

Disagreement is Disappearing on U.S. Cable Debate Shows

S M Mehedi Zaman, Kiran Garimella

Rutgers University, USA
sm.mehedi.zaman@rutgers.edu, kiran.garimella@rutgers.edu

Abstract

Prime-time cable news programs are a highly influential part of the American media landscape, with top-rated opinion shows attracting millions of politically attentive viewers each night. In an era of intense political polarization, a critical question is whether these widely-watched “debate” shows foster genuine discussion or have devolved into partisan echo chambers that deepen societal divides. While these programs claim to air competing viewpoints, no large-scale evidence exists to quantify how often hosts and guests actually disagree. Measuring these exchanges is a significant challenge, as live broadcasts contain overlapping speakers, sarcasm, and billions of words of text. To address this gap, we construct the first speaker-resolved map of agreement and disagreement across U.S. cable opinion programming. Our study assembles over 21,000 episodes from 24 flagship shows on Fox News, MSNBC, and CNN from 2010–2024, segmenting them into host-guest turns and labeling 2.13 million turn-pairs using a high-fidelity large-language-model classifier. We present three findings: (1) the proportion of disagreement/debate on prime time shows a consistent downward trend, dropping by roughly one-third between 2017 and 2024; (2) on-air challenge is partisan and asymmetric—conservatives seldom face push-back on Fox, liberals seldom on MSNBC, with CNN declining toward the midpoint; (3) polarizing issues such as abortion, gun rights, and immigration attract the least disagreement. The work contributes a public corpus, an open-source stance pipeline, and the first longitudinal evidence that televised “debate” is retreating from genuine discussion. By transforming into platforms for partisan affirmation, these shows erode the cross-cutting cleavages essential for a pluralistic society, thereby intensifying affective polarization.

1 Introduction

Televised debate shows—programmes built around a host who interviews or spars with one or more guests about the political story of the day—are advertised as democracy’s front-row seats to an argument. These broadcasts occupy the most valuable real estate in cable schedules: the prime-time block (8p.m.–11p.m. Eastern), when total viewership peaks and advertising rates spike. In 2024, Fox News, MSNBC, and CNN together delivered roughly forty hours of such pro-

gramming each week, and seven of the ten highest-rated cable slots were Fox opinion shows, each drawing well over two million nightly viewers (Littleton 2024). Because the audience is both large and politically attentive, elites treat these discussions as signals—Presidents quote them, lawmakers schedule hearings around them, and social-media influencers amplify their clips (Gertz 2020). If the “debate” has quietly become one-sided affirmation, the entire information chain downstream is fed from an ideological echo.

The puzzle, then, is whether prime-time debate still fulfills its deliberative promise or whether it has degenerated into what researchers call an echo chamber: an environment where political communication overwhelmingly confirms rather than contests the audience’s partisan priors (Löblich and Venema 2021). In such settings, dissenting views are either absent or minimized, producing repeated reinforcement of a shared orientation. While the concept is often applied to social media algorithms (Bakshy, Messing, and Adamic 2015), it is equally relevant to legacy media formats where host-guest exchanges can either foreground disagreement or devolve into choreographed agreement (Jamieson and Cappella 2008). Understanding this dynamic is critical, as the retreat from genuine debate can deepen what political scientists call *affective polarization*—the widening gulf of distrust and dislike between partisans that now rivals disagreement over concrete policy (Iyengar and Hahn 2009).

Scholars have shown that exposure to like-minded news encourages attitude reinforcement (Stroud 2010), while diversified media diets can moderate views (Mutz and Reeves 2005). Yet almost all evidence rests on who watches which channel or how often a partisan guest is booked, not on whether a conservative and a liberal, once seated opposite the same host, actually clash over substance. Without that micro-level record, we cannot know whether televised debate attenuates or accelerates polarization.

Capturing that record is hard for three reasons. First, live television is unscripted and multi-speaker: turns overlap, tone shifts mid-sentence, and topics pivot quickly. Second, genuine disagreement is a semantic property, not a surface cue; a sarcastic “right, exactly” is opposition, not assent. Naïve proxies—assuming every Democrat on Fox is challenged or every negative adjective signals dissent—fail on counter-examples. Third, the scale is daunting. A single year of prime-time opinion programming on the big

three cable networks generates more than a billion words of caption text (Hong et al. 2021). Manual annotation at the speaker-turn level is financially untenable, and existing automated pipelines either collapse all speech into one stream (losing who said what) or ignore rhetorical stance altogether.

Previous studies therefore stop short of the core question. Some sample a handful of episodes and code them by hand, offering depth without generality (Farnsworth and Lichter 2005). Others scrape guest lists or sentiment polarity and infer that a segment was adversarial when the guest's party differs from the host's party (Morris 2016). Such approaches miss intra-party dissent, rhetorical framing, and shifts within a single interview. What we lack is a comprehensive, speaker-resolved measurement of agreement and disagreement across the contemporary cable landscape.

This paper supplies that measurement. We assemble a corpus of 21,000 episodes (approximately 8,000 hours of video) aired between 2010 and 2024 on 24 flagship opinion shows across Fox News, MSNBC, and CNN. An automatic-speech-recognition and speaker-diarization pipeline segments each broadcast into host and guest turns. Adjacent host-guest pairs are then classified as *agreement*, *neutral*, or *disagreement* by a high fidelity stance model which is then applied to classify over 2 million host-guest interaction turns. Finally, we identify the names and political party affiliations of the guests along with the detailed topic being discussed in each episode.

Using this rich dataset, we present five main findings. First, explicit host-guest dissent is rare: across 2017–2024 barely 15% of exchanges register as disagreement, with MSNBC the lowest at roughly 13% and Fox the highest at 17%. Second, dissent is in retreat. On Fox and CNN the share of disagreement turns falls sharply each year, and in the few MSNBC shows with long-run data the baseline is flat but never high, signaling a network-wide drift toward concord. Third, the decline is historically contingent rather than inevitable—archival captions from the early 2010s show *Hannity* operating at nearly double today's disagreement rate, evidence that the format once accommodated sharper clash. Fourth, gatekeeping amplifies the trend: the channels not only book ideologically friendly guests but also moderate them differently once on air—Republican invitees face the stiffest push-back on MSNBC, while Democrats encounter the gentlest treatment there, with the mirror pattern on Fox and a narrowing middle on CNN. Finally, the topics that dominate contemporary culture-war politics—abortion, gun rights, immigration—attract the *least* disagreement, turning segments most likely to mobilize audiences online into echo chambers on air.

Together, these results suggest that prime-time “debate” is increasingly a performance of consensus within partisan boundaries rather than a forum for genuine argumentative contestation. By documenting this shift at scale and at the level of individual utterances, the study clarifies how televised political talk contributes to the broader architecture of affective polarization, and why restoring actual disagreement to mass media remains a pressing democratic task.

2 Related Work

A vast literature links partisan cable news consumption to attitudinal polarization, uncivil discourse, and diminished civic trust. We organize the most relevant findings into four strands: (1) selective exposure and the rise of echo chambers, (2) the evolution of televised debate, (3) democratic consequences of incivility and one-sided programming, and (4) recent computational audits that map ideological drift at scale. Together they illustrate both what we know and the empirical gaps this study fills.

Selective Exposure and the Cable-News Echo Chamber. Once Americans could choose among dozens of channels, many gravitated toward ideologically congenial outlets, producing self-reinforcing “echo chambers.” Laboratory evidence shows Republicans strongly prefer Fox-branded stories while Democrats shun them (Iyengar and Hahn 2009). Longitudinal survey work demonstrates that this sorting hardens attitudes over time (Stroud 2010; Levendusky 2013). Although only a minority of viewers are “active-audience” partisans, their loyalty shapes network programming and, by extension, public debate (Arceneaux and Johnson 2013). Recent large-scale audience data confirm that cable news consumers today almost never stray outside their partisan bubble, with Fox, CNN, and MSNBC audiences sharply segregated (Kim, Lelkes, and McCrain 2022).

Situating cable news as an offline echo chamber is particularly relevant to the ICWSM community. Recent work has begun to challenge the narrative that algorithmic feeds are the primary drivers of polarization (Flaxman, Goel, and Rao 2016; Bakshy, Messing, and Adamic 2015), showing that partisan sorting can be even stronger in offline contexts than online (Vaccari et al. 2016; Brown et al. 2021). Our study contributes directly to this debate by providing a granular, speaker-resolved map of conversational dynamics within a non-algorithmic, host-driven environment. This analysis serves two purposes: first, it provides a reference point for comparing the intensity and mechanisms of human-curated echo chambers versus algorithmically-curated ones; second, it underscores the importance of studying legacy media, as the talking points and video clips generated in these televised debates are the raw material that is subsequently amplified and discussed across social media platforms, shaping the online discourse ecosystem.

Changes in Prime-time Debates. Early-2000s panel shows such as *Crossfire* thrived on raised voices, interruption, and personal digs. Experiments revealed that televised incivility elevates physiological arousal yet erodes trust in government, even when issue content is held constant (Mutz and Reeves 2005). By the mid-2010s, however, genuine left-right sparring had almost vanished; *Crossfire* was canceled again in 2014, symbolizing CNN's retreat from adversarial debate (Forgette and Morris 2006). A decade-long content audit shows that after the 2016 election Fox News booked steadily more conservative guests while CNN and MSNBC booked more liberals, leaving scant room for on-air disagreement (Kim, Lelkes, and McCrain 2022). When dissent does appear, it is often tokenistic or performative rather than deliberative (Sobieraj and Berry 2011).

Consequences for Democratic Deliberation. Incivility and

ideological uniformity carry measurable civic costs. Uncivil exchanges not only entertain but also prime hostility toward political opponents and lower institutional trust (Mutz and Reeves 2005). Because partisan networks rarely present serious cross-cutting dialogue, viewers encounter few competing frames, exacerbating motivated reasoning and affective polarization. The effects cascade beyond the viewing public: elite rhetoric (Gertz 2020) and social-media conversations (Ding, Horning, and Rho 2023) frequently echo televised talking points, amplifying polarization across multiple arenas.

This paper argues that the retreat from on-air contestation is a key, and perhaps underappreciated, driver of affective polarization. Our longitudinal analysis reveals that these prime-time shows are undergoing a fundamental change in their function, shifting from forums of debate into platforms for partisan affirmation. This structural shift has profound implications, a process explained by the political science theory of cross-cutting cleavages (Goodin 1975), which posits that political conflict intensifies when partisan identity becomes the singular, dominant line of social division. By systematically purging substantive disagreement, these programs erode the visibility of other social cleavages and present a world of binary political conflict. A primary mechanism for this is the disappearance of intra-party dissent, which erases the natural diversity of factions within a political coalition. This performance of consensus sends powerful elite cues to viewers that ideological conformity is the norm, framing the in-group’s perspective as a settled truth. Consequently, opposing views are rendered not merely incorrect but illegitimate, hardening the distrust and animosity that define affective polarization.

Computational Mapping of Ideological Slant. Recent work leverages natural-language processing (NLP) and large transcript corpora to quantify network bias with temporal precision. By linking every on-air guest to their campaign-donation ideology score, Kim et al. chart hourly shifts in slant from 2010–2020, showing Fox’s pronounced rightward drift and CNN/MSNBC’s reciprocal move left after 2016 (Kim, Lelkes, and McCrain 2022). Earlier econometric analyses reached similar conclusions using hand-coded tone and carriage data (Martin and Yurukoglu 2017). Such computational audits reveal that partisan skew is dynamic—responsive to electoral cycles, competitive positioning, and programming genre—and they supply the methodological blueprint for the present study.

