Manufacturing *the Enemy*

Anti-Jewish Discourse, Sacralized Violence, and YouTube Policy Violations in Qatar's AJ+ Arabic Series

The show never says, "kill Jews."

It doesn't have to.

Dr. Omar Mohammed
The Antisemitism Research Initiative
Program on Extremism
The George Washington University

2026

Table of Contents

Note on Transliteration

Arabic words in this study are transliterated using Brill's Simple Arabic Transliteration System (sAts), version 1.0 (Pim Rietbroek, December 14, 2010). The system is designed for general scholarly use: one Arabic consonant is transliterated with one Roman letter. Initial hamza is not transliterated, and no distinction is made between *ā* from *alif*, *alif maqṣūra*, or *tāʾ marbūṭa*. Arabic evidence is presented in three forms where relevant: original Arabic script, Brill sAts transliteration, and English translation.

Transliteration Table

Arabic	sAts	Arabic	sAts	Arabic	sAts
ا	a, ā	ط	ṭ	ى	ā
ب	b	ظ	ẓ	ي	ī
ت	t	ع	ʿ	و	ū
ث	ṯ	غ	ġ	ـَ	a
ج	ǧ	ف	f	ـِ	i
ح	ḥ	ق	q	ـُ	u
خ	ḫ	ك	k	يْـَ	ai
د	d	ل	l	وْـَ	au
ذ	ḏ	م	m	ة	a, ah, āh, at, āt
ر	r	ن	n		
ز	z	ه	h		
س	s	و	w, ū		
ش	š	ي	y, ī		
ص	ṣ	ء	ʾ		
ض	ḍ				

1. Executive Summary

This study analyzes the complete corpus of *What Did They Say in Hebrew* (*Šū qālū bi-l-ʿibrī*), an Arabic-language YouTube series produced by AJ+ Arabic, a digital sub-brand of the Al Jazeera Media Network. Launched twenty-three days after the October 7, 2023 attack, the series produced 223 episodes by August 2025, accumulating 83.9 million views, 109,165 comments

from 40,985 unique users across at least 22 countries, and 24,977 transcript segments totaling 843,450 characters. Al Jazeera states that the network is funded by the Qatari government. Scholarship on Al Jazeera as an institution has documented both its role in creating a transnational Arabic public sphere (Lynch, 2006, pp. 1–30) and the tensions between its claims of editorial independence and its function within Qatari state diplomacy (El-Nawawy and Iskandar, 2003, pp. 27–54; Figenschou, 2014, pp. 1–20). AJ+ Arabic, the digital sub-brand that produces the series analyzed here, extends this institutional reach into YouTube's platform ecosystem.

The study introduces Displaced Hostility Analysis (DHA), an original framework for measuring how audience discourse shifts from the editorial layer's political register toward an ethnic-religious register targeting Jews. Its principal findings are as follows:

- The editorial apparatus directs 525 negative evaluations at Israeli actors and zero at Hamas in its own voice, while framing Hamas through legitimation and sacralization.
- The editorial names adversaries through political terms in 95.5% of instances, with 2.0% standalone ethnic naming and 1.0% Zionist terminology. Despite this overwhelmingly political register, the audience independently identifies ethnic-religious targeting in 97.6% of videos in which the editorial contains zero ethnic references (162 of 166 videos).
- Where both editorial and audience ethnic references are present, the audience's ethnic-to-political naming ratio is 7.8 times the editorial ratio (median: 5.2×).
- Audience discourse spans seven thematic domains of anti-Jewish framing — ontological evil, theological condemnation, dehumanization, Holocaust inversion, conspiracy, sacred violence, and essentialization — much of it generated from the audience's own scriptural, eschatological (end-of-times), and devotional repertoires rather than from editorial vocabulary.

- Across all six categories of counter-narrative — including civilian sympathy, rejection of collective blame, and characterization of October 7 as terrorism — the study found zero instances in 109,165 comments.
- The corpus contains material that appears to match categories prohibited by YouTube's Hate Speech Policy and Violent Extremist or Criminal Organizations Policy, including dehumanization, incitement, conspiratorial claims, and praise of Hamas, which the United States designates as a Foreign Terrorist Organization. These instances persist in the comment sections of a channel operating within a Qatari state-funded media ecosystem.
- The Arabic detection gap operates through political-to-ethnic register shifts, Ḫaybar invocations, Quranic dehumanization, scriptural essentialization, pronominal targeting, and devotional ambiguity, rendering much of this content difficult for keyword-based moderation systems to detect.

The study helps define an emerging research agenda on Arabic-language antisemitism in digital ecosystems and provides a corpus-based framework for measuring how formally political discourse becomes ethnic-religious hostility at scale.

2. Introduction

The question that drives this study is deceptively simple: how do you say "Jew" in Arabic?

In Arabic-speaking digital spaces, the word *yahūdī* can function simultaneously as a religious identifier, a political accusation, a pejorative shorthand, and, in certain registers, a term so charged that using it analytically may itself trigger accusations of disloyalty, foreign alignment, or hidden political intent. Across multiple Arabic-speaking contexts, it may also operate as an inherited insult detached from any literal reference to Jewish identity (Mohammed, 2024). This linguistic instability is not incidental to the problem examined here. It is one of the conditions that makes Arabic-language antisemitism difficult to detect, classify, and moderate at scale.

Across many varieties of spoken Arabic, naming practices shift between *al-yahūd* ("the Jews"), *al-Isrāʾīliyyūn* ("the Israelis"), *al-ṣahāyina* ("the Zionists"), and *al-kayān* ("the entity"), collapsing religious, political, ideological, and delegitimizing registers into a single field of reference. The problem is therefore not only what is said, but how naming itself moves across those registers — and how that movement is amplified when it enters digital audience discourse at scale. At stake in this study, then, is not only hostile expression but the instability of naming itself: the movement by which Jews, Israelis, Zionists, and "the entity" become discursively collapsible in Arabic digital discourse.

