

## **Majority Rules: Polarizing Content Dissemination and User Experience on TikTok**

Lexie Matsumoto

Bruno Coehlo

Tobias Lauinger

New York University, USA

Laura Edelson

Northeastern University, USA

Damon McCoy

New York University, USA

An increasing part of public debate and opinion formation occurs on social media. It is therefore crucial to study how different stances about a topic are represented on platforms, both from the perspective of individual users (e.g., concerns about “filter bubbles”) and at the aggregate system level (e.g., whether there is any skew in the dissemination of stances). In this article, we investigate the issue of stance representation at the platform level, and differences in experience at the individual level, in the context of two controversial, mostly two-sided issues: the legality of abortion in the U.S., and the Israel/Hamas conflict. We do so on TikTok, a short-form video platform known for highly viral video content and recommendations that give relatively little weight to the structure of the social graph. In manually annotated representative samples containing over 3.8k videos and spanning a combined 37 weeks, we measure the number of videos available on the platform and the views these videos receive, separated by the stance they represent (Pro Choice vs. Pro Life, and Pro Israel vs. Pro Palestine). We complement these platform-level analyses with a contemporaneous estimate of TikTok

---

Lexie Matsumoto: lb4535@nyu.edu

Date submitted: 2025-09-26

user opinions and a survey of users' content recommendation experiences. Around 34.5% of TikTok users identified as Pro Life and 53.7% as Pro Israel in Pew surveys fielded close to our topic time frames. However, at the platform level, we find that videos labeled as either Pro Life or Pro Israel are markedly underrepresented both in terms of video counts and viewership; we found only 3.2% of abortion-related videos (3.0% of views) represented a Pro Life stance, and 8.4% of Israel/Hamas-related videos (12.0% of views) represented a Pro Israel stance, respectively. Furthermore, our survey found that the user experience differed according to their stance on the topic. While Pro Choice and Pro Palestine respondents tended to be shown mostly videos in line with their opinion, users aligned with the opposing opinions were more likely to report being shown mostly videos they disagreed with, or seeing less on-topic content overall. Our study does not provide evidence for the reason of this imbalance, but it indicates that TikTok's algorithmic recommendation system presents individual users with a distorted view of the debate and of the stances of their fellow users as a whole.

*Keywords:* *TikTok, polarization, content reach, user survey*

## Introduction

Social media has become the primary source of news for many people, surpassing television for the first time in 2025 (Newman, 2025). Social media is a complex system; what users encounter is shaped by a feedback loop involving at least the behavior of content creators (e.g., which content they produce), platform design (e.g., how available content is recommended to available users), and user behavior (e.g., which recommended content users watch or engage with). Ideally, users would be exposed to a variety of opinions through engagement with other users, however, personalization on social media has led to concerns about possible "filter bubbles" where users are served content that predominately aligns with their preexisting viewpoints (Cinelli et al., 2021), which would preclude meaningful debate.

Prior work has often focused on the social network aspects of online debate, i.e. how users interact with their connections and how structurally polarized a debate may be (Serrano et al., 2020; Garimella et al., 2017). An alternative vantage point is how represented different stances of the debate are at an aggregate platform level; for example, whether the content available or viewed on the platform mirrors the opinions of its user base. A skewed distribution would not necessarily imply that the platform is intentionally interfering with a debate, but may indicate that aspects of the overall (eco)system can tilt a debate toward a specific opinion. It is important to be aware of such effects because of their impact on user perceptions and civil discourse on the platform.

This study contributes to a growing body of research auditing political recommendations on TikTok (e.g., (Ibrahim et al., 2025; Tjaden et al., 2025)), but is unique in shifting the focus from broad political partisanship, such as left- and right-wing “camps”, to focusing on stances on specific issues (i.e., how polarizing issues are represented on a platform). If one opinion is disproportionately amplified, it may preclude the ability for users to be exposed to alternative viewpoints, or give users a potentially distorted impression of the dominance of that opinion. In this paper, we take a system-level perspective of this issue to understand the outcome of TikTok’s recommendation system as a whole. Specifically, we aim to measure the dissemination outcomes of a content-driven recommendation system within the context of polarizing issues. We choose two topics that were controversial and prevalent across social media during our study period: the right to abortion access under *Roe v. Wade*, and the Israel/Hamas conflict following October 7th, 2023.

We focus on three main research questions to study TikTok’s platform environment in the context of polarizing debates. First, how are the two opposing opinions on polarizing issues represented and engaged with on TikTok? Does this balance of representation between stances change over time? Finally, does a user’s stance on polarizing issues impact their experience on TikTok in terms of content exposure and alignment?

We choose to observe TikTok, a short-form video sharing platform known for higher virality and lower weight given to structural properties of the social network in recommendations compared to other platforms (Guinaudeau et al., 2022). Using TikTok’s Research API, we collect a total of 1.2M videos from TikTok; 560k videos are related to abortion and 641k

to Israel/Hamas. Our collection periods are centered around major events relating to each issue. For our abortion set, we collect from one month before the leak of the Roe v. Wade Supreme Court draft decision until two months after the overturn was officially announced. For Israel/Hamas, we collect starting the week before the October 7th, 2023 Hamas attack on Israel and the following 4 months. We collect videos by querying the API with a set of curated hashtags for each topic to capture relevant content, which we later validate.

To determine the specific opinions expressed in each video, we rely on manual annotation. For abortion, we classify relevant videos as either *Pro Choice* (supporting the right to have legal access to induced abortion services) or *Pro Life* (opposing induced abortion and supporting its legal prohibition or restriction). Similarly, for the Israel/Hamas data, we classify content as either *Pro Palestine* (sympathize with the Palestinian people) or *Pro Israel* (sympathize with the Israeli people). Our analysis is based on representative samples, for which we manually annotate over 3.8k videos spanning a combined 37 weeks.

Our findings indicate that there is disproportionately little *Pro Life* and *Pro Israel* content available on TikTok, both in terms of available videos (3.2% and 8.4%, respectively) and accrued views (3.0% and 12.0%). We extensively validate our findings to exclude the possibility of a lack of *Pro Life* or *Pro Israel* recall that would skew our results, but we find no evidence to suggest that we failed to acquire a significant amount of content for either stance, indicating our acquisition methodology is robust.

For a better understanding of TikTok’s user base and their experience on the platform, we conduct a User Experience survey of 450 U.S.-based TikTok users. We ask participants for their views of each topic and whether or not they saw content that agreed or disagreed with their opinions. We find that the more extreme of an opinion a user had, the more likely they were to be served content relevant to the topic. When relevant content is served to users, it appears to broadly reflect the platform-wide view distribution, i.e., users aligned with the most popular content reported seeing little disagreeing content, whereas users with opposing views reported disagreeing with content more than they agreed.

Overall, our work contributes a methodology for capturing, validating, and studying the platform-level stance distribution of polarizing topics on TikTok. Using the examples of

abortion access and the Israel/Hamas conflict, we show that both available content and accrued views are extremely skewed towards *Pro Choice* and *Pro Palestine* stances, much more than the opinions of the TikTok user population would suggest. Our results indicate that content dissemination outcomes on TikTok are skewed toward one side of each debate, which overall may present users with a false sense of what their peers think.

### Background

Our study is concerned with differences in content availability and dissemination *in terms of stance* in the context of polarizing debates on TikTok. While we do not study the debates themselves, we provide background information on the two selected topics as it relates to the chosen terminology, topic stances, and key events that happened during the measurement period. (We lay out in the Methodology section how the two topics were selected.)

The first polarizing topic used in our study, abortion access, has been a contentious issue in America for decades. Since the Supreme Court's ruling in 1973 (United States Reports, 1973), access to abortion had been considered a federal constitutional right, making state-level bans on abortion unenforceable. On June 24th, 2022, this precedent was overturned with the Court's decision in *Dobbs v. Jackson Women's Health Organization* (United States Reports, 2022), following a leaked draft opinion earlier that year (Gerstein and Ward, 2022). As of 2026, abortion law varies state by state, with some guaranteeing total access while others impose near-total bans (Guttmacher Institute, 2025). We define the two major stances on the abortion debate to be *Pro Choice* (supporting the right to have legal access to induced abortion services) and *Pro Life* (opposing induced abortion and supporting its legal prohibition or restriction). We name these stances according to the prevalent terminology each side uses to refer to themselves.

The second polarizing topic considered in this study is the Israel/Hamas conflict. While U.S. policy has long aligned with Israel, Americans are divided in their sympathies. Recent Pew surveys show that younger adults and Democrats are more likely to support the Palestinian people, while older adults and Republicans are more likely to side with Israel (Silver, 2024). These generational and partisan divides have fueled conflicts and led to contentious public debate, including widespread campus protests following the Hamas-led attacks on Israel on October 7th, 2023 and Israel's subsequent military response (Associated Press, 2024). For

this topic, we define the major stances as *Pro Palestine* (sympathize with the Palestinian people, which for our labeling purposes does not necessarily imply support for Hamas or for Palestine as a state) and *Pro Israel* (sympathize with the Israeli people). We choose these stance labels based on what the different sides call themselves online, while we name the topic following the terminology used by the Associated Press (Associated Press, 2023) and the U.S. Congress (Sharp and Zanotti, 2024).

### Related Work

To date, it is still an open question how represented the two sides of a polarized debate are on TikTok both in terms of video availability and reach, and how users' experiences differ depending on their stance. Prior work focused either on studying the topics from a content perspective, polarizing debates from a structural social network perspective, user attitudes as a possible outcome of recommendations (and potential intervention by the platform in favor of one stance), or individual user-level recommendations in a controlled setting that does not reveal the aggregate platform-level outcomes of these recommendations.

Recent work by Ibrahim et al. (Ibrahim et al., 2025) uses “sock puppet” accounts across politically diverse U.S. states seeded with partisan content to observe subsequent recommendations made by TikTok. Using only videos with audio transcripts available (26% of recommended TikToks), the authors use human-validated LLMs to classify content collected by the bots. They find that Republican-seeded accounts received significantly more aligned content than Democratic-seeded accounts. Similarly, Tjaden et al. (Tjaden et al., 2025) use behaviorally identical sock puppet accounts to measure the reach of political content in Germany prior to major elections. They find far-right TikToks performed better than moderate alternatives. Unlike our use of API data to understand the platform-level distributions of stanced content, these methods focus on controlled user input to directly observe algorithmic effects.

Previous work has been conducted to understand users' perceptions of polarized content on social media. A mixed methods study reports that newly created accounts show less content unfavorable to the Chinese Communist Party on TikTok compared to other platforms, and a U.S. survey ( $N = 1,214$ ) links heavier TikTok usage to more favorable opinions towards China's human rights record (Titelman et al., 2024). Our study finds a different effect on

topics not directly linked to China, namely a tilt towards a majority stance that is sometimes disconnected from the opinion distribution of TikTok's user base, or much stronger in terms of content availability and reach. In an effort to understand what content a user perceives as political, Greenfield et al. (Greenfield et al., 2025) recruit 358 U.S. participants to annotate over 16,000 videos. Through this study, they find that users label issue-based content (race, gender, policing, economy) as political more often than in terms of institutions (elections, government). User perceptions are important to understand the actual impact of the content that is recommended, an aspect that is often not included in studies of social platforms.