Taken together, these strands demonstrate that *who* watches (selective exposure), *what* they watch (ideological uniformity), and *how* it is presented (incivility) jointly deepen partisan division. Yet we still lack fine-grained evidence on how moment-to-moment argumentative structure within cable segments shapes downstream discourse. By integrating NLP with interaction-level annotations, this paper addresses that gap and specifies the mechanisms linking televised debate to audience polarization.

Furthermore, we connect our empirical findings to the literature on cross-cutting cleavages (Goodin 1975) and elite cue theory (Gilens and Murakawa 2002). While existing work focuses heavily on audience self-sorting into echo

chambers, our contribution demonstrates how the content within these chambers is actively being reshaped to eliminate internal dissent. By documenting this shift, we provide a specific mechanism, the erasure of debate and the performative signaling of conformity, that explains how partisan media can intensify affective polarization beyond simply filtering content.

3 Dataset

Our corpus originates with the Internet Archive Television News Archive,¹ an open repository that mirrors the closed-caption feeds of major U.S. broadcasters. We build on the cleaned release by Laohaprapanon and Sood (2017), but extend it in two important ways. First, debate research requires unambiguous speaker segmentation. Many shows in the Archive omit the `>> speaker:` token that marks a change of voice; others collapse hosts and guests under the same tag. To maximize precision, we retained only shows in which closed captions explicitly annotate both the host (e.g. `>> TUCKER:`) and the guest (`">>"` with no name). This criterion preserved most Fox News content but excluded the bulk of CNN and MSNBC’s schedule, whose captioning conventions lack named turns.

Second, to restore balance across networks we supplemented the Archive with episode transcripts from the shows’ official Apple Podcasts feeds, which supply RSS-accessible .mp3 files which we crawled using a lightweight crawler, to obtain every available episode for shows whose captioning was inadequate in the Archive. We then transcribed and obtained speaker turns these episodes using a large Whisper model (Radford et al. 2023). Next, we then harmonized the formatting to match the Internet-Archive schema. Consistent with the Archive license, all podcast transcripts were already in the public domain; no paywalled or private audio was accessed.

Our final corpus contained 24 flagship debate or opinion show—nine on Fox News, eight on CNN, and seven on MSNBC—broadcast between January 2010 and December 2024. The sample contains over 21,000 episodes, accumulating just over 8,000 hours of video and roughly 194 million captioned words. Across those episodes we identify 2.13 million speaker turns, and on average, about 190 words per show. The earliest transcript dates to 7 July 2009 and the latest to 14 January 2025, giving a span of approximately 5,670 days (15.5 years, though we have gaps in 2014-2017). Table 3 lists every show alongside its coverage window and episode count. For each utterance we have a timestamp, the interlocutor’s role (host or guest), the guest’s inferred party affiliation (Section 4), and a stance label—*agreement*, *disagreement*, or *neutral*—generated by our classifier (Section 4). We also identify the topics discussed in each episode of the show (Section 4).

Dataset completeness. Despite its breadth, the corpus remains incomplete. We focus on the three dominant U.S. cable-news channels—Fox News, CNN, and MSNBC—because they account for the overwhelming share of prime-time debate programming (Pew Research Center

¹<https://archive.org/details/tvarchive>

2023). For each network we obtained every episode for which either the Internet Archive or an active Apple Podcasts RSS feed supplied transcripts. Even so, several shows exhibit intermittent coverage: entire weeks are missing when a programme was pre-empted, went on hiatus, or when its podcast feed was later purged from Apple’s catalogue. Notably, for Hannity and Special Report with Bret Baier, Internet Archive or the Apple Podcasts did not have the data between 2014–2017, possibly due to technical glitches. Table 2 details the principal causes of these gaps. Fox, whose caption files consistently include speaker tags and whose podcast feeds remain intact, therefore contributes more episodes and longer time-series than CNN or MSNBC. This asymmetry constrains some cross-show comparisons, especially for the latter networks, but all analysis we report below explicitly considers episode availability and our results are robust to the missing-data patterns.

Even though television data is widely available through the Internet Archive (Ding, Horning, and Rho 2023) and the Stanford TV data (Hong et al. 2021), data at the level of granularity (including guest names, timestamps, topics, etc) and quality we obtain in this paper is not available previously. We believe the resulting corpus, together with the scraping and alignment procedures documented here, provides a replicable template for future work that seeks speaker-resolved television data at scale. The codebase for dataset construction and our entire pipeline are available on a Github repository.²

4 Methods

Detecting Disagreement

Data Pre-processing For programmes whose closed captions already encode speaker names (predominantly Fox News) we treat every contiguous block attributed to the host as an anchor turn and concatenate the immediately following guest utterances until the host next speaks. Each such *host–guest bundle* constitutes one analytic pair. We further excise material outside the argumentative core of the show by trimming the transcript at the first host utterance and at the final host utterance, thereby discarding lead-in promos, commercial breaks, and closing credits. One possible limitation of our pairing strategy is that, in episodes with multiple guests, consecutive guest turns might be collapsed into a single “guest” block if the host does not intervene. In practice, however, this constellation is extremely rare. Prime-time opinion shows on Fox, CNN, and MSNBC are overwhelmingly organized around a tightly moderated host–guest–host interview rather than open guest–guest debate (Jurkowitz et al. 2013). Prior content audits similarly note that contemporary debate television has shifted away from cross-talk toward “parallel monologues” controlled by the host (Berry and Sobieraj 2013). In our own spot-check of 100 randomly sampled multi-guest episodes, 92% featured a single guest speaking at a time, with fewer than 3% of turns showing two guests in direct succession without host mediation. Thus, while theoretically possible, the loss of distinct

guest reactions in consecutive turns is empirically negligible.

Episodes obtained from Apple Podcasts are .mp3 files without caption metadata. We convert audio to text and speaker diarization using the large-V3 WHISPER model (Radford et al. 2023), producing a token-time-aligned transcript in which turns are labelled generically (SPEAKER 1, SPEAKER 2, ...). To identify the host we apply a two-step heuristic. First, the speaker with the greatest number of turns is presumed the host; every other speaker is provisionally a guest. If we observe multiple guests, occasionally pushing the host below first place, we override the rule and assign the first utterance in the episode to the host, propagating that label to all matching diarization IDs. Residual advertising copy at the head of the file is removed with show-specific regex filters keyed to recurring phrases such as “welcome to *Inside Politics*.” We manually verified every episode where our heuristic did not apply to ensure the dataset was of high quality.

Data Annotation The main analytical backbone of the project is a stance label assigned to every *host–guest bundle*. Because disagreement can flare and subside many times within a single episode, we annotate each contiguous exchange rather than the episode as a whole. Each record therefore contains two text fields—`host_turn` and `guest_turn`—and one categorical target with three possible values: *Agreement*, *Disagreement*, or *Neutral*.

We began by hand-coding five hundred randomly selected pairs from various shows. Two annotators (the authors) independently coded this set, achieving a Cohen’s κ of 0.81, indicating substantial agreement prior to adjudication. Disagreements were then reconciled in adjudication meetings to produce a final gold-standard set with perfect agreement. The following decision rules governed the exercise:

1. **Agreement** is assigned only when the guest endorses, amplifies, or otherwise affirms the host’s preceding claim. *Example*:
Host: “Just to be clear, you agree the Russian invasion was an unprovoked violation of international law?”
Guest: “Absolutely, Moscow crossed a bright red line, and the world has to treat it that way.”
2. **Disagreement** requires an explicit refutation, contradiction, or correction of the host. *Example*:
Host: “Premiums are up twenty–five percent—people hate this law because they’re paying more.”
Guest: “That’s not accurate. Premiums were flat for the two previous years, and the average growth *slowed* after the law passed.”
3. **Neutrality** covers all residual cases: clarifying questions, greetings, news reading, off-topic banter, or exchanges in which stance cannot be inferred.

Manual coding of the 2.13 million host–guest bundles was infeasible, so we adopted a staged pipeline that moves from small, human-verified seeds to large-scale annotation with language models. Our labeling strategy necessarily flattens instances where a guest simultaneously affirms one

²Codebase: <https://github.com/mehedizamane/disagreement-is-disappearing>

part of the host’s statement while rejecting another. However, two factors mitigate its impact. First, prime-time exchanges are typically short and tightly coupled—on average, each host–guest pair contains roughly 90 words, typically addressing a single issue at a time. Second, our classifier is conservative: when a guest response contains both supportive and adversarial elements, the model frequently defaults to the neutral category rather than forcing agreement or disagreement. For *Example*:

Host: “Well, he said he wasn’t aware. He has been specific about that. That’s been one of the things they said. He did not know.”

Guest: “No, I agree. But one of the problems he has in talking about it that way remember he wrote a book and talked about Russia and Iran and Ukraine. And one of the things that the Justice Department is going to do is compare those passages to these documents, and, you know, it’ll either show they were likely used or there’s no mention of them. But it’s a problem for him.”

The model output was neutral, reflecting mixed stance. While this approach sacrifices fine-grained nuance, it ensures that we do not overstate levels of disagreement.

We built the stance classifier through a four-stage, bootstrapped workflow. We began with a small hand-coded seed and found that task-specific encoders (e.g., fine-tuned DeBERTa-v3-base) topped out at barely 55% accuracy, so we pivoted to large language models, using GPT-4 in an interactive loop to expand the seed to 500 high-quality pairs. Two successive rounds of crowd annotation then supplied an additional 1,712 carefully triaged pairs, vetted with gold checks, consensus filtering, and manual audits. At each expansion we re-trained both commercial (GPT-4o) and open-source models (Llama-3 70B 4-bit, DeepSeek-R1-Distill-Llama-8B, Qwen-2.5 32B 4-bit), iteratively pruning noisy items until performance stabilised. The final DeepSeek model attains 89% accuracy on an unseen validation set (comprised of the authors’ manually audited subsets from the crowdsourced datasets), providing the labels used throughout the paper. To ensure the reliability of these labels, two authors independently reviewed all samples in this subset, achieving Cohen’s $\kappa = 0.84$. Disagreements were resolved through discussion until consensus was reached. Full details of the annotation rounds, model variants, and evaluation metrics appear in the Appendix.

This fine-tuned model was then applied to the rest of the 21,000 episodes, generating stance labels for 2.13 million host–guest bundles. For each show, one of the authors manually verified 100 randomly sampled balanced pairs. This spot check was designed as a sanity check and in all shows, model accuracy exceeded 85% (complete evaluation metrics is reported in Table 6 in the Appendix), thus giving us the confidence that our model and predictions are high fidelity and can be relied upon. Overall, we identify 53.6% neutral, 30.4% agreement and 16.0% disagreement labels for all the 21,000 episodes.

Guest Identification and Party Coding

Our procedure for determining who the guests were in each episode, and how those guests leaned politically, unfolded in

two stages.

Stage 1: extracting guest names with GPT 3.5. For every episode we isolated the host’s opening monologue, which typically lists the evening’s interviewees. We then passed that text to the gpt-3.5-turbo model through the OpenAI API in a zero-shot setting. A representative prompt for *Life, Liberty & Levin* was:

```
This is a transcript of the show
`Life, Liberty and Levin`,` hosted
by Mark Levin. He is never a guest.
Extract all invited guests who actually
join the show from the full opening
monologue, and for each give their
political affiliation as one of
[Democrat,Republican,Other,Unknown].
If the guest is a known office-holder,
use their real party. If they are a
journalist, member of a non profit,
an advocate or any other, infer their
closest alignment from your training.
If you absolutely cannot determine
affiliation, use Unknown -- but only
as a last resort. Your reply must be
only a JSON array (no narrative, no
Markdown, no code fences).
```

This step yields, for each episode, a list of guest names paired with the model’s first-pass guess at affiliation. The host’s name is systematically excluded. In total, 1,374 guest affiliations were retrieved in this stage.