This instability of naming is not only a linguistic phenomenon. It is also the mechanism that allows formally political discourse to slide into ethnic-religious hostility without always appearing to do so on the surface.

2.1 Arabic Digital Media and Transnational Reach

Arabic is spoken by more than 400 million native speakers and functions as the liturgical and scriptural language of Islam, a faith with approximately 1.8 billion adherents worldwide (Pew Research Center, 2025). Arabic-language digital content therefore reaches audiences far beyond the Middle East and North Africa (Zayani, 2005): diasporic communities in France, Belgium, the United Kingdom, the United States, Canada, and Australia consume Arabic media as a primary or supplementary information source. Content produced in Doha is consumed in Detroit and

Marseille. Seib (2008, pp. 1–15) describes this reach as the “Al Jazeera effect” — the capacity of Qatar’s media network to shape political perception across borders and reshape the dynamics of international conflict.

This transnational reach has consequences for antisemitism research and for platform governance. Anti-Jewish content circulating in Arabic does not remain confined to societies where Jewish communities are small or absent. It enters European and North American information environments where antisemitic incidents are at historically elevated levels and where Arabic-speaking populations constitute significant minorities. The ecosystem documented in this study is not confined to a regional audience; it is transnational in reach and implication, mediated by the world's largest video-sharing platform.

2.2 The Moderation Asymmetry

The scale of Arabic-language content on major platforms is not matched by the scale of moderation resources devoted to it. Scholarship on content moderation has documented structural asymmetries in which Arabic-language content receives less sophisticated moderation than English-language content, driven by resource allocation, training data availability, dialectal complexity, and the overlap of devotional and political registers (Elswah and Alimardani, 2021).

During the ISIS media ecosystem (2014–2017), this asymmetry was exposed at scale. Arabic-language extremist content — recruitment videos, devotional songs, execution footage — circulated on English-language platforms whose moderation systems were designed primarily for English content. The lesson of that period was not only that extremist content could reach global audiences, but that audience interaction and comment environments formed part of the wider mobilization ecology surrounding Arabic-language extremist content. Arabic, especially when deployed through scriptural, devotional, and historically coded registers, poses interpretive challenges that general-purpose moderation systems are not designed to evaluate contextually (Gillespie, 2018, pp. 1–20; Roberts, 2019, pp. 33–70).

2.3 The Series

On October 30, 2023 — twenty-three days after the October 7 attack — AJ+ Arabic (@ajplusarabi), a digital sub-brand of the Al Jazeera Media Network, launched a YouTube series titled "What Did They Say in Hebrew" (*Šū qālū bi-l-ʿibrī*). The series is hosted by Laylā ʿAbduh, a Palestinian citizen of Israel from Nazareth, and its format is distinctive: the presenter translates and comments on Hebrew-language Israeli media — news broadcasts, Knesset (Israeli parliament) debates, social media posts, military statements — for an Arabic-speaking audience that cannot access the original Hebrew. The network states that it is funded by the Qatari government (Al Jazeera, About Us; see El-Nawawy and Iskandar, 2003; Figenschou, 2014, on the network's institutional structure; Samuel-Azran, 2016, pp. 195–209, on empirical testing of the Qatari–Al Jazeera editorial nexus).

By August 2025, the series had produced 223 episodes, accumulated 83.9 million views, attracted 109,165 comments from 40,985 unique users across at least 22 countries, and maintained a subscriber base of 2.6 million on the parent AJ+ Arabic channel.

The research that produced this study involved months of sustained engagement with the corpus — transcripts, comments, thumbnails, and tags across the full 22-month period. The content documented here is widely accessible across countries through the world's largest video-sharing platform, and it constitutes a discursive environment with radicalizing potential whose scale and persistence this study seeks to document and measure.

2.4 Research Question

Through what mechanisms is a Qatari state-funded Arabic media operation on YouTube associated with measurable anti-Jewish hostility in its audience, and why do these mechanisms evade platform detection?

The question has three components. First, it asks about mechanisms — not merely whether anti-Jewish content is present, but how the editorial apparatus is associated with the conditions under which audience discourse shifts from political to ethnic-religious targeting. Second, it asks about measurability — whether the shift can be quantified rather than merely described. Third, it asks

about platform governance — why such content persists despite appearing to match categories YouTube's published policies prohibit.

2.5 Contribution

This study makes five contributions.

First, it identifies an understudied object of analysis: an Arabic YouTube ecosystem in which a professionally produced editorial layer and a mass audience comment layer function as a linked system.

Second, it assembles a complete production-reception corpus — 223 videos, 109,165 comments, 24,977 transcript segments, 257 thumbnails, 537 content tags — rather than analyzing isolated examples.

Third, it introduces Displaced Hostility Analysis (DHA), an original methodological framework for measuring the shift from political to ethnic-religious targeting across the editorial and audience layers.

Fourth, it documents what this study terms the Arabic detection gap: the specific linguistic, cultural, and theological mechanisms through which anti-Jewish content in Arabic evades automated and human moderation systems designed primarily for English-language content.

Fifth, it maps the corpus evidence onto YouTube's own published policies, demonstrating that the corpus appears to match categories YouTube publicly states it prohibits — across videos, comments, and other platform features — while remaining substantially under-moderated.

The stakes are not only analytical. Such discursive systems can harden sectarian imaginaries, obstruct Arab-Jewish rapprochement, and create reputational and personal risks for researchers who study them from within Arabic-speaking societies.

