

ISD

Powering solutions
to extremism, hate
and disinformation

European **MEDIA AND
INFORMATION** Fund
Managed by
Calouste Gulbenkian Foundation

Narratives of Hate

Post-7 October Antisemitism and Anti-Muslim Hate on Social Media

Hannah Rose and Paula-Charlotte Matlach

About this paper

In the first three months following Hamas' 7 October attack and the subsequent Israel-Gaza war, Jewish and Muslim communities experienced a spike in hate both online and offline.

This research uses innovative research technologies to track the narratives which drove antisemitism and anti-Muslim hate in comments on YouTube videos about the conflict over a three month period.

This research relied on the invaluable support of CASM Technology, including Shaun Ring, Jon Jones and Oli Holmes. Thanks also to ISD colleagues Jacob Davey, Francesca Arcostanzo, Zahed Amanullah and Sid Venkataramakrishnan for their insights, editing and technical support.

The sole responsibility for any content supported by the European Media and Information Fund lies with the author(s) and it may not necessarily reflect the positions of the EMIF and the Fund Partners, the Calouste Gulbenkian Foundation and the European University Institute.

Amman | Berlin | London | Paris | Washington DC

Copyright © Institute for Strategic Dialogue (2024). Institute for Strategic Dialogue (ISD) is a company limited by guarantee, registered office address PO Box 75769, London, SW1P 9ER. ISD is registered in England with company registration number 06581421 and registered charity number 1141069. All Rights Reserved.

www.isdglobal.org

Contents

Introduction	4
Key Findings	5
1. Methodology	6
2. Narratives of Hate	7
2.1 Antisemitism	7
2.1.1 Conspiracy theories	8
2.1.2 Control-based conspiracies	9
2.1.3 Classic antisemitism	10
2.1.4 Conflict-related antisemitism	10
2.1.5 Far-right antisemitism	11
2.2 Anti-Muslim Hate	12
2.2.1 Islam and violence	13
2.2.2 Integration/invasion narratives	13
2.2.3 Cultural incompatibility	14
2.2.4 Slurs	15
2.2.5 Sexual abuse	15
2.2.6 Pro-violence	16
2.2.7 Conflict-related anti-Muslim hate	16
3. Conclusions	17

Introduction

As news of the 7 October Hamas attack reached social media, online hate began to surge before official accounts could provide clear details. This report identifies and analyses both antisemitic and anti-Muslim narratives about the conflict. It uses automated hate speech detection software to track trends over time in YouTube comments.

In the immediate aftermath of the attack, ISD worked with partners at CASM Technology to build bespoke systems to measure these trends. Comparing the three days before and after 7 October, [analysts identified a more than 50-fold rise in the absolute volume of antisemitic comments on YouTube videos about the Israel-Palestine conflict](#); the overall proportion of antisemitic comments tripled. Using a similar classification system, ISD identified an over [40-fold increase in anti-Muslim comments comparing the four days before and after 7 October](#).

Having evidenced increases in the volume of hate speech, we leveraged automated classifiers and topic modelling methods to identify the nature of these waves of hate, to better inform understanding around the ways in which communities are targeted at times of crisis. This analysis identifies the themes and sub-themes which comprised and drove the surge of antisemitism and anti-Muslim hate between October and December 2023.

This report provides a data overview of these themes of hate, and supplements this with a final holistic trend analysis. It aims to inform countermeasures, ranging from government and law enforcement strategies to civil society interventions.

Key Findings

Antisemitism

- **The most common antisemitic narratives relate to conspiracy theories** (32 percent) about a perceived Jewish or Zionist collective acting to subvert the rights and interests of global populations. These include claims that 9/11 was a Zionist plot, Jews are not the 'true' descendants of the 12 tribes, and that 7 October was a 'false flag' operation.
- **Conspiracy theories which allege Jewish or Zionist control of a specific institution** comprised 29 percent of antisemitic comments. Most related posts focus on the US, with additional references to the EU, UK, the media, YouTube and Hollywood.
- While **classic antisemitism** is often considered to be less relevant to contemporary antisemitism, 28 percent of comments relate to this theme. Classic antisemitism references historical medieval antisemitic narratives, often rooted in Christian theology, and includes scriptural or Christian antisemitic narratives, such as labelling Jews the 'Synagogue of Satan' or promoting the deicide myth.
- Debate around definitions of antisemitism often focus around 'grey area' content which relates to where the boundaries of criticism of the state of Israel become antisemitic. However, our analysis found that this type of material – such as Nazi comparisons or the antisemitic denial of Jewish self-determination - only made up 7% of antisemitic comments, with the **vast majority of antisemitic content focusing on non-conflict related themes**.
- Comments which used **overtly extremist language** include assertions that Hitler was right, white supremacist tropes, or that Jews believe in their superiority through use of the term 'goyim' make up only 3% of comments identified.

