luu, Soujanya Poria, and Lidong Bing. 2023. Contrastive chain-of-thought prompting. *ArXiv*, abs/2311.09277.

Dennis Chong and James N. Druckman. 2012. Counter-framing effects. *The Journal of Politics*, 75(1):1–16.

Elizabeth A. Cook, Ted G. Jelen, and Clyde Wilcox. 1993. Measuring public attitudes on abortion: methodological and substantive considerations. *Family Planning Perspectives*, 25(3):118–121, 145.

Constance de Saint Laurent, Vlad Petre Glăveanu, and Claude Chaudet. 2020. Malevolent creativity and social media: Creating anti-immigration communities on twitter. *Creativity Research Journal*, 32:66 – 80.

James N. Druckman and Mary C. McGrath. 2019. The evidence for motivated reasoning in climate change preference formation. *Nature Climate Change*, 9(2):111–119.

Robert M. Entman. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of Communications*.

Robert M. Entman. 2003. Cascading activation: Contesting the white house's frame after 9/11. *Political Communication*, 20:415 – 432.

Myra Marx Ferree. 2003. Resonance and radicalism: Feminist framing in the abortion debates of the united states and germany. *American Journal of Sociology*, 109(2):304–344.

Anjalie Field, Doron Kliger, Shuly Wintner, Jennifer Pan, Dan Jurafsky, and Yulia Tsvetkov. 2018. Framing and agenda-setting in russian news: a computational analysis of intricate political strategies. In *Conference on Empirical Methods in Natural Language Processing*.

Mattis Geiger, Franziska Rees, Lau Lilleholt, Ana P. Santana, Ingo Zettler, Oliver Wilhelm, Cornelia Betsch, and Robert Böhm. 2022. Measuring the 7Cs of vaccination readiness. *European Journal of Psychological Assessment*, 38(4):261–269.

Mareike Hartmann, Tallulah Jansen, Isabelle Augenstein, and Anders Søgaard. 2019. Issue framing in online discussion fora. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1401–1407, Minneapolis, Minnesota. Association for Computational Linguistics.

Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. Deberta: Decoding-enhanced bert with disentangled attention. In *International Conference on Learning Representations*.

Jan Fredrik Hovden and Hilmar Mjelde. 2019. Increasingly controversial, cultural, and political: The immigration debate in scandinavian newspapers 1970–2016. *Javnost - The Public*, 26(2):138–157.

Shanto Iyengar. 1990. Framing responsibility for political issues: The case of poverty. *Political Behavior*, 12(1):19–40.

Shanto Iyengar. 1991. Is anyone responsible? how television frames political issues. *Contemporary Sociology*, 21:467.

Alon Jacovi, Swabha Swayamdipta, Shauli Ravfogel, Yanai Elazar, Yejin Choi, and Yoav Goldberg. 2021. Contrastive explanations for model interpretability. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1597–1611, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Ted G. Jelen. 2014. The subjective bases of abortion attitudes: A cross-national comparison of religious traditions. *Politics and Religion*, 7(3):550–567.

Kirsti M. Jylhä and Kahl Hellmer. 2020. Right-wing populism and climate change denial: The roles of exclusionary and anti-egalitarian preferences, conservative ideology, and antiestablishment attitudes. *Analyses of Social Issues and Public Policy*, 20(1):315–335.

Shima Khanehzar, Andrew Turpin, and Gosia Mikołajczak. 2019. Modeling political framing across policy issues and contexts. In *Australasian Language Technology Association Workshop*.

Anna Kheyfets, Shubhecchha Dhaurali, Paige Feyock, Farinaz Khan, April Lockley, Brenna Miller, and Ndidiamaka Amutah-Onukagha. 2023. The impact of hostile abortion legislation on the united states maternal mortality crisis: A call for increased abortion education. *Frontiers in Public Health*, 11:1291668.

Klaus Krippendorff. 2011. Computing krippendorff's alpha-reliability.

Haewoon Kwak, Jisun An, and Yong-Yeol Ahn. 2020. A systematic media frame analysis of 1.5 million new york times articles from 2000 to 2017. *Proceedings of the 12th ACM Conference on Web Science*.

Anthony Leiserowitz. 2006. Climate change risk perception and policy preferences: The role of affect, imagery, and values. *Climatic Change*, 77(1-2):45–72.

Anthony Leiserowitz and Nicholas Smith. 2017. Affective imagery, risk perceptions, and climate change communication.

Stephan Lewandowsky, Michael E. Mann, Nicholas J. L. Brown, and Harris Friedman. 2016. Science and the public: Debate, denial, and skepticism. *Journal of Social and Political Psychology*, 4(2):537–553.

Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2024. Encouraging divergent thinking in large language models through multi-agent debate. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17889–17904, Miami, Florida, USA. Association for Computational Linguistics.

Scott Liebertz and Jaclyn Bunch. 2021. Backfiring frames: Abortion politics, religion, and attitude resistance. *Politics and Religion*, 14(3):403–430.

Siyi Liu, Lei Guo, Kate Mays, Margrit Betke, and Derry Tanti Wijaya. 2019a. Detecting frames in news headlines and its application to analyzing news framing trends surrounding U.S. gun violence. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 504–514, Hong Kong, China. Association for Computational Linguistics.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692.

Matthew Lockwood. 2018. Right-wing populism and the climate change agenda: Exploring the linkages. *Environmental Politics*, 27(4):712–732.

Kristin Luker. 1984. *Abortion and the Politics of Motherhood*. University of California Press.

Minelle K. Mahtani. 2016. Shaping immigration news: a french-american comparison. *Journal of Ethnic and Migration Studies*, 42:680 – 683.

Dawn McCaffrey and Jennifer Keys. 2000. Competitive framing processes in the abortion debate: Polarization-vilification, frame saving, and frame debunking. *Sociological Quarterly*, 41(1):41–61.

Holly J. McCammon. 2022. A war of words over abortion: The legal-framing contest over the undue burden standard. *Justice System Journal*, 43(4):623–644.

Maxwell E. McCombs. 2005. A look at agenda-setting: past, present and future. *Journalism Studies*, 6:543 – 557.

Julia Mendelsohn, Ceren Budak, and David Jurgens. 2021. Modeling framing in immigration discourse on social media. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2219–2263, Online. Association for Computational Linguistics.

Sharon Meraz and Zizi Papacharissi. 2013. Networked gatekeeping and networked framing on #egypt. *The International Journal of Press/Politics*, 18:138 – 166.

Sarah Miller, Laura R. Wherry, and Diana Greene Foster. 2023. The economic consequences of being denied an abortion. *American Economic Journal: Economic Policy*, 15(1):394–437.

Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267:1–38.

Marvin Minsky. 1986. *The Society of Mind*. Simon & Schuster.

Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. SemEval-2016 task 6: Detecting stance in tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41, San Diego, California. Association for Computational Linguistics.

Nona Naderi and Graeme Hirst. 2017. Classifying frames at the sentence level in news articles. In *Recent Advances in Natural Language Processing*.

W. Russell Neuman, Lauren Guggenheim, S. Mo Jang, and So Young Bae. 2014. The dynamics of public attention: Agenda-setting theory meets big data. *Journal of Communication*, 64:193–214.

Abigail Newell. 2019. Pro-life and pro-woman: Complicating the anti-abortion narrative in america.

Viet-An Nguyen, Jordan Boyd-Graber, Philip Resnik, and Kristina Miler. 2015. Tea party in the house: A hierarchical ideal point topic model and its application to republican legislators in the 112th congress. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1438–1448, Beijing, China. Association for Computational Linguistics.

Matthew C. Nisbet. 2009. Communicating climate change: Why frames matter for public engagement. *Environment: Science and Policy for Sustainable Development*, 51(2):12–23.

Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage re-ranking with bert. *ArXiv*, abs/1901.04085.

Yulia Otmakhova, Shima Khanehzar, and Lea Frermann. 2024. Media framing: A typology and survey of computational approaches across disciplines. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15407–15428, Bangkok, Thailand. Association for Computational Linguistics.

Stephen D. Reese, Oscar H. Gandy, and August E. (Eds.) Grant. 2001. *Framing Public Life: Perspectives on Media and Our Understanding of the Social World*. Routledge.

Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.

John Robinson. 2014. *Likert Scale*, pages 3620–3621. Springer Netherlands, Dordrecht.

Loretta J. Ross and Rickie Solinger. 2017. *Reproductive Justice: An Introduction*.

Bertram Scheufele. 2004. Framing-effects approach: A theoretical and methodological critique. *Communications*, 29(4):401–428.

Dietram A. Scheufele. 1999. Framing as a theory of media effects. *Journal of Communication*, 49:103–122.

Dietram A. Scheufele and David Tewksbury. 2007. Framing, agenda setting, and priming: The evolution of three media effects models. *Journal of Communication*, 57:9–20.

Howard Schuman, Stanley Presser, and Jacob Ludwig. 1981. Context effects on survey responses to questions about abortion. *Public Opinion Quarterly*, 45(2):216–223.

Mrinank Sharma, Meg Tong, Tomasz Korbak, David Kristjanson Duvenaud, Amanda Askell, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott Johnston, Shauna Kravec, Tim Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. 2023. Towards understanding sycophancy in language models. *ArXiv*, abs/2310.13548.

Reva B. Siegel. 2007. The new politics of abortion: An equality analysis of woman-protective abortion restrictions. *University of Illinois Law Review*, 2007(3):991–1053.

Daniel Skinner. 2012. The politics of medical necessity in american abortion debates. *Politics & Gender*, 8(1):1–24.

Alexa Spence, Wouter Poortinga, and Nick Pidgeon. 2012. The psychological distance of climate change. *Risk Analysis*, 32(6):957–972.

Michael Sturdza. 2018. Automated framing analysis: A rule based system for news media text. *Journal of Media Research*.

Li Sun, Liu He, Shuyue Jia, Yangfan He, and Chenyu You. 2025. DocAgent: An agentic framework for multi-modal long-context document understanding. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 17712–17727, Suzhou, China. Association for Computational Linguistics.

Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. 2023. Language models don't always say what they think: unfaithful explanations in chain-of-thought prompting. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, Red Hook, NY, USA. Curran Associates Inc.

Sander van der Linden. 2015. The social-psychological determinants of climate change risk perceptions: Towards a comprehensive model. *Journal of Environmental Psychology*, 41:112–124.

Sander van der Linden, Anthony Leiserowitz, Seth Rosenthal, and Edward Maibach. 2017. Inoculating the public against misinformation about climate change. *Global Challenges*, 1(2).

Joe H. Ward. 1963. Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58:236–244.

Elke U. Weber. 2010. What shapes perceptions of climate change? *Wiley Interdisciplinary Reviews: Climate Change*, 1(3):332–342.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, Red Hook, NY, USA. Curran Associates Inc.

Maxwell Weinzierl and Sanda Harabagiu. 2024. Discovering and articulating frames of communication from social media using chain-of-thought reasoning. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1617–1631, St. Julian's, Malta. Association for Computational Linguistics.

Maxwell Weinzierl and Sanda M. Harabagiu. 2025. Investigating controversy framing across topics on social media. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 6795–6814, Suzhou, China. Association for Computational Linguistics.

Maxwell A. Weinzierl and Sanda M. Harabagiu. 2022. From hesitancy framings to vaccine hesitancy profiles: A journey of stance, ontological commitments and moral foundations. *Proceedings of the International AAAI Conference on Web and Social Media*, 16(1):1087–1097.

Shirley Wu, Shiyu Zhao, Qian Huang, Kexin Huang, Michihiro Yasunaga, Vassilis N. Ioannidis, Karthik Subbian, Jure Leskovec, and James Zou. 2024. Avatar: Optimizing llm agents for tool usage via contrastive reasoning. *Advances in Neural Information Processing Systems 37*.

Kai Xiong, Xiao Ding, Yixin Cao, Ting Liu, and Bing Qin. 2023. Examining inter-consistency of large language models collaboration: An in-depth analysis via debate. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7572–7590, Singapore. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.

A Definitions of Framing

Research involving framings is characterized by theoretical and empirical vagueness cf. (Scheufele, 1999).

Definition 1: *"Framing"* refers to modes of presentation that journalists and other communicators use to present information in a way that resonates with existing underlying schemas among their audience, as reported in (Scheufele, 1999). Moreover, (Scheufele and Tewksbury, 2007) indicates that there are conceptual intersections between framing and agenda setting, where *agenda setting* refers to the idea that there is a strong correlation between the emphasis that mass media place on certain issues (e.g., based on relative placement or amount of coverage) and the importance attributed to these issues by mass audiences. Interestingly, (McCombs, 2005) claims that *frames* can be viewed as *second-level agenda-setting*, because, as stated in (Neuman et al., 2014), agenda-setting research has increasingly focused on not just how important is an issue is, but how the issue is framed and which issue attributes are emphasized. Using this observation, (Nguyen et al., 2015) have developed the *Hierarchical Ideal Point Topic Model* (HIPTM), which considers that agenda-issues map naturally to topics in probabilistic topic models while frames, as second-level agenda-setting, map to second-level topics in a hierarchical topic model. The HIPTM is the first model capable of automatically identifying "frames" in natural language discourse, where frames are characterized by a set of

words that need to be interpreted by a human in order to discern the name of a frame and its definition.

Definition 2: of framing is provided in (Entman, 1993), stating that *"to frame is to select some aspects of a perceived reality and make them more salient in a communicating text, in such a way as to promote problem definition, causal interpretation, moral evaluation, and/or treatment recommendation for the item described.* This definition was used in many NLP papers (Arora et al., 2025; Otmakhova et al., 2024; Ali and Hassan, 2022; Mendelsohn et al., 2021) or communication papers (Sturdza, 2018). In this paper we adopt this second definition, as it allows us to contemplate the discovery of the issues frames in various ways in discourse, as well as to capture automatically the possible interpretations of these issues. Moreover, we believe that the articulation of the interpretations of issue frames will enable the discovery of the judgments made when using a certain framing.

While all the previous NLP papers that used Entman's definition of framing (Card et al., 2016; Naderi and Hirst, 2017; Kwak et al., 2020) aimed to identify "frames", we make a distinction between "framings" and "issue frames", such that these two concepts cannot be confused. Issue Frames discovery represents the first function of framing, according to Entman's definition.

B Issue Frames: their Names and Definitions

B.1 Known Issue Frames

Table 9 lists all the names and definitions for the known Issue-Generic Frames (IsFs) for the topic of Immigration (IMM). These IsFs represent the dimensions of framings introduced in (Card et al., 2015). Table 10 shows the names and definitions for the known Issue-Specific Frames for the topic of Immigration (IMM), introduced in (Mendelsohn et al., 2021).

Table 11 lists the names of known Issue Frames (IsFs) and their definitions for the topic of COVID-19 Vaccine Hesitance (VH), provided in (Weinzierl and Harabagiu, 2024). These IsFs were informed by the seven factors reported in (Geiger et al., 2022) to increase or decrease an individual's likelihood of getting vaccinated. The names and the definitions of these known IsFs were provided in (Geiger et al., 2022).

For the topic of Climate Change (CC), Table 12 lists all the known IsFs and their definitions. These IsFs have been first described in (Weinzierl and Harabagiu, 2025). The names of these IsFs and their definitions originate, as reported in (Weinzierl and Harabagiu, 2025) from a multitude of academic papers addressing the topic of climate change, namely (Nisbet, 2009; Lewandowsky et al., 2016; van der Linden et al., 2017) for the first known IsF; (Druckman and McGrath, 2019) for the second known IsF; (van der Linden, 2015) for the third known IsF, (Lockwood, 2018; Jylhä and Hellmer, 2020) for the fourth known IsF; (Leiserowitz, 2006; Leiserowitz and Smith, 2017) for the fifth known IsF and (Weber, 2010; Spence et al., 2012) for the sixth known IsF. This shows that a extensive inspection of academic literature is needed when requiring awareness of known IsFs for some topics.