2.6 Roadmap

The study proceeds as follows. Section 3 reviews existing scholarship across five bodies of literature — antisemitism research, Arabic language ideology, media and propaganda theory,

parasocial interaction (the one-sided emotional bond audiences form with media figures) and radicalization, and platform governance — and identifies the threefold gap (substantive, methodological, governance-related) this study addresses. Section 4 presents the methodology: research design, corpus structure, the DHA framework, classification instrument, Arabic-specific coding challenges, and limitations. Sections 5–7 present findings across three layers: the editorial apparatus (§5), the audience ecosystem (§6), and the production-reception mechanism linking them (§7). Section 8 interprets the findings within Arabic cultural-religious context and broader theoretical frameworks. Section 9 maps the findings onto YouTube's published platform governance policies. Sections 10–11 offer recommendations and conclusions.

3. Literature Review

This section positions the study within five bodies of scholarship — antisemitism research, Arabic language ideology and scriptural authority, media and propaganda theory, parasocial interaction and radicalization, and platform governance — and identifies the gap the study addresses.

3.1 Antisemitism in Arabic-Language Media

Scholarly understanding of antisemitism in the Arabic-speaking world draws on several overlapping traditions. Lewis (1986) framed an influential early account of anti-Jewish sentiment in Islamic societies as having a long but uneven history, distinct in character from European Christian antisemitism and shaped by different theological, legal, and social dynamics. Webman (2013) surveys the subsequent scholarly debate, noting that the resurgence of antisemitism after the Second Intifada challenged the conventional view that Muslim antisemitism was purely a product of the Arab-Israeli conflict, with newer scholarship contending that it draws on a longer religious and cultural tradition. Wistrich (2010) documented the global resurgence of antisemitism across multiple traditions, identifying both continuities with older forms and distinctly modern developments.

Küntzel (2020) offers the most specific framework for the present study. He identifies "Islamic antisemitism" as a religiously motivated form of modern antisemitism that draws on two distinct sources: anti-Jewish passages within Islamic scripture and tradition, and European antisemitic ideologies transmitted to the Muslim world through specific historical channels — most notably Nazi propaganda broadcasts to Arabic-speaking audiences via Radio Zeesen (1939–1945) and the political alliance between the Jerusalem Mufti Ḥajj Amīn al-Ḥusaynī and the Nazi regime. Küntzel emphasizes that this form of antisemitism is not inherent to Islam as such — Islamic texts include passages that cast Jews positively — but is a specific interpretive tradition that theologizes political conflict by reading Quranic passages about historical Jewish communities as timeless characterizations of Jewish nature. The Hamas Charter (1988), which cites a hadith (a reported saying of the Prophet Muhammad) about Muslims killing Jews at the end of times, exemplifies this fusion of scriptural authority with modern political violence. Crucially, Küntzel demonstrates

that Islamic antisemitism emerged before Israeli statehood — the earliest documented text dates to 1937 — contradicting the claim that it developed as a response to Israeli policy. Complementary archival and historical research has deepened this picture: Herf (2009) documents the content and reach of Nazi Arabic-language radio propaganda; Litvak and Webman (2009) trace how Holocaust denial and inversion became embedded in Arabic political discourse; and Webman (2013) surveys the broader scholarly debate about the origins of Muslim antisemitism. The distinctive feature of the current period, however, is not the persistence of any single tradition but what Herf (2024, pp. 237–238) calls the “era of simultaneity” — the concurrent operation of right-wing, left-wing, and Islamist antisemitism. The AJ+ Arabic corpus documented in this study sits at the intersection of two of these faces: the left-wing anti-Zionist framing that structures the editorial apparatus and the Islamist scriptural-eschatological register that dominates the audience discourse. Across these traditions, Hirschbein and Asfari (2023) demonstrate that the same conspiratorial template — a hermetically sealed theological worldview immune to falsification — operates across white supremacist, Islamist, and other extremist milieus, with Jews consistently positioned as the conspiratorial Other.

Johnson (2024) provides the most comprehensive recent mapping of contemporary left antisemitism, documenting how antizionist discourse has become “an instrument for encrypting hostility to Jews by embedding reference to them in an ideological proxy term” (p. 9) — a formulation directly relevant to the naming instability analyzed in §5.2. The IHRA Working Definition of Antisemitism (2016), with its eleven illustrative examples, provides the operational classification instrument for this study. The definition's explicit distinction between legitimate criticism of Israel and antisemitic content — particularly its identification of Holocaust inversion (Example 9), collective attribution of Israeli state actions to Jews as a people (Example 2), and calls for violence against Jews in the name of religion (Example 1) — structures the coding framework applied in §4.

The "Decoding Antisemitism" project has developed methods for analyzing anti-Jewish discourse in digital media across European languages (Becker et al., 2022; Becker et al., 2024). Their work demonstrates that modern antisemitism frequently operates through indirection — allusion,

irony, visual coding, and contextual implication rather than explicit slurs — a finding directly relevant to the Arabic detection gap documented in §9.4.

This literature establishes the ideological content and historical genealogy of anti-Jewish hostility in Arabic and Islamic contexts, but it does not generally examine how such hostility is produced, circulated, and received within a linked digital media ecosystem at corpus scale.

3.2 Arabic as Sacred Language and Scriptural Authority

The audience discourse documented in this study operates through a language whose sacred status confers a form of authority that secular political language lacks. Arabic is not merely a medium of communication for many Muslim communities; it is the language of divine revelation. The Quran, understood by Muslims as the direct speech of God, is recited, memorized, and liturgically performed in Arabic regardless of the speaker's native language. This sacred status has been analyzed from multiple angles: as a dimension of national identity (Suleiman, 2003, pp. 9–30), as a site of conflict between sacred and secular claims (Suleiman, 2011, pp. 1–22), and as a source of political authority in everyday life (Haeri, 2003, pp. 1–18). The analytical consequence for this study is that when commenters frame political hostility in Quranic vocabulary, the register itself does argumentative work that translation cannot capture. The doctrines of Quranic inimitability (*iʿjāz al-Qurʾān*) and the limits of translation (Abdul-Raof, 2001, pp. 1–25; Suleiman, 2003, pp. 54–71) have contributed to Arabic's distinctive symbolic and religious authority in Muslim communities, an authority that extends beyond theological contexts into political and media discourse. Phrases, formulas, and registers associated with Quranic or devotional Arabic carry a legitimating force that ordinary political vocabulary does not.