There has also been prior work studying our chosen topics on social media. Regarding abortion access, Sharma et al. study discourse on X (previously Twitter). They collect 700k tweets from abortion-related hashtags and classified (in favor, neutral, or against abortion) with 76% accuracy using machine learning classifiers (Sharma et al., 2017). Garimella et al. (Garimella et al., 2017) use a similar data acquisition technique to study more generally how four controversial topics (one of which was abortion) were debated on Twitter in 2011–2016 from an interaction network perspective; Cinelli et al. (Cinelli et al., 2021) expand this work to a comparison of Gab, Facebook, Reddit, and Twitter from an information spreading and polarization perspective. Data acquisition through hashtag-based searches and snowball sampling has been used in other studies as well (Guinaudeau et al., 2022; Vasconcellos et al., 2023; Lu et al., 2015) (not necessarily covering the same topics as our study); Vasconcellos et al. identified hashtags that often co-occur.

As with abortion, online discourse around the Israel/Hamas conflict is highly polarized. Using TikTok's Research API, Edelson et al. collect >280k U.S. posts across 12 three-day windows (Edelson, 2024). They report preliminary findings that Pro Palestine TikToks largely outnumber Pro Israel content and show time-varying amplification periods for each topic. Similarly, a co-hashtag network built from 509 TikToks collected via the hashtag *#israel/#palestine* shows more tightly connected Pro Palestine communities, consistent with a network of activism and a sharp October 2023 spike that decayed quickly afterwards (Hasin, 2025). Outside of TikTok, coordinated actors on X were found to have a measurable impact on organic users' engagement, increasing negative interactions in discussions surrounding Gaza (Dey et al., 2024).

While prior research has demonstrated that TikTok is a source for political discourse and that recommendation algorithms can shape amplification patterns, these studies do not address our research questions. Our work extends this literature by conducting a platform-level analysis of how imbalanced the content and view distributions are for stances of polarizing topics, and complements this with a survey asking respondents about their experience on TikTok. We are especially interested in differing experiences on the platform between majority and minority opinion holders. Taken together, our analysis lends understanding to how TikTok disseminates polarizing content and how that dissemination affects the experiences of its users.

### Methodology

This paper is based on data that was sourced from the TikTok Research API (TikTok, 2023). At a high level, it provides us with a data set of public content originating from the U.S. We constructed two topic-specific data sets for abortion and the Israel/Hamas conflict, each collected during episodes of intense national debate on social media. For abortion, we queried for videos published between April 1st and August 25th, 2022; this window captures both the leak and subsequent overturn of *Roe v. Wade*. For the Israel/Hamas conflict, we captured the period from September 30th, 2023 to January 31st, 2024, which includes the days leading up to and following the Hamas attack on October 7th, 2023.

Due to issues with the TikTok API, we are unable to collect videos published during specific “blackout” timeframes for which the API does not respond to our queries and instead returns an error. We cannot retroactively capture data for the dates October 16th, November 9th-12th, November 23-27th, and November 29-30th, 2023, all of which are within the timeframe for Israel/Hamas. These blackout dates affected queries for all hashtags, including high-volume hashtags such as *#fyp*, and were not within our control. The abortion data is unaffected.

We selected the topics of abortion and Israel/Hamas based on several criteria. To be viable for this study, a topic needed to have sustained and polarized public discourse on TikTok, distinct opinion groups (i.e., clearly distinguishable, binary stances), and the feasibility of building lists of relevant hashtags to retrieve a large portion of the content on the topic with reasonable precision. We considered and ultimately rejected other polarizing issues. The issue of gun control was rejected as there were more nuanced opinions that did not categorize easily

into a binary classification at scale. The Amber Heard v. Johnny Depp trial was not feasible to collect because important time frames would be heavily impacted by the blackout dates in TikTok’s Research API. The Blake Lively v. Justin Baldoni lawsuits lacked the volume of videos necessary to perform meaningful analysis.

We built our initial data sets by querying the TikTok API for hashtags related to each of the topics. In selecting these hashtags, we were challenged to balance precision and recall. For the abortion set, as a starting point, we obtained from unpublished preliminary work by the Algorithmic Transparency Institute<sup>1</sup> a set of hashtags of interest. We revised this set by removing hashtags that led to a high number of videos unrelated to the topic (i.e., false positives). For example, *#christiantiktok* was too general because it is commonly used in video descriptions unrelated to abortion. We proceeded similarly for the Israel/Hamas set, for which we received an initial list of hashtags directly from journalist Mike Scott, who has covered online discourse surrounding the topic (Scott, 2023). After validation of individual hashtags for overall topic relevance, we expanded the remaining sets through a snowball method: we randomly sampled at least 5 videos from each hashtag and added any other related hashtags found in the descriptions of those videos to our list, also to be reviewed. At each step, we only retained hashtags that were clearly relevant to the issues and excluded overly broad or ambiguous hashtags. For example, although *#war* appeared frequently for Israel/Hamas, it was excluded because it captured a wide range of content that was completely unrelated to this specific conflict. The final list of 74 abortion-related and 117 Israel/Hamas hashtags is available in Tables 16 and 17 in the appendix.

We collected our data sets retroactively. We completed acquisition of the abortion data in January 2024 and of the Israel/Hamas data in June 2025. While academic literature on TikTok video lifetimes is scarce, industry analytics tools have reported that a TikTok video receives 72% of its views within 24 hours, and 95% of its views within 35 days (Panigrahi, 2025). Given that more than one year has elapsed between video upload and data collection, it is unlikely that a significant amount of additional engagement has occurred after collection, and we therefore can assume the view/engagement numbers to be “final.” However, our data sets exclude videos that were deleted before our data collection. Overall, the abortion data set

---

<sup>1</sup><https://ati.io/>

compromises of 560k videos from 249k accounts with a total of 3.8b views. For the Israel/Hamas set, we collected 641k videos from 161k accounts for a total of 6.4b views.

While inspecting the abortion data set, we find 8.8% of videos have fewer views than likes, making up 8.5% of total likes and 0.5% of all views. We find similar anomalies in the Israel/Hamas data set, albeit at a much lower rate of 0.6% of videos, accounting for 0.0012% of views and 0.02% of likes. Given the relatively higher weight of likes vs. views in these cases, our best hypothesis is that TikTok might be correcting the amount of views from bots or other malicious accounts. These videos have the same proportions to the global set when analyzed across stance detection, number of accounts, videos with zero views and special types of videos (stitches, duets and reply-to videos, which all allow users to build on content from others). Since the problematic instances are balanced across categories of interest, we keep them in our data set and address the issue by avoiding calculating view totals for small subsets of the data or per-video ratios with views.

### *Stance Detection Methodology*

We assign stance labels to videos so that we can later analyze availability and dissemination as a function of the video's stance. In the context of this study, we rely on sampling and manual stance annotations by the authors. We considered several alternatives, including a hashtag-based methodology and large language models (LLMs) for stance detection, and ultimately settled on human annotation; Section 4.3 justifies our choice.

For each of the two topic data sets, we randomly drew and annotated 3 samples, each meant to measure a different element of the data: a sample to measure distribution of videos, another for distribution of views, and a longitudinal sample meant to capture a temporal analysis of views across our collection period. Video counts are relevant to our analysis because they reflect how much content is available (i.e., how many videos creators upload). To characterize dissemination (i.e., the size of the audience reached by a video), we utilize views as the metric. We favor views over an engagement-based metric, such as the number of likes, because views more directly reflect the audience size as the outcome of the recommendation algorithm as opposed to how engaged the audience is.

To enable analysis of the distribution of stance across available content (video counts), we manually labeled two samples of 500 videos per topic. These samples are uniformly random, ensuring each video was equally likely to be selected. Using these samples, we are able to accurately estimate proportions of stance in the full data sets with an approximated  $\pm 4.4$  percent margin of error at a 95% confidence (Wilson's confidence intervals).

To represent the number of views associated with each stance across topics, we follow a similar procedure, manually labeling two samples of 500 videos per topic. In these samples, we selected TikTok videos using probability proportional to size with replacement (PPSWR), where views were our measure of size. That is, videos with higher views (more reach) were more likely to be included in the sample. Due to this construction of the samples, a stance's proportion of videos in the sample reflects the proportion of views in the full data set. Thus, we have unbiased estimators of a stance's total share of views, and the two samples allow us to estimate view distribution by stance with approximately  $\pm 4.4\%$  error at 95% confidence.

To enable longitudinal analysis while balancing time granularity and annotation effort, we split our entire collection periods into consecutive 7-day intervals (to mitigate day-of-week effects), totaling 21 intervals for abortion data and 16 intervals for Israel/Hamas. Due to the blackout periods in the Israel/Hamas data, we slightly shifted the beginnings of some intervals so that no interval would contain any days within TikTok's API blackout. For each interval, we drew a view-weighted sample of 50 TikTok videos using PPSWR similar to the aggregate sample. Videos were assigned to intervals based on their publication date, and all lifetime views of each video were allocated to the same interval, no matter when the views actually occurred. For our longitudinal analysis, we labeled 1,050 videos for abortion and 800 for Israel/Hamas.

To label the videos in our samples, annotators watched each video and reviewed its metadata such as description and hashtags. Each video was then categorized into six mutually exclusive categories: the two opposing stances plus neutral, unrelated, private or unavailable, as defined in Table 1. Our more detailed annotation guidelines are given in Tables 18 and 19 in the appendix. To assign a category to a given TikTok, annotators considered a combination of available metadata and video content; a video showing general topic-related footage without additional commentary would be classified under the respective stance if the creator's opinion was conveyed through choice of hashtags or description. For example, footage of the Israel/Hamas

conflict with no audio or text commentary and mixed hashtags that used a Palestinian flag emoji visible on screen (without also including an Israeli flag emoji) would be considered Pro Palestine. If we encountered non-English commentary or description, we utilized translation services and relied on stanced visuals (e.g., emojis, flags, icons).

**Table 1:** Annotation categories used for manual stance labeling

Category	Definition
Pro Choice	Supporting the right to legal access to induced abortion services
Pro Life	Opposing induced abortion, supporting its prohibition/restriction
	or
Pro Palestine	Sympathize with the Palestinian people
Pro Israel	Sympathize with the Israeli people
Neutral	Relevant to the topic, but does not express a clear stance
Unrelated	Not relevant to the topic
Private	TikTok displays an error that the account is private
Unavailable	TikTok displays an error that the video has been removed or is otherwise inaccessible

To ensure labeling consistency, two co-authors annotated an initial subset of 100 videos per topic. Inter-annotator agreement was near perfect with a Cohen’s Kappa of 0.949 for abortion and 0.976 for Israel/Hamas. Given high agreement, a single annotator labeled the remaining videos. Overall, we manually labeled a total of 3,850 TikTok videos; we describe all samples in Table 6.

### *Data Set Coverage*

To ensure our findings are representative for the ecosystem of U.S. content on TikTok, we aim to characterize the precision and recall of our hashtag-based data collection methodology for the two topics. We can estimate precision based on our manually annotated random samples. When we combine abortion videos manually labeled as private or unavailable with those not related to abortion, we obtain a precision of 95.2% ( $\pm 1.87$  points at 95% confidence), indicating high

relevancy to abortion. We found for our Israel/Hamas data set, similarly, 85.5% of manually annotated viewable videos were on topic ( $\pm 3.5$ ). Thus we conclude that the vast majority of videos captured in our data sets are related to our topics. (We note that because our analysis is based on manually annotated random samples, the precision metric is not critical.)

