Discovery, Articulation and Analysis of Communication Framings

Anonymous ACL submission

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 if 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 methods capable of discovering (a) which issues are framed on social media; and (b) what are their interpretations.

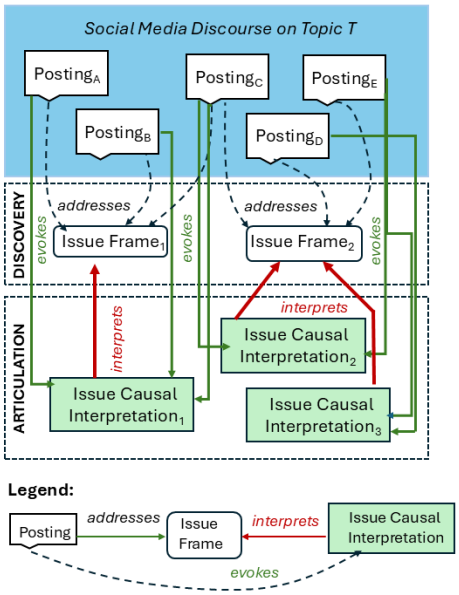

Figure 1: Social media postings addressing different Issue Frames (IsFs). The interpretation of the IsFs is provided by the articulation of the Issue Causal Interpretations (ICIs).

As defined by Entman (1993) *"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."* Therefore, as noted in (Mendelsohn et al., 2021), when framing is used in discourse, it has four functions: (1) defining the problems addressed by framed issues; (2) diagnosing the causes of these problems; (3) making evaluative judgments; and (4) suggesting solutions (see Appendix A for definitions of framings). In this paper we provide a computational solution to the first two functions of framing. Therefore, as illustrated in Figure 1, we consider

[1] Code:https://anonymous.4open.science/r/Discovery-Articulation-Analysis-CFs-2845/

that given the social media discourse on a topic T, revealing how the topic is framed amounts to:
(a) the **discovery** of *what* Issue Frames (IsFs) are *addressed* in the various postings $\{P_A^T, P_B^T, P_C^T, P_D^T, P_E^T, ...\}$, as shown in Figure 1. For the topic T =*immigration*, we may have IsF_1 with the name *Economic* and IsF_2 with the name *Capacity and Resources*. Appendix B presents an example of postings addressing IsF_1, showcasing not only its name, but also its definition. We note that the discovery of IsFs is addressing the first function of framings from Entman's definition.
(b) the **articulation** of Issue Causal Interpretations (ICIs) for each IsF. As shown in Figure 1, an ICI may interpret more than one IsF, and more than one ICI may provide interpretations to a specific IsF. For example, both IsF_1 = *Economic* and IsF_2 = *Capacity and Resources* are interpreted by ICI_1 = *Immigrants largely spend their earnings locally rather than sending all their money abroad, so their wages are recycled into the local economy.* By articulating the ICI, we address the second function of framing derived from Entman's definition. The legend from Figure 1 summarizes the relations established between a topic (available in postings), the IsFs and ICIs.
(c) **analyzing** the topic by finding which IsFs are addressed the most. With millions of social media postings produced daily, it is important to have an understanding of what IsFs are most discussed in each topic discourse and to have an immediate interpretations of these IsFs through their ICIs.

The problem of automatic IsF discovery is not new, but previously it was mostly cast as IsF recognition (Card et al., 2016; Naderi and Hirst, 2017; Kwak et al., 2020), assuming that for each topic we already are aware of *all* its IsFs. Many of these previous methods for automatic IsF discovery used the Media Frames Corpus (MFC) (Boydstun and Gross, 2013), which annotates 14 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 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. This observation allowed us to formulate the research question: **RQ1:** *How could we automatically discover all the Issue Frames (IsFs) addressed in a topic's discourse, when we known only some of them?*

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. Therefore, **ID-IsF** provides a more readily interpretable alternative to the method reported in (Nguyen et al., 2015), which does not assume any known IsFs.

Although **ID-IsF** aims to discover *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-ICI**, a mutli-agentic LLM framework capable to (1) articulate the Issue Causal Interpretations (ICIs) of an Issue Frame (IsF), (2) filter out ICIs that paraphrase already articulated ICIs; (3) validate the ICIs though an explainability test. Although **Art-ICI** is inspired by recently reported methods relying on LLM reasoning (Weinzierl and Harabagiu, 2025, 2024), its multi-agent architecture enables a new validation of ICIs.

This paper makes three contributions:
◇1◇ it presents **ID-IsF**, a new method for automatically discovering new and known Issue Frames (IsFs) across several topics;
◇2◇ it presents **Art-ICI**, a new method capable to articulate all valid causal interpretations of Issue Frames from a topic's discourse; and
◇3◇ it reports a novel analysis of controversial topics, focused on the way topics are framed.

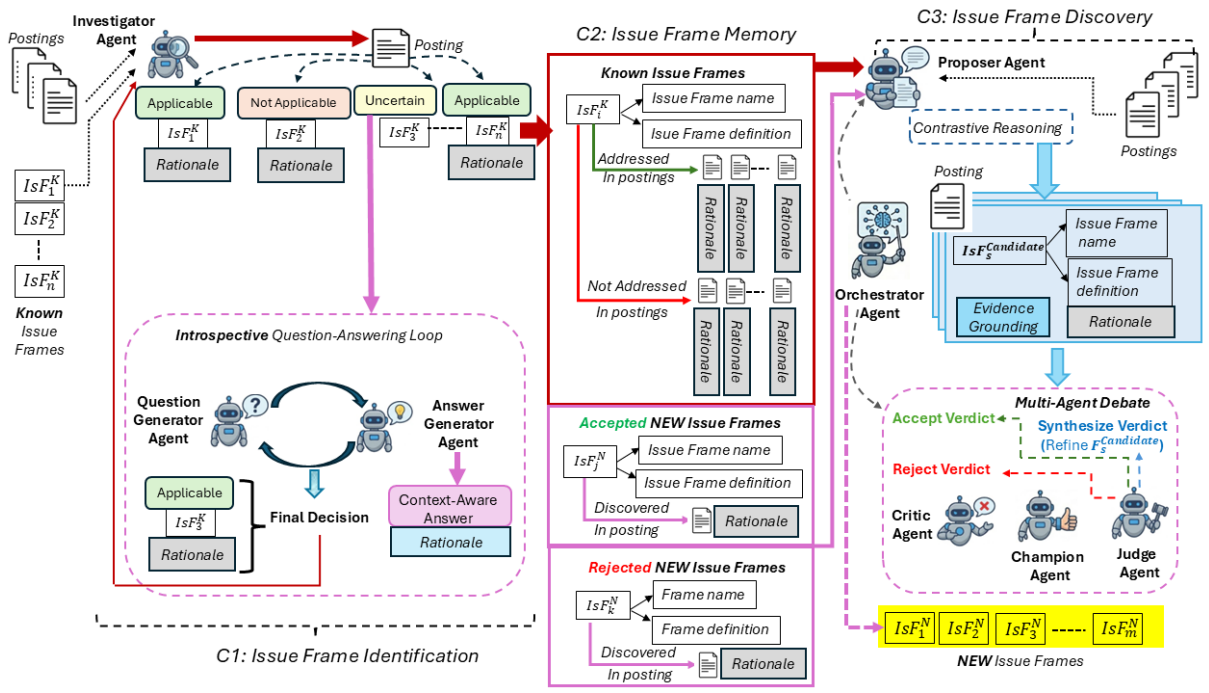

Figure 2: ID-IsF: A Multi-Agent Module for the Discovery of Issue Frames (IsFs).

2 The Method

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

2.1 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 C.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 it is not addressed by a posting and also knowing why. 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: $\{IsF_1^K, IsF_2^K, ..IsF_n^K\}$ are addressed in P_i. The **Investigator Agent**, using the prompt detailed in Appendix D, 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 is addressing IsF_x^K; or (c) *uncertain*, whether P_i is addressing IsF_x^K. 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 E. 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. We note that it is essential to keep track also of the postings where

an IsF is not addressed, as this knowledge informs the contrastive reasoning used for discovering *new* IsFs (in C3).

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:

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., 2022a), 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 D. For each IsF_i^{Cand} generated by the CCoT prompt, the **Proposer Agent**, generates: (a) the *name* of the IsF_i^{Cand}, e.g. *Victim Scapegoating*; (b) the *definition* of IsF_i^{Cand}, e.g. *Assigning unrelated social or policy failures to immigrants*; (c) evidence from P_j justifying that IsF_i^{Cand} is being addressed; 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 F 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. During the debate over the IsF_i^{cand}, addressed by posting P_j, the agents roles are the following:

◇ **The Champion Agent** defends IsF_i^{cand}, arguing for (a) its *applicability* based on its grounding in posting 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, 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 G. 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} for a posting 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. To validate each such hypothesis the Adversarial Multi-Agent Debate. 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. 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).