This sacralized authority has consequences for discourse analysis. When a commenter frames political hostility in Quranic vocabulary — invoking *al-qirada wa-l-ẖanāzīr* (apes and pigs, Quran 5:60) or reciting *ḥasbunā Allāh wa niʿm al-wakīl* (Quran 3:173) with an appended targeting clause — the register itself confers a form of authority that secular political language lacks. The commenter is not merely expressing an opinion; they are speaking within a register that carries the weight of revelation. This is not a claim about what the Quran teaches or what Islamic

doctrine requires; it is an observation about how certain language registers function as legitimation devices in popular digital discourse (cf. Küntzel, 2020, pp. 3–5; Herf, 2009, pp. 15–50; 2024, pp. 67–68, 76; Litvak and Webman, 2009, pp. 165–200). Hirschbein and Asfari (2023, pp. 52–59) provide a complementary framework, showing that conspiratorial antisemitism across traditions operates through theological categories — worship, revelation, gnosticism, theodicy, and eschatology — that render the conspiratorial worldview immune to disconfirmation. The audience discourse documented in this study exhibits several of these features: the Quranic essentialization constitutes a form of revelation (scriptural authority as proof), while the zero counter-narrative finding is consistent with what Hirschbein and Asfari describe as a hermetically sealed worldview that resists falsification (p. 53).

Scholarship on Arabic sociolinguistics and language ideology has documented this sacralized authority in various domains — education, law, media, and public oratory — but has not typically connected it to the analysis of anti-Jewish discourse in digital ecosystems. The result is that one of the most powerful mechanisms through which online antisemitism operates in Arabic — the deployment of scriptural register to authorize ethnic-religious hostility — remains understudied in scholarship on digital antisemitism, Arabic media discourse, and platform governance.

Existing scholarship establishes Arabic as a language of revelation, liturgy, and symbolic authority in Muslim communities, and it has examined Arabic language ideology, scriptural authority, and the social prestige of Quranic register. What it has not adequately examined is how these linguistic and theological properties operate inside contemporary Arabic digital ecosystems to authorize anti-Jewish discourse. In particular, little scholarship has analyzed how "Jews," "Judaism," and "Israel" may become fused within a single theological-linguistic field of naming, in which political reference is absorbed into scriptural, devotional, and eschatological repertoires. The literature establishes the background condition — Arabic's sacralized authority — but the specific mechanism through which that authority appears to authorize anti-Jewish discourse in digital ecosystems is the argument developed and tested in this study through the corpus evidence in §§5–8.

3.3 Media Framing and Propaganda Theory

The editorial apparatus analyzed in this study is best understood through a convergence of framing and propaganda theories.

Entman (1993, p. 52) defines framing as selecting certain aspects of perceived reality and making them more salient in a communicating text, in such a way as to promote a particular problem definition, causal interpretation, moral evaluation, or treatment recommendation. The framing concept, systematized in propaganda analysis (Jowett and O'Donnell, 2019, pp. 1–30), is applied to the corpus evidence in §5.1 and revisited in §8.3.

Ellul (1962/1973, pp. 74–79) distinguishes agitation propaganda, which seeks to provoke immediate action, from integration propaganda, which seeks long-term alignment between the individual and the group's prevailing ideology. The distinction between agitation and integration propaganda provides the framework for interpreting the temporal patterns documented in §6.12 and discussed in §8.3.

Gerbner's cultivation theory (2002, pp. 46–50) predicts that heavy, consistent exposure to media messaging shapes the audience's perception of social reality. The theory's core prediction — that the more television people watch, the more their worldview aligns with what television presents — provides a framework for interpreting the temporal and geographic audience patterns documented in §§6.11–6.12 and discussed in §8.3.

Herman and Chomsky's propaganda model (1988, pp. 1–35) identifies structural filters — ownership, advertising, sourcing, flak, and ideology — through which media systems produce consent. The AJ+ Arabic case inverts one element of this model: rather than manufacturing consent for the powerful, the series manufactures hostility toward a designated enemy. But the structural logic is recognizable: state ownership (Qatar via Al Jazeera; see Samuel-Azran, 2013), advertiser incentives (YouTube monetization), selective sourcing (Hebrew media chosen for its accusatory potential), and ideological coherence (the moral binary). Lynch (2006, pp. 31–75) demonstrates that Al Jazeera's Arabic-language broadcasting played a constitutive role in creating a "new Arab public" — a transnational discursive space in which shared media

consumption generates shared political orientation. The AJ+ Arabic ecosystem documented in this study operates within and extends this transnational public.

Barthes (1957/1972, pp. 109–159) describes mythification as the process by which historically contingent arrangements are presented as natural, inevitable, and eternal. The editorial's structural absence of moral complexity — zero episodes acknowledging competing claims, zero acknowledgment of Israeli civilian victimhood — performs this mythification: the moral binary is not argued but assumed, transforming a political conflict into a timeless confrontation between good and evil.

3.4 Parasocial Interaction, Moral Disengagement, and Identity Consolidation

Horton and Wohl (1956, pp. 215–229) introduced the concept of parasocial interaction to describe the one-sided intimacy that media consumers develop with performers. Although originally applied to television, the concept extends naturally to digital media presenters whose recurring, direct-to-camera address creates the psychological conditions for perceived personal relationship. The corpus evidence for parasocial dynamics is documented in §6.9 and discussed in §8.3.