Anti-Muslim hate

- The majority of anti-Muslim hate speech (55 percent) promote the **narrative that Islam or Muslims are inherently violent, dangerous and a threat to others**.
- The second most common anti-Muslim theme (19 percent) **suggests that Islam or Muslims are an invading force who are unable to integrate into different societies**. Posts containing this theme

reference a wide range of countries including the UK, Canada, Australia, the US, Russia, the Netherlands, Indonesia and France, highlighting the transnational nature of Islamophobic world views.

- **Posts suggesting that Islam is culturally incompatible with perceived Western values** comprise 13 percent of comments. Muslims are frequently accused of being anti-liberal, anti-LGBTQ+, untrustworthy, and seeking to replace Christian sites with Mosques.
- 7 percent of anti-Muslim posts use **slurs** to denigrate or dehumanise Muslim communities.
- Posts referring to the **anti-Muslim narrative that Islam or Muslims are promoters or perpetrators of sexual abuse** make up 3 percent of anti-Muslim comments.
- 1 percent of posts containing anti-Muslim hate speech explicitly **promote violence against Muslim people**.

In the first three months of the 2023 Israel-Gaza war, among both antisemitic and anti-Muslim comments, **established, well-trodden hateful narratives were far more common than new or conflict-specific comments**, suggesting that **the conflict has been an amplifier for entrenched and established hate**, rather than a catalyst for new types of hate.

Many of these posts may meet the threshold for illegal hate speech while others may be legal but banned by YouTube's terms of service. This research shows **continued gaps between the articulation of codes of practice and their enforcement**. Despite diverse legal or platform status, different severities of hate are intrinsically interconnected, and all serve to mainstream and target vulnerable communities.

Findings on the narratives of antisemitic and anti-Muslim hate should be used to **inform and prevention and education measures**. Counter-speech, alternative narratives, or inoculation programming should be data-driven, articulating positive messaging rooted in understandings of the drivers and narratives of hate in the contemporary online context, and thereby targeting the most widespread hateful narratives.

Methodology

To analyse changes in thematic discussions over time, ISD and CASM initially used the YouTube API to collect videos published after 1 October by accounts previously found to have published videos which featured hateful comments. This included 1,493,364 videos from 1,112 accounts that [previously published videos which attracted anti-Muslim comments](#), as well as 732,616 videos from 593 accounts that [published videos which attracted antisemitic comments](#). Keyword filtering was used to extract all videos mentioning Judaism, Islam, Jews or Muslims, as well as the Israel-Hamas war in the title or video description.

Step No.	Description	Anti-Muslim	Antisemitic
1	Initial seed list	1,112 accounts	593 accounts
2	Initial video collection	1,493,364 videos	732,616 videos
3	Thematic filter of videos	307,523 videos	156,723 videos
4	Total comments collected	68,622,297 comments	48,918,039 comments
5	Filtered comments	1,514,018 comments	610,513 comments
6	Classifier annotation	618,660 comments	194,695 comments

Table 1: Breakdown of the dataset per each step of the methodology.

Next, all comments from these videos were collected and screened for English-language content using an English language annotator and then filtered thematically, using the same list of keywords that was initially used to filter the videos. The resulting dataset was then processed using the Language Health Centre classifier - trained on a previous ISD-CASM project - and annotated the comments as either antisemitic or anti-Muslim hate speech.

In a further step, this data was clustered using semantic mapping, an unsupervised technique for organising large quantities of unstructured text data. When applied to social media content, this method creates clusters of messages containing semantically related material. Each document, such as a post, tweet, comment, or image caption, is assigned to a single cluster, in order to group pieces of text together based on their semantic similarity. The antisemitic and anti-Muslim datasets, each consisting of 100,000 posts, were individually processed via semantic mapping. In both cases, this was divided so

that a third consisted of content considered hateful by the relevant classifier, a third was content not considered hateful and a final third was content from imageboard 4chan. This was included for quality control purposes.

Subsequently, the resulting clusters of semantically similar messages were reviewed and labelled by analysts, grouping hateful clusters into themes and sub-themes. The narratives were then evaluated for their accuracy. Due to insufficient accuracy, comments allocated to the anti-Israel antisemitism theme underwent a second round of annotation by the antisemitism classifier. Comments allocated to the conflict-related antisemitism theme underwent additional manual annotation by analysts. Averaged across both ideologies and all themes, the classifier reached 83 percent precision and 80 percent recall scores for detecting hate speech.