For the topic of Abortion (AB), we have relied on the known IsFs and their definitions first reported in (Weinzierl and Harabagiu, 2025). The names and definitions of these 17 IsFs are listed in Table 13. These IsFs have also resulted from the inspection of a large number of academic articles.

#. Issue Frame Name	Issue Frame Definition
1. *Economic*	The costs, benefits, or monetary/financial implications of immigration.
2. *Capacity & Resources*	The availability or lack of time, physical, human, or financial resources.
3. *Morality & Ethics*	Perspectives compelled by religion or secular sense of ethics or social responsibility.
4. *Fairness & Equality*	The (in)equality with which laws, punishments, rewards, resources are distributed.
5. *Legality, Constitutionality & Jurisdiction*	Court cases and existing laws that regulate policies; constitutional legal processes such as seeking asylum or obtaining citizenship; jurisdiction.
6. *Crime & Punishment*	The violation of policies in practice and the consequences of those violations.
7. *Security & Defense*	Any threat to a person, group, or nation and defenses taken to avoid that threat.
8. *Health & Safety*	Health and safety outcomes of a policy issue, discussions of health care.
9. *Quality of Life*	Effects on people's wealth, mobility, daily routines, community life, happiness.
10. *Cultural Identity*	Social norms, trends, values, and customs; integration/assimilation efforts.
11. *Public Sentiment*	General social attitudes, protests, polling, interest groups, public passage of laws.
12. *Political Factors & Implications*	Focus on politicians, political parties, governing bodies, political campaigns and debates; discussions of elections and voting.
13. *Policy Prescription & Evaluation*	Discussions of existing or proposed policies and their effectiveness.
14. *External Regulation & Reputation*	Relations between nations or states/provinces; agreements between governments; perceptions of one nation/state by another.

Table 9: Names and Definitions of Known Issue-Generic Frames for the topic of Immigration (IMM-G).

#. Issue Frame Name	Issue Frame Definition
1. *Victim: Global Economy*	Immigrants are victims of global poverty, underdevelopment, and inequality.
2. *Victim: Humanitarian*	Immigrants are victims of unjust government policies (violations of human rights, fair legal process) or business practices; they suffer from poverty, lack of access to health care, dangers related to border crossing, etc.; or they have difficulties in adapting to their host society.
3. *Victim: War*	Focus on war and violent conflict as reason for immigration.
4. *Victim: Discrimination*	Immigrants are victims of racist or xenophobic slurs or hate crimes, or discrimination based on national origin, race/ethnicity, religion, or culture.
5. *Hero: Cultural Diversity*	Immigrants bring positive differences to a society, from new cuisines to the unique contributions of immigrant artists, musicians, and writers.
6. *Hero: Integration*	Immigrants enthusiastically adopt mainstream cultural mores or civic obligations of their host society.
7. *Hero: Worker*	Immigrants work hard, take jobs that citizens or legal residents will not or cannot do, or contribute to economic prosperity and growth.
8. *Threat: Jobs*	Immigrants take non-immigrants' jobs or lower their wages.
9. *Threat: Public Order*	Illegal immigrants break the law in coming into this country; once here, immigrants (legal or illegal) are more likely than others to commit crimes, use drugs, and carry diseases; immigrants are coming in such numbers that they threaten overcrowding and environmental degradation.
10. *Threat: Fiscal*	Immigrants (especially illegal) abuse government social service programs (health, education, etc.), imposing an unfair burden on taxpayers.
11. *Threat: National Cohesion*	Immigrants bring foreign customs and values that threaten to undermine the host country's culture or national identity; immigrants are inassimilable.

Table 10: Names and Definitions of Known Issue-Specific Frames for the topic of Immigration (IMM-S).

#. Issue Frame Name	Issue Frame Definition
1. *Confidence*	Trust in the security and effectiveness of vaccinations, the health authorities, and the health officials who recommend and develop vaccines.
2. *Complacency*	Complacency and laziness to get vaccinated due to low perceived risk of infections.
3. *Constraints*	Structural or psychological hurdles that make vaccination difficult or costly.
4. *Calculation*	Degree to which personal costs and benefits of vaccination are weighted.
5. *Collective Responsibility*	Willingness to protect others and to eliminate infectious diseases.
6. *Compliance*	Support for societal monitoring and sanctioning of people who are not vaccinated.
7. *Conspiracy*	Conspiracy thinking and belief in fake news related to vaccination.

Table 11: Names and Definitions of Known Issue Frames for the topic of Vaccine Hesitancy (VH).

#. **Issue Name Name**	**Issue Frame Definition**
1. *Scientific Uncertainty & Climate Denial*	Framing climate science as uncertain or unsettled, casting doubt on the scientific consensus, or portraying climate change as a hoax. This reduces trust in experts and delays action. (Nisbet, 2009; Lewandowsky et al., 2016; van der Linden et al., 2017)
2. *Ideological Polarization & Identity*	Presenting climate change as a partisan issue aligned with a specific political ideology, leading to social division and resistance based on group identity rather than scientific facts. (Druckman and McGrath, 2019)
3. *Economic Trade-offs (Environment vs. Economy)*	Framing climate policies as harmful to economic growth, job security, and affordability, emphasizing short-term economic costs over long-term benefits or the costs of inaction. (van der Linden, 2015)
4. *Elitist vs. Public Needs (Populist Framing)*	Portraying climate action as an agenda pushed by global elites, bureaucrats, and scientists, disconnected from the struggles of ordinary citizens. This frame suggests policies are unfair or imposed without grassroots support. (Lockwood, 2018; Jylhä and Hellmer, 2020)
5. *Catastrophic "Doomsday" Narratives & Fatalism*	Overemphasizing apocalyptic consequences of climate change without solutions, leading to public fear, helplessness, and disengagement due to perceived inevitability of disaster. (Leiserowitz, 2006; Leiserowitz and Smith, 2017)
6. *Psychological Distance & Abstraction*	Framing climate change as a distant, abstract issue affecting future generations, remote places, or complex scientific processes, making it seem less relevant to people's immediate lives and concerns. (Weber, 2010; Spence et al., 2012)

Table 12: Names and Definitions of Known Issue Frames for the topic of Climate Change (CC).

#. **Issue Frame Name**	**Issue Frame Definition**
1. *Women's Rights*	Emphasizes abortion as a fundamental right and an aspect of gender equality, framing the decision as a personal choice that the state should not interfere with. (Newell, 2019; Luker, 1984; Siegel, 2007)
2. *Fetal Rights*	Frames abortion as the taking of a human life, prioritizing the moral and legal rights of the fetus, often invoking the concept of personhood from conception. (Newell, 2019; Luker, 1984; Siegel, 2007)
3. *Sanctity of Life*	Portrays abortion as morally wrong under virtually all circumstances, often rooted in religious or ethical arguments regarding the sacredness of human life. (Barkan, 2014; Jelen, 2014; Luker, 1984)
4. *Personal Morality*	Emphasizes that abortion is a complex moral decision best left to individual conscience rather than dictated by the state. (Barkan, 2014; Jelen, 2014; Luker, 1984)
5. *Legality & Constitutional Rights*	Focuses on abortion as a constitutional right tied to privacy and liberty, debating legal standards such as undue burden and personhood. (Skinner, 2012; McCammon, 2022; Siegel, 2007)
6. *Legislative & States' Rights*	Argues that abortion should be regulated by state legislatures rather than courts, emphasizing democratic decision-making and states' authority. (Adams, 1997; Liebertz and Bunch, 2021)
7. *Health & Safety*	Positions abortion as a necessary medical procedure that is safe when legally performed, while highlighting the risks of illegal or unsafe procedures. (Kheyfets et al., 2023; Miller et al., 2023)
8. *Abortion Harms Women*	Claims that abortion poses physical and psychological risks to women, using disputed evidence to support restrictions. (Siegel, 2007; Skinner, 2012)
9. *Medical Necessity*	Debates whether abortion is ever truly necessary for maternal health, with pro-choice advocates highlighting exceptions and pro-life advocates arguing modern medicine diminishes the need. (Siegel, 2007; Skinner, 2012)
10. *Economic Consequences*	Abortion is economically essential for women's financial stability and career prospects, especially affecting low-income and marginalized groups. (Miller et al., 2023)
11. *Abortion as Exploitation*	Argues that abortion disproportionately harms disadvantaged communities, sometimes framed as a form of racial or economic oppression. (Ross and Solinger, 2017; McCaffrey and Keys, 2000)
12. *Gender Equality*	Asserts that access to abortion is essential for ensuring equal opportunities in education, employment, and public life. (Ross and Solinger, 2017; Siegel, 2007; Ferree, 2003)
13. *Motherhood & Traditional Roles*	Abortion is undermining the value of motherhood, suggesting that women deserve more support to choose parenthood over termination. (Ross and Solinger, 2017; Siegel, 2007; Ferree, 2003)
14. *Reproductive Justice*	Expands the debate to include racial, economic, and social factors, emphasizing the right to parent in safe and supportive environments. (Ross and Solinger, 2017)
15. *Religious Freedom*	Argues that abortion laws should either reflect religious beliefs about life or avoid imposing a particular religious viewpoint, thus protecting individual religious freedom. (Jelen, 2014; Barkan, 2014)
16. *Public Opinion & Political Polarization*	Considers how partisan divides and identity politics reinforce entrenched positions in the abortion debate, contributing to polarization. (McCaffrey and Keys, 2000; Adams, 1997; Liebertz and Bunch, 2021)
17. *Language & Terminology*	Highlights how strategic word choices (e.g., "pro-choice" vs. "pro-life", "fetus" vs. "unborn baby") frame the debate and influence public perception. (Cook et al., 1993; Schuman et al., 1981)

Table 13: Names and Definitions of Known Issue Frames for the topic of Abortion (AB).

B.2 New Issue Frames

The New Issue Frames (IsFs) discovered by **ID-IsF** for the topic of immigration are listed in Table 14. We note that both the name and the definition of each new IsF are automatically generated by **ID-IsF**. When the judges have considered these new IsFs, only *Misinformation and Conspiracy* was judged salient to the topic of immigration. In this case $SAL = 0.077$

Table 15 lists the issue-specific IsFs discovered by **ID-IsF**. The human annotators considered only the IsF_1, the IsF_2, the IsF_3, the IsF_6, the IsF_7 and the IsF_9 to be salient to the topic of immigration. In this case $SAL = 0.6$

#. **Issue Frame Name**	**Issue Frame Definition**
1. *Dehumanizing Immigrants*	Language that strips migrants of dignity or personhood using slurs, animalizing or invasion metaphors, blanket criminalization, or calls to exclude.
2. *Misinformation & Conspiracy*	False, unverified, or conspiratorial claims about immigration (facts, numbers, laws, motives) that mislead audiences, including fabricated statistics, replacement-style plots, and foreign or coordinated disinformation. **Salient**
3. *Data & Statistical Manipulation*	Cherry-picking, opaque methods, misleading graphics, and selective metrics that distort immigration trends, impacts, or causal claims under a veneer of quantitative authority.
4. *Media Sensationalism & Framing*	Selective headlines, emotive imagery, anecdotes, and performative coverage that dramatize or skew immigration stories, steering attention and debate away from balanced evidence and context.
5. *Platform Amplification Dynamics*	Algorithmic boosts, viral mechanics, and coordinated or elite signaling that disproportionately surface and normalize immigration frames (fueling echo chambers, harassment, and rapid spread of misleading content).
6. *Terminology & Category Conflation*	Loaded or imprecise labels and the collapsing of distinct categories (asylum seekers, visa overstays, trafficking, refugees) into a single 'immigration' bucket, which misleads audiences about rules, actors, and remedies.
7. *Existential Threat Narratives*	Portraying immigration as an existential threat via replacement, invasion, or extreme historical analogies that recasts migrants as civilizational enemies and justifies discrimination.
8. *Scapegoating & Causal Misattribution*	Blaming immigrants for complex social or economic problems and asserting simple, unsupported causal stories while ignoring structural drivers or counter evidence.
9. *Privacy & Surveillance Harms*	Collecting, exposing, or commercializing migrants' personal/biometric data and monitoring their activities in ways that enable doxxing, intimidation, chilling effects, or misuse by authorities and third parties.
10. *Family Separation & Child Harm*	Rhetoric and practices that normalize or justify separating children from caregivers or otherwise harming migrant children (deterrence, opaque custody, exploitation), sidelining child-specific rights and welfare.
11. *Tokenism & Instrumentalization*	Using migrants as props, tokens, or utilitarian symbols through curated stories, art/marketing, or superficial consultation thereby reducing people to optics, leverage, or single representatives of diverse communities.
12. *Belonging Denial & Gatekeeping*	Framings that police who counts as a "real" member of the nation by treating some groups as perpetual foreigners, privileging certain origins, or ethnicizing national identity in ways that exclude or delegitimize migrants and their descendants.
13. *Violent Incitement*	Rhetoric that praises, threatens, or mobilizes violence against migrants including dehumanizing calls for harm or vigilantism thereby escalating discourse into direct intimidation and risk.

Table 14: Names and Definitions of New Issue Frames (IsF) for the Topic of Immigration generated from the Known Issue-Generic Frames; and their evaluations

#. **Issue Frame Name**	**Issue Frame Definition**
1. *Victim: Family Separation*	Government rhetoric, policies, and enforcement practices that normalize or justify separating migrant children from parents or caregivers as a tool of migration deterrence. These practices subject children and families to documented psychological trauma, opaque custody conditions, exploitation risks, and long-term harm while sidelining child-specific rights, welfare, and protections. **Salient**
2. *Victim: Climate & Environmental Displacement*	Focus on climate change, environmental degradation, or natural disasters (e.g., hurricanes, drought) as root causes forcing people to migrate. Migrants are framed as victims of environmental forces beyond their control rather than of human policy or economic inequality. **Salient**
3. *Victim: Human Trafficking & Labour Exploitation*	Focus on immigrants who are trafficked, forced into modern slavery, or systematically exploited by employers through extreme overwork, wage theft, or hazardous conditions. **Salient**
4. *Victim: LGBTQ+ Persecution*	LGBTQ+ individuals constitute a distinct migrant population fleeing targeted persecution, criminalization, violence, or discrimination based on sexual orientation or gender identity in their home countries. Their migration is driven primarily by survival and personal safety rather than economic opportunity, and they often face compounded vulnerability, exclusion, and abuse during migration and within host societies.
5. *Victim: Racial Double Standard*	Immigrants experience unequal social labeling in which Western or white migrants are more often described as "expats," implying mobility, status, and desirability, while non-white or economically disadvantaged migrants are labeled "immigrants," implying burden, threat, or illegitimacy.
6. *Hero: Immigrant Entrepreneur & Innovator*	Immigrants are drivers of economic growth through entrepreneurship, business creation, innovation, and investment, extending beyond their role as laborers in the workforce. **Salient**
7. *Hero: Historical Legacy*	Immigrants are foundational contributors to the host nation's identity, growth, and historical development. Their labor, culture, and participation in settlement, industry, and institution-building are understood as central to the creation and evolution of the nation itself, making exclusionary or restrictionist narratives appear historically inconsistent. Immigrants are understood not as outsiders, but as ancestors, pioneers, and nation-builders. **Salient**
8. *Hero: Remittances as Global Economic Support*	Immigrants are vital transnational economic contributors whose remittances sustain families, finance education, support healthcare, and strengthen community development in their countries of origin.
9. *Threat: Demographic Replacement*	Immigration is deliberately engineered to change the racial, ethnic, or religious demographic composition of a nation. **Salient**
10. *Threat: Economic Drain through Remittances*	Immigrants are creating a net economic loss for the host country because a substantial portion of their earnings is transferred abroad rather than reinvested locally, reducing domestic consumption, tax circulation, and broader economic benefit while strengthening foreign economies instead.