Estimating recall is more difficult as it would require annotating a large random sample of *all* TikTok content. Instead of numerically estimating recall, we manually validate the primary threat to validity of our findings, specifically whether our data collection methodology failed to acquire minority stance content, which would make any imbalance appear more pronounced than it is in reality. As we will show, our data sets are heavily skewed to Pro Choice and Pro Palestine content, thus we investigate whether Pro Life or Pro Israel content was systematically omitted. We do so by manually searching minority stanced content on TikTok and validating whether it is included in our data set.

### *Abortion*

First, we search for well-known U.S. Pro Life accounts mentioned in the news (Brown, 2021; Hoonhout, 2020; Shugerman, 2021; Taylor, 2023; Wu and Byler, 2022). We find that most accounts mentioned are either already included in our data set or did not publish content during our measurement period, or have been turned private or deleted. Among the anomalies, we find that two out of six videos from a prominent Pro Life advocacy group's account do not contain any related terms in their description, and a third video is not returned by TikTok's API. The remaining three videos are included in our data set.<sup>2</sup>

Furthermore, another popular Pro Life account with 20M all-time views has only one out of 205 videos posted during our measurement period returned by the TikTok API. While we are unsure why the API does not return these videos, it may be related to the account having been previously removed and blocked from advertising (as reported by the organization on their website). However, the total 468k likes of the account's relevant videos make up only 4.0% of

---

<sup>2</sup>All videos and metadata used in this work were publicly available on TikTok. However, TikTok's Research API Terms of Service do not allow us to mention account names or information that could be traced back to the user.

total Pro Life likes in the detected set. In other words, content omitted from well-known Pro Life accounts does not account for a significant share of the Pro Life content that is included in our data set, and does not explain our findings of very little Pro Life content on TikTok.

We also explore whether “algspeak” could explain the scarcity of Pro Life content. Some users utilize nonstandard spellings of relevant terms to evade perceived content moderation (e.g., *#aborshun*, *#ab0rt10n*, *#4b0rtio0n*). We find only 772 videos in total querying the API with these terms, most of which were already discoverable by being accompanied with a more common hashtag such as *#prolife* or *#roevwade*. Therefore, we estimate the impact of algspeak to be negligible.

### *Israel/Hamas*

A limitation of our data set is that we queried the API for content published with a U.S. country code only; any international creators focused on the Israel/Hamas conflict are excluded from our analysis. For example, official government accounts from nations involved in the conflict do not appear in our data set because they would not be classified as a U.S. account.

We manually searched for prominent accounts; among these, most appeared in our data set. However, since there was no forward-facing ability to determine the country code of a user, it was difficult for us to reliably conclude if accounts were excluded due to lack of coverage or because of the country code. For example, we found during our search a popular Pro Palestine creator with more than 230k followers. Despite posting videos in the collection time frame with hashtags contained in our set, none of this content appeared in our data set. Through querying the API, we could confirm that no content (even unrelated to the topic) from this creator was returned, thus we assume they are located outside of the U.S.

Regarding algspeak, we encountered only one instance in the hashtag *#idftherro3ist*, which was included and classified as a Pro Palestine hashtag. We found no other evidence of widespread use of algspeak.

While our validation efforts do not reveal any significant gaps in our data collection methodology within our scope of U.S.-based accounts, the geographical limitation is substantial. Our findings should be interpreted as representing content specifically within the U.S.-based TikTok ecosystem rather than the global conversation about this international conflict.

### *Alternative Stance Detection Methodologies*

We considered alternative methodologies for assigning stance labels to TikTok videos other than manual, human annotation. Using only the hashtags that creators added to their videos to automatically assign stance resulted in poor coverage and limited granularity. We found that large portions of our data could not be given a stance through a hashtag methodology, as creators often used only one or few hashtags that may or may not represent their stance.

We also tested LLMs (GPT-4o, Claude Sonnet 4.0, Gemini Flash 2.5) on both accuracy and inter-rater agreement with a human. LLMs assigned labels to 100 videos each for abortion and Israel/Hamas given hashtags and video descriptions, along with voice-to-text transcriptions when available; these results are in the appendix, in Tables 21 and 20. Our two samples of 100 videos were randomly selected from the set of all human-labeled TikToks per topic. The prompts we used for annotations are provided in Appendix 10 and were written by translating our human annotation codebooks into structured instructions, following guidance for LLM-based text annotation tasks (Gilardi et al., 2023). The models we tested generally were unwilling to infer stance and often defaulted to neutral stances, possibly because of the controversial nature of our topics. Claude Sonnet 4.0 had the highest accuracy of 84.0% for the abortion topic, while GPT-4o performed best for Israel/Hamas videos with 83.0% accuracy. Two human annotators had a Cohen's Kappa of  $\kappa > 0.9$  for each topic; in contrast, models had low to moderate agreement, with the abortion topic performing notably poorly with an average  $\kappa$  of 0.23. Israel/Hamas' average  $\kappa$  was 0.64, which qualifies as moderate to substantial agreement but significantly underperformed when compared to human agreement. While training a custom model was theoretically possible, doing so would have required manually labeling a training set that would have been larger than the samples we ultimately use to approximate the entire population with reasonable confidence. In addition, training an LLM would have introduced an additional source of potential error. Since the analysis needed to answer our research questions is based on the aggregate population and does not require individual labels for all collected

videos, we determined that manual annotation of random samples was the most accurate and efficient approach.

### *TikTok User Surveys*

In addition to data collected from TikTok’s Research API, we conducted a survey to better understand how users perceive polarizing content on TikTok. This survey will allow us to better understand how the available content is served to users of varying stances. We will refer to this survey as our User Experience survey.

We recruited 450 U.S.-based TikTok users through Prolific and administered the survey through Qualtrics in December 2024. The average completion time was 3 minutes and 42 seconds. Participants were compensated at a rate equivalent to approximately \$27/hr. Prior to participating in the study, respondents gave informed consent in accordance with protocols approved by our institution’s IRB. The main questions of interest were how much content related to the two topics respondents recalled seeing on TikTok, and whether it agreed or disagreed with their personal opinion. The question phrasings and possible survey responses can be found in the appendix.

To put respondents’ experience with polarizing content into perspective, we asked participants about their stances on abortion and the Israel/Hamas conflict using the definitions shown in Table 2. Question wording and response options were taken from Pew Research Center surveys contemporaneous with our data collection periods (Gottfried, 2024). We note that for the Israel/Hamas topic, the Pew survey gave respondents the options to express sympathy with both or neither side of the conflict. There were no such options regarding respondents’ opinions on abortion. Respondents were grouped into ordinal ideological categories. For abortion, we use *Strictly Pro Choice*, *Leaning Pro Choice*, *Leaning Pro Life*, and *Strictly Pro Life*. For Israel/Hamas, *Strictly Pro Palestine*, *Leaning Pro Palestine*, *Neutral*, *Leaning Pro Israel*, *Strictly Pro Israel*, and *Neither*.

We use our own User Experience survey responses for self-contained hypothesis testing regarding user experience with polarizing content on TikTok, but not for comparison with our video data sets. Since our survey was fielded over a year from when our collected TikToks were

**Table 2:** Respondent Opinion Classification

<b>Topic</b>	<b>Survey Response</b>	<b>Assigned Stance/Category</b>
<b>Abortion</b>	Legal in all cases	Strictly Pro Choice
	Legal in most cases	Leaning Pro Choice
	Illegal in most cases	Leaning Pro Life
	Illegal in all cases	Strictly Pro Life
<b>Israel/Hamas</b>	Entirely with the Israeli people	Strictly Pro Israel
	Mostly with Israeli people, somewhat with Palestinian people	Leaning Pro Israel
	Equally with both	Neutral / Equally Aligned
	Mostly with Palestinian people, somewhat with Israeli people	Leaning Pro Palestine
	Entirely with the Palestinian people	Strictly Pro Palestine
	Neither group	Neither Side

published, user opinions gathered in this survey are not an accurate measure of the TikTok population’s opinion distribution during our video collection periods.

To gain an understanding of contemporaneous U.S. TikTok user opinions for comparison with our video data sets, we draw on additional public opinion surveys conducted by Pew Research Center. For abortion, we use a study fielded in June-July 2022 following the *Roe v. Wade* overturn (Doherty et al., 2022) (Wave 110), and for Israel/Hamas a survey fielded in February 2024 (Silver et al., 2024) (Wave 143). We will refer to these surveys as the Pew Abortion Opinion survey and Pew Israel/Hamas Opinion survey, respectively. These surveys, downloaded from the publicly available American Trends Panel, allow us to estimate age-specific opinion distributions on each topic using the provided survey weights calibrated to be nationally representative.

These two surveys provide nationally representative opinion distributions measured during the same periods in which the TikTok videos in our data sets were published. Our objective, however, is to approximate the opinion distribution of TikTok users. Because no contemporaneous surveys jointly measure TikTok usage and topic opinions, we use age-based poststratification on the Pew Opinion surveys to estimate opinions of TikTok users. Specifically, we reweight the age-specific opinion estimates to match the age composition of TikTok’s user population, which we infer from Pew Research Center’s Social Media Fact Sheet (Gottfried, 2024) (see Table 3). We weight the User Experience, Pew Abortion Opinion, and Pew Israel/Hamas Opinion surveys with this method. Finally, we aggregate the post-stratified opinion estimates to produce overall benchmarks for U.S. TikTok users. Table 4 contains the resulting opinion distribution.

For the purpose of our analysis, we combine the “leaning” individuals into their shared stances and exclude *Neutral* or *Neither* respondents. This allows us to use stance proportions that correspond more directly to our manual video annotations where we only distinguish between opposing positions with no measure of extremity.<sup>3</sup> Using these aggregated categories

---

<sup>3</sup>While our video annotations do include a “neutral” label, these videos typically correspond to informative news-like reporting that does not reflect an opinion towards the debate. They are qualitatively different from a neutral stance indicated by survey respondents and therefore do not factor into our analysis of polarization or stance representation.

for the Abortion and Israel/Hamas Opinion surveys, and considering only respondents who reported stances, we estimate that 65.5% of U.S. TikTok users identified as Pro Choice and 34.5% as Pro Life. For Israel/Hamas, 46.3% identified as Pro Palestine and 53.7% as Pro Israel.

**Table 3:** TikTok user age distribution taken from Pew’s Social Media Fact Sheet. These margins are applied to User Experience, Pew Abortion Opinion and Pew Israel/Hamas Opinion using poststratification by age.

Category	Group	Share
Age	18–29	37.5%
	30–49	38.5%
	50–64	17.0%
	65+	7.0%

We considered incorporating additional contemporaneous surveys for our topics including those fielded by Gallup, PRRI, and NPR (Montanaro, 2023; Public Religion Research Institute, 2023; Saad, 2022); Gallup and PRRI report U.S. opinions on abortion, while NPR reports the U.S. public’s perception on the Israel/Hamas conflict. Ultimately, we did not incorporate these surveys due to a mismatch in timing and age bucketing. PRRI collected interviews from March to December 2021 and as such capture pre, during, and post-Roe opinions. Gallup’s survey was administered in June, prior to the overturn of *Roe v. Wade*. Furthermore, our post stratification procedure requires age buckets similar to the ones used in Pew’s Social Media Fact Sheet, but none of the additional polls aligned; for example, NPR buckets ages to over and under 45. At a high level, we do not find substantial differences in reported user opinions when comparing the Pew Abortion and Israel/Hamas Opinion surveys to these additional surveys. Therefore, we rely exclusively on Pew Research Center surveys to understand the opinion distribution of U.S. TikTok users. Pew provides consistent strata across surveys for post-stratification, methodological uniformity in fielding and weighting surveys, and timing that aligns more closely with our TikTok data collection periods.