Stage 2: refining affiliation labels with external data.

Journalists, editors, and many subject-matter experts do not have an official party registration, which often led the language model to return Unknown. To improve coverage and cross-check the LLM generated affiliations, we matched the extracted names and affiliations against the DIME database (Bonica 2015), a large-scale compilation of itemized U.S. campaign-finance records. This is a high fidelity resource used in previous research and has been shown to be of high quality, containing a large database of public personas (Kim, Lelkes, and McCrain 2022). To ensure accurate mappings, we normalized names into a consistent “last, first” format and required exact matches. In cases where DIME contained multiple individuals with the same name, we used the “recipient party”, and “contributor occupation” columns to disambiguate. Ambiguous matches (less than 5% of the total matches) with no clear resolution were conservatively excluded rather than risk misclassification. Among the subset of guests who had a party affiliation in both sources, 70.9% of GPT labels matched DIME, while 29.1% diverged. In cases of mismatch, we prioritized the DIME label in the final dataset. For guests marked as Unknown by GPT but present in DIME, we also filled in the affiliation from DIME. This procedure both corrected potential GPT misclassifications and reduced reliance on the Unknown category, yielding a more robust measure of guest partisanship.

Applying this two-step pipeline to the full corpus provided an affiliation label for nearly every recurring guest. Figure 1 shows the final distribution, showing, as expected,

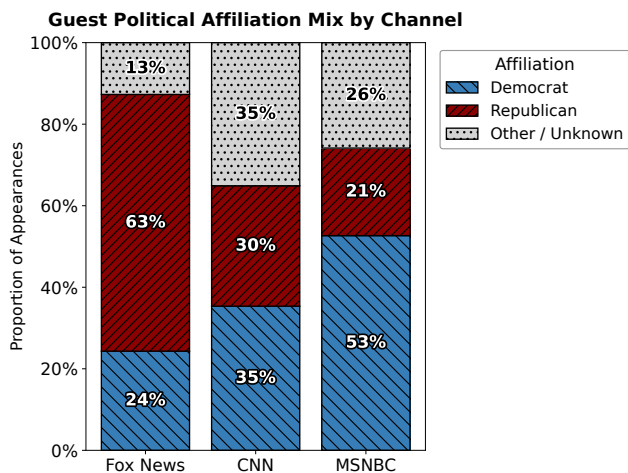


Figure 1: Fraction of guests by political leaning and channel. As expected, we see that Fox News invites mostly Republicans and MSNBC mostly Democrats.

that Fox News features predominantly Republican guests, while MSNBC features predominantly Democratic ones.

Topic Modeling

To recover the topics of each episode we applied BERTopic (Grootendorst 2022) to every episode transcript in the corpus.

Pre-processing. Captions were lower-cased, stripped of non-alphabetic characters, filtered through the NLTK English stop-word list, and purged of “intro-junk” tokens such as host names and sponsor slogans. The cleaned texts were embedded with the `all-MiniLM-L6-v2` sentence-transformer.

Model configuration and tuning. We initially set the vectorizer to `min_df = 2`, `max_df = 0.9`, and `max_features = 30,000`. Dimensionality reduction used UMAP with cosine distance (`n_components = 3`, `min_dist = 0.1`, `n_neighbors = 10`); soft clustering employed HDBSCAN, which yields a membership probability for every document. Hyper-parameters were tuned by grid search on a 5% development slice, optimizing for *codeless* coherence—our best setting achieved a coherence score of 0.689.

Topic consolidation. The tuned model returned 146 raw clusters. Manual inspection revealed substantial lexical overlap, so we merged semantically redundant clusters—using keyword similarity and exemplar phrase matching—down to 80 distinct topics. Finally, two authors grouped those 80 topics into 15 higher-level categories (e.g. ELECTIONS & CAMPAIGNS, CULTURAL POLITICS, COVID-19, etc); any disagreements were resolved by discussion. These 15 top-level labels constitute the topic taxonomy visualized in Figure 8. The list of all 80 topics is added in table 2 in the Appendix.

5 Findings

We organize the findings along three dimensions. First, we present longitudinal disagreement and agreement from 2010 through 2024, documenting how the share of disagreement contracts across most shows and channels (Section 5). Second, we pair every utterance with the guest’s party identification to show how partisan booking decisions shape conversational tone—Democratic invitees receive more agreement on MSNBC, Republicans on Fox, while CNN remains comparatively balanced (Section 5). Third, we disaggregate the data by topic, demonstrating that the lowest levels of disagreement cluster around culture-war flashpoints such as abortion, gun rights, and immigration, a pattern that underscores the echo-chamber character of contemporary debate television (Section 5).

Agreement and Disagreement Over Time

In this analysis, we compute the average (dis)agreement per show, aggregated by month. For each month, we calculate the mean level of (dis)agreement observed across all shows within that period. Alongside the mean, we also present the 95% confidence intervals, which allows us to track how (dis)agreement varies over time.

Figure 2 shows the temporal trajectories of host–guest disagreement for a (non random) sample of the shows in our dataset. Across the full 2017–2024 window, disagreement on average constitutes only a surprisingly small share of airtime—typically between 10% and 15%—and, for most shows, that share is either flat or trending downward. The pattern is clearest on Fox News. Tucker Carlson Tonight begins its run in 2017 with roughly one third of all host–guest exchanges coded as explicit disagreement, yet this fraction drops to about 15% by the time the show goes off-air in April 2023 (Figure 2a). The ordinary-least-squares slope on the monthly series is -3.9 percentage points per year ($p = 0.000$), confirming a systematic retreat from disagreement. Our finding dovetails with the New York Times content audit of 1,100 episodes, which likewise documents a contraction in dissenting perspectives over the same period (Times 2022). Laura Ingraham’s primetime hour shows a similar trend (Figure 2d), albeit from a slightly lower baseline (-1.3 percentage points per year, $p=0.000$), showing that the phenomenon is not idiosyncratic to a single host but characterizes Fox’s broader editorial posture.

CNN exhibits a more heterogeneous picture. Anderson Cooper 360° mirrors Tucker Carlson in direction (Figure 2b), though the glide path is gentler: disagreement falls from roughly 20% in 2019 to 14% by late 2024, a decline of 1.2 percentage points per year ($p = 0.000$). Similarly, the Sunday show State of the Union—which is structurally designed to pit partisan surrogates against one another—has a higher base rate of disagreement (near 30%) in 2018 but goes down to around 20% by 2024 (decline of -2.1 percentage points per year ($p = 0.000$), Figure 2e).

MSNBC’s inventory is the most internally consistent: disagreement begins low and stays low. The Savage Nation and Velshi, for instance, rarely strays above 10-15% disagreement in any month (Figure 2f), and the slope of the

time-trend is statistically indistinguishable from zero. Other MSNBC offerings in the appendix (Figures 10, 11) replicate this flat-line pattern. Taken together, the cross-network comparison reveals that while Fox is actively converging toward a narrow opinion corridor, CNN is bifurcated between a shrinking adversarial prime-time slot and a resilient Sunday shows, and MSNBC has long since embraced homophily as a default.

This long-run flatness is consistent with MSNBC's booking strategy: prime-time programming overwhelmingly features ideologically aligned guests, so baseline opportunities for visible disagreement are rare. Where disagreement does appear, it tends to be intra-coalitional (e.g., tactical disputes within the Democratic camp) rather than cross-partisan clashes, and these are moderated into brief, low-salience exchanges by the host. As a result, even as the network's guest mix varies slightly across shows, the format produces a ceiling effect: disagreement cannot decline much further, because it begins at very low levels. In this sense MSNBC illustrates a homophilous "steady state," contrasting with Fox's dynamic narrowing and CNN's shifting hybrid format.

The complementary analysis of agreement (see Figure 3 for a small sample, Figures 12, 13 in the Appendix for all shows) reinforce these conclusions. Where disagreement contracts, agreement expands. On Tucker Carlson Tonight the proportion of explicit agreement rises by almost 5 percentage points over the study period, a mirror image of the disagreement slide. Neither Anderson Cooper 360° nor The Savage Nation exhibits a statistically significant upward drift in agreement, suggesting that their modest losses in disagreement are being absorbed by neutral or ambiguous exchanges rather than by outright concurrence. Crucially, across all shows the combined share of agreement and neutral statements already exceeds 85% of airtime by 2024, leaving scant room for the deliberative contestation that debate television purports to deliver.

Two patterns emerge with genuine substantive weight. First, the data point to a deliberate editorial pivot away from inviting ideologically dissonant guests: as fewer counter-voices appear, on-air push-back correspondingly evaporates. This retreat is clearest on Fox and CNN, where disagreement falls significantly in 7 out of 9 Fox shows and 5 of 8 CNN shows. MSNBC shows no systematic trend—largely because the available time-series for its line-up is shorter—but its baseline levels of disagreement are already minimal. Second, the erosion of dissent is mirrored by a marked rise in affirmation, especially on Fox, where explicit agreement climbs in 5 of 9 shows.

Although CNN continued to invite a more balanced mix of guests, the incidence of dyadic disagreement still declined over time. This reflects the format of prime-time talk shows, where hosts maintain strong control over the conversation and often avoid sustained cross-cutting debate, instead privileging civility and rapid topic shifts. In line with prior work on cable news as "parallel monologues" rather than genuine argument exchange (Sobieraj and Berry 2011; Jamieson and Cappella 2008), even a heterogeneous guest lineup does not necessarily translate into higher on-air disagreement.

Viewed alongside the topic-level findings in Section 5, the pattern is disquieting. The shrinking space for overt disagreement coincides with sustained—or even heightened—coverage of polarizing policy domains such as abortion, gun rights, and immigration. What remains, therefore, is "echo-chamber prime time": divisive themes are aired, yet almost invariably among like-minded interlocutors who amplify rather than interrogate partisan frames. The evidence is unmistakable: contemporary debate television has not melted into polite conversation; it has shed the adversarial backbone that once separated it from straightforward opinion broadcasting.

Longitudinal case studies To gauge how network-level patterns translate into editorial decisions on individual shows, we isolate three Fox News franchises for which we possess either uninterrupted long-run data or a well-documented format change: *Hannity*, *Special Report with Bret Baier*, and *Gutfeld!*. Each offers a distinct window onto the evolving relationship between host persona, booking strategy, and discursive diversity.

Sean Hannity Sean Hannity's primetime hour has been on the schedule since 2010, but its tone shifted markedly once Donald Trump emerged as the Republican standard-bearer. Hannity not only endorsed Trump on air; he also appeared on stage at a 2018 campaign rally in Cape Girardeau, Missouri, blurring the line between commentator and surrogate (Schwartz and Oprysko 2018). Our data mirror that realignment. The average disagreement fraction drops from 0.115 in the 2010–2013 window to 0.075 between 2017 and 2024—a 35% contraction that is statistically significant at $p < 0.001$ (Figure 4a). Over the same span, explicit agreement rises by nearly six percentage points (Appendix, Figure 12f), confirming that what little debate and discussion remained has largely been replaced by consonant, in-group affirmations. The show thus exemplifies how personal political engagement by a host can accelerate the migration from debate to echo chamber.

Bret Baier By contrast, Bret Baier—long marketed as Fox's "fair and balanced" chief political anchor (Darcy 2023)—shows no comparable drift. Across the full series (2010–2024) the monthly disagreement share hovers near 8%, with early-period (2010–2013) and late-period (2017–2024) means of 0.092 and 0.075 respectively, a difference that is not statistically distinguishable from zero (Figure 4b). Agreement frequencies are likewise flat (Appendix, Figure 13c). In other words, even as Fox's opinion block narrows its viewpoint corridor, the network's flagship news show preserves a modest but steady level of dissent which is slightly comforting for the state of journalism and news reporting.