Definitions

As frameworks for guiding labelling of hateful posts, analysts used the [IHRA Working Definition of Antisemitism](#) that:

“Antisemitism is a certain perception of Jews, which may be expressed as hatred toward Jews. Rhetorical and physical manifestations of antisemitism are directed toward Jewish or non-Jewish individuals and/or their property, toward Jewish community institutions and religious facilities.”

This definition is understood in context with support from its eleven examples of antisemitism.

To assist in the labelling of anti-Muslim hate, analysts used the ISD definition of hate as:

“Activity which seeks to dehumanise, demonise, harass, threaten, or incite violence against an individual or community based on religion, ethnicity, race, sex, gender identity, sexual orientation, disability, national origin or migrant status.”

This definition was interpreted with context from [ISD's Explainer on Islamophobia and anti-Muslim hatred](#).

Both definitions are consistent with previous analysis conducted by ISD on these topics, using the same notes and coding guidebooks. While there are many definitions of both forms of hate and their interpretation across thousands of comments can be diverse, our analysts have sought to build consistent, reliable and detailed definitional frameworks.

Narratives of Hate

This section provides a summary of antisemitic and anti-Muslim themes in relation to the Israel-Gaza war following the 7 October 2023 attack, and their relative prominence over time. Each theme is divided into its constituent sub-themes and their volumes over time.

The data is intended to provide a snapshot of the relative proportion of antisemitic and anti-Muslim narratives, rather than quantifiable measurements. Where methodological constraints, data access issues and the sheer volume of content prohibit views across entire platforms, this analysis omits Y axis volumes in favour of qualitative sub-theme comparisons.

Antisemitism

In the three months after the 7 October attack, five central antisemitic themes emerged in YouTube comments on videos about the conflict. The two most common were both related to conspiracy theories, including manifestations of the myths and tropes of Jewish greed, control and power.

The 'conspiracy theories' theme refers to the diverse manifestation of these narratives. 'Control-based conspiracies' identifies the most common of these myths of Jewish control and influence, while its constituent sub-themes refer to control or influence over a specific entity including states, inter-governmental bodies and media organisations. These findings emphasise the dominance of conspiracy theories in online antisemitism, the vast

majority of which are implicit or coded towards historical events. The central tenet of these conspiracies has remained constant since the early 20th century, the turn towards racialised antisemitism and the dissemination of the Protocols of Zion. While antisemitism is commonly understood as 'old wine in new bottles', the language and narratives identified in these themes are barely repackaged, but rather strongly echo these foundations.

The third most common theme was classic antisemitism, comprising of de-contextualised medieval Christian liturgy and related narratives, such as the narrative that Jews killed Jesus. The prominence of this theme confirms [previous manual thematic analysis of antisemitic comments by ISD analysts](#), which found classic tropes persisted in the online environment. Classic antisemitism in the contemporary landscape receives less attention from scholars of antisemitism and therefore features less prominently in counter-messaging campaigns, perhaps due to assumptions of its relative irrelevance. These findings importantly highlight the persistence of medieval and Christian antisemitic myths and tropes, which continue to proliferate on social media.

A less common theme was conflict or Israel-related antisemitism. While many YouTube comments criticised Israel or discussed the conflict in non-antisemitic ways, this theme included posts which used antisemitic framings or assertions in their criticism of Israel. Anti-Israel antisemitism receives significant attention from campaigners and analysts seeking to define or contest the boundary between hateful and non-hateful rhetoric.




Figure 1: Antisemitic themes over time

This remains an important, and often inconclusive, area of study. This data, however, indicates that while such language does exist, it is far outstripped in volume by covert and widely accepted antisemitic conspiracies and tropes. Online antisemitism is not just borderline, but in its majority, is overtly hateful.

Finally, the least common theme was far-right antisemitism. Whereas many different themes of antisemitism may be expressed by far-right individuals, this theme refers to language originating in far-right online spaces. Far-right networks commonly spread antisemitic hate; however, the relatively low prominence of this theme likely reflects the fact that the data was collected from mainstream rather than far-right channels.

This section will now provide an overview of the constituent sub-themes of each antisemitic theme.

Conspiracy theories

Conspiracy theories was the most common antisemitic theme (32 percent of antisemitic posts relevant to a theme), referring to diverse and generic conspiracies about a Jewish or Zionist collective and its perceived nefarious alternative intentions. By far the common sub-theme were posts referring to a general 'Zionist conspiracy', claiming that a 'Zionist' or 'Jewish' lobby controls world politics, specifically regarding foreign policy in the Middle East.

Criticism of Zionism is not inherently antisemitic and the term has diverse interpretations for different communities. However, it becomes so if 'Zionist' is

simply used as a synonym for 'Jew' in the construction of antisemitic conspiracy theories.