**Table 4:** Estimate of U.S. TikTok users’ opinions towards abortion and Israel/Hamas (based on Pew survey benchmarks calibrated to the age distribution of U.S. TikTok users). Derived from Pew ATP Wave 110 (Abortion) fielded in June-July, 2022, and Wave 143 (Israel/Hamas) fielded in February 2024 using Pew-provided survey weights post-stratified based on age. Aggregated proportions combine “leaning” categories into shared stances. Israel/Hamas aggregates exclude *Neutral* and *Neither*.

Topic	Stance/Category	Percentage
Abortion	Legal in all cases	32.5%
	Legal in most cases	33.0%
	Illegal in most cases	26.3%
	Illegal in all cases	8.2%
	Pro Choice (aggregated)	65.5%
	Pro Life (aggregated)	34.5%
Israel/Hamas	Entirely Pro Palestine	9.4%
	Leaning Pro Palestine	17.2%
	Neutral / Equally Aligned	31.9%
	Leaning Pro Israel	20.3%
	Entirely Pro Israel	10.2%
	Neither Side	10.5%
	Pro Palestine (aggregated, excl. neutral/neither)	46.3%
	Pro Israel (aggregated, excl. neutral/neither)	53.7%

### *Data Set Overview*

In the following, we summarize the data sets, samples, and surveys. We start from two data sets of TikTok videos collected by querying the Research API for topic-specific hashtags, summarized in Table 5. Our analysis is exclusively based on manually annotated samples derived from this data (to estimate the stance distribution of videos and views for each topic, both in aggregate and over time); these samples are described in Table 6. To complement our quantitative analysis of TikTok videos, we rely on multiple surveys to understand user experience and opinions, as shown in Table 7. For comparison to our video samples, we use Pew Opinion surveys to estimate user opinion distributions. Separately, our User Experience survey allows us to analyze exposure and perceived agreement with served content from a user perspective.

**Table 5:** Data sets of videos created by querying TikTok’s Research API.

<b>Name</b>	<b>Size</b>	<b>Publication Dates</b>	<b>Description</b>
Abortion API	560k	Apr 1–Aug 25, 2022	Contains all queried API videos for the abortion topic
Israel/Hamas API	641k	Sep 30, 2023–Jan 31, 2024	All queried API videos for the Israel/Hamas topic

### *Ethics*

To conduct our research, we submitted a proposal to TikTok and were granted access to their Research API. The API returns only public content, and does not specifically disclose any personal information. While public TikTok content may occasionally contain personal information (e.g., names mentioned in videos or descriptions), we do not analyze any category of personal information and do not release it. TikTok users gave general consent to the use of their data for research purposes when accepting TikTok’s privacy policy. We did not obtain separate consent from TikTok users whose videos we collected for this research. Our Institutional Review Board (IRB) approved our study. We acknowledge a risk of negative impacts related to the interpretation of our findings. Our empirical findings characterize the viewpoints on our chosen topics as observed on TikTok, and do not contribute to the debates themselves. However, given the polarized nature of the debates, people might seek to use these findings to support their own

**Table 6:** Random samples used to estimate video counts and views by stance within the full data sets of Table 5.

<b>Name</b>	<b>Size</b>	<b>Metric Estimated</b>
Abortion Video	500	Video counts by stance
Israel/Hamas Video	500	Video counts by stance
Abortion View	500	Video views by stance
Israel/Hamas View	500	Video views by stance
Abortion Weekly	1,050	Longitudinal video views by stance (21 weekly samples of 50 videos each)
Israel/Hamas Weekly	800	Longitudinal video views by stance (16 weekly samples of 50 videos each)

**Table 7:** Surveys used to estimate U.S. TikTok user opinion distributions and study user experience with exposure to polarizing content on TikTok.

<b>Survey Name</b>	<b>Fielded</b>	<b>Usage</b>
Pew Social Media Fact Sheet	Sep–Oct 2024	TikTok user age distribution reference
Pew Abortion Opinion	June–July 2022	Opinion distribution (abortion)
Pew Israel/Hamas Opinion	Feb 2024	Opinion distribution (Israel/Hamas)
User Experience	Dec 2024	User stance vs. content exposure

standpoint, even though this study is not in any way designed to do so. Furthermore, people might hold TikTok responsible for the trends we observe on the platform, even though our study does not address this issue and cannot explain the reason for our findings. We argue that the benefits of raising public awareness for our findings outweigh these risks. Most notably, our results serve to caution TikTok users that the frequency at which they observe differing opinions about abortion or Israel/Hamas on the platform is not representative for TikTok's user base or the population at large. Furthermore, there is a risk that this finding might be more systemic than what we are able to demonstrate in the present work. If our findings can lead to future work to better understand this issue and its causes, the public may benefit through increased awareness and potentially improved platforms.

### Analysis

We begin our analysis with the following question: How much are the opposing stances on polarizing topics represented and consumed on TikTok?

Table 8 presents the proportions of videos belonging to the two stances and the additional categories (neutral, unrelated, and private/unavailable) as manually annotated in the two 500-video samples covering the entire data set duration. On both topics, there is a clear stance represented in most of the available content. When we consider only videos that received a stance classification as the baseline, we find 96.5% to represent Pro Choice, and 88.2% to represent Pro Palestine. In other words, the share of videos representing Pro Life and Pro Israel are only 3.5% and 11.8%, respectively. To put this into perspective, we need to understand the corresponding distribution of user opinions. If almost all users held a specific opinion, then it would not necessarily be surprising to see very little alternative content. Yet, based on the Pew Opinion surveys, we estimate that only 65.5% of U.S. TikTok users consider themselves Pro Choice and 46.3% Pro Palestine, a relatively smaller share compared to the observed content distribution. Consequently, we observe a stark imbalance between content and users; the share of videos on TikTok that represent the Pro Choice and Pro Palestine opinion is much

larger than what the user opinion distribution would suggest.<sup>4</sup> In the following, we explore this imbalance in more detail.

**Table 8:** Distribution of Stances in Manually Annotated Video Samples ( $n = 500$  per topic) using Wilson’s 95% CIs. Aggregated rows are calculated relative only to stanced videos (Pro Choice and Pro Life, Pro Palestine and Pro Israel). Pro Choice and Pro Palestine stances are very dominant in terms of video counts.

Topic	Category	Percentage	95% CI
Abortion	Pro Choice	89.4%	(86.4, 91.8)
	Pro Life	3.2%	(2.0, 5.1)
	Neutral	2.4%	(1.4, 4.1)
	Unrelated	2.8%	(1.7, 4.6)
	Private or Unavailable	2.2%	(1.2, 3.9)
	Pro Choice (stanced only)	96.5%	(94.5, 97.9)
Pro Life (stanced only)	3.5%	(2.1, 5.5)	
Israel/Hamas	Pro Palestine	62.8%	(58.5, 66.9)
	Pro Israel	8.4%	(6.3, 11.1)
	Neutral	10.2%	(7.8, 13.2)
	Unrelated	13.8%	(11.1, 17.1)
	Private or Unavailable	4.8%	(3.2, 7.0)
	Pro Palestine (stanced only)	88.2%	(84.4, 91.2)
Pro Israel (stanced only)	11.8%	(8.8, 15.6)	

<sup>4</sup>In the case of the Israel/Hamas topic, we find that the *less* popular user opinion is the dominant stance on TikTok in terms of available videos. For proper interpretation of this observation, we note that our TikTok user opinion estimate has both sides relatively close at 46.3% vs. 53.7%. Due to our post stratification, both user opinions might be within each other’s error margin. No matter which of the two sides is slightly more popular among users, the content stance distribution is much more skewed at 88.2% Pro Palestine.

When we also include the remaining annotation categories in the baseline, Pro Choice still represents the overwhelming majority of videos (89.4%) within the abortion topic, while only 3.2% are Pro Life. The remainder of content is either neutral (2.4%), unrelated (3.0%), or private/unavailable (2.2%).

For the Israel/Hamas topic, we find that Pro Palestine content comprises most of the videos in the data set (62.8%) compared to just 8.4% Pro Israel. Notably, the minority user opinion in this case represents the majority of content on TikTok. We find that a larger proportion of the content is neutral (10.2%) and unrelated (13.8%) compared to the abortion topic. We hypothesize this is due to the Israel/Hamas conflict having more and sustained news updates as ongoing developments were reported daily compared to Roe v. Wade coverage that primarily centered around just two events. As a result, we observe more content focused on reporting updates or sharing news footage without expressing a stance.

Secondly, we find a higher degree of hashtag “misuse” in the Israel/Hamas data set. Key hashtags for our collection, such as *#gaza* and *#israel*, were used by accounts for view farming or content relevant to the territory itself, but not the conflict (e.g., nature, tourism, history). Despite higher rates of unrelated content, excluding these hashtags from the collection would have greatly impacted our ability to gather relevant videos from TikTok.

Although the total amount of content available provides valuable insight, it is not indicative of what is being served to users; the existence of videos on the platform does not guarantee they are watched. To assess the impact and reach of stanced content, we now examine the number of *views*, which is a more direct measure of how many times a video was delivered to users. Table 9 presents the proportion of views per stance as manually annotated in the two view-weighted random samples of 500 videos per topic, which cover the entire duration of our data sets. We find that, similar to raw video counts, Pro Choice and Pro Palestine content receives far more views than Pro Life and Pro Israel videos.

When we observe the distribution of views across stances, the disparities we found in content availability persist in reach. For the abortion topic, Pro Life videos comprise only 3.0% of all views, whereas Pro Choice accounts for 77.0%. Similarly, Pro Israel videos capture 12.0% of total views, while Pro Palestine makes up over half of all views (50.6%). Notably, for the

**Table 9:** Distribution of Views in Manually Annotated View Samples ( $n = 500$  per topic) using Wilson’s 95% CIs. Aggregated rows are calculated relative only to stanced views (Pro Choice and Pro Life, Pro Palestine and Pro Israel). Pro Choice and Pro Palestine stances are also dominant in terms of video views.

Topic	Category	Percentage	95% CI
Abortion	Pro Choice	77.0%	(73.1, 80.5)
	Pro Life	3.0%	(1.8, 4.9)
	Neutral	3.8%	(2.4, 5.9)
	Unrelated	4.8%	(3.2, 7.0)
	Private or Unavailable	11.4%	(8.9, 14.5)
	Pro Choice (stanced only)	96.3%	(93.9, 97.7)
	Pro Life (stanced only)	3.7%	(2.3, 6.1)
Israel/Hamas	Pro Palestine	50.6%	(46.2, 55.0)
	Pro Israel	12.0%	(9.4, 15.1)
	Neutral	22.2%	(18.8, 26.1)
	Unrelated	12.0%	(9.4, 15.1)
	Private or Unavailable	3.2%	(2.0, 5.1)
	Pro Palestine (stanced only)	80.9%	(76.5, 85.1)
	Pro Israel (stanced only)	19.1%	(15.2, 23.9)

Israel/Hamas topic, Neutral videos account for 22.2% of all views while only making up 10.2% of all content; we hypothesize this is due to the amount of news coverage and live updates relayed during our collection period of ongoing conflict in comparison to the abortion topic, which only had two notable newsworthy events.