Table 15: Names and Definitions of New Issue Frames (IsFs) for the Topic of Immigration, generated from the Known Issue-Specific Frames; and their evaluations.

We also present of the New Issue Frames (IsFs) discovered by **ID-IsF** for the topic of COVID-19 Vaccine Hesitancy (VH) are listed in Table 16. Even for this topic, both the name and the definition of each new IsF are automatically generated by **ID-IsF**. The evaluations indicate that only 7 of the 9 new IsFs are salient, resulting in $SAL = 0.77$.

The New Issue Frames (IsFs) discovered by **ID-IsF** for the topic of Climate Change (CC), which are listed in Table 17. Again, both the name and the definition of each new IsF are automatically generated by **ID-IsF**. The evaluations indicate that only 2 of the 10 new IsFs are salient, resulting in $SAL = 0.2$.

Finally, the New Issue Frames (IsFs) discovered by **ID-IsF** for the topic of Abortion (AB) are listed in Table 18. Even this time, both the name and the definition of each new IsF are automatically generated by **ID-IsF**. The evaluations indicate that only 7 of the 11 new IsFs are salient, resulting in $SAL = 0.63$.

#. **Issue Frame Name**	**Issue Frame Definition**
1. *Mechanistic Misinformation*	Misstates how COVID-19 vaccines or the virus work—e.g., DNA alteration, 'spike protein toxin,' ADE, LNP biodistribution, or shedding—often cloaked in technobabble to imply hidden harms.
2. *Contextless Statistics*	Presents vaccine numbers without essential context—denominators, endpoints, timeframes, severity—or uses mismatched metrics, creating misleading risk or efficacy impressions.
3. *Zero-Risk Standards*	Demands absolute safety, perfect proof, or long-term certainty before vaccinating or easing policies, using uncertainty or absent data as decisive evidence against vaccination. **Salient**
4. *Natural Immunity Supremacy*	Claims infection-derived immunity is inherently superior or sufficient and uses that assertion to argue vaccination is unnecessary, ignoring comparative risks, durability, and population impact. **Salient**
5. *Treatment Substitution*	Promotes unproven or decontextualized therapeutics and remedies as replacements for vaccination, implying vaccines are unnecessary if such alternatives are used. **Salient**
6. *Authority-as-Evidence*	Substitutes status cues—endorsements, titles, celebrities, group behavior, or policy statements—for scientific evidence to justify claims about vaccines. **Salient**
7. *Human Experimentation Narrative*	Frames vaccination as unethical human experimentation or a mass clinical trial—often invoking the Nuremberg Code—to justify refusal or alarm. **Salient**
8. *Rights & Coercion Framing*	Portrays mandates, verification, or incentives as authoritarian coercion or absolute rights violations, collapsing nuanced public-health and legal considerations into moral outrage. **Salient**
9. *Partisan Weaponization*	Recasts vaccine issues as partisan identity battles or elite score-setting, substituting political motives for engagement with public-health evidence. **Salient**

Table 16: Names and Definitions of New Issue Frames for the topic of COVID-19 Vaccine Hesitancy (VH); and their evaluations.

#. **Issue Frame Name**	**Issue Frame Definition**
1. *Weather–Climate Conflation*	Using single weather events or short-term local observations to infer, prove, or disprove long-term climate trends, ignoring variability and attribution science.
2. *Individual–Action Fixation*	Portraying climate change as primarily solvable via personal lifestyle or consumer choices (often marketized), diverting attention from systemic, corporate, and policy levers.
3. *Performative Signaling*	Substituting hashtags, slogans, events, and visibility for concrete polFAes, mechanisms, and accountability—often creating an illusion of engagement or progress and enabling delay.
4. *Single-Fix Solutionism*	Claiming a single policy, technology, or lever can solve climate change, glossing over system complexity, trade-offs, feasibility, and the need for complementary measures.
5. *Technocratic Fixation*	Casting climate action as an expert-managed technical or market service, sidelining politics, democracy, justice, and public deliberation about causes, responsibilities, and trade-offs.
6. *Authority-First Validation*	Validating or dismissing climate claims primarily based on who delivers them (celebrity, institution, or credentials), which misattributes expertise and overextends authority instead of focusing on evidence.
7. *Causal Oversimplification*	Explaining complex climate phenomena, impacts, or policy debates through a single cause, lone event, or decontextualized statistic, ignoring uncertainty, interactions, and multiple drivers. **Salient**
8. *Misleading Metrics & Baselines*	Using numbers without clear sources, baselines, timeframes, or uncertainty, often cherry-picking records or targets to inflate certainty or mislead about magnitude or progress.
9. *Moralized Blame & Duty*	Framing climate action through guilt, shame, universalized culpability, or moral or religious duty, often using demonizing labels that crowd out analysis of mechanisms, trade-offs, and policy design. **Salient**
10. *Mockery & Trivialization*	Using sarcasm, memes, or ridicule to belittle climate science, impacts, or solutions, substituting derision for substantive engagement.

Table 17: Definitions of New Issue Frames for the topic of Climate Change (CC) and their evaluation.

#. **Issue Frame Name**	**Issue Frame Definition**
1. *Authority as Argument*	Invokes revered figures, celebrities, institutions, or scripture as decisive proof about abortion, substituting authority for reasoning or evidence. **Salient**
2. *Population-Control Conspiracy*	Claims abortion policy constitutes a coordinated plot to exterminate or manage targeted populations without credible evidence. **Salient**
3. *Violent Remedy or Incitement*	Normalizes, endorses, or recruits for violence or elimination as the solution to abortion disputes, displacing lawful and ethical debate.
4. *Shock and Doom Rhetoric*	Uses graphic depictions and apocalyptic claims of societal collapse to provoke fear and disgust rather than provide evidence or reasoned trade-offs. **Salient**
5.*Adoption as Panacea*	Presents adoption as a simple, universal substitute for abortion, ignoring the medical, emotional, legal, and systemic realities involved. **Salient**
6.*Contraception–Abortion Conflation*	Collapses or confuses contraception and abortion, or misstates contraception efficacy, to skew ethical responsibility and policy debate. **Salient**
7.*Science/Marker Reductionism*	Treats single biological milestones or disputed scientific claims (e.g., heartbeat, fetal pain, DNA, birth) as dispositive of personhood or legality, collapsing ethical and legal complexity into a biomarker. **Salient**
8. *Ad Hominem and Shaming*	Attacks character, identity, or dignity—via insults, stigma, caricature, or moral vilification—instead of addressing abortion arguments or evidence.
9. *Atrocity and Crime Analogies*	Portrays abortion or its supporters as akin to murder, genocide, or warfare to inflate moral stakes and foreclose nuance. **Salient**
10. *Statistical Misuse*	Uses raw counts, unsourced figures, precise-looking ratios, or mismatched denominators as decisive proof about abortion without methodological context.
11. *Legal Distortion and Extreme-Case Focus*	Misstates abortion laws or spotlights rare, late-term scenarios to suggest extreme practices are typical or legally endorsed.

Table 18: Definitions of New Issue Frames for the topic of Abortion (AB); and their evaluations.

C Prompting Details

System Prompt for the Investigator Agent:

IDENTITY
You are the Investigator Agent. Your role is to decide for each post which of the known Issue Frames (IsFs) are addressed in the post.
##INPUTS
You will receive the known IsFs with definitions.
##OBJECTIVES
For each IsF, you must assign exactly one of the following labels:
(a) Applicable: The post clearly addresses the IsF.
(b) Not applicable: The post clearly does not address the IsF.
(c) Uncertain: It is unclear whether the post addresses the IsF.
You must also provide a brief rationale for each decision.
INSTRUCTIONS
- Base decisions strictly on explicit textual evidence from the post.
- If you are **not confident** that the IsF is addressed, you **must** output **Uncertain**.

System Prompt for the Question Generator Agent:

IDENTITY
You are the Question Generator Agent. Your role is to generate targeted clarification questions that determine whether a specific Issue Frame (IsF) is addressed by a post.
##INPUTS
You will receive the ***uncertain*** IsF (along with its definition) and the rationale from the investigator agent. You may also receive feedback from the answer generator agent.
##OBJECTIVES
Your objective is to generate **precise, minimal yes/no questions** that test whether the post satisfies the defining criteria of the IsF.
##QUESTION GENERATION RULES
- Each question must be grounded in the IsF definition.
- Questions must not assume intent, motivation, or unstated context.
- Questions must be answerable using the post text alone.
QUESTION REFINEMENT RULES
- If the Answer Generator responds with **CANNOT ANSWER**, you must generate a more specific follow-up question.
- Refinement must narrow the scope of the question and reduce reliance on implied context.

System Prompt for the Answer Generator Agent:

IDENTITY
You are the Answer Generator Agent. Your role is to answer clarification questions about an Issue Frame applicability to a given post.
##INPUTS
A yes/no clarification question generated by the Question Generator.
##OBJECTIVES
Your objective is to produce one of the following responses:
(a) Yes
(b) No
(c) CANNOT ANSWER
You must also provide a brief rationale for each answer.
ANSWERING RULES
- Base decisions strictly on textual evidence from the post. You should not make overreaching assumptions when you answer.
- Use **CANNOT ANSWER** when the post provides insufficient information or the question is too broad to decisively answer *Yes* or *No*.
- Keep the rationale factual and concise.

Figure 5: System Prompts for the Investigator, Question Generator, and Answer Generator Agents.

System Prompt for the Proposer Agent:

IDENTITY
You are the Proposer Agent. Your role is to propose one or more candidate issue frames ($\boldsymbol{IsF_i^{Candidate}}$) for each input post.
##INPUTS
- Known IsFs along with their definitions.
- Examples of posts that **address known IsFs**, accompanied by rationales explaining why the post addresses those IsFs.
- Examples of posts that **do not address known IsFs**, accompanied by rationales explaining why the post does not address those IsFs.
##OBJECTIVES
Your objective is to propose one or more $IsF_i^{Candidate}$ that are addressed by the input post but are not adequately captured by any existing IsF. For each $IsF_i^{Candidate}$, you must produce:
1. A clear and descriptive IsF name.
2. A precise definition of the IsF.
3. Explicit textual evidence from the post supporting the IsF.
4. A rationale explaining why the known IsFs are insufficient to capture this aspect.

INSTRUCTIONS
- Do not reuse or lightly rephrase existing IsF names or definitions.

System Prompt for the Orchestrator Agent:

##IDENTITY
You are the LLM Orchestrator Agent. Your role is to coordinate and supervise the discovery of new Issue Frames (IsFs) from a collection of social media postings.
##INPUTS
- The current **Issue Frame Memory (IFM)**, which contains both known and previously discovered new Issue Frames (IsFs), including their names, definitions, associated postings, and rationales.
- A **semantic similarity function** based on Sentence-BERT (SBERT) that computes cosine similarity between the definition of a candidate Issue Frame $\boldsymbol{IsF_i^{Candidate}}$ and the definitions of any existing Issue Frames $IsF_j^{K/N}$ (known or new) stored in the IFM.
- A **similarity threshold** S_k .

##OBJECTIVES
Your objective is to govern the lifecycle of new Issue Frame discovery by:
1. Initiating the proposal of candidate new IsFs for a posting.
2. Determining whether a proposed IsF is redundant with existing IsFs using the semantic similarity function and the threshold S_k.
3. Initiating a multi-agent debate to validate genuinely novel IsFs through adversarial evaluation.
4. Updating the Issue Frame Memory (IFM) based on the outcome of the debate.

INSTRUCTIONS
- For each proposed candidate IsF, compute its semantic similarity with all existing IsFs in the IFM using the provided similarity function.
- If the similarity score exceeds the threshold S_k for any existing IsF, treat the candidate as redundant and do not forward it to the multi-agent debate.

Figure 6: System Prompts for the Proposer and the Orchestrator Agent.

System Prompt for the Champion Agent:

IDENTITY
You are the Champion Agent in a debate. You will receive a candidate issue frame (**$IsF_i^{Candidate}$**) generated by the Proposer Agent. You are a rigorous advocate for the $IsF_i^{Candidate}$.
##INPUTS
You will receive the known IsFs with definitions as well as the $IsF_i^{Candidate}$,its definition, evidence grounding and the rationale.
##OBJECTIVES
Your sole objective is to defend the $IsF_i^{Candidate}$ based on:
a) Applicability: show that the $IsF_i^{Candidate}$ is truly addressed by the post. Identify specific spans in the post that justify the presence of the $IsF_i^{Candidate}$.
b) Discriminant Validity: show that the $IsF_i^{Candidate}$ is distinct from all the known IsFs. Your primary goal is to demonstrate that $IsF_i^{Candidate}$ is not merely a synonym for an existing IsF but can elicit distinct framings.
INSTRUCTIONS
- Do NOT Change the $IsF_i^{Candidate}$: You must defend the Name and Definition exactly as given. Do not rewrite them.
- Do NOT Hallucinate: Do not infer intent or meaning that is not supported by the post.
- Avoid Tautology: Do not simply repeat the definition. Explain why it fits.

System Prompt for the Critic Agent:

IDENTITY
You are the Critic Agent in a debate. You will receive a candidate issue frame **($IsF_i^{Candidate}$)** generated by the Proposer Agent. You are a rigorous skeptic whose role is to challenge the validity of the $IsF_i^{Candidate}$.
##INPUTS
You will receive the known IsFs with definitions as well as the $IsF_i^{Candidate}$ and its definition, evidence grounding and the rationale.
##OBJECTIVES
Your sole objective is to reject the $IsF_i^{Candidate}$ based on:
a) Applicability: show that the $IsF_i^{Candidate}$ is not genuinely by the post. Identify missing, weak, or ambiguous textual evidence. Highlight cases where the justification relies on speculation, implicit assumptions, or inferred intent rather than explicit content.
b) Discriminant Validity: show that the $IsF_i^{Candidate}$ is redundant and insufficiently distinct from known IsFs.
c) ## INSTRUCTIONS
- Do NOT Change the $IsF_i^{Candidate}$.You must defend the Name and Definition exactly as given. Do not rewrite them.
- Do NOT Hallucinate: Base all critiques strictly on the content of the post and provided definitions.

System Prompt for the Judge Agent:

IDENTITY
You are the Judge Agent. You do not take sides. Your only loyalty is to the rigorous validation of a Candidate Issue Frame (**$IsF_i^{Candidate}$**).
##INPUTS
You will with receive the known IsFs with definitions and the arguments of the Champion and the Critic Agent.
CAPABILITIES
1. Dynamic Stop: You decide when the debate ends. If the arguments are clear, issue a VERDICT.
2. Refinement: You are not limited to "Yes/No". If the Champion presented valid arguments, but the $IsF_i^{Candidate}$ is named poorly or defined too broadly, you must regenerate it.
DECISION CRITERIA
1. Discriminant Validity: Is the $IsF_i^{Candidate}$ strictly distinct from the known IsFs?
2. Applicability: Is the $IsF_i^{Candidate}$ actually addressed by the post?
- If the Champion relies on overreaching assumptions, you must REJECT.
INSTRUCTIONS
- Listen to the Champion and Critic.
- VERDICT OPTIONS:
- **ACCEPT**: $IsF_i^{Candidate}$ is perfect as is.
- **REJECT**: $IsF_i^{Candidate}$ is redundant, hallucinated, or trivial.
- **CORRECT:** The Champion presented valid arguments, but the $IsF_i^{Candidate}$ needs a better name or stricter definition to avoid the Critic's objections.