The second most common conspiracy theory alleged that Jews or Israelis are not real Jews, but are [instead descended from the Khazars](#), a nomadic Turkic people. Conspiracy theorists sometimes refer to Jews as the 'Khazarian mafia'. Other groups including [fringe Black Hebrew Nationalists](#) and [British-Israelists](#) use this conspiracy to lay claim to being the 'true' descendants of the Biblical tribes. Many posts claiming that Jews were 'fake' referenced a Book of Revelation verse that the "synagogue of Satan...say they are Jews, but are not".

A third common conspiracy theory, specific to 7 October, claimed that the attacks were a 'false flag' operation – an act where the perpetrator disguises their role in order to pin blame on another actor - and therefore deliberately orchestrated by the Israeli government to lay the groundwork for a Gaza invasion. The false flag narrative is common in radicalised conspiracy networks, using elements of mis- or disinformation to promote extremist anti-establishment narratives. In the case of 7 October, common false flag narratives have relied on [footage of an Apache helicopter wrongly alleged to have shot at Israeli civilians at the Nova festival](#).

False flag conspiracy theories come from a range of political stances; as such, post clusters were only included in this sub-theme where the majority were antisemitic. For example, one comment on the topic stated that "I'M sure THE JEWS KNEW, and like it happening because it gives them more powers to destroy Muslims".

A sub-theme about World War Three summarises




Figure 2: Conspiracy theory sub-themes over time

summarises posts which allege a plot to foment another global conflict orchestrated by Israel, Zionists or an imagined “Illuminati”. One such comment theorised that “The Zionists (Jews in name only) who run the world and desire a NWO [New World Order] one world government need WW3 to shuffle the pieces and depopulate the world so they can rebuild from the ashes a Marxist utopia in honor of their god Satan.” The 7 October attack slots neatly into existing worldviews about Jewish control of populations and power structures.

Two sub-themes were related to country-specific conspiracy theories. Iran-related conspiracies contained posts mixing broader conspiracy theories with references to either modern-day Iran or biblical confrontations between Jews and predecessors to Iran. Turkey-related conspiracy theories contained a range of conspiracies relating to Turkey or Erdogan.

Two sub-themes referred to the events of 9/11, often explicitly spreading the idea that Jews or Israelis orchestrated the 2001 plane attacks. The dancing Israelis sub-theme refers to an [alleged video of Israelis dancing with joy while watching the Twin Towers fall](#), suggesting they were involved with the attacks.

A final sub-theme referenced the early 20th century [Protocols of the Learned Elders of Zion](#), a totally fake account of minutes describing a meeting of Jewish leaders plotting to gain money, control and power.

Control-based conspiracies

The control-based conspiracies category (29 percent of relevant antisemitic posts) refers to theories which

specifically claim that a state or institution is controlled by Jews or Zionists for money or power. Most of these narratives focused on the US government and other political institutions, with many claiming that Biden’s military support for Israel in the immediate aftermath of 7 October stemmed from Jewish or Zionist puppet masters.

One comment labelled under this sub-theme stated that “The US economy, politic, and military [sic] are under the influence of Jewish elite hands. The US politicians are endorsed by the rich American Jews. See how it works!”. Additional sub-themes made similar accusations about the UK and the EU. A comment posted multiple times stated that at the time of the creation of the State of Israel, “Jews were there paying off the crooked politicians of UK, USA , FRANCE ect [sic]”.

A second cluster of sub-themes referred to the conspiracy theory that Jews control the media as a form of information warfare. These sub-themes most commonly focused on news media, with comments such as “The Jews bought off the Western elite and its media.” This was followed by YouTube, perhaps biased by the data collection method. A final cluster referenced the accusation that Jewish people control the film industry and Hollywood.




Figure 3: Control-based conspiracy sub-themes over time

Classic Antisemitism

In the immediate aftermath of the 7 October attack, classic antisemitism was the most common antisemitic narrative (28 percent of antisemitic posts relevant to a theme). As the Israeli offensive in Gaza evolved, classic antisemitism became the third most common form of antisemitism behind conspiracy theories.

Classic antisemitism refers to antisemitic tropes and narratives which derived in the medieval period, often linked to religious Christian texts or sub-cultures, as defined in the [9th example of the IHRA Working Definition of Antisemitism](#). Of the five sub-themes of classic antisemitism, three are rooted in Christian scripture or theology while two stem from medieval-Christian cultural tropes.

The most common sub-theme was the use of the ‘Synagogue of Satan’ as a descriptor of Jews or Israel, a quote lifted from a passage in the Book of Revelation and directed as a term of abuse. The Christian antisemitism sub-theme related to posts which contained other biblical verses calling for the punishment of Jews. One such comment claimed that “The Jewish society have been amassing great wealth for many years in order to take back their control, The bible says that fire shall rain from the skies, the olive branch shall bloom in the desert, the yellow man will rule the earth, these will be signs of the second coming”.