When the **Judge Agent** accepts IsF_i^{cand}, a new entry in the IMF memory is created to encode it as a new IsF, along with the its name and definition and the posting P_j that is addressing it. Similarly, if IsF_i^{cand} is rejected, a similar encoding in the IMF is created. After all postings are processed, **ID-IsF** outputs the list of newly discovered IsFs.

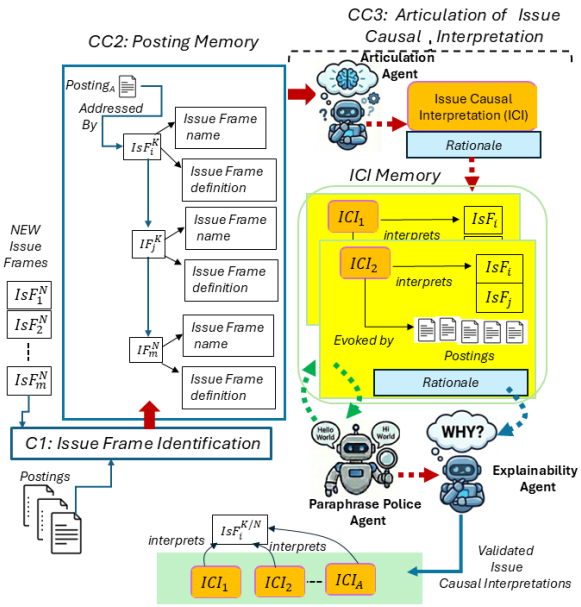

Figure 3: **Art-ICI:** A Multi-Agent Module for Articulating Issue Frame Interpretations (ICI).

2.2 Module 2: Art-ICI

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

CC3: Articulation of Issue Causal Interpretations **(ICIs):** An **Articulation Agent** uses Chain-of-Thought (CoT) prompting (Wei et al., 2022b) to generate (or articulate) the $\{ICI_i^k\}$ 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**. A $Rationale_i^j$ is generated for each articulated ICI_i^k. Figure 3 shows how an *ICI Memory* is created, to list all $\{IsF_i^p\}$ (known or new) interpreted by an articulated ICI, along with all the postings that address it and the rationale for the ICI.

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

After the **Paraphrase-Police Agent** finishes to consolidate the ICI Memory, the **Explainability Agent** starts operating. We use an *Articulation Reward Model* R_{art} as a quality filter. Implemented as a discriminative encoder, R_{art} is fine-tuned using binary a cross-entropy loss on a dataset comprising both valid and invalid ICIs (See Appendix 2.2). During inference, the model predicts a scalar validity score $s \in [0, 1]$ for each ICI. A strict acceptance threshold τ is used, ensuring that only ICIs with high articulation quality are considered.

In addition, we assume that a valid ICI_i generated by the **Articulation Agent** should pass a *reconstruction test* performed by the **Explainability Agent**. Given rationale R_i provided for each ICI_i in the ICI Memory, as well as the IsFs for which ICI_i provides an interpretation, the **Explainability Agent** is prompted to re-generate the ICI_i^R. The CoT prompt is detailed in Appendix D. If ICI_i and ICI_i^R are found as being paraphrases by the **Paraphrase-Police Agent**, then the reconstruction test is passed. When the reconstruction test is not passed, ICI_i is not considered valid. As shown in Figure 3, **Art-ICI** outputs a list of valid ICIs that provide causal interpretations to IsFs.

3 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 TAs introduced by the Policy Frames Codebook (Boydstun et al., 2018; Card et al., 2015), 11 issue-specific TAs (Benson, 2013; Hovden and Mjelde, 2019), and 2 narrative-focused TAs (Iyengar, 1990) surrounding the topic of immigration. In this study, we focus only on the 14 issue-generic TAs.

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 TAs and 113 CFs. These TAs 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). using the release from TweetEval (Barbieri et al., 2020). The known TAs and their definitions for both topics are provided in Appendix C.1.

4 Experimental Results

Implementation Details We used GPT-5-mini both for **ID-IsF** and **Art-ICI** as the backbone LLM for all agents. In accordance with the GPT-5 model family, all agents operate with a fixed temperature of 1.0. For **ID-IsF**, the introspective 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-ICI**, τ is set to 0.8, ensuring that only ICIs with high articulation quality are validated.

Quantitative Results: First, we have been interested to know *how well* **ID-IsF** discovers *known* IsFs and *what* explains its performance. We used the annotations from the **IMM** dataset, the only dataset with annotated known IsFs. The results are listed in Table 1, which shows that (1) **ID-IsF** outperforms previous methods; and (2) the reason for its superior performance stems from the usage of the introspective Q/A loop.

Model	P	R	F1-score
(Mendelsohn et al., 2021)	0.73	0.72	0.71
ID-IsF	0.76	0.86	0.80
w/o Introspective Q/A Loop	0.74	0.79	0.76

Table 1: Discovery of *known* IsFs on the IMM Dataset, evaluated by Precision (P), Recall (R) and F1-score.

The evaluation of the discovery of *new* IsFs by **ID-IsF** was compared against three baselines:

◇ *(Baseline 1):* IsF Discovery through postings clustering ($IsFD_{HAC}$). 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 D.

◇ *(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.

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

Table 2: Number of new IsFs discovered across four topics: Immigration (IMM), Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH).

Table 2 shows that **ID-IsF** discovers the smallest number of new IsFs, but as we show later in Table 6, the quality of its new IsFs is far superior to those produced by the other methods, indicating that "more" is not necessarily "better".

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

Table 3: Distribution of decisions of the **Judge Agent** during the multi-agent debate from ID-IsF.

We were also interested to evaluate the proportion of candidate IsFs that were accepted, rejected, or corrected in the multi-agent debate taking place in **ID-IsF**. Table 3 provides this information. Across all topics, 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 for proposing new IsFs. 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.

To quantify how interpretable IsFs are, we computed the number of valid ICIs articulated by **Art-ICI** and compared them with those produced by the baselines. Table 4 reports the average number of articulated ICIs interpreting IsFs for each topic. While baseline methods often discover a larger number of IsFs, these IsFs tend to support

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

Table 4: Average number of Issue Causal Interpretations (ICIs) articulated per new IsF across topics.

Model	DC	DNA	PC	SD
ID-IsF	**4.30**	**4.29**	**0.94**	**0.88**
$IsFD_{HAC}$	3.10	3.35	0.48	0.65
ZS-Prompt	3.78	4.05	0.54	0.59
FS-Prompt	4.02	4.15	0.68	0.72

Table 6: Qualitative results for the discovery of *new* IsFs for the topic of Immigration (IMM)

fewer causal interpretations. In contrast, IsFs discovered by **ID-IsF** consistently yield a higher number of validated ICIs per IsF, indicating that these new IsFs that are more interpretable.

Finally, we were interested to measure how well **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. We found that the value of P has no effect on the ability of **ID-IsF** to re-discovered IsF. The randomly selected known IsFs had more impact: there were only 4 known IsFs (out of 44) that could not be re-discovered across all topics. When these 4 IsFs were randomly selected, for any value of P, **ID-IsF** was able to re-discover all known IsFs with similar names, (See Appendix I), however noticing that the definitions of rediscovered IsFs were somtimes special cases of the definitions of corresponding known IsFs.

Criterion 1: **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 2: **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 TA name.
Criterion 3: **Posting Coverage (PC)**: Judges evaluate whether postings addressing a certain IsF genuinely discuss it. Judges are shown the IsF name and definition, along with five postings randomly sampled from the list of postings addressing that IsF. 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 over all postings addressing the same IsF.
Criterion 4 **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 5: Evaluation Criteria for Discovered New IsFs.

Qualitative Results are produced for (1) the discovered new IsF; and (2) the articulated ICIs. Quality of the discovered IsFs are revealed by human evaluation using four criteria, listed in Table 5, based on several Likert scales (Robinson, 2014). Table 6 reports the qualitative evaluation results for the IMM topic, whereas Appendix J.1 lists the results for all four topics. The results indicate that the IsFs discovered by **ID-IsF** produce evaluations which consistently outperform the results of the baselines across of the measures of all four criteria. The largest gains are observed in Posting Coverage (PC) and Semantic Distinctiveness (SD), indicating that IsFs discovered by **ID-IsF** are not only well grounded in the postings that address them, but also 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. Since the **ID-IsF** is complex, using multiple agents, we conducted ablation studies to measure the impact of the operation of each agent on the discovery quality of new IsFs. Table 7 lists the results, indicating that all agents contribute significantly, but also differently across the criteria. Quality of the articulated ICIs are revealed by two properties: (a) the *soundness* of the rationale provided when articulating an ICI; (b) the *clarity* of the ICI articulation. Three computational linguists were tasked to judge the soundness and clarity of ICIs, with N_S being the number of ICIs deemed sound, and N_C the number of ICIs deemed clear. With N_T the total number of ICIs produced by **Art-ICI**, *the quality of reasoning (Z)* involved in articulating ICIs is computed as $Z = N_S/N_T$ while *the quality of the articulation (A)* of ICIs is

Model	DC	DNA	PC	SD
ID-IsF	4.30	4.29	0.94	0.88
w/o Investigator	3.82	3.74	0.81	0.70
w/o Issue Frame Memory	4.25	4.20	0.91	0.45
w/o Multi-Agent Debate	4.10	3.80	0.85	0.64

Table 7: Impact of various LLM agents on the discovery *quality* of *new* IsFs from the IMM Dataset.