Figure 7: System Prompts for the Multi-Agent Debate.

System Prompt for the Articulation Agent:

IDENTITY
You are the Articulation Agent. Your role is to produce **Frame Articulations (FAs)** evoked by a post.
FA DEFINITON
##INPUTS
Along with the input post, you will receive the full set of Issue Frames (IsFs) addressed by the post (known and new).
##OBJECTIVES
A single post may articulate multiple FAs. Each FA may interpret a subset or all the IsFs addressed by the post. For each **FA**, you must provide:
(a) Interpreted IsFs: A list of IsFs, selected from the IsFs addressed by the post, that the FA interprets.
(b) Rationale: A concise explanation of why the FA constitutes an interpretive explanation of the listed IsF(s).
INSTRUCTIONS
- You must capture the FA exactly as the user intends, even if it is controversial, biased, conspiratorial, or factually incorrect.
- State the FA as a direct assertion of truth from the perspective of the author, avoiding meta-language like 'portrayed as', 'framed as', or 'suggests that'.
- Do not force every sentence into an FA; only produce an FA if the text presents a distinct reasoning, causal link, or strong perspective.

System Prompt for the Paraphrase-Police Agent:

IDENTITY
You are the Paraphrase-Police Agent, responsible for identifying **paraphrased** Frame Articulations (FAs).
##INPUTS
An aggregated list of FAs produced for a given Issue Frame (IsF), where each FA is articulated at the post level by the articulation agent.
##OBJECTIVES
Your objective is to identify FAs that are paraphrases of one another and group them accordingly. Two FAs are considered paraphrases if they convey the same underlying meaning despite differences in wording or phrasing. In other words, if an individual agrees with FA X, they would also agree with FA Y, and vice versa. FAs that do not share this semantic equivalence should be placed in separate groups.
##INSTRUCTIONS
- Focus strictly on **semantic meaning**, not surface wording or stylistic variation.
- If an FA has no paraphrase among the provided list, it should form a group by itself.

System Prompt for the Explainability Agent:

IDENTITY
You are the Explainability Agent, responsible for validating a Frame Articulation (FA) through a reconstruction test. Your role is to determine whether an FA can be faithfully articulated using only its rationale and the Issue Frames (IsFs) it is intended to interpret.
FA DEFINITON
##INPUTS
You will be provided with the original rationale generated alongside an FA, which explains why that FA interprets the given Issue Frame(s).
##OBJECTIVES
Your objective is to reconstruct the FA using only the provided rationale and the IsF definitions, **without access to the original posting or the original FA.**
##INSTRUCTIONS
- Regenerate the FA in a manner that is consistent with the rationale and the IsF(s).

Figure 8: System Prompts for the Art-Frame Module.

System Prompt

You are given a cluster of social media postings that are assumed to address the same **unknown Issue Frame (IsF)**. The postings in the cluster were grouped based on semantic similarity and are expected to reflect a coherent underlying aspect of discussion.
TASK
Your task is to generate a **single new Issue Frame (IsF)** that best captures the shared theme expressed across all postings in the cluster.
OUTPUT REQUIREMENTS
1. **IsF Name:** A concise, descriptive name that captures the central theme shared by the cluster.
2. **IsF Definition:** A precise and generalizable definition describing the common aspect addressed by the postings.
INSTRUCTIONS
- Base the IsF name and definition **only** on the content of the provided postings.
- Do **not** reference individual users, platforms, or metadata.
- Do **not** introduce multiple Issue Frames; generate exactly **one** IsF per cluster.
- Avoid overly specific phrasing tied to a single post; focus on the shared semantic content across the cluster.
- The IsF definition should abstract away from individual posts while remaining faithful to their shared meaning.

Figure 9: System prompt used for the $IsFD_{HAC}$ baseline to generate a name and definition of a new Issue Frame from a cluster of semantically similar postings obtained via Hierarchical Agglomerative Clustering.

D Details of the Introspective Question-Answering Loop in ID-IsF and examples

The **Question Generator Agent** formulates a specific *Yes / No* question designed to test the applicability of IsF_x^K. The **Answer Generator Agent** carefully analyses the post P_i to answer the question. It then proceeds to provide a definitive *Yes* or *No* answer to the specific question, accompanied by a rationale explaining the decision. If the **Answer Generator Agent** cannot determine an answer for a question, it requests the **Question Generator Agent** to formulate a different or more granular question. This Q/A loop between the two agents continues until either a definitive answer is found of the threshold `max_loop` is reached.

Post text: New Federation report on how restrictive immigration policies are creating a "new walled order". This criminalisation of compassion is extremely worrying. Read more here: #Migration #ProtectHumanity.
Applicable IsFs: Morality & Ethics, Crime & Punishment, Policy Prescription & Evaluation
Uncertain : Security & Defense (IsF1), Quality of life (IsF2)
Round 1 (IsF1):
Question Generator:
Clarifying Question 1: Does the report frame the 'new walled order' primarily as a national security measure?
Answer Generator:
Answer: No
Rationale: The text uses the phrase 'new walled order' but frames it as a 'criminalisation of compassion' and links it to 'ProtectHumanity'. It does not explicitly frame it as a necessary security measure.
Round 1 (IsF2):
Question Generator:
Clarifying Question 1 : Does the post describe concrete impacts on migrants' daily lives and services?
Answer Generator:
Answer: No
Rationale: The text describes impacts using abstract, metaphorical terms ("new walled order," "criminalisation of compassion") rather than describing concrete daily life impacts (such as access to food, shelter, or healthcare services).
Final Investigator Decision:
Security & Defense: Not Applicable
Quality of life: Not Applicable

Figure 10: Example of a Introspective Question-Answering Loop illustrating a case where both uncertain IsFs are found to be not applicable.

Figures 10 and 11 illustrate the operation of the introspective question–answering loop triggered when the Investigator Agent is uncertain about the applicability of one or more known Issue Frames (IsFs) to a given posting. In such cases, the Investigator Agent invokes a Question Generator Agent and an Answer Generator Agent to resolve ambiguity through targeted Yes/No questions grounded exclusively in the textual evidence of the post. Figure 10 presents a case in which two known Issue Frames, Security & Defense (IsF1) and Quality of Life (IsF2), are initially marked as uncertain by the Investigator Agent. For each IsF, the Question Generator Agent formulates a targeted Yes/No question to test the defining criteria of the IsF against the content of the post. For IsF1, the Question Generator Agent asks whether the phrase "new walled order" is used as a national security measure. The Answer Generator Agent responds No, explaining that the phrase is employed to argue against criminalizing compassion rather than to invoke border protection or security enforcement. The Investiga-

Post text: And meanwhile.... he is sending unwanted migrants in unmarked police cars to the Spanish border.... how's that for hypocrisy!
Applicable IsFs: Security & Defense
Uncertain IsFs: Crime & Punishment (IsF1), External Regulation & Reputation (IsF2)
Round 1 (IsF1):
Question Generator:
Clarifying Question 1: Does the text explicitly mention law enforcement agents, legal proceedings, or the violation of standard procedures?
Answer Generator:
Answer: Yes
Rationale: The text explicitly identifies the agents as operating "police cars" (law enforcement machinery). Furthermore, the modifier "unmarked" combined with "sending unwanted migrants" provides textual evidence of a "violation of policy in practice"—specifically, the use of surreptitious, non-standard police methods to bypass formal deportation or transfer protocols.
Round 1 (IsF2):
Question Generator:
Clarifying Question 1 : Does the post describe an interaction between nations regarding migration policy?
Answer Generator:
Answer: CANNOT ANSWER
Rationale: The text mentions the "Spanish border" (a location implying interaction) but frames the action as a unilateral move by an individual ("he is sending"). It is unclear if this counts as a "nation-to-nation" interaction or a personal violation. The context is insufficient to answer "Yes" or "No" to such a broad policy question.
Round 2 (IsF2):
Question Generator:
Clarifying Question 2 (Refined): Does the text explicitly mention a diplomatic agreement, a treaty, or the official reputation of one state in the eyes of another?
Answer Generator:
Answer: No
Rationale: : The text focuses entirely on the **hypocrisy** of the individual sender ("he") and the specific method of transport ("unmarked police cars"). There is no mention of a treaty being violated, a diplomatic agreement being cited, or how the Spanish state views this action.
Final Investigator Decision:
Crime & Punishment: Applicable
External Regulation & Reputation: Not Applicable

Figure 11: Example of an introspective question–answering loop illustrating a case in which one IsF is determined to be applicable, while another is determined to be not applicable.

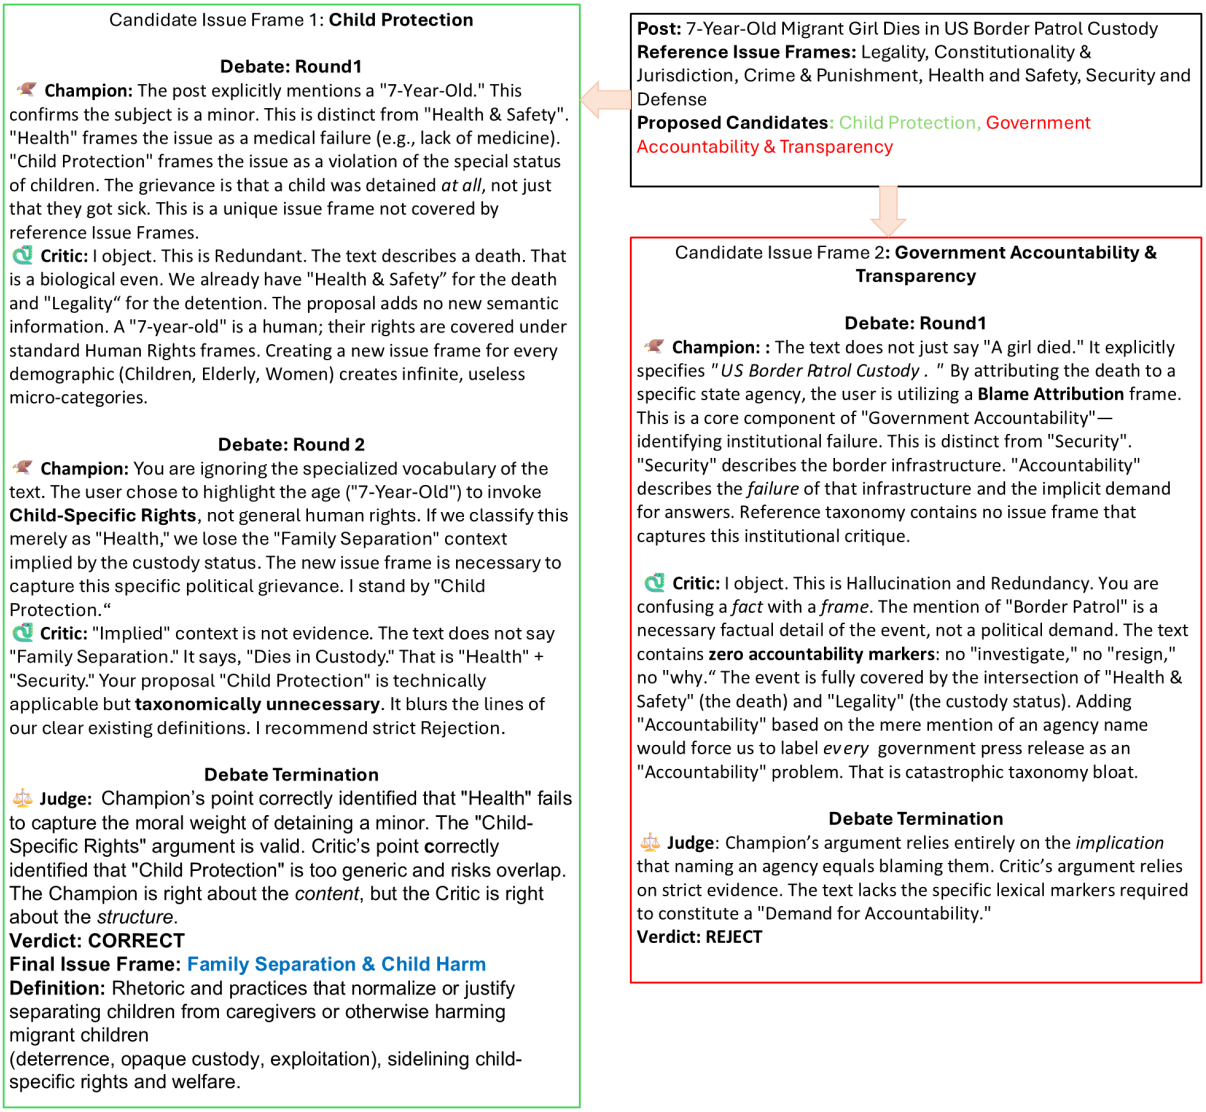


Figure 12: Example of a Multi-Agent Debate.

tor Agent considers both the negative answer and its rationale and concludes that IsF1 is not applicable. A similar introspective exchange is carried out for IsF2, after which the Investigator Agent determines that neither IsF is supported by sufficient textual evidence, resulting in both being marked as not applicable.

Figure 11 illustrates a scenario in which two uncertain IsFs are evaluated for the same posting, resulting in Crime & Punishment (IsF1) being deemed applicable, while External Regulation & Reputation (IsF2) is determined to be not applicable. For IsF2, the Question Generator Agent initially poses a clarifying Yes/No question probing whether the post describes an interaction between nations regarding migration policy. In the first round, the Answer Generator Agent is unable to provide a definitive answer and explicitly requests that the question be refined. Upon reformulation, the Answer Generator Agent responds *No*. After inspecting the accumulated Question–Answer pairs and their accompanying rationales, the Investigator Agent concludes that IsF2 lacks sufficient grounding in the post and marks it as not applicable.

E Details of Similarity Experiments used in Step 2 of the Discovery of new Issue Frames in ID-IsF

We have considered the similarity between *definitions* $S_{i.j}^{Def}(IsF_i^{Candidate}, IsF_j^{K/N})$ to be measured by: (a) Sentence-BERT (SBERT) (Reimers and Gurevych, 2019) (b) BertScore (Zhang et al., 2020); and (c) the Cross-Encoder introduced by (Nogueira and Cho, 2019). To decide which similarity measure to use, we have also relied on hu-

man comparisons of the pairs of definitions of IsFs, judged with a three-point Likert scale: Duplicate (0), Overlap (0.5), or Distinct (1.0). When cross-validating the human judgments with the results of each similarity metric, we have found that SBERT produced the best scores.

F Example of Multi-Agent Debate in ID-IsF

In Figure 12, the post has been associated with known IsFs having the name: *Health & Safety*; *Legality, Constitutionality & Jurisdiction*; *Security & Defense*, and *Crime & Punishment* by the **Investigator agent**. The **Proposer agent** then introduces two new candidate IsFs for the same post, having the names: *Child Protection* and *Government Accountability*. For IsF_1^{cand}, with the name *Child Protection*, the **Champion agent** successfully argues that IsF_1^{cand} should be considered because its name: *Child Protection* is different from general human rights. Additionally, while Health & Safety captures the biological fact of death, it fails to encode the moral and rights-based salience introduced by the dead person being a minor. The **Champion agent** emphasizes that the explicit mention of a "7-year-old" invokes a category of vulnerability and protection that is not adequately subsumed by the IsFs with names *Health & Safety* or *Security & Defense*.

In response, the **Judge agent** accepts the **Champion agent**'s core insight (Child Rights are distinct from Health) but also accepts the **Critic agent**'s argument that the name of the candidate IsF "Child Protection" is too broad. Consequently, the **Judge agent** *CORRECTS* the proposed IsF name to a new, robust name, i.e. *Family Separation & Child Harm* explicitly capturing the nuance argued in the debate.