Scriptural antisemitism refers to posts which claimed that the actions of the Israel government were directly

influenced by the Talmud, implying or stating the primary source of Jewish religious law promotes evil. One such commenter stated that “it’s written into their faith. they’re special. they’re chosen. they rule over all. go read their talmud.”

The two medieval trope sub-themes included the decide myth that Jews killed Jesus and comments about Jewish noses, rooted in medieval caricatures.

Conflict-related antisemitism

Antisemitism specifically relating to the conflict appeared less frequently (7 percent of antisemitic posts relevant to a theme). This reflected ISD’s efforts to establish a high threshold for antisemitism given the many debates around the point at which anti-Israel narratives tip into antisemitism. This article only classified comments under this category if they could not reasonably be interpreted in any other way.

Automated hate speech detection models can have difficulty in detecting the nuanced line between anti-Israel and antisemitic content. This mirrors difficulties in reaching high manual inter-coder reliability, where multiple subject matter experts may disagree on the hateful nature of a post dependent on context. For this reason, a sample of comments in this theme were additionally labelled by subject-matter experts and used to train a conflict-related classifier, which was layered on top of the topic model to generate a 71 percent precision and 95 percent recall score.




Figure 4: Classic antisemitism sub-themes over time




Figure 5: Conflict-related antisemitic sub-themes over time

The most common form of conflict-related antisemitism was comparisons between Israeli actions and the Holocaust, including claims that Jews have become their once-oppressors. Whereas accusations of Israeli actions amounting to genocide were not labelled as antisemitic, specific use of Holocaust language was deemed both as a form of revisionism and the abuse of Jewish trauma for political means.

A second sub-theme denied Jewish self-determination in an antisemitic manner by denying any form of Jewish connection to Israel or accusing Zionists of ‘weaponising’ antisemitism to protect pro-Israel interests. Such posts justified the 7 October attack by asserting that the “Zionist state of isra-hell has been occupying for the past 75 years. This didn’t start on the 7th of october.”

A final sub-theme discussed the persecution of Christians in Gaza and the role of American Christians supporting the war using antisemitic language. For example, one post claimed that “A few days ago Israel destroyed one of the oldest CHURCH in Gaza and region. These filthy Zionists killed Jesus too That’s why Christians should [support] Palestine and not Zionists.”

By contrast to the attention it received in public debate, Israel-related antisemitism was significantly less visible in our research than very overt and conspiratorial forms.

Far-right antisemitism

All antisemitic themes are inherently expressions of extremist worldviews, even when expressed by actors who do not otherwise hold extremist affiliations or views.

In this analysis, far-right antisemitic sub-themes (3 percent of antisemitic posts relevant to a theme) refer to posts whose language is identifiably rooted in far-right framings. This does not preclude that many other themes and sub-themes identified may also fit into far-right ideologies.

The most common far-right sub-theme was overt praise for Adolf Hitler and his actions. This took the form of comments questioning the facts of the Holocaust or claiming that the current Gaza conflict shows that Hitler was right about the ‘Jewish question’. For example, using overtly far-right framings, one comment stated that “If Germany had won the war, none of this would ever have happened. We need a forever & eternal solution to combat the “Jewish” problem.”

The second most common far-right sub-theme was the use of the word ‘goyim’ – a plural of the derogatory Yiddish word for non-Jews, ‘goy’ – as part of a far-right antisemitic narrative accusing Jews of supremacist beliefs. Notably, the word itself is not antisemitic and is commonly used in far-right networks with a sense of irony, presenting challenges for automated detection systems to accurately understand its antisemitic usage at scale. This sub-theme therefore includes posts which were labelled with a high likelihood of antisemitism; implicit antisemitism may well be higher.




Figure 6: Far-right antisemitic sub-themes over time

Anti-Muslim Hate

Seven central themes emerged from the anti-Muslim comments, of which the most common by far was the common trope claiming Muslims or Islam are inherently violent or dangerous. These comments build on the existing anti-Muslim trope that Muslims pose a physical threat, rooted in radical right and counter-jihad movements' attempts to define post 9/11 Islamist violence as an inherent feature of Islam. Sub-themes related to Islam and violence were lexically diverse by shared a strong narrative that the October 7 attacks were emblematic of Islamic violence.

The second most common theme contained comments arguing that Islam or Muslim people are an invading force or refuse to assimilate. This is once again a common anti-Muslim trope, which others Muslim communities and designates them as an enemy. In its more extreme manifestations, it is closely linked to the 'Great Replacement' conspiracy theory, which alleges that a perceived white or European culture is under threat from Muslim invasion. Sub-themes in this category identify where a specific country is mentioned in the context of invasion or integration narratives, including the UK, Russia, Australia, the US, Canada, Indonesia, the Netherlands, France, and mentions of Europe and the West.