Bandura (1999, pp. 193–209) identifies moral disengagement as the set of cognitive mechanisms through which individuals neutralize self-censure for harmful conduct: moral justification, euphemistic labeling, advantageous comparison, displacement of responsibility, diffusion of responsibility, dehumanization, and attribution of blame to the victim. The application of this framework to the corpus evidence is developed in §8.3.

Tajfel and Turner (1986, pp. 7–24) describe social identity theory, in which individuals derive self-esteem from group membership and maintain positive group identity through intergroup comparison. The AJ+ Arabic ecosystem constructs a clear in-group (the *umma*, the resistance, the Palestinian people) and an equally clear out-group (Israel, Jews, Zionists), with the editorial apparatus providing the daily reinforcement that maintains the boundary. The geographic patterns documented in §6.11 are interpreted through this framework in §8.3.

Radicalization research has further theorized the process through which individuals move toward extremism, identifying escalation pathways (Moghaddam, 2005, pp. 161–169) and cognitive mechanisms of political radicalization (McCauley and Moskalenko, 2008, pp. 415–433). This literature explains how individuals are recruited into hostile worldviews, but it does not typically measure the editorial-to-audience transmission chain through which political framing becomes ethnic-religious targeting in Arabic digital discourse.

3.5 Platform Governance and Arabic Content Moderation

A growing body of research documents the asymmetry between platform governance systems designed primarily for English-language content and the challenges of Arabic-language moderation. Gillespie (2018, pp. 1–20) provides the foundational analysis of how platforms make content moderation decisions, while Roberts (2019, pp. 33–70) documents the human labor infrastructure behind automated systems.

Elswah and Alimardani (2021), analyzing the case of #SaveSheikhJarrah, identify what they term "digital orientalism" in content moderation — a documented asymmetry in which Arabic-language content received less sophisticated moderation than English-language content in that specific platform context, driven by resource allocation, training data availability, and linguistic complexity. While their analysis is case-specific, the structural factors they identify — dialectal diversity, morphological richness, and the overlap of devotional and political registers — are directly relevant to the moderation challenges documented in this study. Arabic's dialectal diversity (at least nine major national varieties), its morphological richness, and the prevalence of devotional language that overlaps with extremist registers all compound the detection problem.

The ISIS media ecosystem (2014–2017) remains the most important precedent for Arabic-language extremism on Western platforms. Berger and Morgan (2015) estimated 46,000–90,000 ISIS-supporting Twitter accounts. Platform responses — hash-matching, visual classifiers, and later industry coordination through the Global Internet Forum to Counter Terrorism (GIFCT) — were designed primarily for overt extremist content: black flags, execution footage, and

organizational insignia. Recent research suggests that jihadist online ecosystems have adapted to a more fragmented platform environment rather than disappearing (Ayad, 2025, *CTC Sentinel*), raising the question of whether moderation systems designed for recognizable extremist content can address the implicit, coded, and audience-generated forms documented in this study.

Recent scholarship on content moderation increasingly distinguishes between overtly violative content and borderline or implicit extremist content that carries harmful meaning through coded language, cultural reference, and audience interpretation rather than explicit threats or insignia. General-purpose automated systems struggle to interpret extremist coded language consistently, and effective moderation appears to require hybrid systems combining technical screening with expert human contextual judgment (Rogers, 2025). This challenge is amplified in Arabic, where the devotional register documented in §4.5 — *ḥasbunā Allāh* appearing in 6,094 comments, the vast majority legitimate prayer — illustrates the false-positive problem that makes keyword-based moderation ineffective.

This literature identifies the Arabic moderation gap and the difficulty of governing borderline extremist content, but it does not show, with corpus-level evidence, how a formally compliant editorial layer can coexist with and plausibly help structure a hateful audience layer.

3.6 The Gap This Study Fills

Existing scholarship has established several important foundations: the historical and theological dimensions of antisemitism in Islamic and Arabic contexts; the role of media framing and propaganda in shaping political perception; the importance of parasocial and identity-based mechanisms in digital radicalization; and the structural weaknesses of platform moderation in Arabic-language environments. Yet these literatures remain largely disconnected.

Scholarship on antisemitism in Arabic contexts has not generally been linked to corpus-scale studies of digital audience behavior. Scholarship on propaganda and framing has rarely been brought into direct conversation with Arabic-language antisemitism. Scholarship on platform governance has identified moderation asymmetries in Arabic, but has not typically examined how

hostility is generated across a linked editorial-audience system rather than in isolated content items. The result is a threefold gap: substantive, methodological, and governance-related.

Substantively, there remains limited scholarship on Arabic-language antisemitism as it operates in contemporary digital ecosystems. While Wodak and Reisigl (2001, pp. 31–72) developed discourse-analytical methods for European antisemitism, and the "Decoding Antisemitism" project (Becker et al., 2022) has extended these to digital contexts, Arabic-language digital ecosystems remain largely unexamined at corpus scale. The closest precedent is Rouhana (2023), who analyzes 23,457 Al-Jazeera Arabic articles and 125,501 user comments on Syrian war coverage using a CDA-guided topic modeling method, finding that Al Jazeera used sect-based language in its content and that articles reporting violence received the most sect-based comments. The present study extends this approach in three directions: from sectarian (Sunni–Shi'a) discourse to anti-Jewish hostility, from a news website to a YouTube comment ecosystem, and from topic modeling to a full displacement measurement framework (DHA). Methodologically, there is little work that links editorial production, audience uptake, and platform consequences within a single corpus. From a governance perspective, there is insufficient analysis of systems in which formally compliant editorial content is associated with large-scale audience-layer hate that accumulates in comments rather than in the primary media object itself.