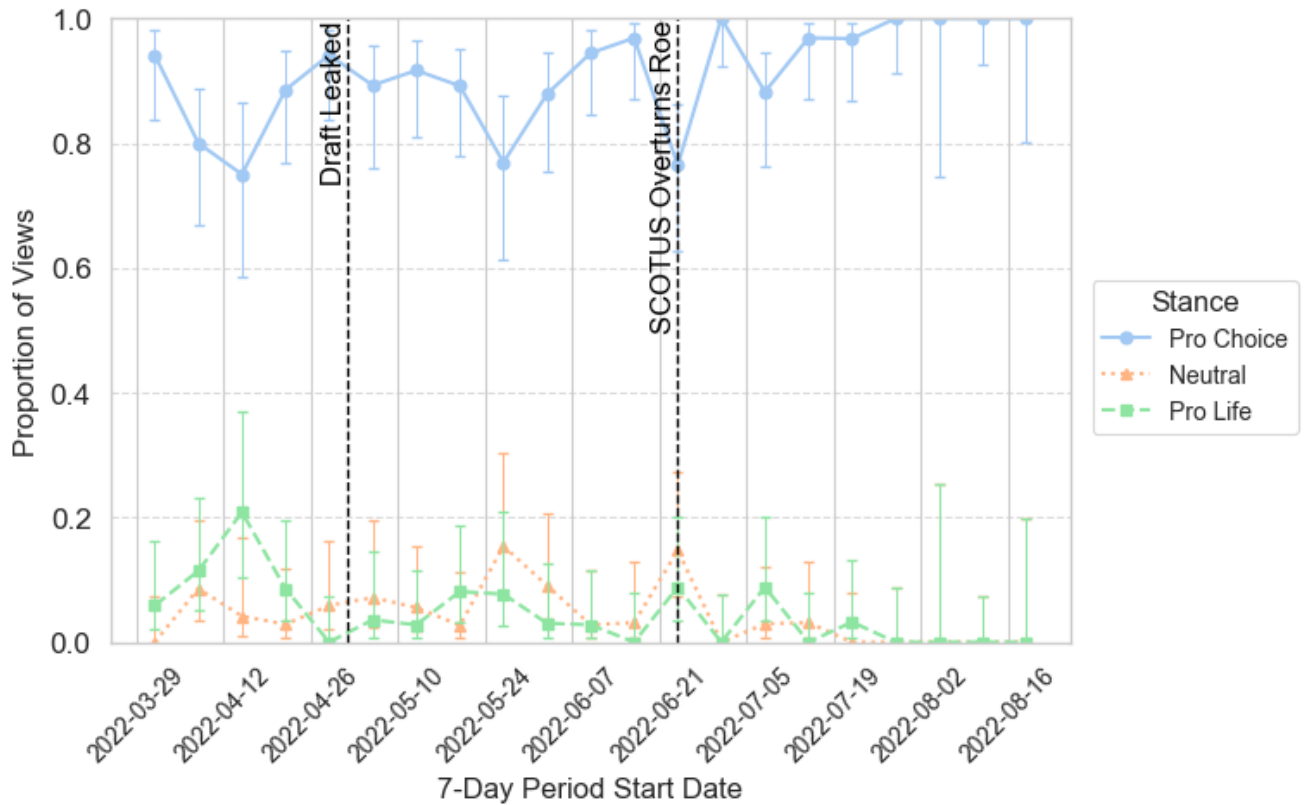
Video counts reflect the supply of content, while views better capture the reach a video was able to achieve on TikTok, and subsequently, the ability to impact users. Pro Life and Pro Israel videos are underrepresented in both production and viewership compared to the estimated opinion distribution of TikTok users.

### *Temporal Analysis*

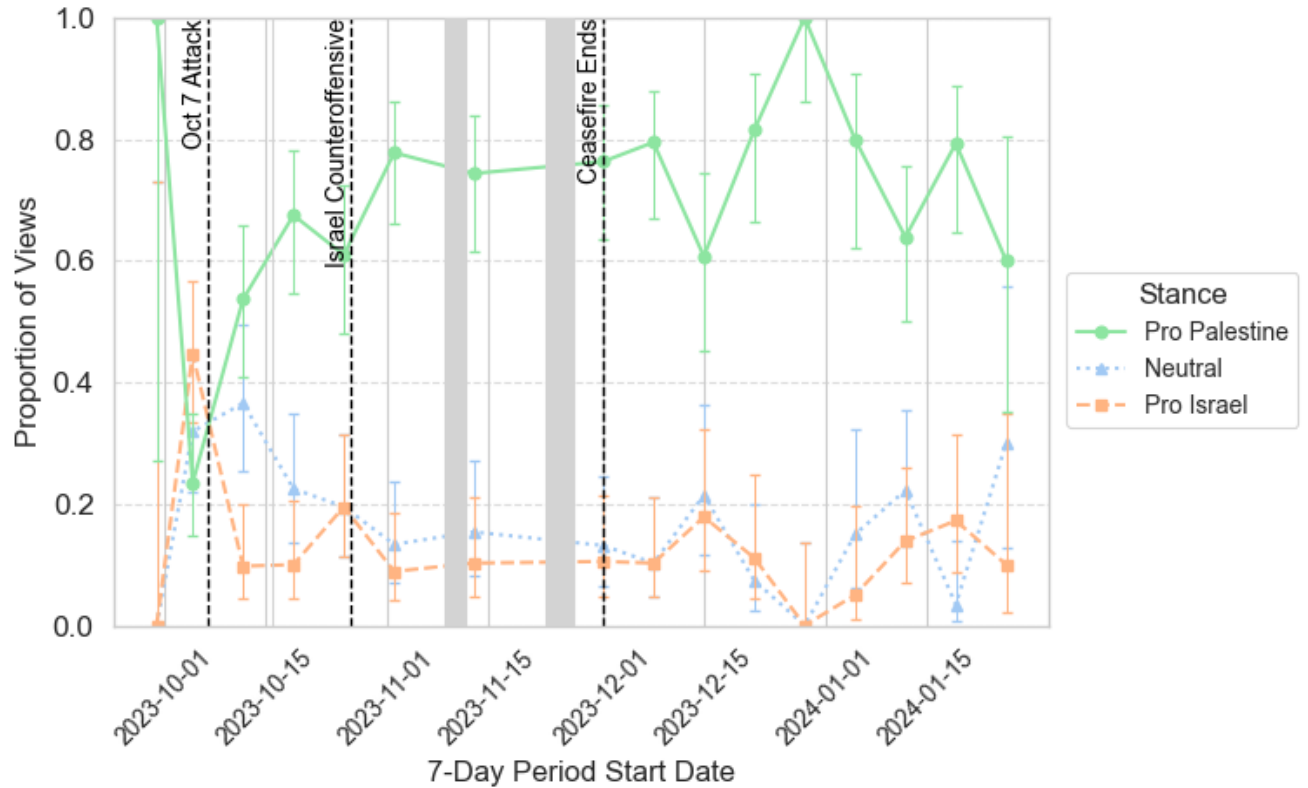
Given that we observe such a drastic difference between the two stances in aggregate, we now analyze this phenomenon longitudinally. More specifically, we investigate whether the aggregate effect we have observed is consistent over time. The following longitudinal analysis is based on manually annotated, view-weighted random samples of consecutive 7-day intervals across our observed time periods for each topic. We chose to base our analysis on views rather than video counts because a video's view count indicates that a user has been exposed to the content, and is a better measure of the actual reach achieved by a TikTok. In the following, we compute view shares relative to all views of stanced or neutral content, i.e., excluding videos labeled as unrelated or private/unavailable.

Figure 1 shows the measured proportion of views in the abortion topic from videos labeled with the Pro Choice, Pro Life, or Neutral stance each week of our collection period. Pro Choice content is clearly dominant over the entire collection period, consistently accounting for the vast majority of views. For the last 6 weeks, Pro Choice makes up over 97% of weekly views, effectively a complete domination of user reach. Notably, during the week following the Supreme Court overturn, Pro Choice videos account for 99.0% of all views. Neutral content retains a small percentage of the total views across the entire period, with its highest share of 15.4% in the last week of May (25 days post leak), and another peak (14.7% of views) during the week of the overturn decision.

Figure 2 shows the corresponding plot for the Israel/Hamas topic. We observe a similar dynamic of the Pro Palestine opinion accounting for the vast majority of views throughout



**Figure 1.** Abortion: Proportion of views over time (relative to Pro Choice, Pro Life, and Neutral content only) in 7-day intervals. Error bars indicate Wilson 90% confidence intervals. Pro Choice accounts for over 80% of views for most of the collection period, while Pro Life has its highest share of 21.8% 2 weeks before the draft was leaked.



**Figure 2.** Israel/Hamas: Proportion of views over time (relative to Pro Palestine, Pro Israel, and Neutral content only) in 7-day intervals. Grey areas indicate periods for which the API returned errors (intervals shifted accordingly). Error bars indicate Wilson 90% confidence intervals. After October 7, 2023, Pro Palestine quickly surges and maintains the majority proportion over the collection period.

most of our observation window. The first two 7-day periods of this data set, however, exhibit large variation among the three labels and are to be interpreted with caution. During the first 7-day period, which was entirely before Hamas' attack on October 7, 2023, there was little relevant content in absolute terms; we collected only 2,019 videos for this interval (in contrast to 54,701 TikToks in the week following October 7th). This is simply the result of the topic not being discussed widely on TikTok, and yields an unstable estimate with a very wide confidence interval, which means that no real conclusions should be drawn from this period. In the second 7-day period, the week of October 7th, Pro Israel (44.7%), Pro Palestine (23.4%), and Neutral (31.9%) content all accounted for a relatively even share of views (with overlapping confidence intervals). This initial balance (which may partially be due to the alignment of the 7-day period) was short lived as in the week following, Pro Palestine content accounted for 53.7% of views and Pro Israel only 9.8%. Within another week, Pro Palestine surpassed 67.5% of all weekly views and remained above 60% for the rest of the collection period, while Pro Israel never recovered beyond 19.5%.

Overall, the longitudinal results suggest that stance distributions were consistent over the time of our data sets. In both cases, a single opinion (Pro Choice, Pro Palestine) remained dominant over time once it emerged and achieved substantially more reach than any other stance. That is, our aggregate finding does not appear to be the result of a temporal anomaly.

### *User Experience*

Given this consistent skew in the view share between stances, which was much stronger than the opinion distribution of TikTok users, we ask what the implications are for users' experience on the platform: Does TikTok show users of both stances similar levels of on-topic content, and does the content users see align with their own stance? The samples of API data that we have studied so far provide an ecosystem-wide aggregate, but do not reveal the experiences of individual users. (For example, the API data cannot replicate a user's For You Page on TikTok.) Therefore, we turn to our User Experience survey to answer these questions. The survey is disconnected from the API data sets in important ways (e.g., temporally), thus it should be seen as an independent and not directly comparable vantage point.

To explore whether users' views on a topic predict exposure to on-topic content, we ran a bidirectional logistic regression using two predictors: (1) stance extremity (the "distance" from Neutral) and (2) stance direction (e.g., Pro Palestine or Pro Israel). Our dependent variable was whether or not a user reported seeing any content on a topic. Full statistical results can be found in Tables 15 and 14 in the appendix. Both predictors were statistically significant for the Israel/Hamas topic; users who self-identified as holding more extreme views were more likely to see content about the conflict, and Pro Palestine users were 2.3 times more likely to report exposure than Pro Israel respondents ( $p < 0.008$ ; Cramer's  $V = 0.359$ ). For the abortion topic, we similarly found that users with more extreme opinions (e.g., users who reported Pro Choice vs. reporting Leaning Pro Choice) were more likely to report exposure. However, there was not a statistically significant difference between Pro Choice and Pro Life users for likelihood of exposure. These results are subject to a number of caveats. Respondents across all opinions reported lower exposure to abortion-related content than Israel/Hamas-related content, likely because our survey was administered after abortion had receded from public debate, while the Israel/Hamas conflict remained ongoing during our survey period. Furthermore, while we asked users if they had seen content related to our topics, we did not ask if they were interested in receiving content about abortion or Israel/Hamas. This is a confounding factor (independent from the user's opinion on the topic) that could impact the likelihood of a user being recommended content.