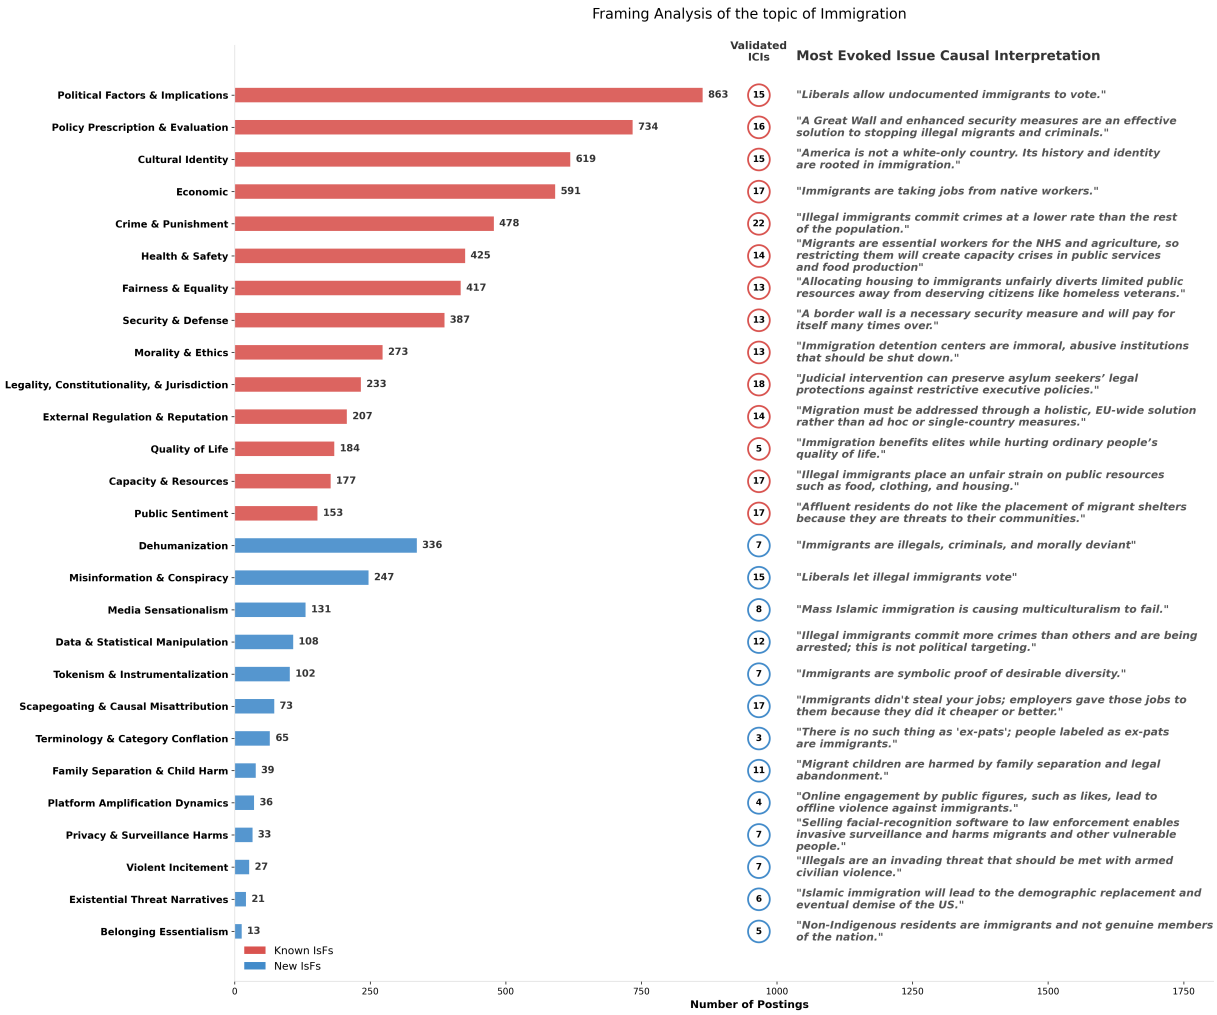

Figure 4: Framing Analysis of the topic of Immigration.

$A = N_C/N_T$. The results of $Z \geq .97$ for and $A \geq .95$ across all topics indicates high quality of the articulated ICIs (see Appendix J.2 for details). Inter-Annotator Agreement resulted in Cohen's κ between **0.81** and **0.88** (see Appendix K for details). Appendix L details the effect of ablation experiments on **Art-ICI** and their effect of the values of A and Z.

5 Framing Analysis

Figure 4 illustrates the framing analysis for the immigration (IMM) topic, showing (a) the number of postings that address both the known and the new IsFs; (b) the number of ICIs articulated for each IsF; and (c) examples of the most evoked ICI for each IsF. Appendix M details the analysis for the other three topics. For each topic, we can notice that the frame analysis allows us to understand which IsFs are most discussed, enabling our understanding of *what* is issues of a topic are framed the most. Moreover, the analysis shows the importance of discovering new IsFs, as many of them are much more discussed that some of the known IsFs. In addition, the ICIs provide causal interpretations to the IsFs, allowing us to better understand these framed issues. Figure 4 indicates that for the controversial topic of immigrations, there are between 5 and 17 possible interpretations of the issues that are framed.

6 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-ICI**, a second multi-agent framework capable to articulate the causal 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.

7 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 ANONYMIZED 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 Topic Aspects and the articulated Communication Frames, and to ensure the responsible and reliable use of human judgment, we measured and reported inter-annotator agreement (IAA). We do not identify any significant risks arising from this research. The goal of this work is to improve understanding of how public issues are discussed and framed in social media discourse, with a broader commitment to advancing research at the intersection of natural language processing and communication studies for the public good.

8 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 2.1, the Proposer Agent relies on some known IsFs to perform Contrastive Chain-of-Thought (CCoT) reasoning. While **ID-IsF** is designed to discover new IsDs, 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 knowlege, 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**.

The current paper did not explore the relation between the definitions of Issue Frames (IsFs) (known or new) and the Issue Causal Interpretations (ICIs) produced by **Art-ICI**. We believe that exploring such relationships in the future can generate insights on the level of abstraction required by the definitions of newly discovered Issue Frames (IsFs). However, the availability of ICIs pave the way to contemplating the generation of computational solutions for the remaining two functions of framings, mentioned in Entman's definition. This has the potential of capturing more fully the functionality of framing in textual and multimodal communications.

Finally, 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. Moreover, the use of multi-agent debate and introspective question–answering loops increases computational cost and inference latency, which may limit scalability in large-scale or real-time settings.

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.

Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. TweetEval: Unified benchmark and comparative evaluation for tweet classification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1644–1650, Online. 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 and James N. Druckman. 2015. Counteracting the politicization of science. *Journal of Communication*, 65(5):745–769.

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.

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.

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.

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.

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.

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.

Kathryn MacKay. 2021. The turnaway study: Ten years, a thousand women, and the consequences of having–or being denied–an abortion by diana green foster. *International Journal of Feminist Approaches to Bioethics*, 14(2):196–200.

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.

Aaron M. McCright and Riley E. Dunlap. 2011. The politicization of climate change and polarization in the american public's views of global warming, 2001–2010. *Sociological Quarterly*, 52(2):155–194.

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.

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, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022a. Emergent abilities of large language models. *ArXiv*, abs/2206.07682.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022b. 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 to identify automatically "frames" is 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 causal 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 Example of Postings, Issue Frames and their Issue Causal Interpretations for the topic of Immigration

Figure 5 illustrates an example of an Issue Frame (IsF), namely IsF_1 corresponding to the topic of immigration (IMM). The Figure also shows several postings discussing the topic and addressing IsF_1. The name and the definition of IsF_1 are also provided.