This study addresses those gaps in five ways. First, it identifies an understudied object of analysis: an Arabic YouTube ecosystem in which a professionally produced editorial layer and a mass audience comment layer function together. Second, it assembles a complete linked corpus of videos, transcripts, tags, thumbnails, and comments rather than analyzing isolated examples. Third, it introduces an original methodological framework — Displaced Hostility Analysis (DHA) — to measure the shift from political framing to ethnic-religious targeting. Fourth, it shows how Arabic scriptural, eschatological, and devotional repertoires shape the form of audience hostility in ways not captured by existing moderation models. Fifth, it demonstrates that platform risk may reside not only in overtly extremist content, but in production-reception systems where formally compliant editorial material is associated with large-scale audience-layer hate.

Existing scholarship explains anti-Jewish scriptural essentialization, implicit online antisemitism, Arabic-language moderation asymmetries, and the governance difficulty of borderline extremist content. What it does not yet provide is a linked Arabic-language production-reception analysis capable of measuring how a formally political editorial discourse is transformed, at scale, into ethnic-religious hostility in the audience layer. This study addresses that gap through Displaced Hostility Analysis (DHA), a corpus-based framework designed to measure that shift across editorial, audience, and platform-governance levels. DHA is introduced here as the study's methodological contribution; its full operationalization appears in §4. DHA sits at the intersection of framing analysis, discourse classification, and platform-governance analysis, allowing the study to measure not only hostile language itself but the shift from formally political editorial discourse to ethnic-religious audience uptake.

In this sense, the study is intended not simply to document one case, but to help define an emerging research agenda on Arabic-language antisemitism in digital ecosystems.

4. Methodology

4.1 Research Design

This study employs a non-experimental, retrospective, longitudinal design (Kumar, 2014, pp. 122–136) analyzing naturally occurring data across a 22-month period (October 2023–August 2025). The design is non-experimental because the researcher did not intervene in the production or reception process; retrospective because the complete corpus was collected after the content had been produced and consumed; and longitudinal because it tracks change across four conflict phases, enabling measurement of temporal escalation rather than a single cross-sectional snapshot.

The analytical approach is mixed-method, combining quantitative content analysis (frequency counts, ratio measurements, temporal trend analysis) with qualitative discourse analysis (mechanism identification, contextual interpretation, cultural-linguistic analysis). This design builds on recent computational approaches to Arabic media discourse, particularly Rouhana's (2023, pp. 904–906) CDA-guided topic modeling method, which demonstrated the feasibility of corpus-scale analysis of Al Jazeera Arabic content and audience comments simultaneously. Kumar (2014, pp. 14–20) argues that research objectives should determine the methodological approach and that most real-world research combines both methods. In this study, the quantitative dimension provides the measurements — the displacement ratios, the evaluative vocabulary asymmetry, the temporal escalation — while the qualitative dimension provides the interpretive framework that explains how those measurements are produced.

AI-assisted analysis. The scale of this corpus — 109,165 comments in Arabic across at least nine national dialects, 24,977 transcript segments, 257 thumbnails, and 537 content tags — exceeds what a single researcher could process, code, and systematically analyze through manual methods alone within a feasible research timeline. This study deployed Anthropic’s Claude AI as a core analytical and coding instrument throughout the research process, and does so with full transparency: the use of AI is not an incidental tool choice but a deliberate methodological decision that this study discloses and defends.

Claude AI was deployed across the following tasks. *Data processing:* the systematic byte-level decoding chain (MacRoman → UTF-8) that restored all Arabic text from garbled encoding was developed and executed through Claude AI; frequency computation, pattern extraction, and cross-tabulation across the full 109,165-comment corpus were performed through Claude's processing capacity, which excels at large-scale structured text analysis. *Coding:* the classification of comments across the seven thematic domains, the counting of editorial naming instances, the computation of per-video DHA ratios, and the identification of policy-matching violations were conducted through Claude AI under the direct supervision and iterative guidance of the lead researcher. The coding procedure operated as follows: the researcher established the codebook, operational definitions, and decision rules (Appendix A); Claude AI applied these rules systematically across the corpus, flagging ambiguous cases for researcher adjudication; the researcher reviewed flagged instances, refined decision rules where boundary cases revealed ambiguity, and directed re-coding where necessary. This iterative human-AI coding loop — systematic AI application of researcher-defined rules, with human override on all contextual and ambiguous decisions — enabled consistent application of the codebook across 109,165 comments at a speed and scale that manual coding alone could not achieve. *Cross-layer analysis:* linking the editorial and audience layers at the per-video level — matching transcript naming counts to comment naming counts across 220 videos to compute displacement ratios — was performed through Claude AI. *Visualization:* all figures and data visualizations in this study were generated through Claude AI. *Analytical structuring:* Claude AI assisted in organizing the multi-layered evidence across the study's editorial, audience, and governance analyses, and in developing the theoretical housing for the DHA framework — including the four named dimensions of displacement and the comparative European framing in the discussion.

The researcher's irreplaceable contribution lies in the domain-specific competence that no AI system currently possesses: native proficiency across Arabic dialects, reading knowledge of Hebrew, deep familiarity with Islamic devotional language and sacred-historical memory, and the contextual judgment required to distinguish prayer from weaponized prayer, political naming from ethnic targeting, and scriptural citation from scriptural essentialization. All final interpretive decisions — particularly the contextual coding challenges documented in §4.5 (devotional

register disambiguation, dialectal variant identification, and false-positive adjudication) — were made by the lead researcher. Claude AI applied the rules; the researcher wrote the rules, adjudicated the exceptions, and bears responsibility for the interpretive judgments.