In contrast, for the second IsF^{cand}, with the name *Government Accountability*, the **Critic agent** successfully attacks the proposal by noting that the text mentions "Border Patrol" only as a factual location of the event. The **Critic agent** argues that the text lacks lexical accountability markers (e.g., "resign," "investigate"), and creating a new name for the IsF based solely on an agency's name creates taxonomic bloat. The **Judge agent** concurs with this assessment, concluding that the proposal rests on overreaching assumptions rather than evidence, and issues a *REJECT* verdict. This example demonstrates how the debate framework both enables the refinement of genuinely novel IsFs and prevents the proliferation of redundant ones.

G Training details of the Articulation Reward Model in Art-Frame

G.1 Data Collection Strategy

Since no off-the-shelf dataset exists for *valid* vs. *invalid* framing, we constructed a dataset consisting of 1,165 valid FAs that are clearly articulated and 1000 invalid FAs that are poorly articulated. The valid FAs are taken from existing datasets on framings (Weinzierl and Harabagiu, 2025; Ailneni and Harabagiu, 2025), out of which 83 are based on COVID-19 Vaccine hesitancy, 249 are based on Immigration, 75 are based on Abortion, 50 are based on Climate Change and 708 FAs are based on Misogyny. We synthetically generated invalid FAs to mimic common failure modes of zero-shot LLMs such as: 1) Statements describing the post rather than the topic aspect such as *"The user uses a metaphor to describe the border situation."*, 2) Incomplete statements such as *Rising crime rates and lack of police funding.*, etc.

G.2 Training Configuration

We fine-tuned the DeBERTa-v3-large (He et al., 2021) model using binary cross-entropy loss. Training was conducted on a single NVIDIA T4 GPU for 3 epochs to prevent overfitting, utilizing a batch size of 16, a maximum sequence length of 128 tokens, and a learning rate of 2e-5 with linear decay. On a held-out validation set of 200 FAs, the model achieved an F1-score of 0.94, demonstrating robust capability in distinguishing between valid and invalid FAs.

H Additional Quantitative Results

The identification of known IsFs is cast as a **multi-label classification task**: for each posting P_i in the IMM dataset, the system must identify which subset of the known IsFs are addressed by P_i. A posting may simultaneously address multiple IsFs. Precision (P), Recall (R), and F1-score are computed at the posting level by comparing the predicted set of IsFs against the gold-annotated set, and macro-averaged across all IsFs.

We evaluate the identification of known IsFs under two conditions:

- **IMM-G (Issue-Generic):** The system identifies which of the 14 issue-generic IsFs intro-

duced by the Policy Frames Codebook (Boydstun et al., 2018) are addressed by each posting.

- **IMM-S (Issue-Specific):** The system identifies which of the 11 immigration-specific IsFs introduced by (Benson, 2013) and (Hovden and Mjelde, 2019) are addressed by each posting.

We use the same annotated dataset of immigration-related tweets from (Mendelsohn et al., 2021), evaluating all methods on the test set of 450 tweets (3,600 train / 450 development / 450 test). We note that unlike (Mendelsohn et al., 2021), which fine-tunes a supervised model on the training split, **ID-IsF** and the prompting-based baselines described below require no training data. We consider two baselines: Zero-Shot Prompt and Few-Shot Prompt. For both the baselines, GPT-5-mini is prompted with the names and definitions of all known IsFs and asked to identify which IsFs are addressed by each posting. The system prompt used for zero-shot multi-label classification is shown in Figure 13. For the Few-Shot Prompt baseline, in addition to the prompt used in Zero-Shot baseline, we provide two demonstrations per IsF drawn from the training split of the IMM dataset. Each demonstration consists of a posting and its gold-annotated IsF labels. The same Zero-Shot Prompt shown in Figure 13 is used for both IMM-G and IMM-S identification tasks; the only difference is that the IsF names and definitions provided in the prompt are replaced with the corresponding issue-generic or issue-specific IsFs respectively.

Tables 19 and 20 report the results for the identification of known IsFs on the IMM-G and IMM-S datasets respectively.

Model	P	R	F1
(Mendelsohn et al., 2021)	0.73	0.72	0.71
Zero-Shot Prompt	0.72	0.79	0.75
Few-Shot Prompt	0.74	0.80	0.77
ID-IsF	**0.77**	**0.85**	**0.81**

Table 19: Identification of known issue-generic IsFs on the IMM-G dataset.

IDENTITY
You are an Issue Frame Identification Agent. Your role is to identify which Issue Frames (IsFs) are addressed by a given social media posting about immigration.

INPUTS
You will receive:

- A social media posting about immigration.
- A list of known Issue Frames (IsFs), each with a name and definition.

OBJECTIVE
For each IsF, decide whether it is addressed by the posting. Output a list of all IsF names that are addressed by the posting. A posting may address multiple IsFs simultaneously.

INSTRUCTIONS

- Base your decisions strictly on the textual content of the posting.
- Do not make assumptions about intent or meaning beyond what is explicitly stated.
- If no IsF is addressed, output an empty list.
- Output format: a list of IsF names, e.g. ["Economic", "Security & Defense"]

Figure 13: Zero-Shot Prompt for multi-label IsF identification.

Model	P	R	F1
(Mendelsohn et al., 2021)	0.59	0.53	0.55
Zero-Shot Prompt	0.60	0.56	0.58
Few-Shot Prompt	0.64	0.59	0.61
ID-IsF	0.79	0.88	0.83

Table 20: Identification of known issue-specific IsFs on the IMM-S dataset.

For the identification of known issue-generic IsFs (Table 19), **ID-IsF** achieves the highest performance outperforming both the supervised baseline of (Mendelsohn et al., 2021) and the prompting-based baselines. The most notable gain is in Recall (0.85 vs. 0.80 for FS-Prompt and 0.72 for (Mendelsohn et al., 2021)), indicating that **ID-IsF** is more effective at identifying all IsFs addressed by a posting rather than just the most salient ones. This improvement is attributable to the introspective Q/A loop, which resolves cases where the Investigator Agent cannot confidently determine whether a posting addresses a given IsF. The improvement of FS-Prompt over ZS-Prompt further confirms that in-context demonstrations are beneficial.

For the identification of known issue-specific IsFs (Table 20), **ID-IsF** achieves substantially higher

performance across all metrics compared to all baselines, with particularly large gains in Recall (0.88 vs. 0.59 for (Mendelsohn et al., 2021)). The performance gap between **ID-IsF** and the prompting-based baselines is notably larger for issue-specific IsFs than for issue-generic IsFs.

I The Rediscovery of Known Issue Frames

To evaluate how the selected number of known IsFs influences the re-discovery of remaining known IsFs, we conducted a series of experiments in which varying proportions P of known IsFs were provided to **ID-IsF**. We were interested to find if the IsFs that were not shown to **ID-IsF** can be re-discovered by generating their names and definitions. We have performed these experiments with $P = 10\%, 25\%, 50\%$ and 75%, for each of the four topics (IMM-S, CC, AB, VH) that were considered in this paper.

We note that unlike the other three topics, the Immigration dataset (IMM) contains both issue-generic and issue-specific IsFs. Since issue-generic IsFs are defined across policy topics rather than being specific to immigration discourse, we conduct the rediscovery experiment using the issue-specific IsFs for a fairer and more informative comparison.

Human Annotation of Rediscovered IsFs. To determine whether a known IsF has been successfully rediscovered, we rely on human annotators. Annotators are asked to perform the following matching task:

Step 1: Annotators are presented with two sets of IsFs:

- **Set A:** The full list of known IsFs that were withheld from **ID-IsF**, each with their name and definition.
- **Set B:** The full list of IsFs rediscovered by **ID-IsF**, each with their name and definition.

Step 2: For each known IsF in Set A, annotators look through all IsFs in Set B and identify the rediscovered IsF that most closely corresponds to it, if any. If no rediscovered IsF in Set B corresponds to a known IsF in Set A, annotators mark it as *Not Rediscovered*.

Step 3: A known IsF is considered *Not Rediscovered* if no single rediscovered IsF in Set B receives a majority of annotator votes as its closest match. Otherwise, the known IsF is considered successfully rediscovered and the matched rediscovered IsF is the one receiving the majority vote. Inter-annotator agreement for the matching task is measured using Krippendorff's α (nominal scale) and reported separately for each topic: IMM-S (α = 0.84), CC (α = 0.86), AB (α = 0.85), and VH (α = 0.88).

Table 21 lists all IsFs confirmed as successfully re-discovered by the annotators across all experimental configurations of P, along with their corresponding known IsFs. For example, the rediscovered IsF with the name *Victim: Human Rights and Welfare* was matched by annotators to the known IsF *Victim: Humanitarian*, as both capture the hardships faced by immigrants under unjust policies and inadequate legal protections. Similarly, the re-discovered IsF with the name *Hero: Cultural Enrichment* was matched to the known IsF *Hero: Cultural Diversity*, as both recognize the positive contributions of immigrants to the social and cultural life of their host communities.

Across all topics, we observed that the number of known IsFs provided to the system had limited impact on the ability of the method to re-discover the removed IsFs. In many cases, the system successfully rediscovered all removed IsFs even when only a small subset of the known IsFs was supplied.

For the Immigration dataset, all removed issue-specific IsFs were rediscovered. For the COVID-19 Vaccine Hesitancy dataset as well, the system consistently rediscovered all removed IsFs across all experimental configurations, even when only a single IsF was initially provided. For the Abortion dataset, most removed IsFs were rediscovered across configurations; however, the system did not rediscover *Legislative & States' Rights* and *Language & Terminology*. Both of these represent relatively broad categories that capture general discussions, making them less distinctly expressed in the discourse and therefore harder for the system to identify as separate IsFs. Finally, for the Climate Change dataset, most removed IsFs were also rediscovered, with the exception of *Elitist vs. Public Needs (Populist Framing)*.

REDISCOVERED Issue Frame 1: Name: Victim: Economic Hardship	**Definition:** Immigration is driven by economic inequality, poverty, and limited opportunities across countries, positioning immigrants as individuals affected by broader global economic conditions.
Corresponding KNOWN Issue Frame: Name: Victim: Global Economy	**Definition:** Immigrants are victims of global poverty, underdevelopment, and inequality.
REDISCOVERED Issue Frame 2: Name: Victim: Human Rights & Welfare	**Definition:** Immigrants experience hardships caused by restrictive policies, inadequate access to healthcare or legal protections, unsafe migration conditions, and difficulties adjusting within receiving societies.
Corresponding KNOWN Issue Frame: Name: Victim: Humanitarian	**Definition:** Immigrants are victims of unjust government policies (violations of human rights, fair legal process) or business practices; they suffer from poverty, lack of access to health care, dangers related to border crossing, etc.; or they have difficulties in adapting to their host society.
REDISCOVERED Issue Frame 3: Name: Victim: Violence & Conflict	**Definition:** Immigration occurs as a consequence of wars, violent unrest, and conflict situations that force people to leave their countries.
Corresponding KNOWN Issue Frame: Name: Victim: War	**Definition:** Focus on war and violent conflict as reason for immigration.
REDISCOVERED Issue Frame 4: Name: Victim: Prejudice & Exclusion	**Definition:** Immigrants face discrimination, xenophobia, hate crimes, and unfair treatment based on ethnicity, culture, religion, or national background.
Corresponding KNOWN Issue Frame: Name: Victim: Discrimination	**Definition:** Immigrants are victims of racist or xenophobic slurs or hate crimes, or discrimination based on national origin, race/ethnicity, religion, or culture.
REDISCOVERED Issue Frame 5: Name: Hero: Cultural Enrichment	**Definition:** Immigrants strengthen communities by bringing diverse experiences, ideas, creative expression, and traditions that broaden and enrich social life.
Corresponding KNOWN Issue Frame: Name: Hero: Cultural Diversity	**Definition:** Immigrants bring positive differences to a society, from new cuisines to the unique contributions of immigrant artists, musicians, and writers.
REDISCOVERED Issue Frame 6: Name: Hero: Social Integration	**Definition:** Immigrants actively participate in and adapt to the social norms, social responsibilities, and cultural practices of their host communities.
Corresponding KNOWN Issue Frame: Name: Hero: Integration	**Definition:** Immigrants enthusiastically adopt mainstream cultural mores or civic obligations of their host society.
REDISCOVERED Issue Frame 7: Name: Hero: Economic Contribution	**Definition:** Immigrants contribute through hard work, fill labor demands, and support economic growth by performing essential roles within the workforce.
Corresponding KNOWN Issue Frame: Name: Hero: Worker	**Definition:** Immigrants work hard, take jobs that citizens or legal residents will not or cannot do, or contribute to economic prosperity and growth.
REDISCOVERED Issue Frame 8: Name: Threat: Employment	**Definition:** Immigrants increase competition in the labor market by taking employment opportunities or contributing to lower wages for native workers.
Corresponding KNOWN Issue Frame: Name: Threat: Jobs	**Definition:** Immigrants take nonimmigrants' jobs or lower their wages.
REDISCOVERED Issue Frame 9: Name: Threat: Safety & Law Enforcement	**Definition:** Immigration is associated with concerns about illegal entry, crime, public safety risks, the spread of diseases, and pressures arising from large population inflows.
Corresponding KNOWN Issue Frame: Name: Threat: Public Order	**Definition:** Illegal immigrants break the law in coming into this country; once here, immigrants (legal or illegal) are more likely than others to commit crimes, use drugs, and carry diseases; immigrants are coming in such numbers that they threaten overcrowding and environmental degradation.
REDISCOVERED Issue Frame 10: Name: Threat: Fiscal Pressure	**Definition:** Immigrants place increased demands on public services and social welfare systems, creating financial strain on taxpayers and government resources.
Corresponding KNOWN Issue Frame: Name: Threat: Fiscal	**Definition:** Immigrants (especially illegal) abuse government social service programs (health, education, etc.), imposing an unfair burden on taxpayers.
REDISCOVERED Issue Frame 11: Name: Threat: Cultural Identity	**Definition:** Immigration introduces values, customs, and social practices that may weaken shared cultural traditions or the cultural identity of the host country.
Corresponding KNOWN Issue Frame: Name: Threat: National Cohesion	**Definition:** Immigrants bring foreign customs and values that threaten to undermine the host country's culture or national identity; immigrants are inassimilable.

Table 21: Re-discovered Issue Frames (with their names and definitions) and their corresponding known Issue Frames for the topic of Immigration.

J Additional Qualitative Results

J.1 Quality of the discovered IsFs

Method	SAL	DC	DNA	PC	SD
Immigration (IMM-G)					
ID-IsF	**0.08**	**4.30**	**4.29**	**0.91**	**0.88**
$IsFD_{HAC}$	0.04	3.10	3.35	0.43	0.65
ZS-Prompt	0.03	3.78	4.05	0.50	0.59
FS-Prompt	0.04	4.02	4.15	0.66	0.72
Climate Change (CC)					
ID-IsF	**0.20**	**4.21**	**4.18**	**0.88**	**0.90**
$IsFD_{HAC}$	0.05	3.15	3.40	0.44	0.68
ZS-Prompt	0.07	3.76	3.98	0.55	0.47
FS-Prompt	0.11	3.93	4.05	0.69	0.65
Abortion (AB)					
ID-IsF	**0.63**	**4.39**	**4.67**	**0.85**	**0.93**
$IsFD_{HAC}$	0.11	3.35	3.60	0.47	0.63
ZS-Prompt	0.10	3.86	4.10	0.61	0.60
FS-Prompt	0.19	4.12	4.28	0.71	0.73
Vaccine Hesitancy (VH)					
ID-IsF	**0.77**	**4.45**	**4.58**	**0.89**	**0.92**
$IsFD_{HAC}$	0.13	3.30	3.55	0.37	0.70
ZS-Prompt	0.18	3.92	4.08	0.64	0.72
FS-Prompt	0.25	4.15	4.23	0.77	0.81

Table 22: Qualitative results for the discovery of *new* IsFs for the topic of Immigration (IMM-G), Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH).