Greg Gutfeld The third case illustrates how structural changes in scheduling can rewire host–guest interaction. Fox announced in April 2021 that *The Greg Gutfeld Show*, formerly a weekend panel show, would move to weeknights at 11 p.m. under the new title *Gutfeld!* (imd 2021).

The shift coincided with a pronounced jump in disagreement: mean dissent climbs from 0.031 in 2016–2021 to

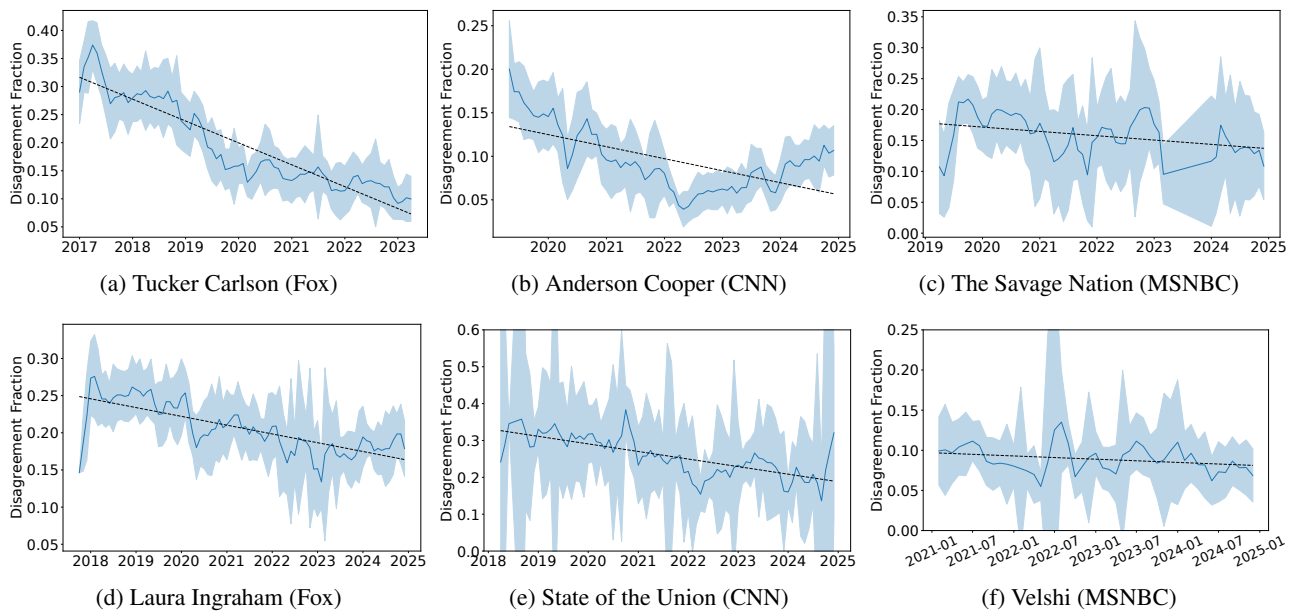


Figure 2: Trends in mean disagreement fraction over the years for a sample of the shows in our dataset.




Figure 3: Trends in mean agreement fraction over the years for a sample of the shows in our dataset.

0.122 in 2021–2024 (Figure 4c). Agreement, by contrast, contracts slightly (Appendix, Figure 12e). One plausible interpretation is that the late-night talk-show format rewards performative sparring, yielding more frequent but still low-stakes disagreement; yet even after the quadrupling, adversarial exchanges fill barely 12% of airtime, leaving the show well within the echo-chamber regime typified by its prime-time siblings.

Taken together, these longitudinal profiles reveal the heterogeneous pathways through which debate shows devolve (or resist devolving) into an ideological monologue. Hannity’s alignment with partisan activism precipitates a steady erosion of dissent; Bret Baier’s adherence to straight-news conventions stabilizes it; and *Gutfeld!*’s format reboot jolts the metric upward without restoring balance. Each trajectory, however, converges on the same broader outcome documented in Sections 5 and 5: a shrinking space for genuine contestation even as coverage gravitates toward the most polarizing policy arenas of American politics.

Channel-level trajectories: yearly aggregates, 2017–2024

Figure 5 collapses the show-specific results into annual means of disagreement fractions for each channel. Three facts stand out. First, the direction of movement in disagreement is clear: between 2017 and 2024 the average share of explicit host–guest disagreement falls across the board—by roughly one-third on Fox News, a quarter on CNN, and a fifth on MSNBC. The decline is gradual rather than abrupt, but its consistency across eight successive years suggests a structural shift in booking or production practices rather than ephemeral news-cycle shocks. Second, the ordering of the bars is counter-intuitive (especially for the Overall averages (right most set of bars in Figure 5)). Fox, long portrayed in the literature as an archetypal echo chamber (Jamieson and Cappella 2008), nonetheless retains the highest level of dissent throughout the series, finishing 2024 at just over 0.10—still above CNN’s 0.08 and MSNBC’s 0.07. One plausible explanation is our sample’s composition: since we have more shows from Fox which also include panel, opinion and news shows, whereas for CNN and MSNBC we primar-




Figure 4: Disagreement over time for Hannity, Baier and Gutfeld.

ily have only opinion shows. Third, although the networks start from different baselines, the slopes are convergent; the gap between Fox and its competitors shrinks by almost two percentage points over the eight-year window, signaling a sector-wide contraction in adversarial dialogue.

The mirror image appears in Figure 6. Average agreement climbs steadily on all three channels, with Fox again leading the pack: its mean agreement fraction rises from just under 0.40 in 2010 to about 0.55 in 2024, a gain of fifteen percentage points. CNN and MSNBC follow the same upward path, though from lower starting points and at a slightly slower pace. By the end of the period, the combined share of agreement plus neutrality exceeds 90% of airtime for every network—a level at which meaningful counter-argument becomes the exception rather than the rule.

Taken together with the show-level trajectories documented above, the yearly aggregates show clearly a troubling pattern where networks are shedding disagreement more quickly than it is gaining it. The data thus corroborate a core claim of selective-exposure scholarship that modern cable news tends to narrow, rather than widen, the interpretive frame of its audiences (Stroud 2011) but they also sharpen it: the narrowing is not merely a function of audience choice across outlets; it is being actively reinforced within the outlets themselves.




Figure 5: Yearly trends in disagreement for the 3 channels.

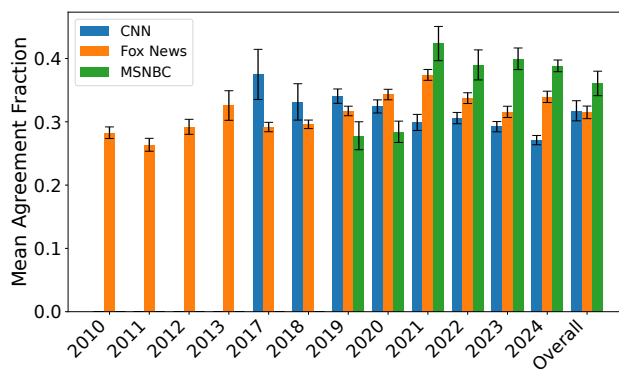


Figure 6: Yearly trends in agreement for the 3 channels.

Disagreement by Speakers Affiliation

The most straightforward pathway to ideological consonance is simply to curate the guest list. Figure 1 confirms that strategy in practice: Republicans dominate Fox, Democrats crowd MSNBC, while CNN remains roughly balanced, mirroring earlier content-analysis work on cable line-ups (Kim, Lelkes, and McCrain 2022). Yet composition alone cannot explain on-air dynamics; what matters is how those guests are treated once the cameras roll. Figure 7 disaggregates the data by the guest’s party label and the network to which they have been invited, allowing us to test whether “friendly” visitors are rewarded with deference and “outsiders” confronted with push-back.

On Fox, Republican guests do enjoy marginally warmer receptions: the average share of explicit agreement for Republicans significantly higher than for Democrats, while disagreement is identical for both groups. Statistical significance tested using a Welch’s t-test assuming unequal variance. $p < 0.01$ (Figure 7a). The substantive effect, however, is well under one percentage point suggesting that Fox’s echo-chamber reputation is driven more by *who* appears on screen than by large differences in tone once the invitation is extended. Put differently, the ideological filter operates at the booking stage; once a Democrat is allowed through that filter, the host’s deference scarcely deviates from the base-

line accorded to Republican regulars.

MSNBC displays a clearer in-group bias. Democratic guests receive a 0.40 agreement share, four points higher than Republicans (0.36), and experience slightly less disagreement (0.12 vs. 0.13). Both differences are statistically significant at $p < 0.01$ (Figure 7b). Although the raw differences may look small, recall that disagreement is already exceedingly rare; a one-point reduction in dissent represents nearly 8% of all hostile exchanges on the network. The result tallies with scholarship arguing that liberal-leaning outlets enforce ideological solidarity more tightly than their conservative counterparts once the guest roster is set (Sobieraj and Berry 2011).

CNN, by contrast, shows virtually no partisan differential. The absence of a treatment effect is consistent with CNN’s self-branding as a “down-the-middle” outlet and resonates with earlier findings that its conflict cues flow more from issue framing than from guest identity (Arceneaux and Johnson 2013).

Disagreement by Topic

Table 1 ranks the five topics that elicit the least and most host–guest disagreement on each channel. Topics with least disagreement indicate issues where there was little counter voice or debate. Three patterns emerge:

First, the same topics surface in the “least disagreement” column across otherwise opposed outlets. Polarizing topics such as Abortion, Immigration, Gun violence and Policing appear among the five topics on most of channels. These are precisely the issues that animate contemporary culture-war politics, yet they rarely provoke open discussion on the news channels. This pattern strongly supports our argument that these topics are not being used as subjects for debate, but rather as vehicles for narrative reinforcement and in-group consensus-building. Second, the topics that generate what most disagreement are largely procedural, event-driven or news related: the economy, elections & campaigns, COVID-19, etc. Third, there are interesting topics with little disagreement specific to channels such as discussions on Culture wars on both Fox News and MSNBC. Similarly, shows related to the investigations and impeachment of Donald Trump were among the least disagreed on MSNBC.

The implication from the findings in the previous two sections is that ideological filtering operates first at the guest-selection stage and then again at the topic-selection stage, leaving only narrow windows in which cross-cutting viewpoints make it to air. Taken together with the longitudinal evidence above, the topic analysis shows a central irony of partisan television. The shows devote ever more airtime to the controversies that polarize American politics (see Figure 9), yet they purge those very segments of meaningful disagreement.

6 Discussion

Summary of principal findings. Our speaker–turn analysis reveals a debate ecosystem that is strikingly placid by historical and cross-network comparison. Across the

Channel	Least Disagreement	Most Disagreement
Fox News	Abortion	COVID
	Cultural Politics	Economy
	Immigration	Election campaigns
	Gun Violence	Judiciary
	Race & Policing	Trump Investigations
MSNBC	Abortion	Judiciary
	Trump Investigations	Election campaigns
	Cultural Politics	Economy
	Immigration	Foreign Policy
	Race & Policing	COVID
CNN	Immigration	Economy
	Judiciary	Election campaigns
	Gun Violence	COVID
	Foreign Policy	Cultural Politics
	Abortion	Trump Investigations

Table 1: Topics with the most & least disagreement per Channel.