This human-AI research partnership is itself a methodological contribution of this study. It demonstrates that AI-augmented research designs can enable individual researchers or small teams to conduct corpus-scale investigations that would otherwise require large research groups, extended timelines, and substantially greater funding. The speed, consistency, and analytical reach that Claude AI provided in this study — processing 109,165 comments, computing displacement ratios across 220 videos, generating all visualizations, and systematically applying a complex codebook — compressed what would have been years of manual work into months. This has direct implications for the research gap documented in §3: the scarcity of corpus-scale studies of Arabic-language digital antisemitism is partly a function of the resource constraints that AI-augmented methods can now overcome. The study therefore argues not only for a specific set of findings but for a research model: that AI-assisted corpus analysis, directed by researchers with the requisite linguistic and cultural competence, can produce policy-relevant evidence at a speed and scale that the current threat environment demands.

The study treats the video as the baseline case unit, while specific analyses draw on the relevant sub-corpora available for each layer: 220 transcribed videos for editorial-language analysis, 223 videos for comment-level analysis, and 221 for tag analysis. Each video carries a measurable set of editorial features (naming practice, evaluative vocabulary, endorsed quotation, Hamas framing) and a measurable set of audience-response indicators (ethnic targeting rate, thematic domain activation, violation density). Process tracing across the production and reception layers links the two analytically, while multi-source triangulation — transcript data, comment data, thumbnail data, and tag data converging on the same mechanisms — supports triangulated interpretation of the production and reception layers.

4.2 The Corpus

The study population is the complete corpus of all publicly available content produced by the series "What Did They Say in Hebrew" (*Šū qālū bi-l-ʿibrī*) between its launch on October 30, 2023 and the scraping date of August 4, 2025. No sampling was performed; the study analyzes the full collected corpus rather than a subset (Krippendorff, 2018, pp. 83–112; Neuendorf, 2017, pp. 117–154), avoiding the sampling decisions that introduce selection bias in smaller-scale studies (Kumar, 2014, pp. 228–240).

Data was collected through automated YouTube scraping on August 4, 2025, using Bright Data's web data collection infrastructure. Bright Data is a commercial data collection platform that provides structured access to publicly available web content, enabling large-scale harvesting of YouTube metadata, transcripts, comments, and associated fields in compliance with the platform's publicly accessible data. The scraping yielded five linked datasets, which were supplemented by open-source intelligence (OSINT) methods conducted by the lead researcher: manual verification of comment content against live YouTube pages, cross-referencing of user geographic self-identifications, thumbnail collection and visual analysis, and contextual verification of Arabic-language content that automated collection alone could not interpret. The combination of automated data harvesting and researcher-conducted OSINT was methodologically essential: Bright Data provided the scale (109,165 comments, 24,977 transcript segments), while OSINT provided the contextual verification that ensured data integrity and enabled the culturally informed coding documented in §§4.4–4.5. Five linked datasets were obtained:

Table 1: Corpus Datasets

Dataset	Records	Content	Analytical function
Video metadata	223 videos	URL, title, upload date, duration (39–747 sec), views (51K–4M; total 83.9M), likes (1.9K–113K), comment count (97–17K), 537 unique tags, thumbnail URL	Series reach; editorial priority mapping
Transcripts	220 videos; 24,977 timed segments; 843,450 characters	Bilingual Arabic-Hebrew text with per-segment start time, end time, duration	Editorial framing analysis (production layer)
Comments	109,165 records; 40,985 unique users; 223 videos	Comment text, YouTube comment ID, username, user ID, like count (1–1,500), reply count (0–100), parent video URL	Audience response measurement (reception layer)

Thumbnails	257 images	Custom-designed video preview images (129 JPG, 128 PNG)	Visual preconditioning analysis
Content tags	537 unique tags across 221 videos	Editorial team's own content classifications	Editorial priority taxonomy

Encoding correction. Comment text was stored in MacRoman encoding by the scraping platform, rendering all Arabic as garbled Latin sequences. A systematic byte-level decoding chain (MacRoman → UTF-8) restored all Arabic text and usernames. The restoration was verified against a random sample of 200 comments by manual comparison with their YouTube originals.

Temporal structure. The 22-month corpus was divided into four conflict phases:

Table 2: Conflict Phases

Phase	Period	Videos	Comments	Views	Key events
I	Oct–Dec 2023	27	16,857	28.8M	Series launch; immediate post-Oct 7
II	Jan–May 2024	65	34,620	21.7M	Ground offensive; hostage negotiations
III	Jun–Dec 2024	68	27,171	17.0M	Sinwar killed; Lebanon front; Nasrallah
IV	Jan–Aug 2025	61	30,007	16.4M	Ceasefire periods; continued production

Note on comment dating. Individual comment posting dates are available only in relative form ("4 days ago" from the scraping date of August 4, 2025), preventing precise temporal assignment of individual comments to specific conflict events. Temporal analysis therefore relies on the upload date of the parent video to assign comments to conflict phases. Two videos (510 comments) fall on phase-boundary dates and are excluded from the phase table; their comments remain in all non-temporal analyses.

Corpus mismatches. Three discrepancies between the video count (223) and the sub-dataset counts require explanation. Two videos returned crawl errors during transcript extraction, yielding 220 transcribed videos out of 223; the two missing videos are excluded from editorial-layer analysis but their comments remain in the audience dataset. Two additional videos lacked

tag metadata, yielding 221 tagged videos out of 223. The thumbnail count (257) exceeds the video count because 34 episodes were re-uploaded or had variant thumbnails captured during scraping; for analysis purposes, the most recent thumbnail per episode is used.

4.3 Variables

The study identifies two sets of variables linked through a central measurement.