Linked to these narratives is 'cultural incompatibility', which uses a number of tropes to argue that Muslims are fundamentally at odds with a perceived Western culture or way of life. In sub-themes, Muslims are claimed to

be illiberal or anti-LGBTQ+, in line with broader tropes that ultra-conservative viewpoints are universal to expressions of Islam. In this theme, posts also claimed that Muslims are dishonest and instructed by the Qur'an to lie to non-Muslims. Similar to antisemitic posts, niche verses of religious texts are deliberately decontextualised by those seeking to weaponise them.

Due to their prominence, stereotypes that Muslims promote sexual abuse or paedophilia constituted their own theme. Such posts accused sexual violence of being inherent to the behaviours of Muslim communities, closely linked to the idea of Islamic barbarism and posts which accused the Prophet Mohammed of paedophilia. This has often manifested in recent years through the trope of the 'Muslim grooming gang', where Muslim people's criminality is seen as motivated by their faith, ethnicity, or migrant status. Comments in this theme link this trope of sexual abuse to the abuses suffered by Israeli women on October 7.

Aside from existing anti-Muslim tropes, the 'slurs' theme incorporates posts which used basic slurs against Muslim people, including that they are 'cowards', 'brainwashed', 'animals', 'satanic', 'a curse', 'cancer', 'evil', or like 'pigs'.

Similarly to the antisemitism data, specifically conflict-related anti-Muslim hate, such as posts discussing the events in Gaza, Western universities or the hostages were less numerous. While support for Israel is of course not inherently Islamophobic, some posts used anti-Muslim narratives to promote support for Israel. Given that the vast majority of anti-Muslim comments did not mention




Figure 7: Anti-Muslim themes over time

these sub-themes, it is clear that the conflict is used as a jumping-off point for the promotion of existing beliefs.

All themes are tied together by their broad essentialisation of Muslim communities as monolithic, assigning tropes or presuming characteristics of Muslim individuals and collectives based on nothing more than their religion.

A clear omission from this data is the common anti-Muslim trope that Islam is inherently misogynistic and oppresses women. This trope is identified by [Facing History](#) as one of five key Islamophobic tropes, and manifests in discussions of women's use of the burqa, niqab or other face coverings. The omission of this topic from this data may be for a multitude of reasons, both methodological and analytical. It may well be that keywords relevant to this theme were not semantically similar enough to be identified by the topic model. It could also be the case that this theme was not as prominent in conflict-related data as in other anti-Muslim datasets.

Islam and Violence

Clustered into this theme (55 percent of anti-Muslim posts relevant to a theme) were posts promoting the idea that Muslims or Islam are inherently violent, terroristic or dangerous to other faiths. This included specific links to Hamas' actions in the post-7 October conflict. The most common sub-theme was the accusation that all Muslims support Hamas, linking Islam to violence and terrorism. Others included a belief that all Muslims are dangerous, responsible for wars or genocide, antisemitic, supporters of violence, or terrorists. It is this stereotyping which

serves to dehumanise Muslim people.

One common sub-theme focused on references to the 'religion of peace'. This phrase, originally used by activists against anti-Muslim hate following the September 11, 2001 terrorist attacks, has been ironically repurposed to denote Islam as violent and Muslims as disingenuous. This includes, for example, a comment that "the religion of peace sure looks violent." Similarly, the 'jihad' sub-theme claimed that all Muslims aim to commit violent jihad, while the 'death cult' sub-theme accused Islam of being a violent cult over a religion.

The above sub-themes are semantically diverse but thematically extremely similar, united in their perception of Muslims as engaging in or supporting violence. This theme dominated anti-Muslim discourse in October 2023, comprising 59 percent of relevant anti-Muslim posts that month.

Invasion/integration narratives

The second most common theme among relevant anti-Muslim comments (19 percent) involved narratives about Muslims failing to integrate or invading other countries. These narratives were often linked with anti-migrant beliefs, claiming that there were too many Muslims in a specific geography, and called for reduced migration or even expulsion. One comment claimed that "All Muslim immigration is a threat to your country. Muslims do not belong in the West."

Sub-themes included the application of these narratives




Figure 8: Islam and violence sub-themes over time




Figure 9: Integration/invasion sub-themes over time

to specific geographical contexts, most notably the UK; other posts referred to Canada, Australia, the USA, Russia, Netherlands, Indonesia and France, as well as Europe and 'the West'. Posts in individual sub-themes often referred to specific geo-political events or figures, varying from praise for Dutch politician Geert Wilders' anti-Muslim statements to discussion of Russia's attitude towards Muslim populations in the Caucasus.