2017–2024 window fewer than one-in-six host–guest exchanges contains explicit dissent; MSNBC is lowest at 13% and Fox highest at 17% (Figure 7). More troubling is direction of travel: disagreement declines year-on-year for almost every Fox and CNN show and is flat (but never high) on the MSNBC shows for which long-run data exist (Figures 10, 11). The trajectory is not inevitable: archival captions from the early 2010s show *Hannity* hovering near 25% dissent, a level now unmatched by any prime-time hour. Guest composition compounds the effect. Networks already skew their invitation lists toward ideologically sympathetic voices, yet our turn-level labeling shows that even once a cross-partisan guest appears the dialogue tilts toward conformity: Republican invitees receive gentler treatment on Fox, Democrats on MSNBC, with CNN occupying a narrowing center. Finally, the very topics that polarize the electorate—abortion, gun rights, immigration—attract the least disagreement, converting ostensibly “hot” segments into performances of in-group consensus.

What do these patterns mean? Two implications stand out. First, the findings operationalize the “spectacle without contestation” hypothesis advanced by selective-exposure theory: cable outlets dramatize ideological battle lines through topic choice and affective framing yet withhold the clash of arguments that would expose audiences to competing rationales. That configuration can intensify *affective* polarization—warmth toward one’s own side and hostility toward the other—while leaving factual beliefs untested, a dynamic known to entrench misperceptions and policy gridlock. Second, the disappearance of dissent is not merely a product of audience self-sorting across channels; it is reinforced *within* channels by editorial practices that book, frame, and moderate conversations so as to minimize overt conflict. The data therefore complicate optimistic accounts that place the onus solely on viewer choice: even citizens who seek out disagreement may be offered little once they tune in.




Figure 7: Average agreement/disagreement by political party of guests invited on a channel. * indicates statistical significance at $p < 0.01$. *n.s.* indicates no significance.

Why the decline? Several mechanisms are plausible. Audience analytics reward affective reinforcement more reliably than cognitive challenge; hosts cultivate personal brands tied to partisan authenticity; and the economics of booking favor a stable roster of predictable allies. The COVID-19 pandemic further normalized remote interviews, making it easier to screen out wild-card guests. Our data cannot pinpoint causal weightings, but the cross-network symmetry—in which left-leaning and right-leaning channels mirror each other—suggests structural, not idiosyncratic, forces.

Is disagreement the right metric? Raw frequency of dissent is a blunt instrument. Some disagreements are performative, others substantive; some are civil, others demagogic. Still, frequency matters for at least three reasons. First, democratic theory assigns deliberation a central role: citizens form better-justified preferences when exposed to counter-arguments. Second, low baseline dissent limits the potential for “constructive conflict,” because intensity cannot compensate for absence. Third, our classifier is conservative: sarcasm, hedged rebuttals, and rapid-fire interruptions often register as *neutral*, implying that true argumentative engagement may be even rarer than we report. Future work should layer additional dimensions—tone, evidence quality, emotional valence—onto the stance signal to distinguish shallow sparring from reason-giving debate.

How Do Offline Debates Mirror Online Echo Chambers?

Unlike algorithmically mediated feeds, the cable-news environment is tightly moderated by hosts, limiting opportunities for organic cross-cutting exchange (Hosseinmardi et al. 2025). In this sense, our findings highlight an offline pathway to echo chamber effects (Ash et al. 2024): the absence of disagreement in televised debates can serve as a functional analogue to algorithmic curation online. Situating our work in this way underscores its relevance: understanding echo chambers requires analyzing both digital and legacy media environments that together shape citizens’ partisan worldviews (Muisse et al. 2022).

Limitations. The study has four chief constraints. *Coverage:* speaker-label quality forced reliance on podcast feeds

for many MSNBC and CNN shows, shortening their temporal span relative to Fox. *Measurement error:* ASR inaccuracies, diarization drift, and stance-classifier mistakes (accuracy=0.89) introduce noise, though audits indicate no directional bias large enough to change headline trends. *Affiliation coding:* DIME donations are an imperfect proxy for ideological leaning, especially for non-U.S. guests or low-profile experts. *Missing modalities:* our analysis ignores visuals, tone of voice, and real-time fact-checking, all of which shape viewer perception of conflict.

Future directions. Integrating sentiment arcs and evidence quality could distinguish “heat” from “light” in televised argument. Linking caption-level disagreement to real-time audience metrics (e.g., minute-by-minute ratings, social-media engagement) would clarify whether viewers reward or punish dissent. Finally, extending the method to streaming talk shows and YouTube political channels would test whether the retreat from contestation is a cable-specific artifact or a broader feature of contemporary political media.

Overall, the erosion of on-air disagreement documented here challenges the longstanding notion of cable news as a public forum for competing ideas. If deliberation is an essential democratic input, then reviving formats that encourage genuine clash—not just spectacle—remains an urgent task for journalists, regulators, and scholars alike.

References

- 2021. Fox News Shifts ‘Greg Gutfeld Show’ to Weeknights. *IMDb*. Accessed: 2025-05-16.
- Arceneaux, K.; and Johnson, M. 2013. *Changing minds or changing channels?: Partisan news in an age of choice*. University of Chicago Press.
- Ash, E.; Galletta, S.; Pinna, M.; and Warshaw, C. S. 2024. From viewers to voters: Tracing Fox News’ impact on American democracy. *Journal of Public Economics*, 240.
- Bakshy, E.; Messing, S.; and Adamic, L. A. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239): 1130–1132.
- Berry, J. M.; and Sobieraj, S. 2013. *The outrage industry:*

- Political opinion media and the new incivility*. Oxford University Press.
- Bonica, A. 2015. Database on ideology, money in politics, and elections (DIME). *Harvard Dataverse*, 2.
- Brown, J. R.; Enos, R. D.; Feigenbaum, J.; and Mazumder, S. 2021. Childhood cross-ethnic exposure predicts political behavior seven decades later: Evidence from linked administrative data. *Science Advances*, 7(24): eabe8432.
- Darcy, O. 2023. Fox News anchor Bret Baier’s reputation takes hit after text messages reveal what he said in wake of 2020 election. *CNN Business*. Accessed: 2025-05-16.
- Ding, X.; Horning, M.; and Rho, E. H. 2023. Same words, different meanings: Semantic polarization in broadcast media language forecasts polarity in online public discourse. In *ICWSM*.
- Farnsworth, S. J.; and Lichter, R. S. 2005. *The mediated presidency: Television news and presidential governance*. Rowman & Littlefield Publishers.
- Flaxman, S.; Goel, S.; and Rao, J. M. 2016. Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly*, 80(S1): 298–320.
- Forgette, R.; and Morris, J. S. 2006. High-conflict television news and public opinion. *Political Research Quarterly*.
- Gertz, M. 2020. Study: Trump sent 657 live tweets of fox programming in 2019. *Media Matters for America*.
- Gilens, M.; and Murakawa, N. 2002. Elite cues and political decision-making. *Research in micropolitics*, 6(1): 15–49.
- Goh, H. W.; Tkachenko, U.; Mueller, J.; and Cleanlab, C. C. 2022. Utilizing supervised models to infer consensus labels and their quality from data with multiple annotators. *CoRR*.
- Goodin, R. E. 1975. Cross-cutting cleavages and social conflict. *British Journal of Political Science*, 5(4): 516–519.
- Grootendorst, M. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794*.
- Hartman, R.; Moss, A. J.; et al. 2023. Introducing Connect by CloudResearch: Advancing online participant recruitment in the digital age.
- Hong, J.; Crichton, W.; Zhang, H.; Fu, D. Y.; Ritchie, J.; Barenholtz, J.; Hannel, B.; Yao, X.; Murray, M.; Moriba, G.; et al. 2021. Analysis of faces in a decade of us cable tv news. In *KDD’21: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*.
- Hosseinmardi, H.; Wolken, S.; Rothschild, D. M.; and Watts, D. J. 2025. Unpacking media bias in the growing divide between cable and network news. *Scientific Reports*, 15(1).
- Iyengar, S.; and Hahn, K. S. 2009. Red media, blue media: Evidence of ideological selectivity in media use. *Journal of communication*, 59(1): 19–39.
- Jamieson, K. H.; and Cappella, J. N. 2008. *Echo chamber: Rush Limbaugh and the conservative media establishment*. Oxford University Press.
- Jurkowitz, M.; Mitchell, A.; Santhanam, L. H.; Adams, S.; Anderson, M.; and Vogt, N. 2013. The Changing TV News Landscape. Technical report, Pew Research Center. Report.
- Kim, E.; Lelkes, Y.; and McCrain, J. 2022. Measuring dynamic media bias. *Proceedings of the National Academy of Sciences*, 119(32): e2202197119.
- Laohaprapanon, S.; and Sood, G. 2017. Archive.org TV News Closed Caption Corpus. https://github.com/notnews/archive_news_cc/. Accessed May 5, 2025.
- Levendusky, M. S. 2013. Why do partisan media polarize viewers? *American journal of political science*, 57(3).
- Littleton, C. 2024. Cable News Ratings 2024: Fox Gains, MSNBC and CNN In Post-Election Dropoff. *Deadline*. Accessed 2025-05-16.
- Löblich, M.; and Venema, N. 2021. Echo Chambers. *Digital roots: Historicizing media and communication concepts of the digital age*, 4: 177.
- Martin, G. J.; and Yurukoglu, A. 2017. Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9): 2565–2599.
- Morris, L. 2016. Effect of Media Bias on Credibility of Political News. *Xavier Journal of Undergraduate Research*.
- Muise, D.; Hosseinmardi, H.; Howland, B.; Mobius, M.; Rothschild, D.; and Watts, D. J. 2022. Quantifying partisan news diets in Web and TV audiences. *Science advances*.
- Mutz, D. C.; and Reeves, B. 2005. The new videomalaise: Effects of televised incivility on political trust. *American Political Science Review*, 99(1): 1–15.
- Pew Research Center. 2023. Cable News Fact Sheet. <https://www.pewresearch.org/journalism/fact-sheet/cable-news/>. Accessed: 2025-09-08.
- Radford, A.; Kim, J. W.; Xu, T.; Brockman, G.; McLeavey, C.; and Sutskever, I. 2023. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, 28492–28518. PMLR.
- Schwartz, J.; and Oprysko, C. 2018. Fox News raps Hannity for appearing onstage at Trump rally. *POLITICO*. Accessed: 2025-05-16.
- Sobieraj, S.; and Berry, J. M. 2011. From incivility to outrage: Political discourse in blogs, talk radio, and cable news. *Political Communication*, 28(1): 19–41.
- Stroud, N. J. 2010. Polarization and partisan selective exposure. *Journal of communication*, 60(3): 556–576.
- Stroud, N. J. 2011. *Niche news: The politics of news choice*. Oxford University Press.
- Times, N. Y. 2022. Inside the Apocalyptic Worldview of ‘Tucker Carlson Tonight’ (Published 2022) — nytimes.com. <https://www.nytimes.com/interactive/2022/04/30/us/tucker-carlson-tonight.html>. [Accessed 29-04-2025].
- Vaccari, C.; Valeriani, A.; Barberá, P.; Jost, J. T.; Nagler, J.; and Tucker, J. A. 2016. Of echo chambers and contrarian clubs: Exposure to political disagreement among German and Italian users of Twitter. *Social media+ society*, 2(3).
- Wu, Y.; Zhou, Y.; et al. 2024. KnowComp at DialAM-2024: Fine-tuning Pre-trained Language Models for Dialogical Argument Mining with Inference Anchoring Theory. In *Proceedings of the 11th Workshop on Argument Mining (ArgMining 2024)*, 103–109. ACL.