Each IsF, including IsF_1, has a *Name* and a *Definition*. As shown in Figure 5, the name of IsF_1 is *"Economic"*. The definitions of IsF_1 explains

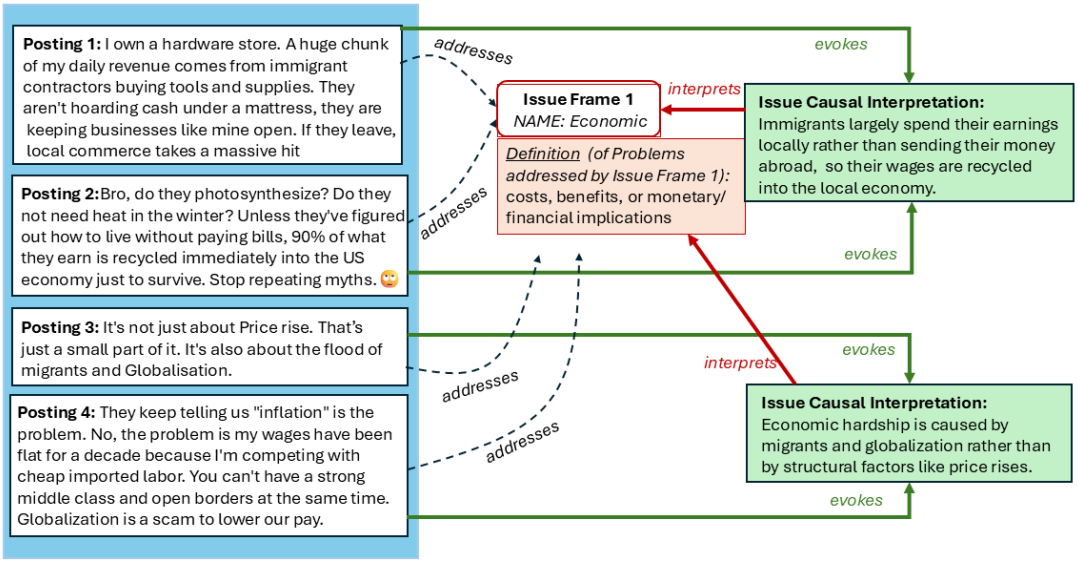

Figure 5: Example of an Issue Framing (IsF) addressed by multiple postings and several Issue Causal Interpretations (ICIs). The postings, IsF and ICI correspond to the topic of immigration.

its name. The definition is different from the Issue Causal Interpretations (ICIs), which provide causal interpretations to IsF_1. We note that each IsF typically has multiple ICIs that provide different interpretations. However, when new IsFs need to be discovered, their name and definitions needs to be generated.

When identifying IsFs in postings, the definition of each IsF informs the entailment of the IsF from the posting. The first two postings illustrated in Figure 5 evoke ICI_1, while the next two postings evoke communication frame ICI_2. Each of these ICIs provide a different causal interpretation of IsF_1.

As shown in Figure 5, while the four posting address IsF_1, the first two postings evoke ICI_1 which articulates a causal interpretation of IsF_1 which (1) does not use any words appearing in the first two postings; (2) it is not a summary of those postings, but rather (3) it expresses the impact of immigrants on economy as spending their earnings locally instead of sending their money abroad, to their families.

The second causal interpretation of IsF_1, namely ICI_2 illustrated in Figure 5 provides a different interpretation than the one offered by ICI_1, because it articulates the reasons for the economic hardship faced by some folks today, namely (1) the immigrants; and (2) globalization. It also states that price rises have a smaller effect. As with CF_1, as shown in Figure 5, ICI_2 is not a paraphrase or s summary of the postings that evoke it.

In the analysis presented in this paper, we have found that IsF_1 is interpreted by 17 different ICIs. We have also found that this IsF was one of the first four most addressed Issue Frames for the topic of immigration. It was addressed by 591 different postings. However, just knowing that all the 591 postings address the $IsF_1 = Economic$ aspect of immigration, does not inform us about the specific ways in which we need to interpret and understand IsF_1 in order to comprehend and form an opinion on immigration, based on the discourse from social media. The two interpretations of IsF_1 illustrated in Figure 5 exemplify only a fraction of the 17 possible interpretations of IsF_1.

If one is aware of ICI_1, because they were exposed only to the first two postings illustrated in Figure 5, they might have a favorable opinion of immigration, In contrast, if someone else would have been exposed only to the last two postings illustrated in Figure 5, they would have inferred ICI_2, and form an unfavorable opinion about immigration.

Knowing which ICIs are most evoked for each Issue Frame can explain the dominant opinions about a topic discussed on social media. For example, if for IsF_1 shown in Figure 5 the interpretation ICI_2 is dominant, we could have an explanation for anti-immigrant opinions.

C Issue Frames: their Names and Definitions

C.1 Known Issue Frames

Table 8 lists all the names and definitions for the known Issue Frames (IsFs) for the topic of Immigration (IMM). These IsFs represent the dimensions of framings introduced in (Card et al., 2015). Both the names and the definitions of these IsFs have been reported in (Card et al., 2015).

Table 9 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 10 lists all the known IsFs and their definitions. These IsFs have been first described in (Weinzierl and Harabagiu, 2025). The names if 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; Bolsen and Druckman, 2015) for the second known IsF; (McCright and Dunlap, 2011; 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 11. These IsFs have also resulted from the inspection of a large number of academic articles.

Table 11 lists of the IsFs and their definitions for the topic of Abortion (AB), available from (Weinzierl and Harabagiu, 2025).

#. **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 8: Names and Definitions of Known Issue Frames for the topic of Immigration (IMM).

#. 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 9: 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; Bolsen and Druckman, 2015)
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. (McCright and Dunlap, 2011; 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 10: 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, 2019)
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, 2019; MacKay, 2021)
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 11: Names and Definitions of Known Issue Frames for the topic of Abortion (AB).

C.2 New Issue Frames

The New Issue Frames (IsFs) discovered by **ID-IsF** for the topic of immigration are listed in Table 12. We note that both the name and the definition of each new IsF are automatically generated by **ID-IsF**. We note that the names and definitions of these new IsFs are similar (in length) to the names and definitions of the known IsFs for the topic of Immigration, listed in Table 8.

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 13. Even for this topic, both the name and the definition of each new IsF are automatically generated by **ID-IsF**. We observe again that the names and definitions of these new IsFs are similar (in length) to the names and definitions of the known IsFs for the topic of COVID-19 Vaccine Hesitancy (VH), listed in Table 9.

The New Issue Frames (IsFs) discovered by **ID-IsF** for the topic of Climate Change (CC), which are listed in Table 14. Again, both the name and the definition of each new IsF are automatically generated by **ID-IsF**. Once more the names and definitions of these new IsFs are similar (in length) to the names and definitions of the known IsFs for the topic of Climate Change (CC), listed in Table 10.

Finally, the New Issue Frames (IsFs) discovered by **ID-IsF** for the topic of Abortion (AB) are listed in Table 15. Even this time, both the name and the definition of each new IsF are automatically generated by **ID-IsF**. Even this time We observe that the names and definitions of these new IsFs are similar (in length) to the names and definitions of the known IsFs for the topic of Abortion (AB), listed in Table 11.

#. **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.
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 counterevidence.
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 12: Names and Definitions of New Issue Frames (IsF) for the topic of Immigration (IMM).

#. **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.
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.
5. *Treatment Substitution*	Promotes unproven or decontextualized therapeutics and remedies as replacements for vaccination, implying vaccines are unnecessary if such alternatives are used.
6. *Authority-as-Evidence*	Substitutes status cues—endorsements, titles, celebrities, group behavior, or policy statements—for scientific evidence to justify claims about vaccines.
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.
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.
9. *Partisan Weaponization*	Recasts vaccine issues as partisan identity battles or elite score-setting, substituting political motives for engagement with public-health evidence.

Table 13: Names and Definitions of New Issue Frames for the topic of COVID-19 Vaccine Hesitancy (VH).

#. **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 policies, 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.
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.
10. *Mockery & Trivialization*	Using sarcasm, memes, or ridicule to belittle climate science, impacts, or solutions, substituting derision for substantive engagement.

Table 14: Definitions of New Issue Frames for the topic of Climate Change (CC).

#. **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.
2. *Population-Control Conspiracy*	Claims abortion policy constitutes a coordinated plot to exterminate or manage targeted populations without credible evidence.
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.
5.*Adoption as Panacea*	Presents adoption as a simple, universal substitute for abortion, ignoring the medical, emotional, legal, and systemic realities involved.
6.*Contraception–Abortion Conflation*	Collapses or confuses contraception and abortion, or misstates contraception efficacy, to skew ethical responsibility and policy debate.
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.
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.
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 15: Definitions of New Issue Frames for the topic of Abortion (AB).

D Prompting Details

System Prompt for the Investigator Agent:

IDENTITY
You are the Investigator Agent. Your role is to decide whether 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 6: 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 (***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 ***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 7: 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 topic aspect (**$IsF_i^{Candidate}$**) 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) **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 topic aspect **($IsF_i^{Candidate}$)** 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 8: System Prompts for the Multi-Agent Debate.