Cultural incompatibility

The third most common anti-Muslim theme was that Muslims are culturally incompatible with Western societies or values (13 percent of anti-Muslim comments relevant to a theme).

The sub-theme stereotyping Muslims as untrustworthy links to the slur sub-theme of dishonesty and the accusation that Muslims are instructed to lie by their

faith. Such comments called for people to "never trust Muslims". This often included the common far-right phrase "without lies, Islam dies". The 'lies' sub-theme included the concept of taqiyya, [an Islamic concept about concealing the truth for the greater good which has been seized upon by anti-Muslim actors](#) to suggest all believers are inherently untrustworthy.

The 'anti-LGBTQ' sub-theme stereotyped Islam or Muslims as hateful or even violent towards LGBTQ+ communities based on perceptions of Sharia law. These narratives go hand in hand with others which frame Islam as illiberal and incompatible with perceived Western values. Similarly, the 'anti-liberal' sub-theme contained posts claiming that Islam is incompatible with liberal values.

The broad sentiments of the incompatibility of Muslims with Western cultures or values is summarised in the




Figure 10: Cultural incompatibility sub-themes over time

'Muslims/Islam is a problem' sub-theme, where Muslims are presented as inherently different and antithetical to the West.

The 'false accusations' sub-theme focuses on Muslims 'playing the victim card', either denying victimhood or accusing them of manufacturing Islamophobia for political gain. Posts also claimed that it is reasonable to fear Muslims. For example, one commenter posted that "People who fear Islam are not phobic. Islam is currently engaged in about 50 conflicts around the world, and in nearly all of them, the Muslims are the aggressors, the occupiers, and the persecutors." Some commenters discussed the literal translation of the word 'Islamophobia' as a 'fear of Muslims', thereby claiming that the stereotypes they used were not Islamophobic as they were not scared of Muslims.

The Qur'an sub-theme related to messages applying stereotypes about Islam (including paedophilia or lying) to the teachings of the Qur'an. While it is not anti-Muslim to bring theological challenges to the teachings of the Qur'an, this cluster of messages transplanted known Islamophobic narratives onto perceptions of the Qur'an. Finally, the Mosques sub-theme discussed alleged Muslim plans to build Mosques on other religious sites, drawing parallels between the location of al-Aqsa Mosque on Temple Mount and the perceived replacement of Jewish or Christian religious sites with Muslim ones in the West. The replacement of Christian places of worship with Muslim ones is a common far-right mobilising trope, seen as representative of the alleged erasure of Western culture by Islam.

Slurs

The fourth most common theme among relevant anti-Muslim posts was the use of slurs (7 percent).

Sub-themes clustered around slurs unique to anti-Muslim hate include posts slurs relating to Islam's rule not to consume pork. This includes references to Pakistan, a Muslim-majority country, as "Porkistan", or using references to pigs as a targeted term of abuse, such as "Your allah is a pig".

Other common slur words, which clustered into sub-themes, include references to Muslims as cowards, brainwashed or stupid, animals, satanic, a curse, and evil.

Sexual abuse

The theme of sexual abuse centred on claims that Muslims or Islam support or promote paedophilia, rape or sexual abuse, or are sexually promiscuous (3 percent of anti-Muslim posts relevant to a theme). This is a well-established anti-Muslim trope, evident in references to 'grooming gangs' in Europe where perpetrators' ethnic or religious identity is presented as motivation for their criminality.

The first sub-theme of 'paedophilia' made highly offensive claims that the Prophet Mohammed was a paedophile, which was often linked to the accusation that all Muslims therefore support paedophilia. A second sub-theme mocked Muslims for perceived beliefs that when they die, there are 72 virgins waiting for them in heaven. This claim is also used to suggested that Muslims aim to die by violent means to receive this. Finally, the 'rape' sub-theme claims that Islam promotes rape or that Muslims are rapists. Comments linked the documented sexual abuses by Hamas on 7 October to characteristics inherent to Muslim people or Islam.




Figure 11: Anti-Muslim slur sub-themes over time




Figure 12: Sexual abuse sub-themes over time

Pro-violence

A small cluster of posts promoted violence against Muslim people (1 percent of anti-Muslim posts relevant to a theme). The pro-violence theme contained only one sub-theme of calls to ‘destroy Islam’. Posts called to “wipe out Islam”, “finish Islam from this world” or that “Islam shud [sic] be completely wiped out from the world map after this attack”.