Paper Checklist

1. For most authors...
 - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
 - (b) Do your main claims in the abstract and introduction accurately reflect the paper’s contributions and scope? **Yes**
 - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes**
 - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes**
 - (e) Did you describe the limitations of your work? **Yes**
 - (f) Did you discuss any potential negative societal impacts of your work? **Yes**
 - (g) Did you discuss any potential misuse of your work? **No**
 - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **No**.
 - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes**
2. Additionally, if your study involves hypotheses testing...
 - (a) Did you clearly state the assumptions underlying all theoretical results? **NA**
 - (b) Have you provided justifications for all theoretical results? **NA**
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **NA**
 - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **NA**
 - (e) Did you address potential biases or limitations in your theoretical framework? **NA**
 - (f) Have you related your theoretical results to the existing literature in social science? **NA**
 - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **NA**
3. Additionally, if you are including theoretical proofs...
 - (a) Did you state the full set of assumptions of all theoretical results? **NA**
 - (b) Did you include complete proofs of all theoretical results? **NA**
4. Additionally, if you ran machine learning experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Yes**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **Yes**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **NA**
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **Yes**
 - (f) Do you discuss what is “the cost” of misclassification and fault (in)tolerance? **Yes**
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
 - (a) If your work uses existing assets, did you cite the creators? **Yes**
 - (b) Did you mention the license of the assets? **NA**
 - (c) Did you include any new assets in the supplemental material or as a URL? **NA**
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **Yes**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **Yes**
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? **Yes**
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? **NA**
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
 - (a) Did you include the full text of instructions given to participants and screenshots? **Yes**
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **No**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **No**
 - (d) Did you discuss how data is stored, shared, and de-identified? **Yes**

A Appendix

The complete list of shows used in the paper are shown in Table 3. Due to hiatus or simply not being able to scrape data from the Internet Archive platform or even Podcast, there are some gaps in the data. They are noted in table 2.

Detecting Agreement/Disagreement

Our stance-labeling workflow unfolded in four stages, moving from small, manually curated seeds to a production-scale LLM annotator.

Show	Coverage Gaps or Anomalies
<i>Hannity</i>	No data from 2014 to 2016.
<i>The Rachel Maddow Show</i>	Only one episode in September 2022, all labeled as agreement or neutrality because it was just Rachel and a reporter talking through the day’s news. Weekly airing format may reduce disagreement frequency.
<i>Outnumbered</i>	Very few episodes before September 2015; analysis starts from October 2015.
<i>Special Report with Bret Baier</i>	No episode available from November 2013 to December 2016.
<i>Life, Liberty & Levin</i>	Some months contain only 2–3 episodes with limited speaker turns; often results in 0 recorded disagreements.
<i>The Savage Nation</i>	Did not air from March 2023 through December 2023.
<i>Erin Burnett OutFront</i>	Did not air from September 2021 to December 2021 and again from October 2023 to December 2023.

Table 2: Data coverage gaps and broadcast anomalies for selected shows.

Stage 1: baseline with task-specific encoders. We first adopted the DialAM-2024 framework of Wu, Zhou et al. (2024), fine-tuning DEBERTA-v3-BASE on 160 of 200 hand-labeled bundles from *Tucker Carlson Tonight*. Test accuracy reached only 55%, and augmenting the training set with the original DialAM corpus drove accuracy down to 48.5%. These results ruled out conventional encoder-classifier models for our domain.

Stage 2: iterative GPT-seed expansion. Turning to instruction-tuned LLMs, we grew the seed set to 500 bundles via an interactive loop with ChatGPT (GPT-4). The model labelled batches of 20 pairs in zero-shot mode; authors corrected its mistakes and fed the revised items back as demonstrations before issuing the next batch. A subsequent 80/20 train-test split yielded 73% accuracy for zero-shot GPT-4 and 79% after fine-tuning GPT-3.5-TURBO. These figures established a workable floor but left room for improvement.

Stage 3: crowd round 1 and consensus filtering. To enlarge the training pool we recruited 150 U.S. participants on CloudResearch (Hartman, Moss et al. 2023). The participants were eligible if they were located in the U.S., aged ≥ 18 , had an historical approval rate $\geq 95\%$, and passed a short English comprehension screener. These were all set by the CloudResearch platform and we just had to upload our task website and collect the data from the participants. We sampled 1,000 bundles drawn equally from *Tucker Carlson Tonight*, *Hannity*, and *The Rachel Maddow Show*. Each 20-pair bundle was triple-annotated on a purpose-built website featuring guideline examples and embedded gold checks. We applied quality screening based on gold-check performance and implausibly short completion times, removing 12% of submissions. Inter-annotator reliability among the crowdsourced annotators was modest (Fleiss $\kappa = 0.302$; Krippendorff’s $\alpha = -0.005$), so to mitigate noise we used CLEANLAB (Goh et al. 2022) to infer consensus labels, retaining 459 high-confidence pairs. Combining these with the 500 GPT-seed pairs (959 total) and re-training multiple models produced the accuracies reported in Table 4.

Stage 4: crowd round 2, manual triage, and final model. Because model performance improved with training-set size, we commissioned a second CloudResearch round for another 1,000 bundles, and did the same cleaning with

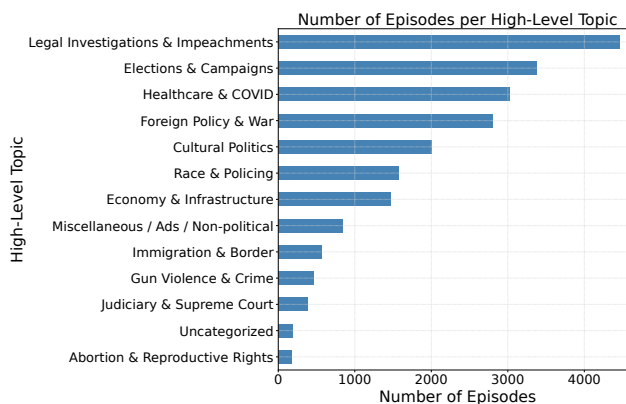


Figure 8: Episodes per high level topic.

CLEANLAB, to finally get 951 pairs of raw labels. Unlike in Stage 3, however, incorporating these raw labels directly reduced performance (Table 5). Manual audits by one of the authors revealed a systematic source of error: crowd workers frequently labeled backchannel acknowledgments (e.g., “mm-hm,” “yeah”) as “agreement,” contrary to our stricter definition of substantive stance. We therefore hand-verified every ambiguous triple, retained 712 balanced, high-quality pairs, and merged them with the earlier 959 from stage 3, yielding a 1,671-pair gold corpus. Fine-tuning GPT-4o on this set restored accuracy to 84%; DeepSeek-R1-Distill-Llama-8B, fine-tuned on the same data, achieved 89% (Table 6).

Table 7 reports the 100 manual sample checks mentioned in Section 4, including if the show’s dataset was taken from Archive or Apple podcast. The samples were not balanced across the 3 classes.

Topic Analysis

Figure 9 shows the high level topics per channel over time. We see clear trends like the increase in coverage in COVID/Race & policing in 2020, consistent coverage for foreign policy content, etc.

The full list of topics are shown below.

- 2012 GOP debate on Pakistan and national security post-bin Laden

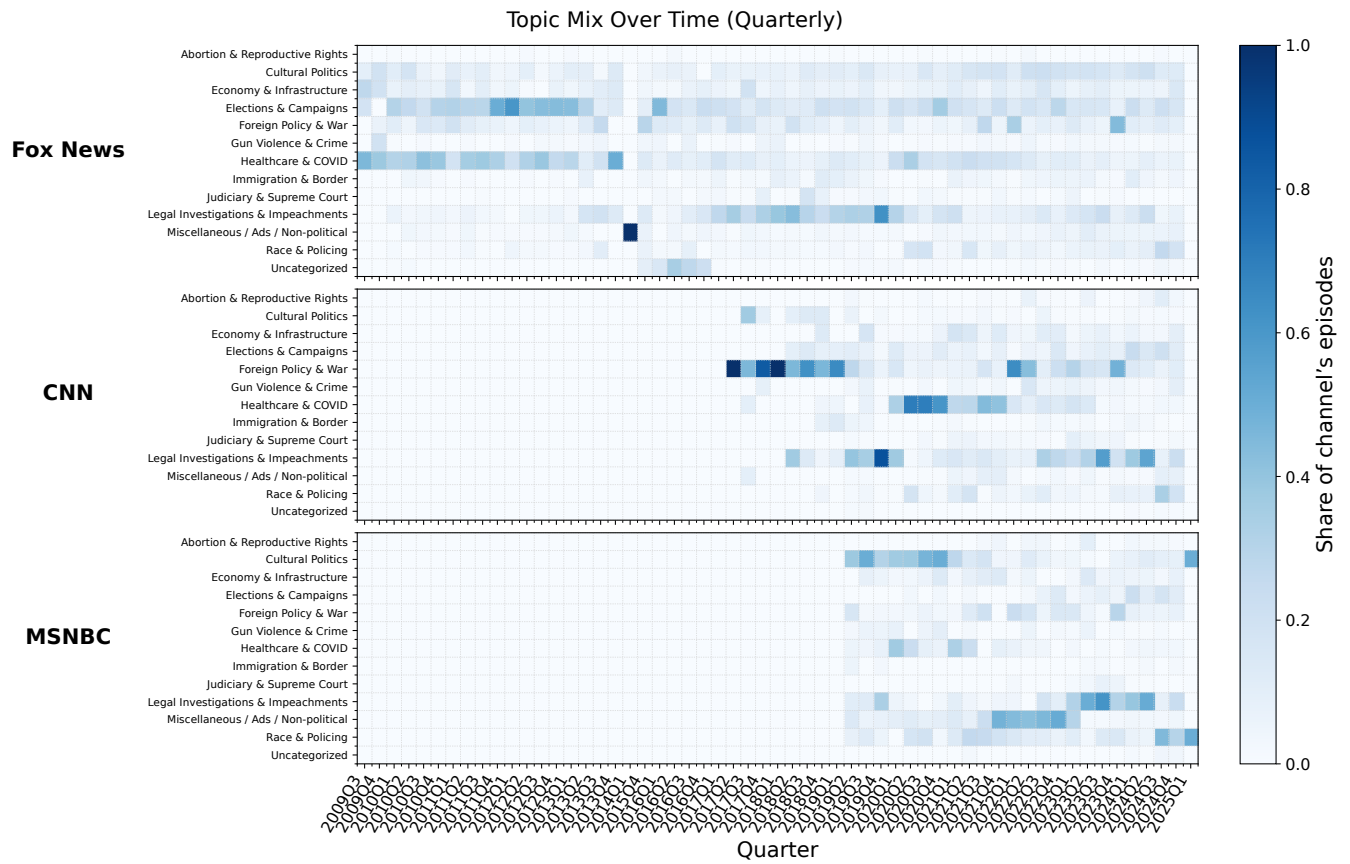


Figure 9: Topics per channel over time.