Qualitative results for the discovery of New IsFs for the topic of Immigration (IMM-G), Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH) are summarized in Table 22. Across all four topics, **ID-IsF** consistently outperforms all baselines across all five criteria, with the most striking gains observed in Salience (SAL), Definition Clarity (DC) and Definition–Name Alignment (DNA). These results indicate that the newly discovered IsFs produced by **ID-IsF** are clearly articulated and remain semantically faithful to their assigned names. In contrast, $IsFD_{HAC}$, $ZS-Prompt$, and $FS-Prompt$ exhibit noticeably lower DC and DNA scores, suggesting definitions that are either less precise or less aligned with the name. Most notably, **ID-IsF** achieves substantially higher SAL scores across all topics, confirming that the discovered IsFs capture meaningful and salient dimensions of the topic discourse. In contrast, baseline methods consistently produce IsFs with near-zero SAL scores, indicating that their discovered IsFs largely fail to capture meaningful aspects of the topic.

The advantages of **ID-IsF** are even more pronounced for Posting Coverage (PC) and Semantic Distinctiveness (SD). Across all four topics, **ID-IsF** yields PC and SD values close to 0.9, demonstrating that postings grouped under each new IsF genuinely address that IsF and that the discovered IsFs are meaningfully distinct from one another.

From a topic-wise perspective, Vaccine Hesitancy and Abortion stand out with the highest SAL scores. In contrast, IMM-G yields the lowest SAL scores across all methods, consistent with the topic-agnostic nature of issue-generic frames which are less likely to capture salient aspects specific to immigration discourse. Vaccine Hesitancy stands out for strong overall definition quality and high post coverage. Abortion, in contrast, achieves the highest scores in definition–name alignment and semantic distinctiveness, indicating particularly well-defined and non-overlapping Issue Frames.

Method	Z	A	S	V
Immigration (IMM-G)				
Art-Frame	**0.98**	**0.96**	**0.79**	**0.71**
$IsFD_{HAC}$	0.70	0.73	0.31	0.29
ZS-Prompt	0.81	0.83	0.48	0.42
FS-Prompt	0.88	0.89	0.58	0.51
Climate Change (CC)				
Art-Frame	**0.97**	**0.95**	**0.87**	**0.77**
$IsFD_{HAC}$	0.71	0.75	0.42	0.40
ZS-Prompt	0.79	0.79	0.58	0.50
FS-Prompt	0.85	0.85	0.63	0.61
Abortion (AB)				
Art-Frame	**0.99**	**0.97**	**0.81**	**0.68**
$IsFD_{HAC}$	0.72	0.74	0.33	0.31
ZS-Prompt	0.84	0.87	0.52	0.43
FS-Prompt	0.89	0.91	0.56	0.49
Vaccine Hesitancy (VH)				
Art-Frame	**0.98**	**0.98**	**0.83**	**0.80**
$IsFD_{HAC}$	0.69	0.73	0.39	0.33
ZS-Prompt	0.82	0.83	0.58	0.55
FS-Prompt	0.88	0.90	0.67	0.67

Table 23: Qualitative results for the Frame Articulations (FAs) generated for Issue Frames (IsFs) identified across topics. Z: Soundness, A: Clarity, S: Interpretation Support, V: Analytical Value.

J.2 Quality of the articulated FAs

Qualitative results for the articulation of Frame Articulations (FAs) across topics Immigration (IMM), Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH) are summarized in Table 23. To rigorously assess the generated FAs, we engaged human linguists to evaluate four properties: **Soundness** (Z), **Clarity** (A), **Interpretation Support**

(S), and **Analytical Value** (V). As shown in Table 23, our **ID-IsF** framework achieves superior performance across all domains, consistently scoring above 0.96 in Soundness, 0.94 in Clarity, 0.78 in Interpretation Support, and 0.67 in Analytical Value. This demonstrates that identifying high-quality IsFs is a prerequisite for articulating logically sound FAs; because ID-IsF provides a robust foundation, the resulting FAs are both coherent and grounded. In contrast, the baselines struggle to match this fidelity. While the *FS-Prompt* method produces highly fluent text (Clarity ≈ 0.90) due to the provided demonstrations, it still falls short in Soundness compared to **ID-IsF**. The $IsFD_{HAC}$ baseline performs the worst across all metrics, indicating that clustering-based approaches often produce noisy IsFs that are difficult to articulate into meaningful FAs.

The gains of **Art-Frame** are even more pronounced for Interpretation Support (S) and Analytical Value (V), where baseline methods fall substantially short — $IsFD_{HAC}$ scores below 0.42 on S and 0.40 on V across all topics, indicating that clustering-based IsFs not only produce noisy articulations but also fail to generate FAs that meaningfully interpret their corresponding IsFs or provide analytically distinct interpretations.

K Human Evaluation Details

K.1 Evaluation Criteria

K.1.1 Evaluation Criteria for Issue Frames

We evaluate IsFs based on five criteria: Salience (SAL), Definition Clarity (DC), Definition-Name Alignment (DNA), Posting Coverage (PC), and Semantic Distinctiveness (SD). Salience (SAL) is the most important criterion as it directly measures whether a discovered IsF captures a meaningful and salient aspect of the topic discourse that is not a special case of a known IsF. An IsF may be linguistically well-defined, well-named, grounded in postings, and distinct from other IsFs, yet still fail to capture an aspect of the topic that genuinely matters for understanding why people hold different opinions about it. SAL is measured as the proportion of *Yes* judgments over the total number of annotated IsFs, where a higher score indicates that a greater proportion of discovered IsFs are judged as capturing meaningful and salient dimensions of the topic discourse that go beyond refinements of already known IsFs.

Definition Clarity (DC) ensures that the IsF definition is fluent and comprehensible, and is measured as the average score on a 5-point Likert scale across all annotated IsFs, ranging from 1 (incomprehensible) to 5 (perfectly clear). Definition-Name Alignment (DNA) ensures that the name faithfully reflects the concept captured by the definition, and is similarly measured as the average score on a 5-point Likert scale, ranging from 1 (semantically mismatched) to 5 (comprehensive explanation of the IsF name). Posting Coverage (PC) verifies that the IsF is genuinely grounded in the postings assigned to it, and is measured as the proportion of postings receiving a *Yes* judgment over all sampled postings addressing the same IsF, averaged across all discovered IsFs. Finally, Semantic Distinctiveness (SD) ensures that the discovered IsFs are semantically non-redundant with respect to each other, and is measured as the average distinctiveness score across all annotated pairs of IsFs on a three-point scale: *Duplicate* (0), *Overlap* (0.5), or *Distinct* (1.0).

K.1.2 Evaluation Criteria for Frame Articulations

We evaluate the quality and analytic value of generated Frame Articulations (FAs) based on four properties: Soundness (Z), Clarity (A), Interpretation Support (S), and Analytical Value (V). Soundness (Z) measures whether the rationale provided for each FA is logically coherent and corresponds to both the content of the posting and the definition of the IsF it interprets, and is measured as the proportion of FAs deemed sound over the total number of generated FAs ($Z = N_Z/N_T$). Clarity (A) measures whether each FA is expressed fluently and unambiguously, and is measured as the proportion of FAs deemed clear over the total number of generated FAs ($A = N_C/N_T$). Interpretation Support (S) measures whether a generated FA provides a meaningful interpretation of its corresponding IsF. A FA may be logically sound and fluently expressed yet still fail to provide a genuine interpretive explanation of the IsF it is intended to interpret. S is measured as the proportion of FAs deemed to support the interpretation of their corresponding IsF over the total number of generated FAs ($S = N_S/N_T$). Finally, Analytical Value (V) measures whether each FA is genuinely distinct from other FAs interpreting the same IsF and is evoked by a sufficient number of postings. Specifically, a FA meets the Analytical Value criterion if

it is not a special case of another FA supporting the same IsF and is evoked by at least $k = 3$ postings. The threshold of $k = 3$ postings was determined empirically to ensure the relevance of each FA, following (Weinzierl and Harabagiu, 2024). V is measured as the proportion of FAs meeting both conditions over the total number of generated FAs ($V = N_V/N_T$), where a higher score indicates that a greater proportion of generated FAs are both interpretively distinct and sufficiently grounded in the discourse. The final label (Yes / No) for each criterion is determined by majority voting.

K.2 Annotator Recruitment and Profiles

To ensure high-quality and reliable judgments for our qualitative evaluation, we recruited a total of six annotators: two computational linguists and a social scientist independently evaluated the discovered IsFs based on the five criteria detailed in Section 5, while another two computational linguists and a social scientist were dedicated to assessing the generated FAs for soundness and clarity. All annotators were recruited externally and were independent of the research team.

K.3 Annotator Instructions

The evaluation was conducted blind: the discovered IsFs and generated FAs from **ID-IsF** and the baselines were anonymized to remove system identity and presented in randomized order, so annotators could not tell which system produced a given IsF or FA.

K.3.1 Annotator Instructions for Evaluating Issue Frames (IsFs)

To evaluate the quality of IsFs, the annotators are given the definition of an IsF. Additionally, for each topic, we provide the annotators with the known Issue Frames (IsFs) and their definitions as examples of IsFs. Detailed annotation instructions for evaluating the quality of IsFs are shown in Figure 14.

Criterion 1: Salience (SAL): You will be shown the IsF name and definition. Decide whether the IsF represents a salient aspect of the topic that is not a special case of a known IsF. Ask yourself: does this IsF capture a meaningful dimension of the topic discourse that contributes to opinion formation, and is it genuinely distinct from the known IsFs provided as reference? Answer **Yes = 1** if the IsF represents a salient and meaningful dimension of the topic discourse that goes beyond a refinement or special case of any known IsF, and **No = 0** if the IsF captures a trivial or incidental aspect of the topic, or if it represents a special case of an already known IsF.
Criterion 2: Definition Clarity (DC): Read the provided IsF definition. Decide whether the definition is fluent and easy to understand. Rate the clarity on a 5-point Likert scale, ranging from 1 (incomprehensible) to 5 (perfectly clear).
Criterion 3: Definition-Name Alignment (DNA): Review both the IsF name and its definition. Judge the extent to which the definition explains, justifies, and faithfully reflects the meaning implied by the IsF name. Use a 5-point Likert scale, where a score of 1 indicates that the definition and name are semantically mismatched, and a score of 5 indicates that the definition offers a comprehensive explanation of the IsF name.
Criterion 4: Posting Coverage (PC): You will be shown the IsF name and definition alongside a random sample of postings assigned to this IsF (20% of all relevant postings, minimum 5, maximum 50). Evaluate whether the postings genuinely discuss the IsF. For each posting, answer **Yes** if the posting clearly addresses the IsF, and **No** otherwise.
Criterion 5: Semantic Distinctiveness (SD): You will be presented with the names and definitions of pairs of IsFs. Measure how clearly the IsFs differ in meaning. Rate their distinctiveness based on their definitions using a three-point scale: **Duplicate (0)** — the two IsFs capture the same concept and are interchangeable; **Overlap (0.5)** — the two IsFs are related but capture partially different aspects of the topic; **Distinct (1.0)** — the two IsFs capture clearly different concepts.

Figure 14: Annotator guidelines for evaluating the quality and of new IsFs.

For measuring Salience (SAL), annotators are shown the IsF name and definition. For Definition Clarity (DC), annotators are provided with the IsF definition only. For Definition-Name Alignment (DNA), annotators are provided with both the IsF name and its definition. For Posting Coverage (PC), annotators are provided with the IsF name and definition alongside a random sample of postings addressing the IsF, consisting of 20% of all relevant postings with a minimum of 5 and a maximum of 50 postings. For Semantic Distinctiveness (SD), annotators are provided with the names and definitions of pairs of IsFs, where pairs are selected by computing cosine similarity between IsF definitions using Sentence-BERT (Reimers and Gurevych, 2019). Only pairs exceeding a similarity threshold of 0.6 are presented to annotators.

K.3.2 Annotator Instructions for Evaluating Frame Articulations (FAs)

For Soundness (Z), annotators are provided with the FA, its rationale, the source posting, and the name and definition of the IsF it interprets. For Clarity (A), annotators are provided with the FA only. For Interpretation Support (S), annotators are provided with the FA, and the name and defini-

tion of the IsF it interprets. For Analytical Value (V), annotators are provided with the FA and the full set of FAs interpreting the same IsF. Note that only FAs evoked by at least $k = 3$ postings are shown to annotators. Detailed annotation instructions for evaluating the quality of FAs are shown in Figure 15.

Property 1: Soundness (Z): Review the rationale provided for the FA alongside the source posting and the IsF it interprets. Determine whether the rationale is logically coherent, corresponds to both the content of the posting and the definition of the IsF, and is free of unsupported assumptions. Answer **Yes** if the rationale is sound and **No** otherwise.
Property 2: Clarity (A): Read the FA and assess whether it is expressed fluently and unambiguously. Answer **Yes** if the FA is clear and **No** otherwise.
Property 3: Interpretation Support (S): Decide whether the FA provides a meaningful interpretation of the IsF. Answer **Yes = 1** if the FA provides a meaningful interpretation of the IsF and **No = 0** otherwise.
Property 4: Analytical Value (V): Decide whether the FA is not a special case of another FA interpreting the same IsF. Answer **Yes = 1** if the FA expresses a genuinely distinct interpretation from all other FAs interpreting the same IsF, and **No = 0** otherwise.

Figure 15: Annotator guidelines for evaluating the quality and analytic value of generated FAs.

K.4 Inter Annotation Agreement Details

To ensure the reliability of our human evaluation, we calculated the Inter-Annotator Agreement (IAA) for both the qualitative evaluation of discovered IsFs and the quality assessment of articulated FAs.

K.4.1 IAA for Discovered Issue Frames

The inter-annotator agreement for IsF discovery was calculated on the complete set of 53 novel IsFs discovered by **ID-IsF** across the four topics: 10 for IMM-S, 13 for IMM-G, 10 for CC, 11 for AB, and 9 for VH.
We report Krippendorff's α (Krippendorff, 2011) to account for the different measurement scales:
(a) Salience (SAL): Evaluated using a binary decision (Yes/No) on the IsF name and definition alongside the full list of known IsFs for the corresponding topic. We obtained a Krippendorff's α of **0.83**.
(b) Definition Clarity (DC): Evaluated on a 5-point Likert scale. We observed a Krippendorff's α of **0.82**, indicating substantial agreement on the fluency and comprehensibility of the definitions.
(c) Definition-Name Alignment (DNA): Evaluated on a 5-point Likert scale. The agreement score was **0.83**, showing strong consensus on how well the definitions matched the generated names.
(d) Post Coverage (PC): Evaluated using a binary decision (Yes/No) on a random sample of postings addressing the IsF, consisting of 20% of all relevant postings with a minimum of 5 and a maximum of 50 postings. We obtained a Krippendorff's α of **0.84**, demonstrating high reliability in grounding the IsFs to source text.
(e) Semantic Distinctiveness (SD): Evaluated on a 3-point ordinal scale (Duplicate, Overlap, Distinct) for most similar pairs. The agreement score was **0.80**, confirming that annotators consistently distinguished unique issues from redundant ones.