Conflict-related anti-Muslim hate

As with Israel-related antisemitism, it was more difficult for the automated detection software to determine the hateful nature of conflict-related anti-Muslim comments, generating a low accuracy score. Due to the relatively small thematic dataset, analysts manually labelled comments to return a high degree of accuracy. The Gaza-related sub-theme applied diverse anti-Muslim

narratives to the Gaza strip, called for attacks against Muslims or used anti-Muslim invasion narratives in response to 7 October (1 percent of anti-Muslim posts relevant to a theme). Similarly, posts in the pro-Israel sub-theme stated their support for Israel using anti-Muslim slurs or stereotypes.

Posts in the hostages sub-theme applied anti-Muslim stereotypes to discussions about the hostages kidnapped by Hamas on 7 October. One such post stated that “I support Israel. It’s about standing up to ISLAM. Islam is what is causing these wars.”, while another stated that “the hostages are likely already dead. Muslims have no honor or respect for human life”. Finally, the universities sub-theme used anti-Muslim narratives to discuss antisemitism on university campuses in the UK and US, where, even before the recent encampments started, both pro- and anti-Israel movements received attention.




Figure 13: Conflict-related anti-Muslim sub-themes over time

Conclusions

This research brief has presented the findings of topic modelling analysis of antisemitic and anti-Muslim comments on YouTube videos about the conflict from 1 October to 24 December 2023. While empirical conclusions cannot be drawn from experimental data, it provides insight into the relative prevalence of themes and sub-themes, with reference to different geographies and individuals.

Antisemitic narratives on videos about the Israel-Gaza conflict majority comprised of conspiracy theories, applying conspiratorial tropes about Jewish money, power, and influence over different events, bodies and geographies. The most common anti-Muslim narrative categorised Muslims as uniquely violent, dangerous or threatening either physically or socially. The prevalence of both these themes demonstrates how the conflict has been used to promote narratives which long pre-date it, with deep-rooted and widespread hateful narratives using the conflict as a springboard for the promotion of existing ideas.

The vast majority of antisemitic and anti-Muslim themes and sub-themes in online discourse about the conflict are familiar, even if different narratives achieved new prominence. This finding is highly relevant to ongoing attempts to delegitimise antisemitism or anti-Muslim hate as merely legitimate criticism of Israel or Palestine. Rather than generating new narratives, such conflicts are used as a jumping-off point for the promotion of old ones, often not even bothering to repackage them in covert or liberal language. The consistency of these narratives across events permits the construction of campaigns which, modified to suit current events, can rely on consistent counter-narratives.

Geographical data, available through the control-based conspiracies antisemitic sub-theme and the integration/invasion anti-Muslim sub-theme, permits comparison of the contexts most discussed in hateful comments. In the case of anti-Muslim hatred, the UK was the most prominent country while for antisemitism, the USA was most commonly mentioned. Given the nature of local debates and politics, countries such as the Netherlands and France arose in anti-Muslim discussions but were not identified in antisemitic comments.

While this research was sparked by the evidenced upsurge in online antisemitism and anti-Muslim hate following the 7 October Hamas attack and subsequent Israeli

offensive in Gaza, the online hate landscape demands rigorous, systematic and live analysis. This is prohibited by poor data access and the difficulties in identifying hate – much of which is borderline in nature – at scale. Current data access provisions have permitted this study on YouTube data, but similar research on other platforms is necessary to gain a zoomed-out understanding of trends across social media. Additionally, a research gap in both multi-lingual and multi-modal analysis needs to be filled.

The diverse nature of hateful online content means that they sit across different areas of legality for borderline content, hate speech and incitement to violence. Although many platforms' Terms of Service now prohibit hateful content, some sub-themes may not reach these thresholds or fit into the specifics of policies. With the introduction of regulatory frameworks in the UK and EU, platforms will be held to their own Terms of Service on the moderation of hateful content.

Antisemitic and anti-Muslim themes evidenced in this research should be used to inform prevention measures and the construction of intervention strategies such as counter-narratives. Through a greater understanding of the narratives of hate, civil society, governments, law enforcement and platforms can work to better understand its drivers. With surging polarisation, hate and extremism, these measures are more important than ever to establish appropriate and robust policies to protect the rights and freedoms of Jewish and Muslim communities.

Amman | Berlin | London | Paris | Washington DC

Copyright © Institute for Strategic Dialogue (2024). Institute for Strategic Dialogue (ISD) is a company limited by guarantee, registered office address PO Box 75769, London, SW1P 9ER. ISD is registered in England with company registration number 06581421 and registered charity number 1141069. All Rights Reserved.

www.isdglobal.org

PO Box 7814, London
United Kingdom, W1C 1YZ
www.isdglobal.org

European | **MEDIA AND
INFORMATION** | Fund

Managed by
Calouste Gulbenkian Foundation

ISD | Powering solutions
to extremism, hate
and disinformation