Show	Channel	Type	Time span	Episodes
The Tucker Carlson Show	Fox News	Weekday	Jan 2017 – Jun 2023	1491
Hannity	Fox News	Weekday	Jul 2009 – Jan 2014 and Jan 2017 – Dec 2024	2441
Special Report with Bret Baier	Fox News	Weekday	Jul 2009 – Oct 2013 and Jan 2017 – Dec 2024	2672
Outnumbered	Fox News	Weekday	Oct 2015 – Dec 2024	1644
The Ingraham Angle	Fox News	Weekday	Oct 2017 – Dec 2024	1584
Gutfeld!	Fox News	Weekday	Jan 2017 – Dec 2024	1001
Jesse Watters Primetime	Fox News	Weekday	Jan 2022 – Dec 2024	634
The Story with Martha MacCallum	Fox News	Weekday	May 2017 – Dec 2024	1639
Life, Liberty & Levin	Fox News	Weekend	Feb 2018 – Dec 2024	481
Anderson Cooper 360°	CNN	Weekday	May 2019 – Dec 2024	1203
Erin Burnett OutFront	CNN	Weekday	Jun 2019 – Dec 2024	1206
Inside Politics	CNN	Weekday	Jul 2019 – Dec 2024	1210
The Lead with Jake Tapper	CNN	Weekday	Jun 2019 – Dec 2024	1056
The Source with Kaitlan Collins	CNN	Weekday	Jul 2023 – Dec 2024	334
Laura Coates Live	CNN	Weekday	Oct 2023 – Dec 2024	283
State of the Union	CNN	Weekend	Apr 2018 – Dec 2024	316
Fareed Zakaria GPS	CNN	Weekend	Apr 2017 – Dec 2024	353
The Beat with Ari Melber	MSNBC	Weekday	Jan 2023 – Dec 2024	288
Inside with Jen Psaki	MSNBC	2 Weekday	Jun 2023 – Dec 2024	130
Alex Wagner Tonight	MSNBC	Weekday	Jan 2024 – Dec 2024	179
Velshi	MSNBC	Weekend	Feb 2021 – Dec 2024	340
The Sunday Show with Jonathan Capehart	MSNBC	Weekend	Jan 2023 – Dec 2024	117
The Savage Nation	MSNBC	Weekday	Apr 2019 – Dec 2024	648
The Rachel Maddow Show	MSNBC	Weekly	May 2022 – Dec 2024	126

Table 3: Broadcast shows, overall time span, and total episode counts.

2. 2012 GOP primaries and healthcare policy debates
3. 2020 Democratic Primary: Sanders vs. Bloomberg, Super Tuesday, and Media Scandals
4. 2020 Democratic primary candidates
5. 2020 Trump Foreign Policy
6. 2024 Republican Primary Campaigns in Iowa and New Hampshire
7. Abortion Rights and the Overturning of Roe v. Wade
8. Ads
9. Afghanistan Withdrawal and Evacuation Crisis at Kabul Airport
10. Andrew Cuomo's Resignation Amid Nursing Home and Sexual Harassment Scandals
11. Black Lives Matter, socialism, and cultural politics
12. Black Voter Representation, Kamala Harris Campaign, and Tim Walz Commentary
13. Brett Kavanaugh Confirmation Hearings and Sexual Assault Allegations
14. Brexit
15. British Royal Family Coverage: Queen Elizabeth's Legacy and King Charles's Ascension
16. COVID origins, Christianity, immigration, and media figures
17. Campus Encampments, Free Speech, and Debates on Antisemitism in Higher Education
18. Chinese Spy Balloon Incidents and U.S. Airspace Surveillance Controversies
19. Collective bargaining and emergency authority
20. Conservative Commentary on Marijuana, Free Speech, and Cultural Politics
21. Covid-19
22. Critiques of the Chinese Communist Party and US-China Strategic Rivalry
23. DOJ and FBI Leadership During Mueller Investigation and Trump-Era Controversies
24. Debates Over Tax Reform, Immigration, and Progressive Policy Priorities
25. Debt Ceiling Debates and Elon Musk's Influence in Political Discourse
26. Democrats
27. Diplomatic Relations and Summit Talks with North Korea
28. Early Trump administration power struggles and turnover
29. Egyptian Revolution and Mubarak's Fall Amid Muslim Brotherhood Rise
30. Electoral Strategy and Media Framing in Battleground States Like Wisconsin
31. Gabby Petito case coverage
32. George Floyd Protests and Police Response in Minneapolis and Beyond
33. George Floyd's Murder and the Prosecution of Derek Chauvin
34. Georgia Senate Elections and Runoff Campaigns
35. Gutfeld! Comedy
36. Hurricane Coverage: Storm Surge, Flooding, and Landfall Impacts
37. Hurricanes and FEMA Disaster Response in Hawaii, Texas, and the Southeast
38. IRS Fraud

Model	Dataset (Balanced classes)	Train, Test	Accuracy	Precision	Recall	F1-score
DeepSeek-R1-Distill-Llama-8B	High quality subsets from the crowdsourced dataset	(339, 120)	66.97%	65.87%	69.12%	66.98%
DeepSeek-R1-Distill-Llama-8B	High quality subsets from the crowdsourced dataset + our own annotated 500 pairs	(498, 184)	70.64%	70.31%	68.27%	70.01%
Llama-3.3 70b 4bit	High quality subsets from the crowdsourced dataset	(339, 120)	69.64%	67.98%	65.75%	67.43%
Llama-3.3 70b 4bit	High quality subsets from the crowdsourced dataset + our own annotated 500 pairs	(498, 184)	74.98%	70.56%	73.21%	71.97%
Qwen-2.5 32b 4bit	High quality subsets from the crowdsourced dataset	(339, 120)	64.08%	65.00%	64.78%	64.88%
Qwen-2.5 32b 4bit	High quality subsets from the crowdsourced dataset + our own annotated 500 pairs	(498, 184)	65.22%	65.01%	64.89%	64.95%
GPT-4o	High quality subsets from the crowdsourced dataset	(339, 120)	79.43%	79.41%	78.90%	79.10%
GPT-4o	High quality subsets from the crowdsourced dataset + our own annotated 500 pairs	(498, 184)	80.43%	81.00%	80.02%	80.55%

Table 4: Comparison of fine-tuned model accuracy using different combination of datasets (phase 1)

Model	Dataset (Balanced classes)	Train, Test	Accuracy	Precision	Recall	F1-score
DeepSeek-R1-Distill-Llama-8B	High quality subsets from the crowdsourced datasets	(528, 245)	62.45%	61.98%	59.90%	60.42%
DeepSeek-R1-Distill-Llama-8B	High quality subsets from the crowdsourced datasets + our own annotated 500 pairs	(861, 344)	68.90%	69.02%	68.86%	69.00%
Llama-3.3 70b 4bit	High quality subsets from the crowdsourced datasets	(528, 245)	62.04%	62.49%	62.11%	62.37%
Llama-3.3 70b 4bit	High quality subsets from the crowdsourced datasets + our own annotated 500 pairs	(861, 344)	69.00%	68.43%	69.49%	69.12%
Qwen-2.5 32b 4bit	High quality subsets from the crowdsourced datasets	(528, 245)	64.08%	64.07%	63.99%	64.00%
Qwen-2.5 32b 4bit	High quality subsets from the crowdsourced datasets + our own annotated 500 pairs	(861, 344)	62.50%	60.99%	61.78%	62.34%
GPT-4o	High quality subsets from the crowdsourced datasets	(528, 245)	72.24%	73.00%	71.44%	71.49%
GPT-4o	High quality subsets from the crowdsourced datasets + our own annotated 500 pairs	(861, 344)	66.87%	67.08%	67.00%	66.09%

Table 5: Comparison of fine-tuned model accuracy using different combination of datasets (phase 2)

- | | |
|--|---|
| <ul style="list-style-type: none"> 39. ISIS-Linked Terror Attacks in Europe: Paris, Manchester, and Brussels 40. Inflation, Energy, and Economic Slowdown 41. Israel-Gaza Conflict and Ceasefire Discussions 42. January 6th 43. Jeffrey Epstein Case and Legal Accountability of Labor Secretary Acosta 44. Kevin McCarthy 45. Kyle Rittenhouse Trial, Verdict, and Broader Jury Trials in Media Focus 46. Las Vegas Mass Shooting and Debates Over Gun Violence and Background Checks 47. Leaks, executive oversight, and antisemitism in political media discussions 48. Manchin's role in Democratic infrastructure and vaccine policy negotiations 49. Mar-a-Lago Classified Documents Case 50. Mass School Shootings and Mental Health Crisis in the US 51. Matt Gaetz Ethics Scrutiny and Cabinet Appointments During Presidential Transition 52. Mueller, Obamacare, and political controversies 53. National identity, foreign influence, and Second | <ul style="list-style-type: none"> Amendment debates in conservative media 54. Nationalist cultural commentary and anti-communist rhetoric 55. Nikki Haley's 2024 Campaign, Southern Primaries, and IVF Policy Debates 56. Obamacare 57. Paul Pelosi Attack Incident and Related Legal Proceedings 58. Police Brutality and Public Outcry Following the Deaths of Floyd, Taylor, and Brooks 59. Postal service, shutdowns, and COVID-era political coverage 60. Religion, patriotism, and cultural conflict in U.S. media (Christianity, NFL, anthem protests) 61. Republican National Convention Coverage Amid Assassination Attempt and Political Tensions 62. Right-Wing Critiques of Race, Immigration, and Progressive Politicians 63. Right-Wing Framing of the Democratic Party and Radical Leftist Ideologies 64. Rudy Giuliani's Legal Troubles: Election Worker Defamation and Mueller-Era Ties 65. Russia-ukraine war and nato's strategic response |
|--|---|

Model	Dataset (Balanced classes)	Train, Test	Accuracy	Precision	Recall	F1-score
GPT-4o	High quality subsets from the crowdsourced datasets (manually checked)	(459, 253)	84.09%	83.96%	84.00%	83.98%
DeepSeek-R1-Distill-Llama-8B	High quality subsets from the crowdsourced datasets (manually checked)	(459, 253)	89.61%	88.50%	87.91%	88.35%

Table 6: Final model used

Show	Source	Accuracy	Precision	Recall	F1-score
The Tucker Carlson Show	Archive	89.0%	87.5%	88.2%	87.8%
Hannity	Archive	88.6%	87.0%	87.9%	87.4%
Special Report with Bret Baier	Archive	88.2%	86.5%	87.3%	86.9%
Outnumbered	Archive	88.0%	86.2%	87.0%	86.6%
The Ingraham Angle	Archive	87.8%	86.0%	86.8%	86.4%
Gutfeld!	Archive	87.5%	85.8%	86.6%	86.2%
Jesse Watters Primetime	Archive	87.2%	85.5%	86.3%	85.9%
The Story with Martha MacCallum	Archive	88.1%	86.3%	87.2%	86.7%
Life, Liberty & Levin	Archive	88.4%	86.7%	87.5%	87.1%
Anderson Cooper 360°	Podcast	86.5%	84.5%	85.3%	84.9%
Erin Burnett OutFront	Podcast	86.2%	84.3%	85.0%	84.6%
Inside Politics	Podcast	86.0%	84.0%	84.8%	84.4%
The Lead with Jake Tapper	Podcast	85.8%	83.9%	84.7%	84.3%
The Source with Kaitlan Collins	Podcast	85.6%	83.7%	84.5%	84.1%
Laura Coates Live	Podcast	85.5%	83.6%	84.4%	84.0%
State of the Union	Podcast	85.7%	83.8%	84.6%	84.2%
Fareed Zakaria GPS	Podcast	85.9%	84.0%	84.8%	84.4%
The Beat with Ari Melber	Podcast	86.3%	84.4%	85.2%	84.8%
Inside with Jen Psaki	Podcast	86.1%	84.2%	85.0%	84.6%
Alex Wagner Tonight	Podcast	85.8%	83.9%	84.7%	84.3%
Velshi	Podcast	86.0%	84.1%	84.9%	84.5%
The Sunday Show with Jonathan Capehart	Podcast	85.7%	83.8%	84.6%	84.2%
The Savage Nation	Podcast	86.2%	84.3%	85.1%	84.7%
The Rachel Maddow Show	Podcast	86.4%	84.5%	85.3%	84.9%