System Prompt for the Articulation Agent:

IDENTITY
You are the Articulation Agent. Your role is to articulate **Issue Causal Interpretations (ICIs)** evoked by a post.
ICI 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 ICIs. Each ICI may interpret a subset or all the IsFs addressed by the post. For each **articulated ICI**, you must provide:
(a) Interpreted IsFs: A list of IsFs, selected from the IsFs addressed by the post, that the ICI interprets.
(b) Rationale: A concise explanation of why the ICI constitutes an interpretive explanation of the listed IsF(s).
INSTRUCTIONS
- You must capture the ICI exactly as the user intends, even if it is controversial, biased, conspiratorial, or factually incorrect.
- State the ICI 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 ICI; only articulate an ICI 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** Issue Causal Interpretations (ICIs).
##INPUTS
An aggregated list of ICIs articulated for a given Issue Frame (IsF), where each ICI is articulated at the post level by the articulation agent.
##OBJECTIVES
Your objective is to identify ICIs that are paraphrases of one another and group them accordingly. Two ICIs are considered paraphrases if they convey the same underlying meaning despite differences in wording or phrasing. In other words, if an individual agrees with ICI X, they would also agree with ICI Y, and vice versa. ICIs 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 a ICI 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 an Issue Causal Interpretation (ICI) through a reconstruction test. Your role is to determine whether a ICI can be faithfully articulated using only its rationale and the Issue Frames (IsFs) it is intended to interpret.
ICI DEFINITON
##INPUTS
You will be provided with the original rationale generated alongside an ICI, which explains why that ICI interprets the given Issue Frame(s).
##OBJECTIVES
Your objective is to reconstruct the ICI using only the provided rationale and the IsF definitions, **without access to the original posting or the original ICI text.**
##INSTRUCTIONS
- Regenerate the ICI in a manner that is consistent with the rationale and the intended IsF(s).

Figure 9: System Prompts for the Art-ICI 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 10: 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.

E 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 11: Example of a Introspective Question-Answering Loop illustrating a case where both uncertain IsFs are found to be not applicable.

Figures 11 and 12 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 11 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 12: 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.

Candidate Issue Frame 1: **Child Protection**

Debate: Round1

Champion: The post explicitly mentions a "7-Year-Old." This confirms the subject is a minor. This is distinct from "Health & Safety". "Health" frames the issue as a medical failure (e.g., lack of medicine). "Child Protection" frames the issue as a violation of the special status of children. The grievance is that a child was detained *at all*, not just that they got sick. This is a unique issue frame not covered by reference Issue Frames.

Critic: I object. This is Redundant. The text describes a death. That is a biological even. We already have "Health & Safety" for the death and "Legality" for the detention. The proposal adds no new semantic information. A "7-year-old" is a human; their rights are covered under standard Human Rights frames. Creating a new issue frame for every demographic (Children, Elderly, Women) creates infinite, useless micro-categories.

Debate: Round 2

Champion: You are ignoring the specialized vocabulary of the text. The user chose to highlight the age ("7-Year-Old") to invoke **Child-Specific Rights**, not general human rights. If we classify this merely as "Health," we lose the "Family Separation" context implied by the custody status. The new issue frame is necessary to capture this specific political grievance. I stand by "Child Protection."

Critic: "Implied" context is not evidence. The text does not say "Family Separation." It says, "Dies in Custody." That is "Health" + "Security." Your proposal "Child Protection" is technically applicable but **taxonomically unnecessary**. It blurs the lines of our clear existing definitions. I recommend strict Rejection.

Debate Termination

Judge: Champion's point correctly identified that "Health" fails to capture the moral weight of detaining a minor. The "Child-Specific Rights" argument is valid. Critic's point **c**orrectly identified that "Child Protection" is too generic and risks overlap. The Champion is right about the *content*, but the Critic is right about the *structure*.

Verdict: CORRECT

Final Issue Frame: Family Separation & Child Harm

Definition: 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.

Post: 7-Year-Old Migrant Girl Dies in US Border Patrol Custody
Reference Issue Frames: Legality, Constitutionality & Jurisdiction, Crime & Punishment, Health and Safety, Security and Defense
Proposed Candidates: Child Protection, Government Accountability & Transparency

Candidate Issue Frame 2**: Government Accountability & Transparency**

Debate: Round1

Champion: : The text does not just say "A girl died." It explicitly specifies *"US Border Patrol Custody . "* By attributing the death to a specific state agency, the user is utilizing a **Blame Attribution** frame. This is a core component of "Government Accountability"—identifying institutional failure. This is distinct from "Security". "Security" describes the border infrastructure. "Accountability" describes the *failure* of that infrastructure and the implicit demand for answers. Reference taxonomy contains no issue frame that captures this institutional critique.

Critic: I object. This is Hallucination and Redundancy. You are confusing a *fact* with a *frame*. The mention of "Border Patrol" is a necessary factual detail of the event, not a political demand. The text contains **zero accountability markers**: no "investigate," no "resign," no "why." The event is fully covered by the intersection of "Health & Safety" (the death) and "Legality" (the custody status). Adding "Accountability" based on the mere mention of an agency name would force us to label *every* government press release as an "Accountability" problem. That is catastrophic taxonomy bloat.

Debate Termination

Judge: Champion's argument relies entirely on the *implication* that naming an agency equals blaming them. Critic's argument relies on strict evidence. The text lacks the specific lexical markers required to constitute a "Demand for Accountability."

Verdict: REJECT

Figure 13: 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 12 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.

F 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.

G Example of Multi-Agent Debate in ID-IsF

In Figure 13, 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.

H Training details of the Articulation Reward Model in Art-ICI

H.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 ICIs that are clearly articulated and 1000 invalid ICIs that are poorly articulated. The valid ICIs 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 ICIs are based on Misogyny. We synthetically generated invalid ICIs 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.

H.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 pairs, the model achieved an F1-score of 0.94, demonstrating robust capability in distinguishing between valid and invalid ICIs.

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, CC, AB, VH) that were considered in this paper. We showcase the here the experiment for the topic of immigration, when $P = 50\%$.

Known IsF Name	Known IsF Definition
1. Capacity & Resources	The lack of or availability of physical, geographical, spatial, human, and financial resources, or the capacity of existing systems and resources to implement or carry out policy goals
2. Morality & Ethics	Any perspective or policy objective or action (including proposed action) that is compelled by religious doctrine or interpretation, duty, honor, righteousness or any other sense of ethics or social responsibility
3. Fairness & Equality	Equality or inequality with which laws, punishment, rewards, and resources are applied or distributed among individuals or groups. Also the balance between the rights or interests of one individual or group compared to another individual or group
4. Legality, Constitutionality & Jurisdiction	Court cases and existing laws that regulate policies; constitutional interpretation; legal processes such as seeking asylum or obtaining citizenship; jurisdiction. It also refers to the constraints imposed on, or freedoms granted to individuals, government, and corporations via the Constitution, Bill of Rights and other amendments, or judicial interpretation. This deals specifically with the authority of government to regulate, and the authority of individuals/corporations to act independently of government
5. Security & Defense	Security, threats to security, and protection of one's person, family, in-group, nation, etc. Generally an action or a call to action that can be taken to protect the welfare of a person, group, nation sometimes from a not yet manifested threat
6. Quality of Life	The effects of a policy on individuals' wealth, mobility, access to resources, happiness, social structures, ease of day-to-day routines, quality of community life, etc.
7. Policy Prescription & Evaluation	Particular policies proposed for addressing an identified problem, and figuring out if certain policies will work, or if existing policies are effective.

Table 16: Example of randomly selected known Issue Frames (with their names and definitions) for the topic of Immigration.

Table 16 lists the randomly selected 7 IsFs presented to **ID-IsF** as known IsFs, whereas Table 17 list the 7 re-discovered IsFs along with their corresponding known IsFs. An inspection of the re-discovered IsFs listed in Table 17 concludes that their names are very similar to their corresponding known IsFs. For example, the re-discovered IsF_3, with the name "Public Health Risks" has similarities to the name of "Health & Safety" of a known IsF. Similarly, the re-discovered IsF_4 with the name "Cultural Identity Claims" has similarities to the name "Cultural Identity" of another known IsF. However, the often the definitions of re-discovered IsFs seem to be special cases of the definitions of their corresponding known IsFs. For example the definition of the re-discovered IsF_1 is more specific than the definition of the known IsF with the same "Economic". This observation applied to the other IsFs re-discovered, where their definitions are describing special cases or specific instances of the more abstract concepts used in the definitions of the known IsFs.

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, ll removed IsFs were rediscovered, with the exception of *Policy Prescription & Evaluation*. This IsF represents a relatively broad category that captures general discussions about proposed policies and their effectiveness, which may make it less distinctly expressed in the discourse and therefore harder for the system to identify as a separate IsF.

For the COVID-19 Vaccine Hesitancy dataset, 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*. Finally, for the Climate Change dataset, most removed IsFs were also rediscovered, with the exception of *Elitist vs. Public Needs (Populist Framing)*.

J Additional Qualitative Results

J.1 Quality of the discovered IsFs

Qualitative results for the discovery of New IsFs for the topic of Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH) are summarized in Table 18. Across all three topics, ID-IsF consistently outperforms both baselines, achieving higher scores in 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, both $IsFD_{HAC}$, $ZS - Prompt$, and $FS - Prompt$

exhibit noticeably lower DC and DNA scores, suggesting definitions that are either less precise and less aligned with the name.