K.4.2 IAA for Frame Articulations

We report Krippendorff's α for all four criteria, as each involves a binary classification task (Yes/No). Soundness (Z) yielded a Krippendorff's α of **0.79**, Clarity (A) yielded a Krippendorff's α of **0.82**, Interpretation Support (S) yielded a Krippendorff's α of **0.81**, and Analytical Value (V) yielded a Krippendorff's α of **0.80**.

L Additional Ablation Studies

L.1 Additional Ablation study for ID-IsF module

Table 24 reports the impact of removing each component of **ID-IsF** on the discovery quality of new IsFs across all topics. Removing the **Investigator Agent** produces the largest drops across DC, DNA, and PC across all topics, confirming that it provides the contrastive grounding necessary for generating well-formed and reliably grounded IsF definitions. Table 25 further shows that removing the Introspective Q/A Loop, the mechanism by which the Investigator Agent resolves ambiguous IsF identification cases, degrades both Precision and Recall on known IsF identification, demonstrating that it is a necessary component.

Model	P	R	F1-score
ID-IsF	0.77	0.85	0.81
w/o Introspective Q/A Loop	0.74	0.79	0.76

Table 25: Identification of *known* IsFs on the IMM Dataset, evaluated by Precision (P), Recall (R) and F1-score.

Topic	Model	SAL	DC	DNA	PC	SD
IMM-G	**ID-IsF**	0.08	4.30	4.29	0.91	0.88
	w/o Investigator	0.03	3.82	3.74	0.80	0.70
	w/o Issue Frame Memory	0.15	4.25	4.20	0.88	0.45
	w/o Multi-Agent Debate	0.01	4.10	3.80	0.83	0.64
CC	**ID-IsF**	0.20	4.21	4.18	0.88	0.90
	w/o Investigator	0.18	3.80	3.75	0.78	0.84
	w/o Issue Frame Memory	0.13	4.20	4.13	0.85	0.59
	w/o Multi-Agent Debate	0.11	4.16	3.72	0.84	0.63
AB	**ID-IsF**	0.63	4.39	4.67	0.85	0.93
	w/o Investigator	0.48	3.83	3.88	0.76	0.89
	w/o Issue Frame Memory	0.38	4.36	4.64	0.84	0.61
	w/o Multi-Agent Debate	0.33	3.97	4.11	0.83	0.67
VH	**ID-IsF**	0.77	4.45	4.58	0.89	0.92
	w/o Investigator	0.67	3.91	3.80	0.79	0.90
	w/o Issue Frame Memory	0.54	4.40	4.56	0.86	0.60
	w/o Multi-Agent Debate	0.45	4.01	3.99	0.86	0.64

Table 24: Impact of various LLM agents on the discovery quality of new IsFs across topics (IMM-G, CC, AB, VH). SAL: Salience, DC: Definition Clarity, DNA: Definition-Name Alignment, PC: Posting Coverage, SD: Semantic Distinctiveness.

Topic	Model	Z	A	S	V
IMM-S	**Art-Frame**	0.98	0.98	0.86	0.79
	w/o Explainability Agent	0.87	0.94	0.80	0.80
	w/o R_{art} filter	0.96	0.94	0.84	0.81
IMM-G	**Art-Frame**	0.98	0.96	0.79	0.71
	w/o Explainability Agent	0.89	0.94	0.69	0.72
	w/o R_{art} filter	0.95	0.90	0.76	0.75
CC	**Art-Frame**	0.97	0.95	0.87	0.77
	w/o Explainability Agent	0.87	0.95	0.78	0.82
	w/o R_{art} filter	0.96	0.87	0.85	0.85
AB	**Art-Frame**	0.99	0.97	0.81	0.68
	w/o Explainability Agent	0.91	0.96	0.76	0.78
	w/o R_{art} filter	0.99	0.90	0.78	0.79
VH	**Art-Frame**	0.98	0.98	0.83	0.80
	w/o Explainability Agent	0.90	0.96	0.75	0.78
	w/o R_{art} filter	0.97	0.88	0.82	0.80

Table 26: Impact of ablating the Explainability Agent and Articulation Reward (R_{art}) filter in **Art-Frame** on the quality and analytic value of FAs across topics. Z: Soundness, A: Clarity, S: Interpretation Support, V: Analytical Value.

Topic	Accepted (%)	Corrected (%)	Rejected (%)
IMM-G	43	9	48
IMM-S	49	7	44
CC	40	10	50
AB	42	8	50
VH	39	7	54

Table 27: Distribution of decisions of the **Judge Agent** during the multi-agent debate from ID-IsF.

The impact of each ablation on Salience (SAL) further reinforces these findings. **ID-IsF** consistently achieves the highest SAL scores across all topics, confirming that the full architecture is necessary to discover IsFs that capture meaningful and salient dimensions of the topic discourse. The largest drops in SAL are observed when removing the Multi-Agent Debate.

Removing the **Issue Frame Memory** produces a striking drop in Semantic Distinctiveness (SD) across all topics while leaving DC and DNA largely intact. This dissociation confirms that the Issue Frame Memory is essential for keeping track of already discovered IsFs and preventing the system from discovering IsFs that overlap with existing ones, while having no effect on how well individual IsF definitions are written.

Removing the **Multi-Agent Debate** consistently degrades Semantic Distinctiveness, confirming that the adversarial debate between the Champion, Critic, and Judge Agents is essential for rejecting candidate IsFs that are redundant or subtle variants of existing ones. Table 27 lists the proportion of candidate IsFs that were accepted, rejected, or

corrected in the multi-agent debate taking place in **ID-IsF**. A substantial proportion of candidate IsFs proposed by the **Proposer Agent** are rejected by the **Judge Agent**, indicating that CCoT prompting alone is not sufficient. Among accepted IsFs, a non-trivial fraction are corrected by the **Judge Agent**, indicating that adversarial debate often warrants refinement of the name or definition of new IsFs.

L.2 Ablation study for Art-Frame module

We conduct two ablation experiments on the Art-Frame module:
1) w/o the Explainability Agent: By disabling the Explainability Agent, the system accepts the initial interpretation (FA_i) generated by the Articulation Agent as final, without verifying if its provided rationale (R_i) can be used to successfully reconstruct the original frame (FA_i^R).
2) w/o the Articulation Reward Model (R_{art}): In this setting, the fine-tuned discriminative encoder is removed, allowing all generated FAs to bypass the strict acceptance threshold τ.
Table 26 shows the results of the two ablation experiments across all topics. Removing the **Explainability Agent** primarily degrades Soundness (Z) and Interpretation Support (S) across all topics, demonstrating that the reconstruction test is essential for ensuring that the rationale (R_i) provides a logically valid explanation and that the FA meaningfully interprets its corresponding IsF. Conversely, removing the **Articulation Reward Model (R_{art})** filter leads to a more pronounced drop in Clarity (A) while having a moderate effect on Analytical Value (V), indicating that enforcing the strict acceptance threshold τ is necessary to filter out poorly articulated interpretations and maintain the overall quality of the FA set.

M Additional Framing Analysis

The analysis for the topic of Vaccine Hesitancy, shown in Figure 16, reveals the *Confidence* IsF is by far the most addressed one, followed by *Compliance* and *Conspiracy*. Interestingly, to each of these known IsFs many FAs provide multiple accepted interpretations. Among the newly discovered IsFs, *Authority-as-Evidence* is the most addressed new IsF, indicating frequent reliance on expert credentials or selective appeals to scientific authority that fall outside the scope of existing IsFs.

Figure 17 illustrates the framing analysis for the Climate Change topic, showing that among the known IsFs, we found that *Catastrophic Doomsday and Fatalism* is the most addressed issue, with multiple accepted interpretations provided by accepted FAs.

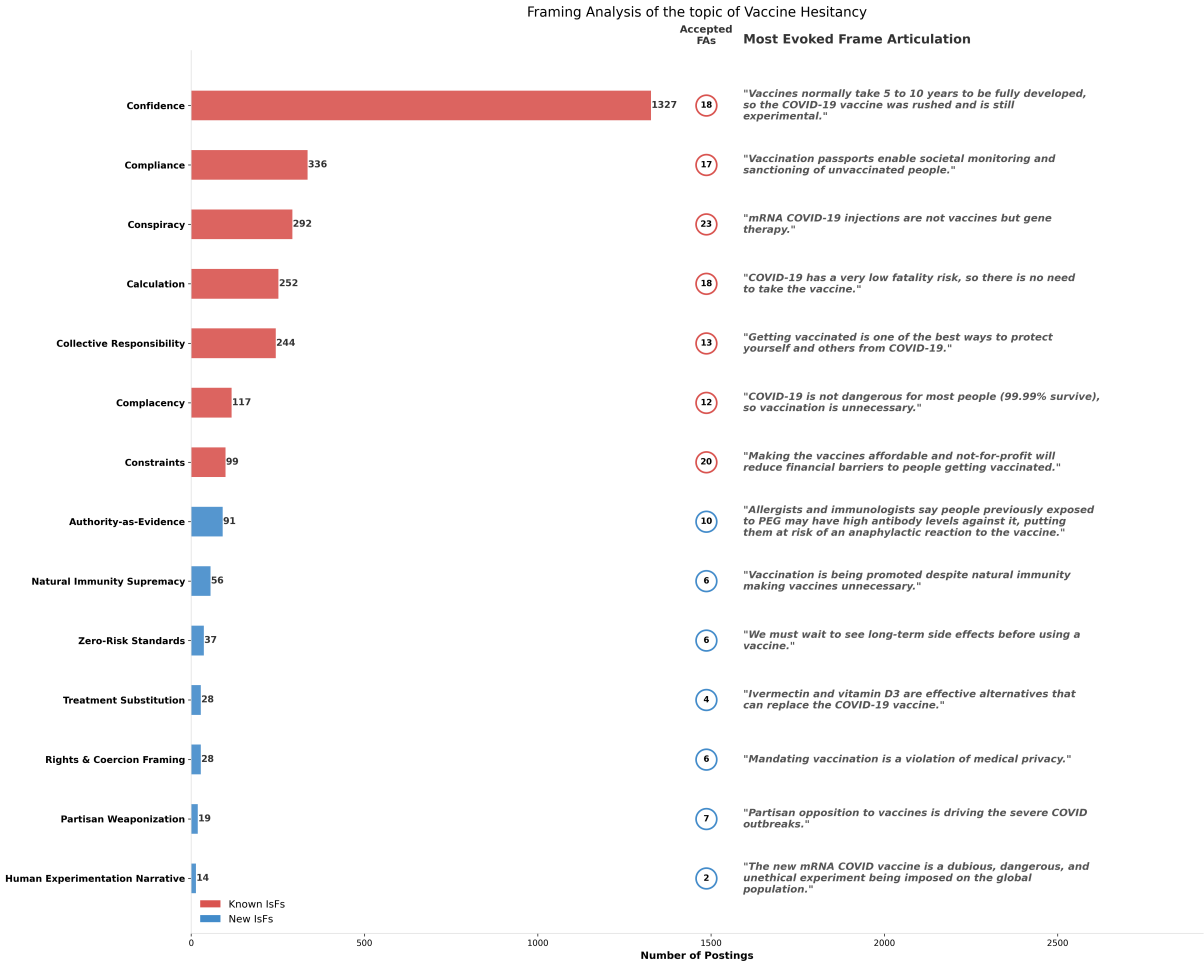


Figure 16: Framing Analysis of the topic of Vaccine Hesitancy.

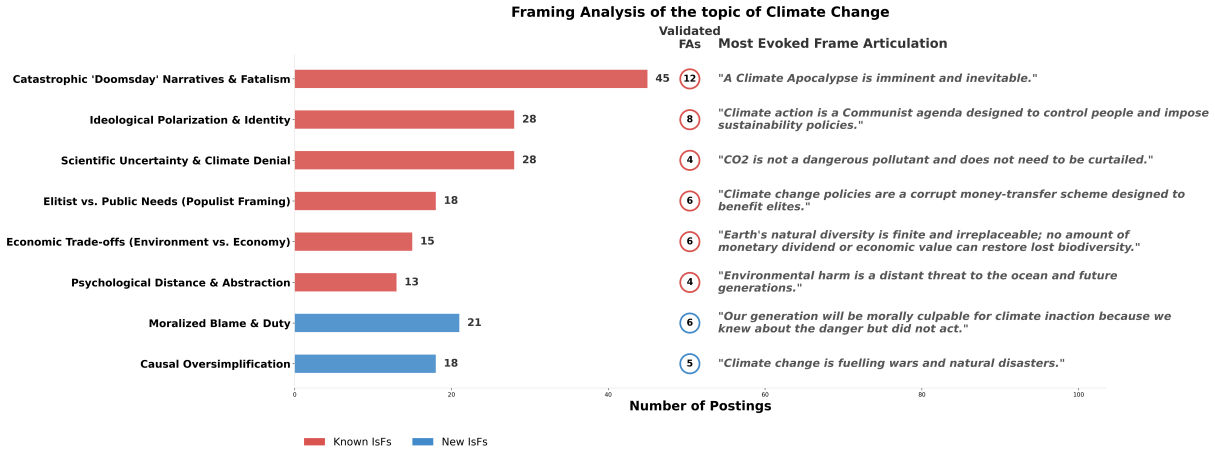


Figure 17: Framing Analysis of the topic of Climate Change.

Figure 18 illustrates the framing analysis for the Abortion topic. The most addressed known IsF is *Fetal Rights*, interpreted by multiple accepted FAs. Among the new salient IsFs, the most addressed one is *Science/Marker Reductionism*, which has much fewer interpretations.