Table 7: Performance metrics of different datasets after final model labeling

66. San Francisco Homelessness
67. Saudi Arabia’s Crown Prince and the Killing of Jamal Khashoggi
68. Scandals and Public Figures (e.g., Avenatti, Bernie)
69. Syrian Civil War and U.S. Policy on Assad, Chemical Weapons, and ISIS
70. Tea Party-Era Health Reform Debates
71. Titan Submersible Disaster and Ocean Exploration Catastrophe Investigation
72. Trump Administration Infighting, Mueller-Era Investigations, and DACA Policy Debates
73. Trump Hush Money Trial Featuring Cohen and Daniels
74. Trump Impeachment Hearings and Whistleblower Testimony
75. U.S. Drone Strike on Qasem Soleimani and Escalating Tensions with Iran
76. US-Mexico border security and cartel-driven immigration
77. Urban Crime and Policing in Central Park and Manhattan with NYPD Coverage
78. VP Kamala Harris and Democratic Campaign Strategies on Abortion

79. Virginia and New Jersey Gubernatorial Races with National Culture War Overtones

Disagreement by Profession

To assess whether a guest’s profession predicts on-air disagreement, we first needed to ensure our comparisons were statistically stable. We established a minimum sample size for each occupation by defining a desired level of precision: a 95% confidence interval no wider than ± 3 percentage points around the global mean disagreement rate (16%). The calculation determined that the minimum number of labeled pairs required to meet this criterion is 575. Therefore, we restricted our analysis to occupations with at least 575 samples, providing a robust basis for comparison. The occupation labels themselves were sourced from the DIME dataset and normalized to create consistent categories (e.g., attorney, county attorney, and attorney general were all merged into lawyer/attorney). The results in Table 8 show that legislators—both members of Congress (19.9%) and senators (18.6%)—are among the most frequently challenged guests, along with managers and entertainers (20%). Lawyers, executives, and CEOs show slightly lower rates (16–17%), while

educators, researchers, and students are the least disagreed with (<15%). These patterns suggest that disagreement is more common when guests hold formal decision-making authority or represent political institutions, but is still broadly distributed across professional categories. Importantly, the overall network-level decline in disagreement is not driven by occupational composition, as all high-frequency groups show parallel declines over time. However, it is important to note that while the DIME dataset includes occupation fields, approximately 35% of the guests in our corpus could not be reliably matched to an occupation entry due to missing or incomplete data in the DIME dataset. As a result, the occupation-based analysis should be interpreted with caution, as it necessarily reflects a subset of the full guest population.

Agreement/Disagreement Plots for All Shows

Figures 10, 11 and 12, 13 show the disagreement and agreement plots for all shows, respectively.

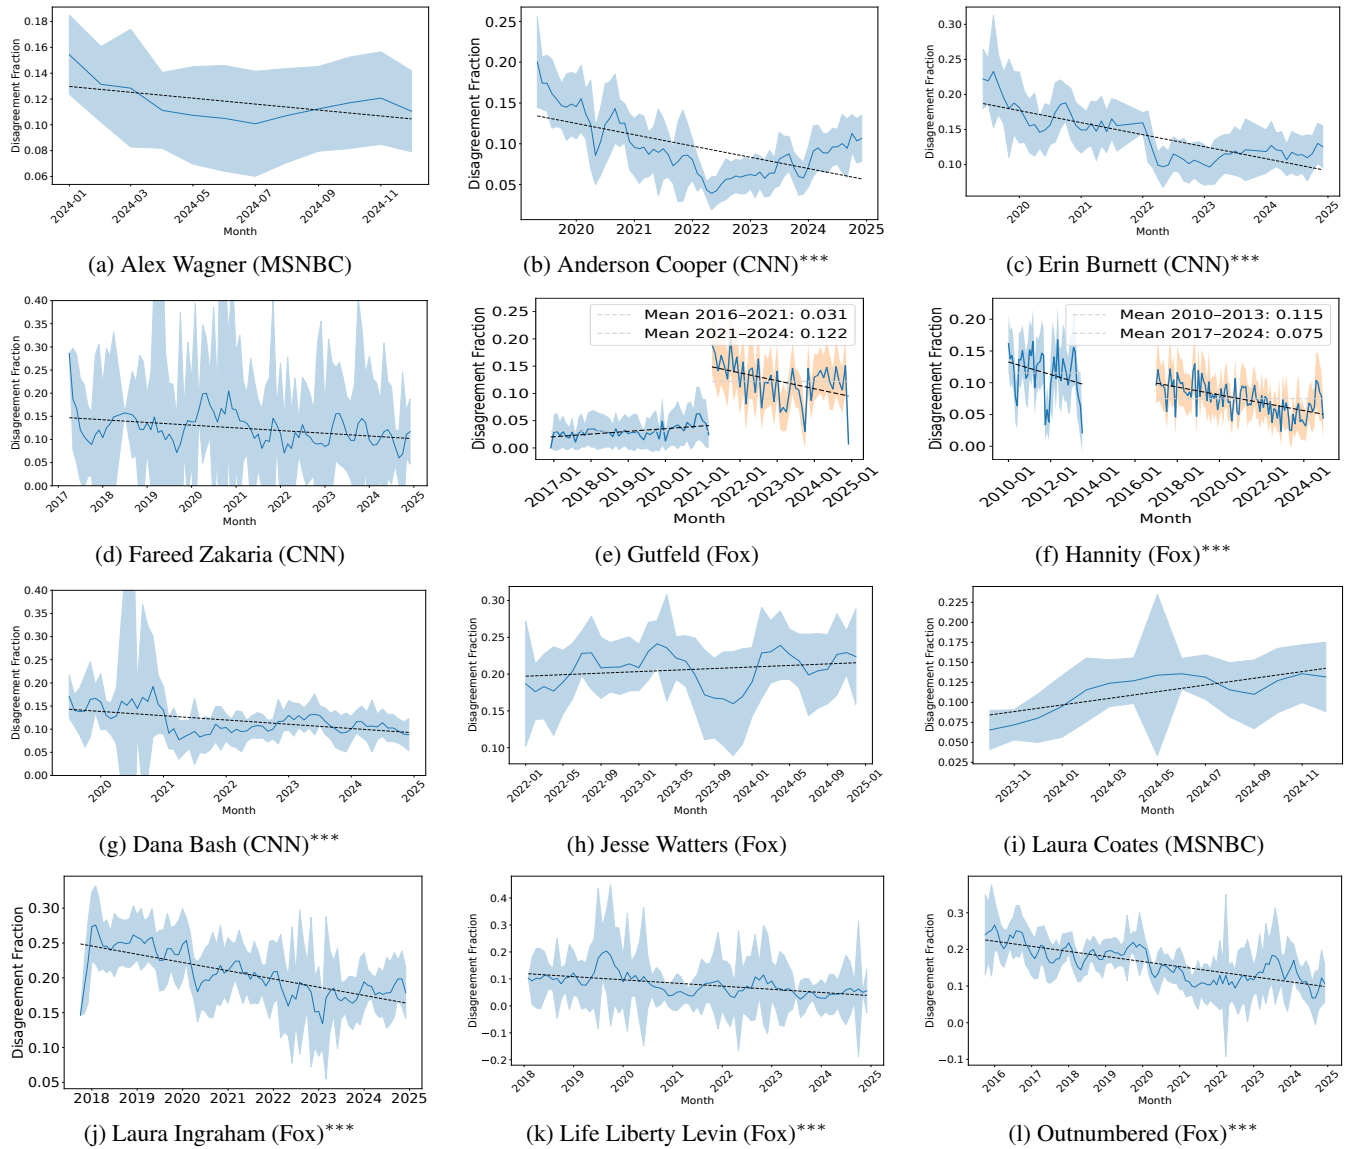




Figure 10: Disagreement patterns over time (half of the shows). *** indicates a significant decline ($p < 0.001$) in slope.

Occupation	Disagree Count	Total Pairs	Disagreement Rate
Entertainer	393	1,929	0.204
Manager	552	2,752	0.201
Member of congress	1,598	8,036	0.199
Rancher	170	855	0.199
Insurance	132	689	0.192
Director	1,059	5,588	0.190
Therapist	231	1,221	0.189
Disabled	301	1,597	0.188
Senator	1,599	8,590	0.186
Airline pilot	124	687	0.180
Artist	547	3,041	0.180
Principal	348	1,952	0.178
Architect	326	1,884	0.173
Lawyer/attorney	6,431	37,245	0.173
Farmer	276	1,601	0.172
Homemaker	424	2,488	0.170
Banker	249	1,462	0.170
Pastor	151	894	0.169
Retired	5,212	30,897	0.169
Not employed	4,992	29,599	0.169
Consultant	1,680	10,011	0.168
Teacher	605	3,611	0.168
Legislator	269	1,607	0.167
Executive	1,425	8,517	0.167
CEO/president	2,906	17,424	0.167
Self employed	683	4,116	0.166
Author	140	847	0.165
Chairman	464	2,809	0.165
Engineer	633	3,848	0.165
Political scientist	98	598	0.164
Contractor	209	1,280	0.163
Public relations	135	828	0.163
Physician	909	5,585	0.163
Accountant	152	955	0.159
Firefighter	224	1,410	0.159
Partner	381	2,426	0.157
Economist	181	1,154	0.157
Professor/writer	1,667	10,659	0.156
Investor	147	946	0.155
Psychologist	118	764	0.154
Real estate	574	3,727	0.154
Educator	172	1,126	0.153
Programmer	122	802	0.152
Analyst	88	584	0.151
Clerk	106	711	0.149
Administrator	87	586	0.148
Founder	317	2,155	0.147
Producer	112	779	0.144
Finance	134	935	0.143
Student	274	1,916	0.143
Research scientist	143	1,030	0.139
Designer	104	833	0.125
Unemployed	125	1,042	0.120


Table 8: Disagreement rates by occupation




(a) Rachel Maddow (MSNBC)




(b) Saturdays Sundays with Jonathan Capehart (MSNBC)




(c) Special Report with Bret Baier (Fox) ***




(d) State of the Union (CNN) ***




(e) The Beat with Ari Melber (MSNBC)




(f) The Lead with Jake Tapper (CNN) ***




(g) The Savage Nation (MSNBC)



(h) The Source with Kaitlin Collins (CNN)




(i) The Story with Martha MacCallum (Fox) ***



(j) Tucker Carlson (Fox) ***




(k) Velshi (MSNBC)




(l) Inside with Jen Psaki (MSNBC)


Figure 11: Disagreement patterns over time (rest of the shows). *** indicates a significant decline ($p < 0.001$) in slope.




(a) Alex Wagner Tonight (MSNBC)




(b) Anderson Cooper (CNN)




(c) Erin Brunett Outfront (CNN)




(d) Fared Zakaria GPS (CNN)




(e) Gutfeld (Fox)




(f) Hannity (Fox) ***




(g) Inside Politics with Dana Bash (CNN)




(h) Jesse Waters (Fox)




(i) Laura Coates Live (MSNBC)



(j) Laura Ingraham (Fox) ***




(k) Life Liberty Levin (Fox) ***




(l) Outnumbered (Fox) ***


Figure 12: Agreement patterns over time (half of the shows). *** indicates a significant increase ($p < 0.001$) in slope.




(a) Rachel Maddow (MSNBC)




(b) Saturdays Sundays with Jonathan Capehart (MSNBC)




(c) Special Report with Bret Baier (Fox)




(d) State of the Union (CNN)




(e) The Beat with Ari Melber (MSNBC)




(f) The Lead with Jake Tapper (CNN)




(g) The Savage Nation (MSNBC) ***




(h) The Source with Kaitlin Collins (CNN)




(i) The Story with Martha MacCallum (Fox)



(j) Tucker Carlson (Fox) ***



(k) Velshi (MSNBC)



(l) Inside with Jen Psaki (MSNBC) ***

Figure 13: Agreement patterns over time (rest of the shows). *** indicates a significant increase ($p < 0.001$) in slope.