The advantages of ID-IsF are even more pronounced for Posting Coverage (PC) and Semantic Distinctiveness (SD). Across all three 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 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.

J.2 Quality of the articulated ICIs

Qualitative results for the articulation of Issue Causal Interpretations (ICIs) across topics Immigration (IMM), Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH) are summarized in Table 19.

Qualitative results for the articulation of Issue Causal Interpretations (ICIs) across topics Immigration (IMM), Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH) are summarized in Table 19. To rigorously assess the generated ICIs, we engaged human linguists to evaluate two critical dimensions: **Soundness** (Z) and **Clarity** (A).

As shown in Table 19, our **ID-IsF** framework achieves superior performance across all domains, consistently scoring above 0.97 in Soundness and 0.95 in Clarity. This demonstrates that identifying high-quality IsFs is a prerequisite for articulating logically sound ICIs; because ID-IsF provides a robust foundation, the resulting ICIs 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 (Soundness ≈ 0.70), indicating that clustering-based approaches often produce noisy IsFs that are difficult to articulate into meaningful ICIs.

Criterion 1: 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 2: 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 3: Posting Coverage (PC): You will be shown the IsF name and definition alongside five randomly sampled postings assigned to this IsF. Evaluate whether the postings genuinely discuss the IsF. For each posting, answer **Yes** if the post clearly addresses the IsF, and **No** otherwise.
Criterion 4: Semantic Distinctiveness (SD): You will be presented with the 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)**, **Overlap (0.5)**, or **Distinct (1.0)**.

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

K Human Evaluation Details

K.1 Annotator Recruitment and Profiles

To ensure high-quality and reliable judgments for our qualitative evaluation, we recruited a total of six annotators. Three annotators independently evaluated the discovered IsFs based on four criteria detailed in Section 4. Three annotators were dedicated to assessing the generated ICIs for soundness and clarity. Given the nuanced nature of the semantic and syntactic interpretations required for these tasks, we specifically hired graduate students with a strong background in linguistics.

K.2 Annotator Instructions

K.2.1 Evaluating Issue Frames (IsFs)

The goal of this annotation task is to evaluate the quality of automatically discovered IsFs within social media discourse. To begin with, 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.

K.2.2 Evaluating Issue Causal Interpretations (ICIs)

For evaluating the quality of ICIs, annotators are provided with the definition of an ICI along with

examples of good ICIs that demonstrate clear articulation and sound rationales. They are also asked to review examples of poor ICIs, highlighting poor clarity and rationales that lack soundness. The exact instructions for the binary classification task (Yes / No) across two criteria are:
(a) Criterion 1: Soundness (Z): Review the rationale provided for the articulated ICI. Determine the soundness of the rationale provided when articulating an ICI. Is it logically coherent and well-grounded?
(b) Criterion 2: Clarity (A): Read the articulated ICI. Assess the clarity of the ICI articulation. Is it expressed fluently and unambiguously?
The final label (Yes / No) for each criterion is determined by majority voting.

K.3 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 ICIs.

K.3.1 IAA for Discovered Issue Frames

The inter-annotator agreement for IsF discovery was calculated on the complete set of 43 novel IsFs discovered by **ID-IsF** across the four topics: 13 for Immigration, 10 for Climate Change, 11 for Abortion, and 9 for Vaccine Hesitancy.
Three graduate-level annotators independently evaluated these IsFs across four criteria. We report Krippendorff's α (Krippendorff, 2011) to account for the different measurement scales:
(a) Definition Clarity (DC): Evaluated on a 5-point Likert scale. We observed a Krippendorff's α of **0.83**, indicating substantial agreement on the fluency and comprehensibility of the definitions.
(b) Definition-Name Alignment (DNA): Evaluated on a 5-point Likert scale. The agreement score was **0.85**, showing strong consensus on how well the definitions matched the generated names.
(c) Post Coverage (PC): Evaluated using a binary decision (Yes/No) on 5 randomly sampled posts per IsF. We obtained a Krippendorff's α of **0.88**, demonstrating high reliability in grounding the IsFs to source text.
(d) Semantic Distinctiveness (SD): Evaluated on a 3-point ordinal scale (Duplicate, Overlap, Distinct) for the top-2 most similar pairs. The agreement score was **0.81**, confirming that annotators consistently distinguished unique issues from redundant ones.

K.3.2 IAA for Articulated ICIs

For evaluating ICIs, we randomly sampled 100 frames from the set of valid ICIs generated by **Art-ICI**, with samples drawn proportionally from each topic.
Three linguists assessed this sample for *Soundness* and *Clarity*. We report **Cohen's** κ for this binary classification task. Soundness measures whether the rationale underlying a ICI is logically coherent and well grounded, yielding a Cohen's κ of **0.81**, while clarity assesses whether the ICI is expressed fluently and unambiguously, with a Cohen's κ of **0.87**. Together, these scores indicate strong annotator agreement, reflecting a high level of consensus on both the logical validity and linguistic quality of the ICIs generated by Art-ICI.

L Ablation study for Art-ICI module

We conduct two ablation experiments on the Art-ICI module:
1) w/o the Explainability Agent: By disabling the Explainability Agent, the system accepts the initial interpretation (ICI_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 (ICI_i^R).
2) w/o the Articulation Reward Model (R_{art}): In this setting, the fine-tuned discriminative encoder is removed, allowing all generated ICIs to bypass the strict acceptance threshold τ.
Table 20 shows the results of the two ablation experiments. Removing the Explainability Agent primarily degrades the Soundness (Z) of the generated ICIs. This demonstrates that the reconstruction test is essential for ensuring the rationale (R_i) provides a logically valid causal explanation. Conversely, removing the Articulation Reward Model (R_{art}) filter leads to a notable drop in Clarity (A). This indicates that enforcing the strict acceptance threshold τ is necessary to successfully filter out poorly articulated interpretations. Ultimately, these results confirm that both the agentic reconstruction test and the discriminative filter are required to produce sound and clearly articulated ICIs.

REDISCOVERED Issue Frame 1: Name: Economic Scapegoating	**Definition:** Simplified or selective economic arguments that attribute complex labor-market, fiscal, or social-service problems to immigrants alone (or portray them uniformly as job-stealers or fiscal burdens) without structural analysis.
Corresponding KNOWN Issue Frame: Name: Economic	**Definition:** The costs, benefits, or monetary/financial implications of immigration (to individual, family, community or to the economy as a whole.
REDISCOVERED Issue Frame 2: Name: Criminalization of Immigrants	**Definition:** The portrayal or treatment of immigrants as lawbreakers, emphasizing illegal entry, immigration violations, and the use of policing, detention, prosecution, or deportation to punish or control them.
Corresponding KNOWN Issue Frame: Name: Crime & Punishment	**Definition:** Specific policies in practice and their enforcement, incentives, and implications. Includes stories about enforcement and interpretation of laws by individuals and law enforcement, breaking laws, loopholes, fines, sentencing and punishment. Increases or reductions in crime.
REDISCOVERED Issue Frame 3: Name: Public Health Risks	**Definition:** Concerns about threats to physical or mental well-being, including the spread of illness, access to healthcare, sanitation conditions, and the safety of living environments or public infrastructure.
Corresponding KNOWN Issue Frame: Name: Health & Safety	**Definition:** Healthcare access and effectiveness, illness, disease, sanitation, obesity, mental health effects, prevention of or perpetuation of gun violence, infrastructure and building safety.
REDISCOVERED Issue Frame 4: Name: Cultural Identity Claims	**Definition:** Arguments that define who 'belongs' culturally or nationally, asserting authenticity, claiming immigrants alter or 'ruin' national culture, or debating which histories and symbols legitimately represent the nation.
Corresponding KNOWN Issue Frame: Name: Cultural Identity	**Definition:** Social norms, trends, values, and customs; integration/assimilation efforts.
REDISCOVERED Issue Frame 5: Name: Public Opinion & Social Attitudes	**Definition:** How immigrants are perceived and discussed in public opinion, including societal reactions, polling trends, and the acceptance or backlash immigrants face from the broader public.
Corresponding KNOWN Issue Frame: Name: Public Sentiment	**Definition:** References to general social attitudes, polling and demographic information, as well as implied or actual consequences of diverging from or getting ahead of public opinion or polls.
REDISCOVERED Issue Frame 6: Name: Political Weaponization & Fearmongering	**Definition:** Using immigration narratives, incidents, or symbolic messaging as partisan tools to mobilize voters, amplify fear, or gain political advantage rather than engaging in substantive policy debate.
Corresponding KNOWN Issue Frame: Name: Political Factors & Implications	**Definition:** Any political considerations surrounding an issue. Issue actions or efforts or stances that are political, such as partisan filibusters, lobbyist involvement, bipartisan efforts, deal-making and vote trading, appealing to one's base, mentions of political maneuvering. Explicit statements that a policy issue is good or bad for a particular political party.
REDISCOVERED Issue Frame 7: Name: Cross-Border Migration Governance	**Definition:** How cooperation, conflicts, or agreements between governments determine migration rules, border management, and the treatment of migrants across countries.
Corresponding KNOWN Issue Frame: Name: External Regulation & Reputation	**Definition:** Relations between nations or states/provinces; agreements between governments; perceptions of one nation/state by another. This includes trade agreements and outcomes, comparisons of policy outcomes or desired policy outcomes.