**Table 10:** Perceived Alignment of Abortion Content on TikTok by User Stance

User Stance	Agreeing	Neutral	Disagreeing	No/Little Content
Pro Choice	55.7%	9.6%	0.5%	34.2%
Leaning Pro Choice	31.2%	20.0%	5.7%	43.1%
Leaning Pro Life	10.1%	17.2%	21.6%	51.1%
Pro Life	19.9%	10.8%	29.9%	39.5%

We then turned to the question of content alignment, i.e., whether the user reported agreeing or disagreeing with most of the content that they saw. Tables 10 and 11 note the reported agreement per stance on the two topics. For abortion, a chi square test revealed a highly significant association between stance and reported agreement with content ( $\chi^2(9) = 97.5, p < 0.00001$ ; Cramer's  $V = 0.269$ ). Specifically, 55.7% of Pro Choice users and 31.2%

**Table 11:** Perceived Alignment of Israel/Hamas Content on TikTok by User Stance

User Stance	Agreeing	Neutral	Disagreeing	No/Little Content
Pro Palestine	58.0%	15.5%	14.5%	12.0%
Leaning Pro Palestine	39.0%	32.6%	16.1%	12.3%
Neutral	12.0%	33.7%	9.9%	44.4%
Leaning Pro Israel	18.8%	28.1%	23.3%	29.9%
Pro Israel	24.9%	12.5%	44.1%	18.5%
No Opinion	2.8%	35.4%	8.0%	53.8%

of Leaning Pro Choice users said most content they recalled seeing agreed with their opinion, while only 0.5% and 5.7%, respectively, said it disagreed. In contrast, only 19.9% of Pro Life users reported seeing mostly affirming content, while 29.9% saw mostly disagreeing content. We found the same relationship between stance and alignment for the Israel/Hamas topic ( $\chi^2(5) = 56.5, p < 0.00001$ ; Cramer's  $V = 0.359$ ). Pro Palestine users were the most likely to report seeing content that mostly agreed with their stance (58.0%), while Pro Israel users were nearly double as likely to report seeing mostly disagreeing content (44.1%) rather than mostly agreeing content (24.9%).

Unfortunately for our interpretation of these results, we do not know the stance distribution of video views on TikTok during the (unspecified) time period that survey respondents recalled in their answers. The survey was fielded 10-28 months after the time span represented by our manually annotated video samples. In the context of the Israel/Hamas conflict, Radway and Edelson found that the Pro Palestine stance still dominated both available content and accrued views on TikTok in the beginning of September 2025 (Radway and Edelson, 2025). For this topic, we can therefore assume that the “content majority” on TikTok continued to be Pro Palestine during the period covered by our survey. In this light, our User Experience survey suggests that the user experience on TikTok is neither ideologically neutral nor a perfect filter bubble. It appears that a stance that is dominant on the platform in terms of available videos or reach can also be the dominant stance shown to individuals who disagree with the stance, although this does not appear to hold true for all users. As a result, users whose opinions differ from the content-majority stance are more likely to be exposed to mostly opposing views,

while users aligned with the content majority tend to be mostly exposed to affirming content. Furthermore, we find that ideological extremity may increase exposure to the polarizing topics we studied, particularly when a user opinion aligns with the majority of TikTok content. Due to the limitations of our User Experience survey, these interpretations require additional study to be confirmed and refined.

### Discussion

Our study, in the context of two polarizing debates, abortion access and the Israel/Hamas conflict, examined how much reach each of the two sides achieved on TikTok. For both debates, we found a strong imbalance between the stance distribution of videos and views on TikTok on the one hand, and the estimated opinion distribution of TikTok's U.S. user base on the other hand. In the case of abortion access, videos and views on TikTok skewed much more towards Pro Choice than the opinions of TikTok's user base. In the case of the Israel/Hamas conflict, Pro Palestine became dominant in terms of content availability and reach, whereas we estimated user opinions to be close together (with Pro Palestine slightly in the minority). This effect was consistent across the duration of our data sets.

Theoretical concepts such as filter bubbles (Pariser, 2011), echo chambers (Cinelli et al., 2021), and the spiral of silence (Noelle-Neumann, 1974) have been introduced to characterize polarization from the perspective of individuals. In contrast, our work was primarily concerned with the platform-wide situation, which has not previously been studied. In both topics, we found one stance to be very underrepresented in terms of content availability and views. We do not know whether low availability or low views are the cause of this effect. One conceivable hypothesis that we cannot presently prove or disprove is that creators might be unwilling to invest into content that will reach only a small audience. This would be conceptually related to, but different from, a *spiral of silence*, which theorizes that individuals will withhold their opinions if they believe they are in the minority to avoid social isolation. Our User Experience survey can be interpreted as being consistent with a spiral of silence effect, but it was not designed to understand the effect of recommendations on user behavior. More research is necessary to understand whether an effect akin to a spiral of silence contributes to the skewed stance distribution observed on TikTok.

While methodologically more limited, our User Experience survey provides some insight into what effects individuals encounter on TikTok. Within both topics, most users of one opinion recalled seeing mostly agreeing content, whereas most users of the opposing opinion recalled seeing mostly disagreeing content. *Filter bubbles* are understood as users consuming ideologically aligned content (either through homogeneous social networks or hyper-personalized recommendations) that reinforces their current belief. While our User Experience survey leaves open whether a filter bubble effect applies to the users who recalled seeing mostly agreeing content,<sup>5</sup> it appears to rule out a filter bubble effect for the users of the opposing opinion, since they recalled seeing mostly disagreeing content. Therefore, filter bubbles do not seem to be the correct framework to generalize to all users' experiences on TikTok in the context of the two polarizing topics that we have studied. Additional research is required to understand in detail how the platform-level effects we have observed in our study translate to the experience of individual users.

In broad terms, these platform-level effects may reduce the ability and willingness of TikTok users to meaningfully understand and contribute to discourse. We measure substantially different distributions of user opinions and reach/availability of content, and as such, users' perception of agreement (or disagreement) among peers could be distorted. This is an important caveat for users to be aware of when they (dis)engage with debates on the platform.

It is important to note that our study has measured the *outcome* of TikTok's recommendation system as a whole, but does not reveal any specific cause for these imbalances. Rather, a recommendation system can be understood as a large feedback loop with many components (e.g., user behavior, creator behavior, platform design), all of which could potentially contribute to the effect we have observed in our measurement. We refrain from speculating which of these components might be more to "blame." Regardless of the cause, our findings underscore how a recommendation system can make a polarized debate even more polarized.

---

<sup>5</sup>Our methodology does not allow us to demonstrate the *presence* of filter bubbles. Since our survey asked participants only for the most prevalent stance in their feeds, it does not reveal to which degree other stances were or were not represented.

### *Limitations*

Our results are limited by the time period of analysis and the U.S.-focused data acquisition. The restriction on content uploaded by U.S. creators is partly due to restrictions of TikTok's Research API (for example, it does not return Canadian content), and partly due to practical considerations (for example, the necessity of cultural and language background within our research team). U.S. TikTok users can be recommended content from international creators, which is not included in our data sets. We anticipate this limitation to be less of an issue for the abortion topic, since it revolves around U.S. policy, but more so for the international Israel/Hamas conflict. (We note, though, that we observed the stronger skew towards one stance in the abortion topic.) Per the Research API guidelines, our work is also constrained to only currently available content from adult accounts. This means videos that have been turned private or are unavailable (either due to deletion or TikTok's content moderation) between their creation and our data acquisition are not included. Using hashtags to query videos allows us to collect relevant videos. However, we do not collect videos that may be relevant but do not contain hashtags. Prior work has also detected various issues with the TikTok Research API. For example, Pearson et al. found differences in video metadata between the front end and API, especially around engagement metrics (Pearson et al., 2025); Corso et al. encountered periods of "blackout" dates where the API did not return videos (Corso et al., 2024). As detailed in the methodology section, similar issues arose during our data collection, which we addressed as much as feasible. While our extensive validation did not unearth any additional inconsistencies, we cannot independently establish absolute certainty as to the accuracy of the data supplied by TikTok.

Our User Experience survey is instrumental in understanding the user experience on TikTok, but we acknowledge its limitations. Our survey was fielded after public debate had reached a crescendo for each of our topics, not concurrent with the collection periods from the API. Therefore, the results from the User Experience survey should not be directly compared to the API-based data analysis. Our survey results are based on respondents' recall of their feeds, not the actual feed itself, and thus subject to distortion. Furthermore, respondents' perceptions of content stance may differ from our definitions; for example, users may vary in what they consider "neutral" content. Lastly, we asked participants what they saw and if they

agreed with the content, but we did not ask if they were *interested* in seeing it, which could also potentially impact the reported visibility of topic-related content in their feeds.

With abortion access and the Israel/Hamas conflict, we have studied two topics that differ in various aspects (e.g., different time period, domestic vs. international issue). While these differences establish some degree of diversity, two topics are not sufficient to ascertain whether the skewed stance distribution we have observed is an exception or the norm in the TikTok ecosystem.

### **Conclusion**

Our research provides the first large-scale analysis of the skewed reach of polarizing content on TikTok. We used platform-level measurements to show that available and viewed videos were strongly imbalanced in favor of Pro Choice and Pro Palestine stances, whereas TikTok users' personal opinions were more balanced (and even slightly in favor of Pro Israel). By asking users directly, we found that participants identifying as Pro Choice or Pro Palestine more often encountered mostly relevant and agreeable content, while holders of the opposing opinions reported less exposure and more disagreement.

Taken together, our analysis lends understanding to how TikTok disseminates polarizing content and how that dissemination affects the experiences of its users. While our study does not attempt to identify causality, it does demonstrate the importance of measuring the outcome of how platforms disseminate content, especially in the context of polarizing debates. Future work is needed to investigate whether this finding is limited to the two topics we selected, how opinions are represented in debates with more than two major stances, what polarizing debates look like on other platforms, and more generally, which factors cause this imbalance.

### **Acknowledgments**

We thank TikTok for access to their Research API, without which this study would not have been possible. This study was funded in part by the National Science Foundation (NSF) under grant 2344939.