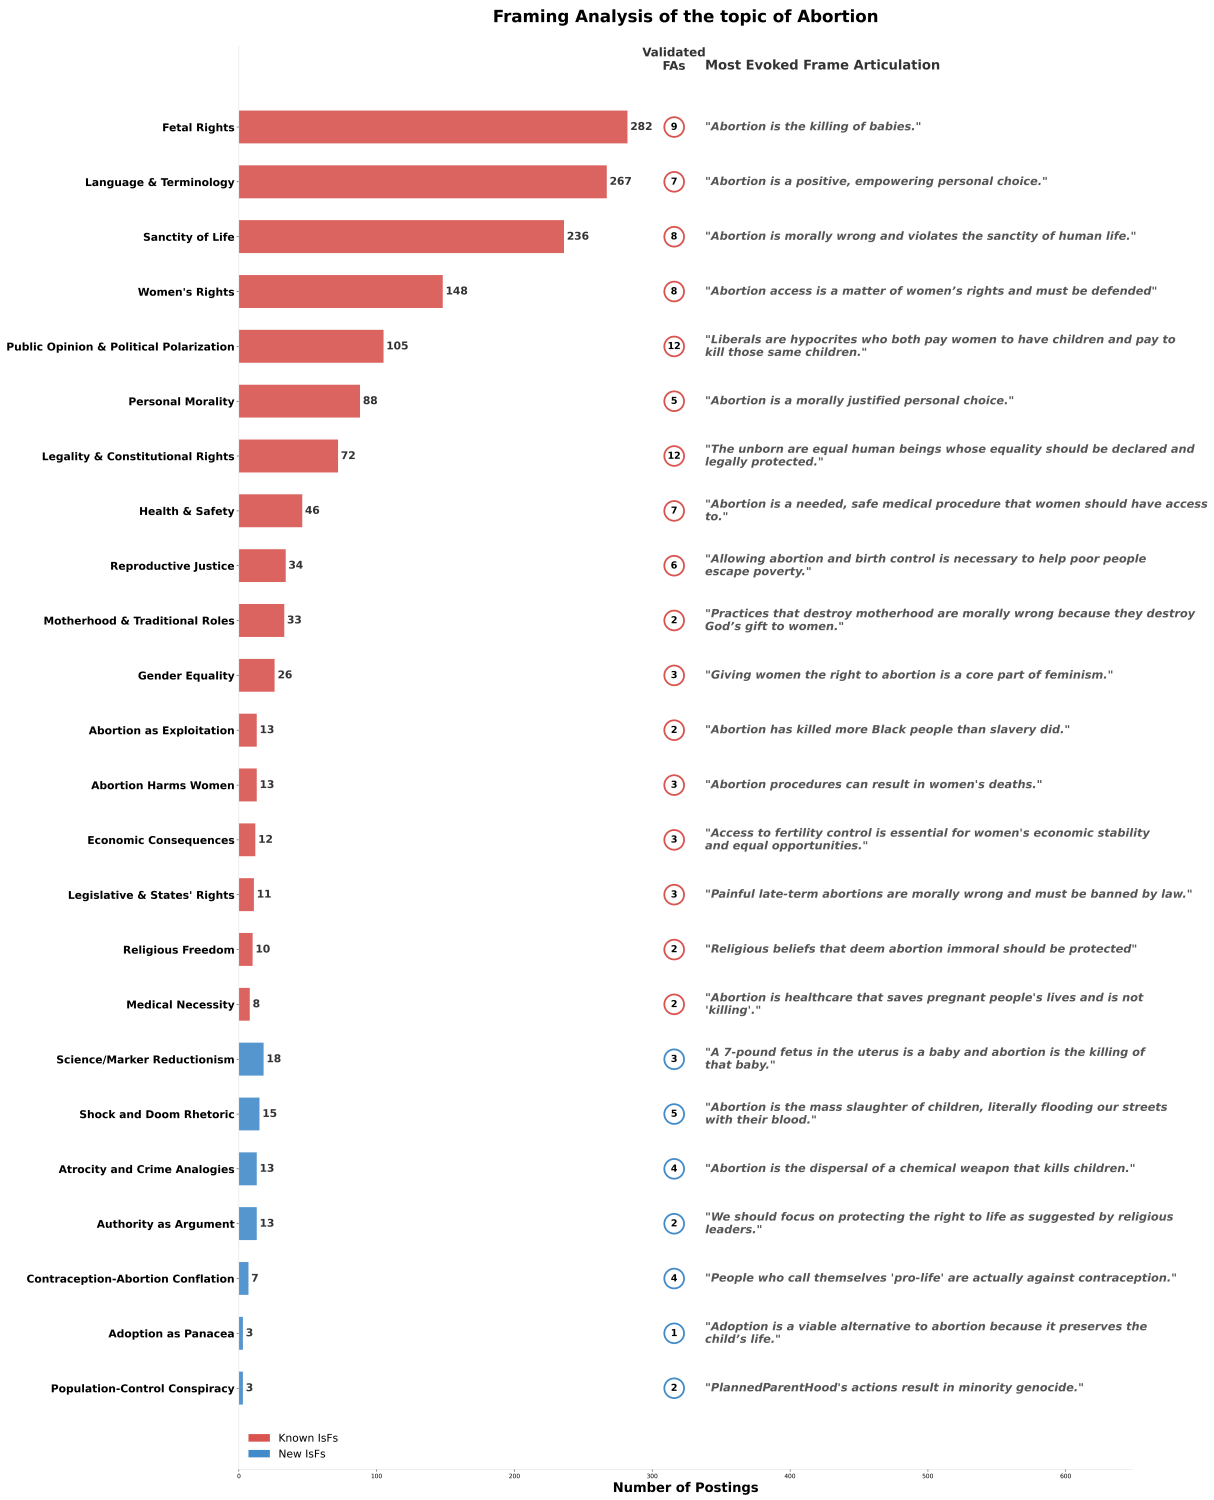


Figure 18: Framing Analysis of the topic of Abortion.

Finally, In Figure 19 we show the analysis for the topic of immigration, in which we show both the known and new IsFs, nmaking the distinction between generic and issue-specific IsFs. For this topic, we show that the analysis reveals the impact on framing of the gneric IsFs, which dominate the framing of this topic. However, from Figure 19 we may conclude that the new generic IsF *Misinformation and Conspiracy* is more addressed than some of the issue-specific IsFs. The analysis also shows the role of the newly discovered IsFs, that capture the long tail of issues framed for this topic. To note is also the fact for the topic of immigration, the generic IsFs are interpreted by more accepted FAs than the issue-specific IsFs, which is not surprising. However, some of the new issue-specific IsFs are interpreted by more

accepted FAs than some of the known issue-specific IsFs.

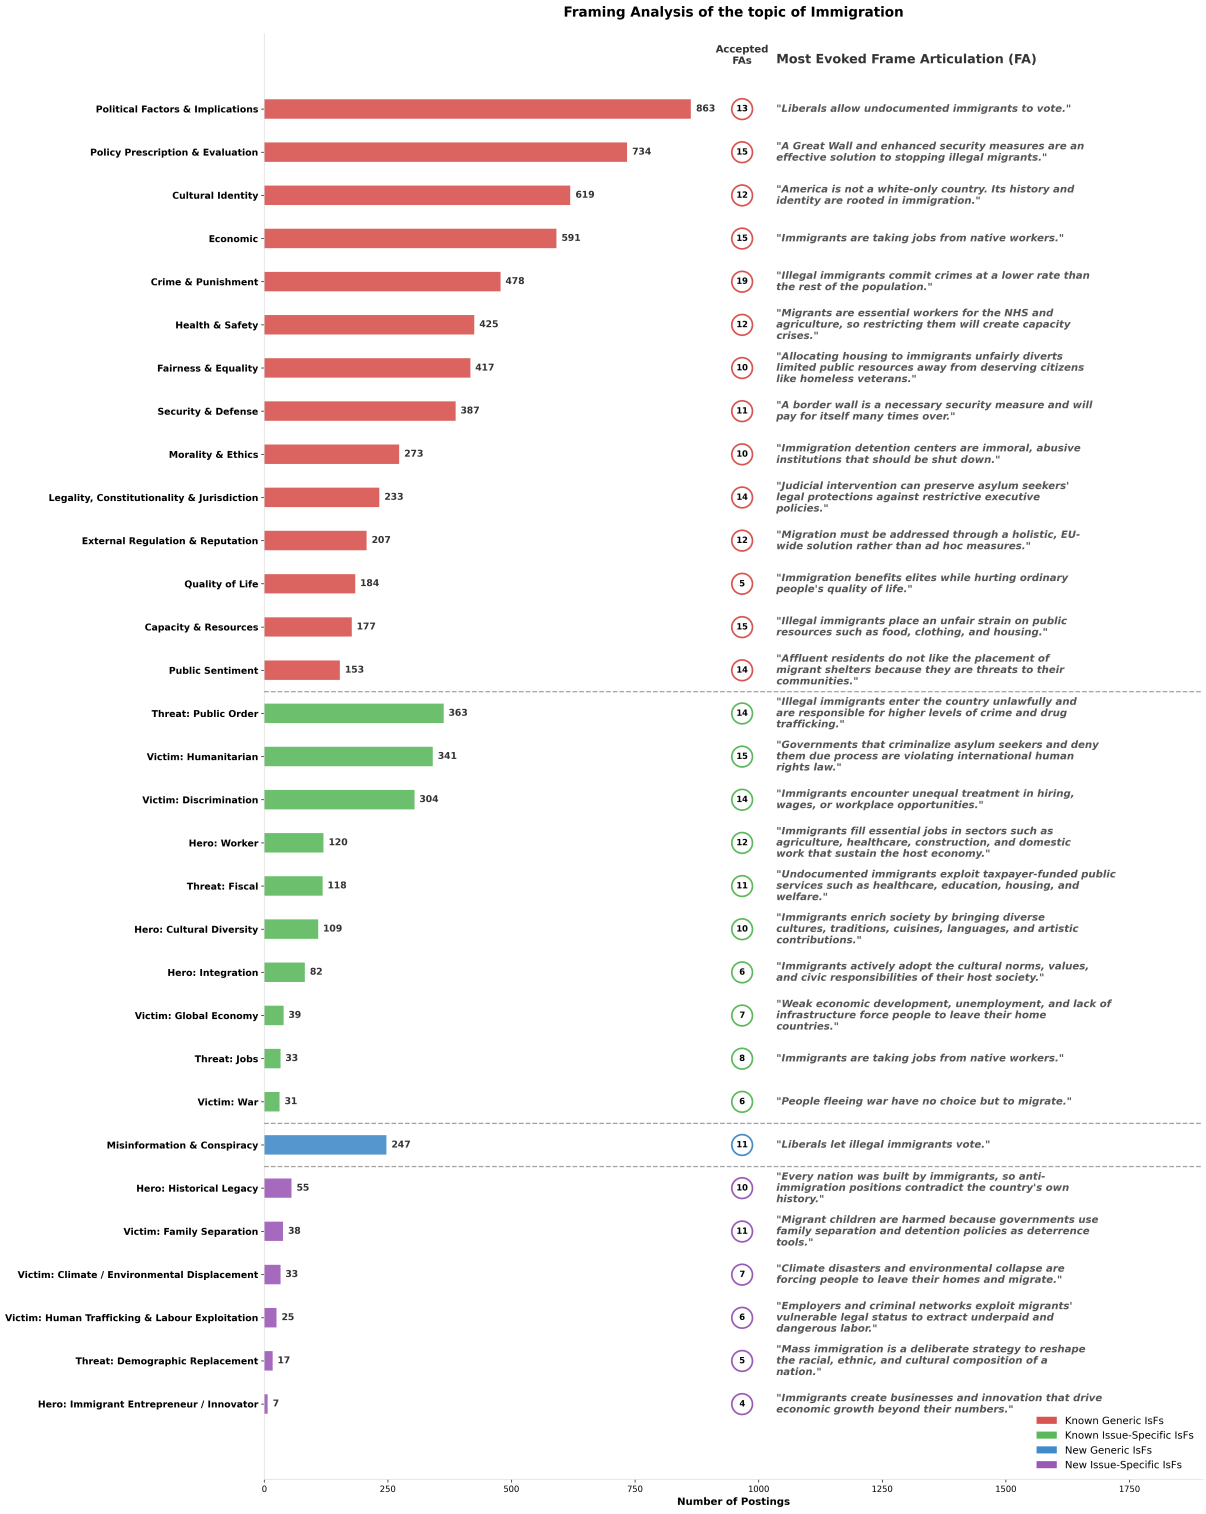


Figure 19: Framing Analysis of the topic of Immigration.

N Examples of Frame Articulations

Frame Articulations (FAs) that interpret Issue Frames provide an important understanding of how a topic is framed. In Table 28 we show the FAs generated by **Art-Frame** for one known, generic IsF used in the framing of immigration. Table 28 also shows the evaluation results for these FAs, according to the properties listed in Table 6. All accepted FAs are sound, clear, support the interpretation of the IsF and

provide analytical value. We show in Table 28 that the accepted FAs are also relevant, as they are evoked by significant number of postings, ranging from 17 up to 104 postings, which is not surprising, given that the IsF of *Political Factors and Implications* is a generic one, originating in the Media Frames.

Table 29 lists the FAs for the newly discovered generic IsF for the topic of immigration: *Misinformation & Conspiracy*. The Table also lists the evaluation results of these FAs as well as their coverage in terms of the postings that evoke them. Interestingly, the number of postings for accepted FAs range from 4 to 104 postings.

#	Frame Articulation (FA)	No. of Posts	Accepted
1	Liberals allow undocumented immigrants to vote.	104	Yes
2	Brexit was driven in part by political dissatisfaction with EU migration policies and freedom of movement rules.	99	Yes
3	Political leaders who cater to constituencies that reject universal human rights enable and normalize institutional cruelty toward migrants at the border.	98	Yes
4	National sovereignty over immigration is undermined by international institutions such as the UN and EU, along with political elites, through migration obligations imposed without sufficient democratic consent.	98	Yes
5	Immigration enforcement tools, including the census, are used to create fear and discourage immigrants from participating in civic and government processes.	65	Yes
6	The U.S. government attributes poor conditions and humanitarian harms at the border onto political opponents who refuse to fund detention infrastructure.	65	Yes
7	Trump spreads fake news about the migrant caravan travelling north through Mexico to manipulate public opinion ahead of the midterm elections.	60	Yes
8	Anti-immigration parties portray themselves as defenders of national sovereignty while indirectly reinforcing economic systems that depend on vulnerable and easily exploitable undocumented labor.	49	Yes
9	The Syrian refugee crisis intensified political disputes within the EU over burden-sharing and migration responsibility.	48	Yes
10	Merkel's government is suppressing media coverage of migrant-related crimes to limit public backlash and preserve political support for its immigration policies.	43	Yes
11	Russia and other external actors deliberately amplify immigration as a wedge issue to deepen political polarization and interfere in Western electoral outcomes.	35	Yes
12	Governments worsen shortages of essential workers, including doctors, by removing qualified immigrants over minor infractions to appeal to anti-immigration constituencies.	34	Yes
13	The EPP's opposition to the Migration Compact prevented European governments from reaching consensus, leaving migration governance in deliberate political paralysis.	29	No
14	Widespread public dissatisfaction toward migrants is a major force shaping voting behavior in Italy and pushing mainstream political parties toward increasingly restrictive immigration positions.	22	Yes
15	France's support for General Haftar over Libya's internationally recognized government is contributing to regional instability and migration pressures while avoiding shared responsibility for managing their consequences.	17	No

Table 28: Frame Articulations (FAs) for the Issue Frame: *Political Factors and Implications*, sorted by number of posts.

#	Frame Articulation (FA)	No. of Posts	Accepted
1	Liberals allow undocumented immigrants to vote.	104	Yes
2	Russia and other external actors deliberately amplify immigration as a wedge issue to deepen political polarization and interfere in Western electoral outcomes.	35	Yes
3	Democrats intentionally encourage illegal immigration to manufacture social chaos that they can then exploit for political gain.	23	Yes
4	Immigration is inherently a net economic burden on every country it affects.	14	No
5	The United Nations Migration Pact is a secret globalist instrument that surrenders national sovereignty over immigration to unelected international bodies, a threat most citizens have never been told about.	13	Yes
6	Child separation at the border is a politically manufactured controversy created to attack the government rather than a genuine humanitarian crisis.	12	Yes
7	ISIS and other terrorist organizations deliberately embed operatives within refugee and migrant flows entering Europe in order to infiltrate countries and carry out future attacks.	11	No
8	The European migrant crisis is not the result of humanitarian and geopolitical conditions, but a deliberate and coordinated effort under the Kalergi Plan to reshape Europe through engineered demographic change.	10	Yes
9	Pharmaceutical companies, global health organizations, and governments are deliberately using migrant populations as unmonitored carriers of disease, while suppressing information on migrant-linked transmission to avoid public alarm and scrutiny.	9	Yes
10	Migrants receive financial incentives and benefits unavailable to citizens in order to encourage migration.	8	No
11	Immigration statistics are systematically falsified by governments and international organizations to conceal the true scale of arrivals and suppress public resistance.	8	Yes
12	All drugs entering the United States are coming through Mexico via tunnels and immigrant caravans, making migration the primary source of the country's drug problem.	7	No
13	Daily rapes across Sweden and Europe are being committed by Muslim immigrants, a fact governments and media are deliberately concealing from the public.	6	Yes
14	Western societies are being deliberately destabilized through a coordinated combination of manufactured wars, false-flag operations, and engineered mass migration designed to erode democratic institutions.	5	Yes
15	George Soros funds and organizes migrant caravans, paying for transport, supplies, and protest salaries to engineer mass migration into Western countries.	4	Yes

Table 29: Frame Articulations (FAs) for the Issue Frame: *Misinformation & Conspiracy*, sorted by number of posts.