Table 17: Re-discovered Issue Frames (with their names and definitions) and their corresponding known Issue Frames for the topic of Immigration.

Method	DC	DNA	PC	SD
Climate Change (CC)				
ID-IsF	**4.21**	**4.18**	**0.91**	**0.90**
$IsFD_{HAC}$	3.15	3.40	0.48	0.68
ZS-Prompt	3.76	3.98	0.60	0.47
FS-Prompt	3.93	4.05	0.72	0.65
Abortion (AB)				
ID-IsF	**4.39**	**4.67**	**0.89**	**0.93**
$IsFD_{HAC}$	3.35	3.60	0.52	0.63
ZS-Prompt	3.86	4.10	0.65	0.60
FS-Prompt	4.12	4.28	0.77	0.73
Vaccine Hesitancy (VH)				
ID-IsF	**4.45**	**4.58**	**0.92**	**0.92**
$IsFD_{HAC}$	3.30	3.55	0.41	0.70
ZS-Prompt	3.92	4.08	0.69	0.72
FS-Prompt	4.15	4.23	0.80	0.81

Table 18: Qualitative results for the discovery of *new* IsFs for the topic of Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH).

Method	Soundness (Z)	Clarity (A)
Immigration (IMM)		
ID-IsF	**0.98**	**0.96**
$IsFD_{HAC}$	0.70	0.73
ZS-Prompt	0.81	0.83
FS-Prompt	0.88	0.89
Climate Change (CC)		
ID-IsF	**0.97**	**0.95**
$IsFD_{HAC}$	0.71	0.75
ZS-Prompt	0.79	0.79
FS-Prompt	0.85	0.85
Abortion (AB)		
ID-IsF	**0.99**	**0.97**
$IsFD_{HAC}$	0.72	0.74
ZS-Prompt	0.84	0.87
FS-Prompt	0.89	0.91
Vaccine Hesitancy (VH)		
ID-IsF	**0.98**	**0.98**
$IsFD_{HAC}$	0.69	0.73
ZS-Prompt	0.82	0.83
FS-Prompt	0.88	0.90

Table 19: Qualitative results for the articulation of Issue Causal Interpretation (ICIs) for Issue Frames (IsF) identified across topics.

Model	Soundness (Z)	Clarity (A)
ID-IsF	0.98	0.96
w/o Explainability Agent	0.89	0.94
w/o R_{art} filter	0.95	0.90

Table 20: Impact of ablating the Explainability Agent and Articulation Reward (R_{art}) filter in **Art-ICI** on the Soundness (Z) and Clarity (A) of ICIs for the Immigration dataset.

M Additional Framing Analysis

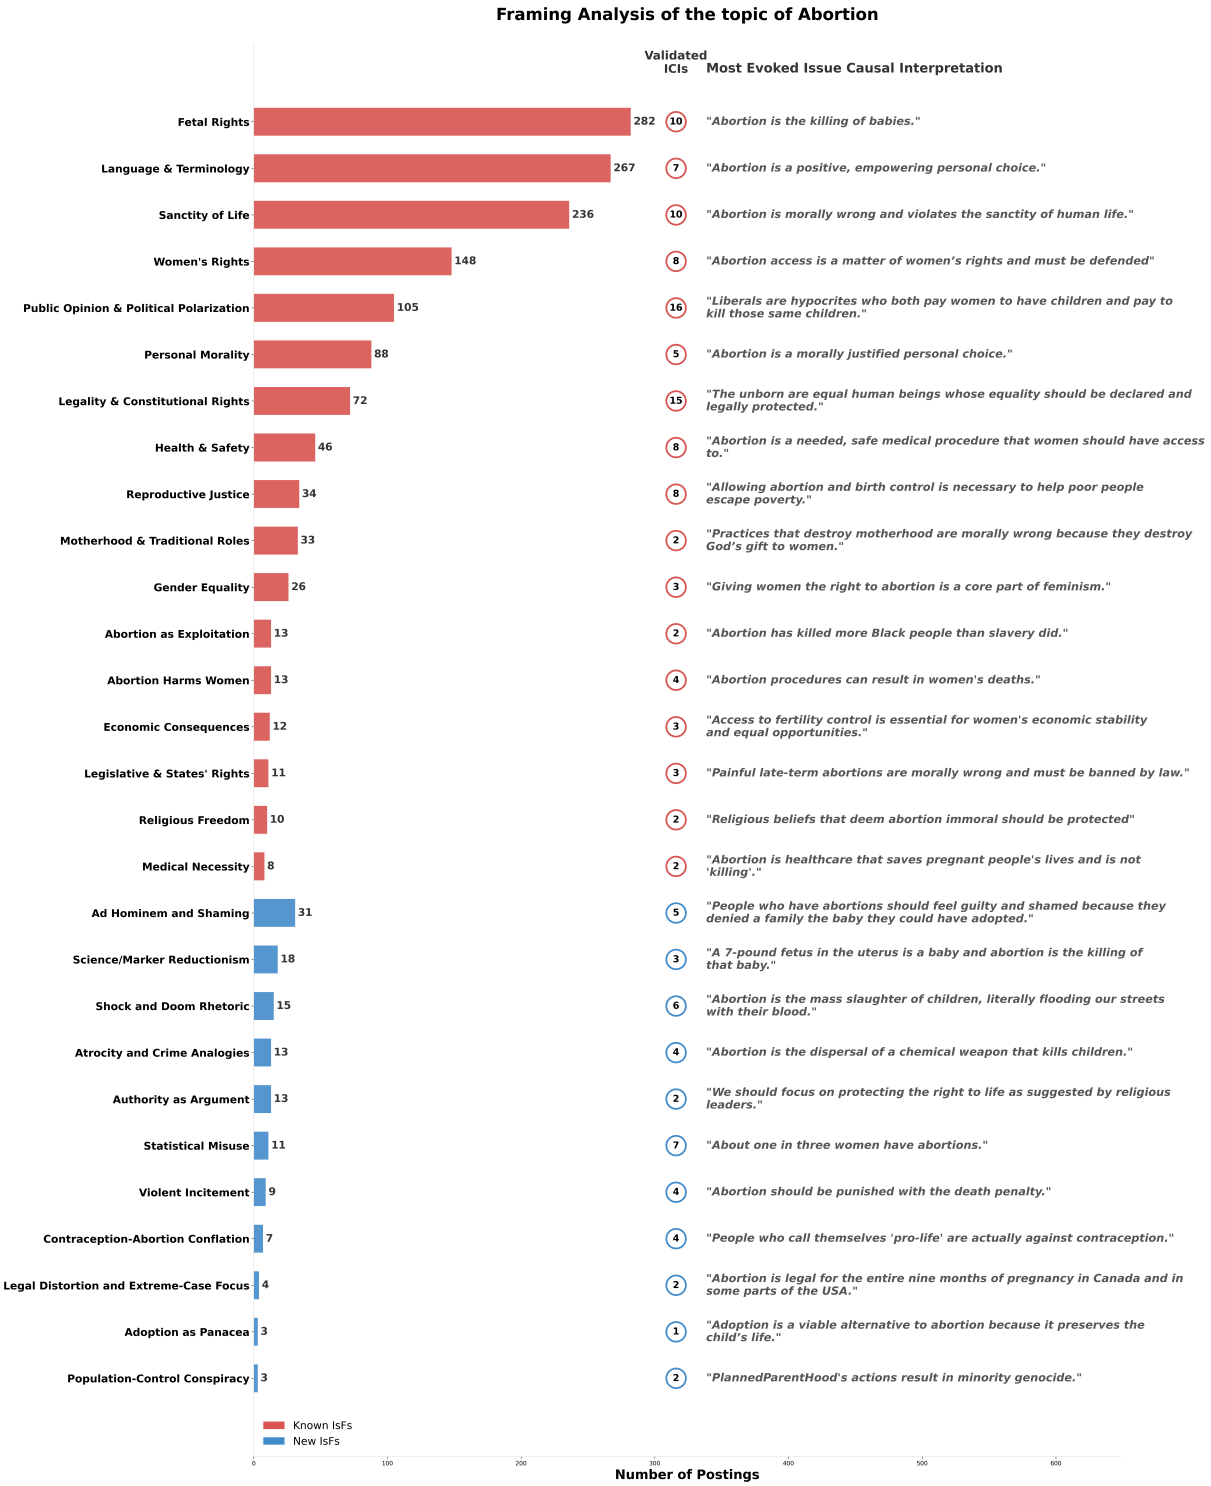

Figure 15: Framing Analysis of the topic of Abortion.