## References

- Associated Press (2023). Israel-Hamas war. AP News. <https://apnews.com/hub/israel-hamas-war>.
- Associated Press (2024). A look at the Gaza war protests that have emerged on US college campuses. AP News. <https://apnews.com/article/gaza-war-campus-protests-966eb531279f8e4381883fc5d79d5466>.
- Brown, L. (2021). Gen Z is spreading the Pro-Life message in a whole new way on TikTok. NCR. <https://www.ncregister.com/news/gen-z-is-spreading-the-pro-life-message-in-a-whole-new-way-on-tiktok>.
- Cinelli, M., Morales, G. D. F., Galeazzi, A., Quattrociocchi, W., and Starnini, M. (2021). The Echo Chamber effect on social media. In *Proceedings of the National Academy of Sciences*, volume 118, page e2023301118.
- Corso, F., Pierri, F., and Morales, G. D. F. (2024). What we can learn from TikTok through its Research API. In *Proceedings of the Web Science Conference*.
- Dey, P., Luceri, L., and Ferrara, E. (2024). Coordinated activity modulates the behavior and emotions of organic users: A case study on tweets about the Gaza conflict. <https://arxiv.org/abs/2402.05873>.
- Doherty, C., Kiley, J., Oliphant, B., Daniller, A., Hartig, H., Borelli, G., Dunn, A., Van Green, T., Gómez, V., and Salzer, R. (2022). Majority of public disapproves of Supreme Court's decision to overturn Roe v. Wade. Pew Research Center. <https://www.pewresearch.org/politics/2022/07/06/majority-of-public-disapproves-of-supreme-courts-decision-to-overturn-roe-v-wade/>.
- Edelson, L. (2024). Getting to know the TikTok Research API. Cybersecurity for Democracy. <https://cybersecurityfordemocracy.org/getting-to-know-the-tiktok-research-api>.
- Garimella, K., Morales, G. D. F., Gionis, A., and Mathioudakis, M. (2017). The effect of collective attention on controversial debates on social media. In *Proceedings of the Web Science Conference*, page 43–52.

- Gerstein, J. and Ward, A. (2022). Exclusive: Supreme Court has voted to overturn abortion rights, draft opinion shows. Politico. <https://www.politico.com/news/2022/05/02/supreme-court-abortion-draft-opinion-00029473>.
- Gilardi, F., Alizadeh, M., and Kubli, M. (2023). ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences of the United States of America*, 120(30):e2305016120.
- Gottfried, J. (2024). Americans' social media use. Pew Research Center. <https://www.pewresearch.org/internet/2024/01/31/americans-social-media-use/>.
- Greenfield, J., Wang, S., and Metaxa, D. (2025). Perceiving politics on TikTok: A user-centered approach to understanding political content on TikTok. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*.
- Guinaudeau, B., Munger, K., and Votta, F. (2022). Fifteen seconds of fame: Tiktok and the supply side of social video. *Computational Communication Research*, 4(2):463–485.
- Guttmacher Institute (2025). Interactive map: US abortion policies and access after Roe. Guttmacher Institute. <https://states.guttmacher.org/policies/>.
- Hasin, R. (2025). Unveiling voices: A co-hashtag analysis of TikTok discourse on the 2023 Israel-Palestine crisis. <https://arxiv.org/abs/2501.07182>.
- Hoonhout, T. (2020). TikTok reinstates Pro-Life org Live Action, says initial ban for 'multiple community guidelines violations' was 'result of a human error'. National Review. <https://www.nationalreview.com/news/tiktok-bans-pro-life-org-live-action-for-unspecified-multiple-community-guidelines-violations/>.
- Ibrahim, H., Jang, H. D., Aldahoul, N., Kaufman, A. R., Rahwan, T., and Zaki, Y. (2025). Tiktok's recommendations skewed towards Republican content during the 2024 U.S. presidential race. <https://arxiv.org/abs/2501.17831>.
- Lu, H., Caverlee, J., and Niu, W. (2015). Biaswatch: A lightweight system for discovering and tracking topic-sensitive opinion bias in social media. In *Proceedings of the International Conference on Information and Knowledge Management*, page 213–222.

- Montanaro, D. (2023). Americans are split over Israel's response in its war with Hamas. NPR. <https://www.npr.org/2023/11/15/1212913674/poll-israel-hamas-war-biden-democrats-republicans>.
- Newman, N. (2025). Digital news report 2025. Reuters Institute for the Study of Journalism. <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2025>.
- Noelle-Neumann, E. (1974). The spiral of silence a theory of public opinion. *Journal of Communication*, 24(2):43–51.
- Panigrahi, S. (2025). What engagement data reveals about content lifespan in TikTok vs Instagram. Measure Studio. <https://www.measure.studio/post/tiktok-vs-instagram>.
- Pariser, E. (2011). *The Filter Bubble: What the Internet is hiding from you*. Penguin Press.
- Pearson, G. D., Silver, N. A., Robinson, J. Y., Azadi, M., Schillo, B. A., and Kreslake, J. M. (2025). Beyond the margin of error: A systematic and replicable audit of the TikTok research API. *Information, Communication & Society*, 28(3):452–470.
- Public Religion Research Institute (2023). Abortion attitudes in a post-Roe world: Findings from the 50-state 2022 American values atlas. PRRI. <https://prri.org/research/abortion-attitudes-in-a-post-roe-world-findings-from-the-50-state-2022-american-values-atlas/>.
- Radway, E. and Edelson, L. (2025). Updating our Israel/Gaza analysis 2 years on. Cybersecurity for Democracy. <https://cybersecurityfordemocracy.org/updating-our-israelgaza-analysis-2-years-on>.
- Saad, L. (2022). 'Pro-Choice' identification rises to near record high in U.S. Gallup. <https://news.gallup.com/poll/393104/pro-choice-identification-rises-near-record-high.aspx>.
- Scott, M. (2023). Does social media favor Palestine over Israel? Politico Europe. <https://www.politico.eu/newsletter/digital-bridge/does-social-media-favor-palestine-over-israel/>.
- Serrano, J. C. M., Papakyriakopoulos, O., and Hegelich, S. (2020). Dancing to the partisan beat: A first analysis of political communication on TikTok. In *Proceedings of the Web Science Conference*, page 257–266.

- Sharma, E., Saha, K., Ernala, S. K., Ghoshal, S., and De Choudhury, M. (2017). Analyzing ideological discourse on social media: A case study of the abortion debate. In *Proceedings of the International Conference of The Computational Social Science Society of the Americas*, pages 1–8.
- Sharp, J. and Zanotti, J. (2024). Israel and Hamas conflict in brief: Overview, U.S. policy, and options for congress. Congressional Research Service. <https://www.congress.gov/crs-product/R47828>.
- Shugerman, E. (2021). Meet the teen who got the Pro-Choice movement on TikTok. The Daily Beast. <https://www.thedailybeast.com/meet-the-teen-who-schooled-the-pro-choice-movement-on-tiktok>.
- Silver, L. (2024). Younger Americans stand out in their views of the Israel-Hamas war. Pew Research Center. <https://www.pewresearch.org/short-reads/2024/04/02/younger-americans-stand-out-in-their-views-of-the-israel-hamas-war/>.
- Silver, L., Alper, B., Keeter, S., Lippert, J., and Mohamed, B. (2024). Majority in U.S. say Israel has valid reasons for fighting; fewer say the same about Hamas. Pew Research Center. <https://www.pewresearch.org/2024/03/21/majority-in-u-s-say-israel-has-valid-reasons-for-fighting-fewer-say-the-same-about-hamas/>.
- Taylor, L. (2023). Pro-abortion radicals want me to be ‘raped’ and ‘killed’ just because I’m a Pro-Life college student. Fox News. <https://www.foxnews.com/opinion/pro-abortion-radicals-want-raped-killed-just-because-pro-life-college-student>.
- TikTok (2023). Research API | TikTok for Developers. TikTok. <https://developers.tiktok.com/products/research-api/>.
- Titelman, L., Brundidge, J., et al. (2024). Understanding TikTok’s political role: A mixed-methods user study of news and political content consumption. *Frontiers in Social Psychology*, 2.
- Tjaden, J., Wolfgram, J., Philipp, A., Weißmann, S., Bobzien, L., Kohler, U., and Verwiebe, R. (2025). Does the TikTok feed lean right? Exposure to political party content among non-partisan users during regional and federal elections in Germany. [https://osf.io/preprints/socarxiv/7vdex\\_v1](https://osf.io/preprints/socarxiv/7vdex_v1).

United States Reports (1973). U.S. Reports: Roe v. Wade, 410 U.S. 113 (1973). Library of Congress. <https://www.loc.gov/item/usrep410113/>.

United States Reports (2022). Dobbs v. Jackson Women’s Health Organization 2022. Library of Congress. [https://www.supremecourt.gov/opinions/21pdf/19-1392\\_6j37.pdf](https://www.supremecourt.gov/opinions/21pdf/19-1392_6j37.pdf).

Vasconcellos, P. H. S., Lara, P. D. D. A., and Marques-Neto, H. T. (2023). Analyzing polarization and toxicity on political debate in Brazilian TikTok videos transcriptions. In *Proceedings of the Web Science Conference*, page 33–42.

Wu, Y. and Byler, D. (2022). Opinion | What we found when analyzing 1,000 viral TikToks on #Abortion. Washington Post. <https://www.washingtonpost.com/opinions/interactive/2022/tiktok-abortion-debate-gen-z/>.

## Appendix

### *Additional Tables*

Tables 12 and 13 combine the weekly samples from our longitudinal analysis as an additional validation of our aggregate findings in Table 9. Because of their construction, these pooled weekly samples do not directly replicate the aggregate samples, and we caution that the pooled weekly samples are not to be considered representative of TikTok’s long-term view distribution. For example, the total numbers of published videos, accrued views and their stance distributions differ each week, which explains why the distributions in the pooled samples deviate slightly from our aggregate samples. Since the distributions in the samples of both types are broadly consistent at a high level, we gain additional confidence about our aggregate Abortion View and Israel/Hamas View samples, which are representative.

**Table 12:** Aggregated views per stance distribution from pooled longitudinal PPSWR samples for abortion (21 weekly samples combined; includes only Pro Choice, Pro Life, and Neutral labels). We calculate this distribution to validate our findings in comparison to our Abortion View sample (Table 9) and find it to be consistent. However, this is not a single random sample and should not be interpreted as representative.

Category	View Percentage	95% CI
Pro Choice	91.9%	(88.9, 94.1)
Pro Life	3.6%	(2.2, 5.8)
Neutral	4.5%	(2.9, 7.0)

**Table 13:** Aggregated views per stance distribution from pooled longitudinal PPSWR samples for Israel/Hamas (16 weekly samples combined; includes only Pro Palestine, Pro Israel, and Neutral labels). We calculate this distribution to validate the findings from our Israel/Hamas View sample (Table 9) and find it to be consistent. However, this is not a single random sample and should not be interpreted as representative.

Category	View Percentage	95% CI
Pro Palestine	59.7%	(54.9, 64.2)
Pro Israel	14.2%	(11.1, 17.8)
Neutral	26.2%	(22.2, 30.6)

**Table 14:** Bidirectional logistic regression predicting self-reported exposure to Israel/Hamas content using stance extremity and direction.

	Coefficient	Odds Ratio	<i>p</i> -value
Intercept	0.376	1.46	0.038
Distance from Neutral	0.523	1.69	0.010
Pro Palestine (vs. Pro Israel)	0.851	2.34	0.008
Observations		385	
Pseudo $R^2$		0.079	

**Table 15:** Bidirectional logistic regression predicting self-reported exposure to abortion-related content using stance extremity and direction.

	Coefficient	Odds Ratio	<i>p</i> -value
Intercept	-0.011	0.99	0.947
Distance from Neutral	0.450	1.57	0.026
Pro Choice (vs. Pro Life)	0.211	1.24	0.303
Observations		449	
Pseudo $R^2$		0.012	

*Survey Questions*

1. In what country do you currently reside? <sup>6</sup>
  - USA
  - Other
  
2. Which of the following social media sites do you use on a regular basis (at least once a month)? Choose any that apply.<sup>7</sup>
  - Facebook
  - Youtube
  - Twitter
  - LinkedIn
  - Pinterest
  - Google Plus
  - Tumblr
  - Instagram
  - Reddit
  - VK (ВКонтакте)
  - Flickr
  - Vine.co
  - Meetup
  - ask.fm
  - Snapchat
  - TikTok
  - Medium
  - N/A, Rather not say, or None of the above

---

<sup>6</sup>If a user selects "Other", they do not continue the survey.