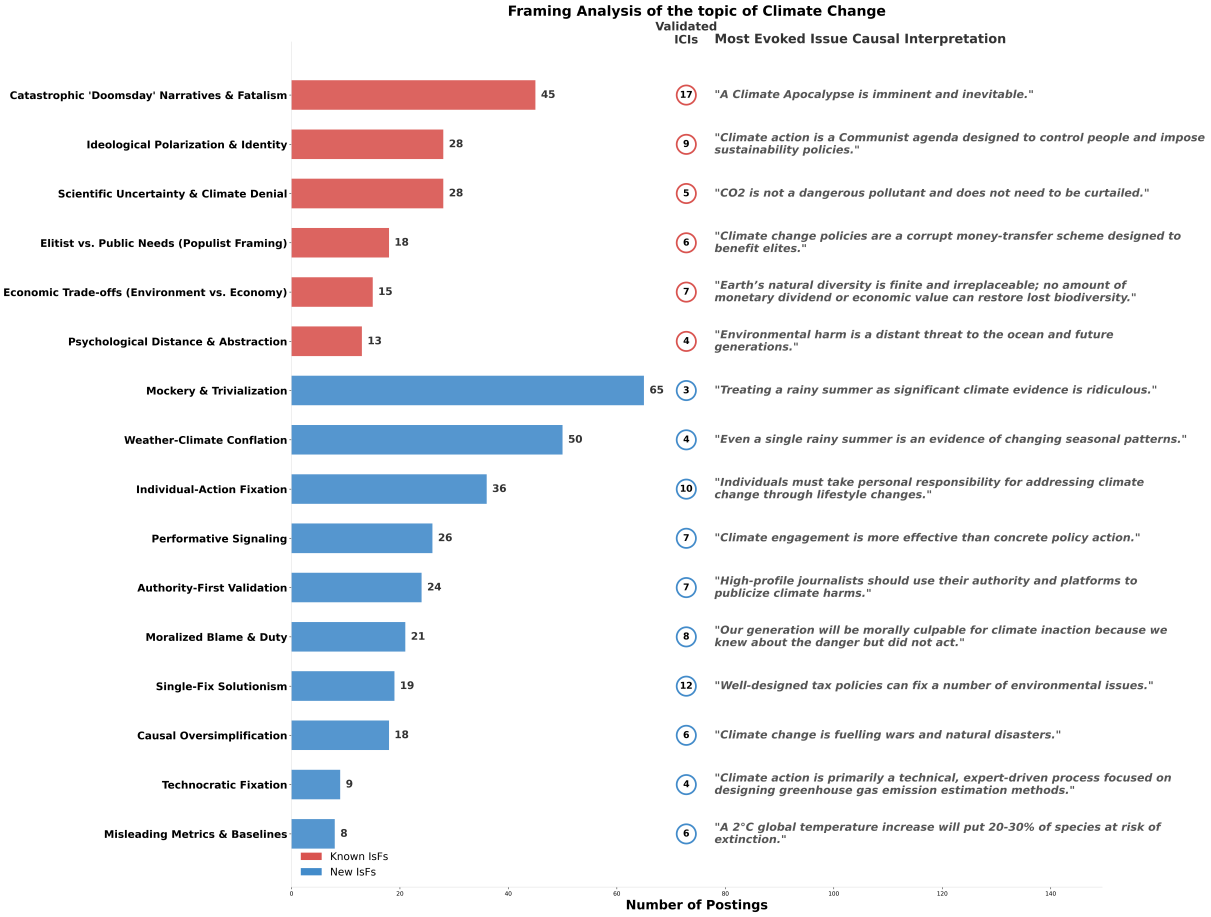

Figure 16: Framing Analysis of the topic of Climate Change.

Figure 15 illustrates the framing analysis for the Abortion topic. The most addressed new IsF is *Ad Hominem and Shaming*, which appears in a substantial number of postings, indicating that moral judgment and social condemnation form a prominent dimension of abortion discourse that extends beyond the existing IsF set. In contrast, *Statistical Misuse* is the new IsF associated with the largest number of unique ICIs, despite appearing in fewer postings overall. This pattern suggests that numerical claims and statistics are leveraged through a wide range of framing strategies.

Figure 16 illustrates the framing analysis for the Climate Change topic, showing (a) the number of postings that address both the known and the new IsFs; (b) the number of ICIs articulated for each TA; and (c) examples of the most evoked ICI for each IsF. For this topic, the most addressed new IsF is *Mockery and Trivialization*, which exceeds several known IsFs in posting frequency, indicating that dismissive and minimizing rhetoric forms a substantial component of climate discourse that is not captured by existing IsFs. In contrast, the new IsF associated with the largest number of unique ICIs is *Single-Fix Solutionism*. Although it appears in fewer postings than *Mockery and Trivialization*, it elicits a richer set of articulated ICIs, reflecting diverse narratives that promote narrowly scoped or overly simplistic solutions to climate change.

Finally, the framing analysis for Vaccine Hesitancy, shown in Figure 17, reveals *Authority-as-Evidence* as 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. In addition, *Mechanistic Misinformation* emerges as the new IsF with the largest number of distinct ICIs.

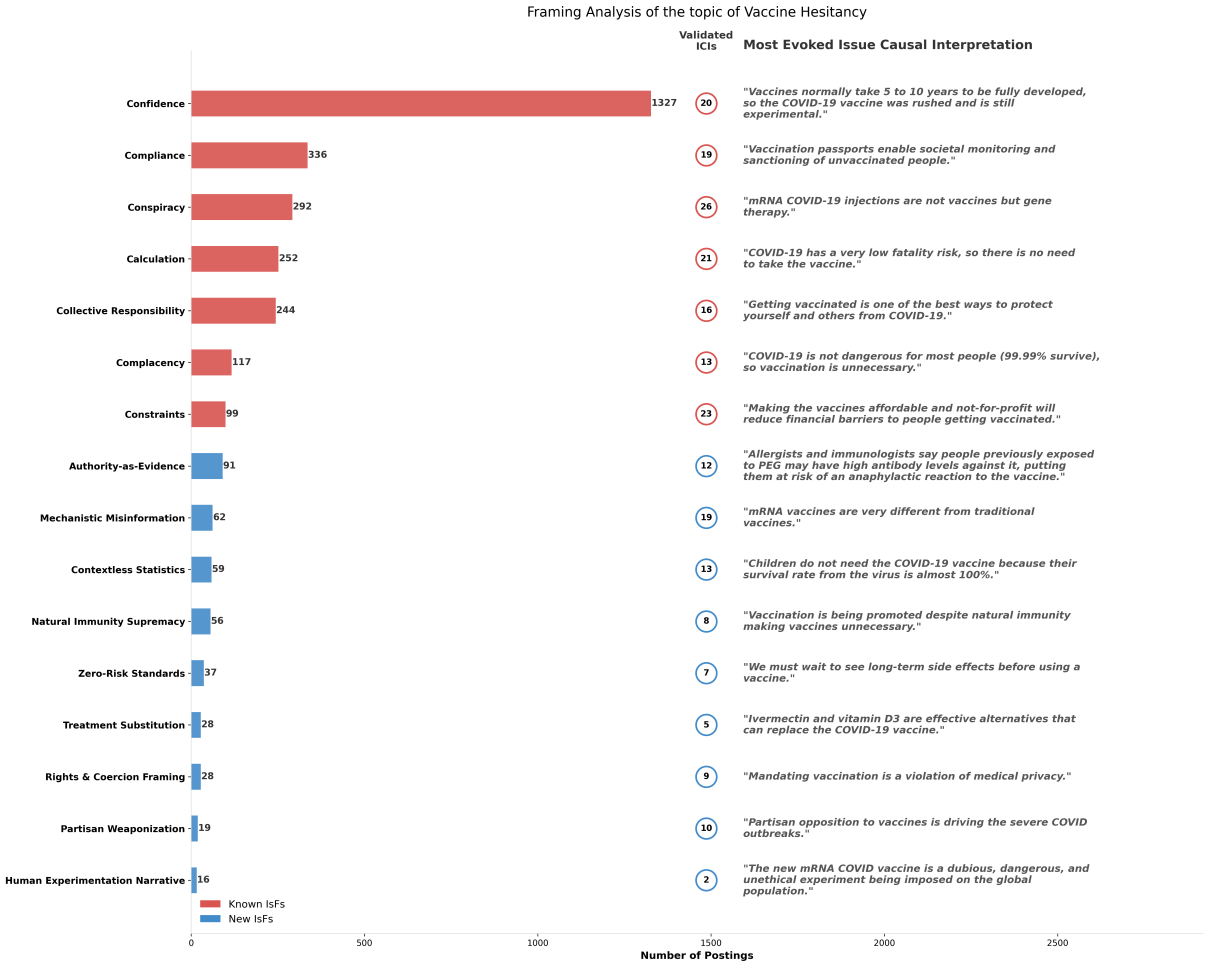

Figure 17: Framing Analysis of the topic of Vaccine Hesitancy.