<sup>7</sup>If a user does not select "TikTok", they do not continue the survey.

3. Do you think abortion should be

- Legal in all cases
- Legal in most cases
- Illegal in most cases
- Illegal in all cases

4. On TikTok, did you see a lot of content about abortion?

- Yes and most of it disagreed with my views
- Yes and most of it neither agreed nor disagreed with my views
- Yes and most of it agreed with my views
- No, I saw none or very little content about abortion
- I do not remember

5. Would you say your sympathies lie

- Entirely with the Palestinian people
- Mostly with the Palestinian people but also somewhat with the Israeli people
- Mostly the Israeli people but also somewhat with the Palestinian people
- Entirely with the Israeli people
- Equally with both of them
- With neither of them

6. On TikTok, did you see a lot of content about the Israel-Hamas War?

- Yes and most of it disagreed with my views
- Yes and most of it neither agreed nor disagreed with my views
- Yes and most of it agreed with my views
- No, I saw none or very little content about Israel-Hamas War
- I do not remember

7. What is your age (in years)
8. How much time do you spend on TikTok on a daily basis (in hours)?
9. Do you describe yourself as a man, a woman, or in some other way?
  - Man
  - Woman
  - Some other way
10. Are you of Hispanic, Latino, or Spanish origin, such as Mexican, Puerto Rican or Cuban?
  - Yes
  - No
11. Which of the following describes your race?
  - White
  - Black or African-American
  - Asian or Asian-American
  - Native American/American Indian/Alaska Native
  - Native Hawaiian/Other Pacific Islanders
  - Some other race
12. In politics today, do you consider yourself a
  - Republican
  - Democrat
  - Independent
  - Something else

**If “Independent” or “Something else” is chosen:**
13. As of today do you lean more towards
  - Lean Republican

- Lean Democrat

**The following Attention Check Questions were distributed through the survey. Failure to answer correctly disqualified a user from the survey.**

14. The following question is straightforward, when asked about your preferred movie you must answer ‘Groundhog Day.’ This question serves as an attention check. Given the text above, what is your preferred movie?

- Pirates of the Caribbean
- Dune
- Shrek
- Groundhog Day
- The Godfather

15. The following question is simple, when asked for your food, you must select “pizza” and ignore all other alternatives. This serves as a check of your attentiveness. Based on the above, what should you enter?

- Ice-cream
- Pizza
- Cake
- Pie

**Table 16:** Hashtags used for initial discovery of abortion content.

---

Hashtags
<p> #abortionisahumanright #abortionisessentialhealthcare #abortionishealthcare #abortionisnormal  #bansoffourbodies #clinicdefender #forcedbirth #forcedbirthextremists #mybodymychoice  #plannedparenthood #pro-choice #proabortion #prochoice #prochoiceisprolife  #prochoiceisprowomen #prochoicewithheart #prolifeisajoke #prolifeisalie #prolifelogic  #reproductivejustice #reproductiverights #righttochoose #riseup4abortionrights  #roevswadeneedstostay #womensreproductiverights #abortionisevil #abortionisnotbirthcontrol  #abortionismurder #abortioniswrong #abortionkills #antiabortion #antichoice #endabortion  #endabortion2022 #endabortionchristian #endabortionforever #equippedforlife  #humanityoftheunborn #makeabortionillegal #marchforlife #overturnroe #postroegeneration  #pro-life #prolife #prolife4life #prolifeapologetics #prolifegen #prolifegenerationb  #prolifeisprowomen #prolifetiktok #prolifeyouth #righttolife #roevwadehasgottogo  #stopabortion #theyshouldhaveachoice #unbornlivesmatter #youarebabykillers #4b0rti0n  #ab0rt1on #aborshun #abortion #abortionban #abortiondebate #abortionlaw  #abortionrights #prochoicearguments #prochoicers #prolifeprochoice #prolifevsprochoice  #roe #roevswade #roevwade #roevwadeovertun #whatisabortion </p>

---

**Table 17:** Hashtags used for initial discovery of Israel/Hamas content.

---

Israel/Hamas Hashtags
<p>#israel #Israeli #Jewish #freepalestine #free_palestine #freepalistine🇵🇸 #zionist #gaza  #hamas #westandwithpalestine #gazaunderattack #palestineprotest #PrayforPalestine  #StopGenocide #SavePalestine #IDF #Zionism #westandwithisrael #standwithisrael  #standwithpalestine #istandwithisrael #istandwithpalestine #prayforisrael #chidrenofgaza  #parentsofgaza #israelhostagerelapse #jewishpride #proudjew #prayingforisrael #israelpalestine  #israelunderattack #saveisrael #helpisrael #christiansunitedwithisrael #usisrael #israelarmy  #israelzionist #idf3ist #idfsoldier #jewishlivesmatter #proisrael #hamsisis  #freegazafromhamas #netanyahu #marchforisrael #palestinianstatehood #palestine truth  #palestinereality #rivertotheseapalestinewillbefree #weareallpalestinians #vivapalestina  #endtheoccupationofpalestine #boycottisrael #boycottisraelproducts #stoparmingisreal  #palestinianlivesmatter #gazagenocide #israeliblockade #stopgazagenocide  #stopthegenocideingaza #savejournalistingaza #savechildrenpalestine #savegazachildren  #childrenofgaza #gazastarving #stopzionism #stopzionistterrorism #stopzionistoccupation  #independentpalestine #weargreenforpalestine #savepalestine #wewantfreepalestine  #gazapalestine #gazatoday #palestinetiktok #palestinechild #palestinewillbefree  #israeliwarcrimes #palestinianliberation #israelwarongaza #palestineforever #saverafahpalestine  #justiceforpalestine #savegazapalestine #savenorthgaza #dontstoptalkingaboutpalestine  #stopgenocideingaza #worldstandwithpalestine #gazaholocaust #gazahospital #endzionism  #endzionistapartheid #peaceforpalestine #gazachildren #supportpalestine #endtheoccupation  #shutitdown4palestine #endisraeliapartheid #endisraeliwarcrimes #jerusalemisthecapitalofpalestine  #palestinesolidarity #fromtherivertothesea #unitedwithisrael #israelgaza #supportisrael  #israelunderfire #marchforisrael #israelhamaswar #hamasisis #hamasterrorattack  #free_falestine #palestine #terrorismohamas #helpgaza #freepalestinefromhamas #savepalestine🍉</p>

---

**Table 18:** Guidelines for the manual annotation of abortion content on TikTok.

<b>Label</b>	<b>Definition</b>
Pro Choice	Advocating for the right to have legal access to induced abortion services. Includes criticism of the Roe vs Wade overturn. The use of Pro Choice hashtags may indicate stance. When in doubt, other videos from the same account were queried.
Pro Life	Opposing induced abortion and support for its legal prohibition or restriction. Includes support of the Roe vs Wade overturn. The use of Pro Life hashtags may indicate stance. When in doubt, other videos from the same account were queried.
Neutral	Abortion related video that does not indicate the authors position or does not argue for one specific side. Includes abortion related news and comedy content.
Not related	The content is deemed to not be related to abortion, and its hashtag usage is considered to not indicate stance.
Private	Error message “This account is private” or similar.
Unavailable	Error message “Video currently unavailable” or similar.

**Table 19:** Guidelines for the manual annotation of Israel/Hamas content on TikTok.

Label	Definition
Pro Palestine	Expresses support or sympathy for Palestine, Palestinians, or their actions. Criticizes Israel or expresses sympathy for opposing perspectives. The use of Pro Palestine hashtags may indicate stance. When in doubt, other videos from the same account were queried.
Pro Israel	Expresses support or sympathy for Israel or its actions. Criticizes Hamas, Palestine, Palestinians, or their actions. The use of Pro Israel hashtags may indicate stance. When in doubt, other videos from the same account were queried.
Neutral	Discusses or clearly implies relevance to the Israel/Hamas conflict but does not support or pass clear judgment on either Israel, Hamas, Palestine, or the Palestinian People. Largely includes news sources that post updates on the conflict.
Not related	The content is deemed to not be related to the Israel/Hamas conflict, and its hashtag usage is not considered to indicate stance.
Private	Error message “This account is private” or similar.
Unavailable	Error message “Video currently unavailable” or similar.

**Table 20:** Cohen’s Kappa agreement between human labels and LLM predictions.

Topic	GPT-4o	Claude Sonnet 4.0	Gemini Flash 2.5
Abortion	0.25	0.41	0.29
Israel/Hamas	0.68	0.62	0.63

**Table 21:** Classification accuracy of large language models relative to human-annotated ground truth labels.

Topic	GPT-4o	Claude Sonnet 4.0	Gemini Flash 2.5
Abortion	71%	84%	73%
Israel/Hamas	83%	79%	79%

### Alternative Methodology LLM Annotation Prompts

#### *Abortion Prompt*

You are a classifier tasked with annotating a data point with a label. You must use the codebook below to assign a label related to abortion.

Inputs:

- video\_description: caption text written by creator
- hashtags: list of hashtags used by creator
- voice\_to\_text: transcript of audio in the video, if available

Use only these inputs to make a decision.

Codebook:

- Pro Choice: advocating for the right to have legal access to induced abortion services; includes criticism of the Roe v. Wade overturn; use of Pro Choice hashtags may indicate stance
- Pro Life: opposing induced abortion and supporting its legal prohibition or restriction; includes support of Roe v. Wade overturn; use of Pro Life hashtags may indicate stance
- Neutral: abortion-related video that does not indicate author's position or argue for one specific side; includes abortion-related news and comedy content
- Unrelated: the content is deemed not to be related to abortion, and it's hashtag usage is not considered to indicate stance

Rules:

- Only output EXACTLY one of the four labels from the codebook: "Pro Choice", "Pro Life", "Neutral", or "Unrelated"
- Base your decision solely on the provided inputs, and do not use outside knowledge.
- Do not explain your answer.

### *Israel/Hamas Prompt*

You are a classifier tasked with annotating a data point with a label. You must use the codebook below to assign a label related to the Israel/Hamas conflict.

Inputs:

- video\_description: caption text written by creator
  - hashtags: list of hashtags used by creator
  - voice\_to\_text: transcript of audio in the video, if available
- Use only these inputs to make a decision.

Codebook:

- Pro Palestine: expresses support or sympathy for Palestine, Palestinians, or their actions; criticizes Israel or expresses sympathy for opposing perspectives; use of Pro Palestine hashtags may indicate stance
- Pro Israel: expresses support or sympathy for Israel or its actions; criticizes Hamas, Palestine, Palestinians, or their actions; use of Pro Israel hashtags may indicate stance
- Neutral: discusses or clearly implies relevance to the Israel/Hamas conflict but does not support or pass clear judgment on either Israel, Hamas, Palestine, or the Palestinian People; largely includes news sources that post updates on the conflict
- Unrelated: the content is deemed not to be related to Israel/Hamas, and its hashtag usage is not considered to indicate stance.

Rules:

- Only output **EXACTLY** one of the four labels from the codebook: "Pro Palestine", "Pro Israel", "Neutral", or "Unrelated"
- Base your decision solely on the provided inputs, and do not use outside knowledge.
- Do not explain your answer.