The production layer is measured through nine editorial features: evaluative vocabulary applied to Israeli actors (525 instances) and Hamas (0 in editorial voice), naming practice (ethnic-to-political ratio), endorsed quotation frequency, directive speech acts (445 instances), emotional modeling (sarcastic laughter, 102 instances across 63 videos), structural absences (empathy, criticism, complexity), thumbnail design (257 images), and content tags (537 unique tags). The reception layer is measured through six audience-response indicators: ethnic/political naming ratio per video, thematic domain activation across seven domains, parasocial engagement (2,312 intimacy markers), temporal escalation across four conflict phases, counter-narrative frequency (0 across all six categories), and hostile discourse density (26,620 domain-classified term instances aggregated across the seven thematic domains in §§6.2–6.8 and Hamas sanctification in §6.9. Because comments may be coded under more than one domain (§A.2), this figure represents term-level instances rather than unique comments; the number of unique comments containing at least one domain-classified term is lower. Approximately 18,400 unique comments contain at least one domain-classified term, produced by approximately 12,300 unique users — indicating that hostile discourse is distributed across roughly 30% of the unique user base rather than concentrated in a small prolific subset. Of these, 3,438 are individually mapped to specific YouTube policy categories in §§9.1–9.2). Full operationalization details, coding rules, and decision protocols for all variables appear in Appendix A (§§A.1–A.6).

Theoretical basis. DHA rests on a convergence of four theoretical traditions introduced in §3. Framing theory (Entman, 1993, p. 52) establishes that editorial choices — selection, emphasis, exclusion — shape audience interpretation. Cultivation theory (Gerbner, 2002, pp. 46–50)

predicts that sustained, consistent exposure to a media frame progressively aligns audience perception with that frame. Social identity theory (Tajfel and Turner, 1986, pp. 7–24) predicts that media-constructed in-group/out-group boundaries harden through repeated reinforcement. Moral disengagement theory (Bandura, 1999, pp. 194–196) predicts that structural absences — the erasure of out-group victimhood, the absence of moral complexity — neutralize the audience's self-censure regarding hostile expression toward the out-group. Together, these frameworks predict a measurable consequence: an audience operating within a sustained, morally binary editorial frame will shift from the editorial's register toward a more hostile register drawn from the audience's own cultural, theological, and historical repertoire. The ethnic-to-political naming ratio captures this shift because it measures the point at which political framing is transformed into ethnic-religious targeting — the observable indicator that the theoretical processes are operating.

Four dimensions of displacement. DHA identifies four dimensions of displacement, each supported by distinct evidence in the corpus. First, target displacement: the shift from political to ethnic target, measured by the per-video ethnic-to-political naming ratio (the editorial says "Israel"; the audience says "the Jews"). Second, register displacement: the shift from secular-political to sacred-theological vocabulary, measured by the audience's activation of scriptural, eschatological, and devotional registers absent from the editorial (§§6.2–6.8). Third, amplification displacement: the audience's vocabulary exceeds the editorial's in scale and intensity, measured by the amplification factors in §7.2 (Panel B: up to 88× for muqāwama, 52× for mujrimīn). Fourth, authorization displacement: the audience draws on sources of authority — Quranic scripture, prophetic tradition, eschatological prophecy — that the editorial does not deploy, measured by the audience-only terms in §7.2 (Panel A: mujāhidīn, Ḫaybar, laʿnat Allāh, all with zero editorial instances). The four dimensions are analytically distinct but empirically co-occurring: a single comment may perform target displacement (naming Jews), register displacement (using Quranic vocabulary), amplification displacement (exceeding the editorial's intensity), and authorization displacement (invoking scriptural authority) simultaneously.

Terminological note. The term "displaced hostility" has an independent history in experimental psychology, rooted in the frustration-aggression hypothesis (Dollard et al., 1939; Marcus-Newhall et al., 2000). In that tradition, displaced aggression describes an individual-level process in which a person redirects hostility from an unreachable provoking source to a substitute target. The framework introduced here operates at a different level of analysis: it measures a media-system-level process in which an editorial apparatus's political register is transformed into the audience's ethnic-religious register at corpus scale. The "displacement" in DHA is a register shift — from political naming to ethnic naming — not an aggression redirection from one target to another. However, a structural analogy connects the two: the editorial creates hostility toward a political target (Israel, the occupation) that the audience cannot act upon directly; the audience then extends that hostility to the broader ethnic-religious category (Jews), drawing on its own scriptural and cultural repertoire. The psychological tradition's core insight — that hostility migrates from a provoking source to a proxy target — thus operates in DHA at the media-system level rather than the individual-experimental level. This study does not claim continuity with the experimental displaced-aggression literature; it acknowledges the terminological overlap and the structural analogy while marking DHA as an original framework for corpus-based media analysis.

Central measurement — the Displaced Hostility Analysis (DHA) factor:

The DHA factor is the ratio of ethnic-to-political naming in audience comments divided by the same ratio in editorial transcripts, calculated per video and averaged across the corpus. It measures the degree to which the audience shifts from the editorial's political register to an ethnic-religious register.

A DHA factor of 1.0 would indicate that audience and editorial naming proportions are identical. Across the 54 videos where both editorial and audience ethnic references are present, the mean DHA value is 7.8× (median: 5.2×). The distribution is right-skewed, with an interquartile range of approximately 2.6× to 10.4× and a maximum of 26.5×. Because the DHA ratio divides by the editorial ethnic-to-political ratio, it is sensitive to small-integer variation in the editorial denominator: a video with one editorial ethnic reference will produce a higher DHA than a video with two, even if the audience's ethnic naming is identical. The mean and median are therefore

reported together, and the median (5.2×) is the more robust central-tendency measure. In the 166 videos where the editorial contains zero ethnic references, the audience introduces ethnic targeting in 97.6% of cases (162 of 166). This binary measure does not depend on ratio calculation and is therefore not subject to the small-denominator sensitivity.

Zero-denominator rule. In 166 of 220 videos, the editorial transcript contains zero ethnic references, making the standard ratio undefined. These videos are therefore analyzed separately through a binary measure: whether the audience introduces ethnic targeting despite zero editorial ethnic naming. This yields the 97.6% finding (162 of 166 videos). Across the 54 videos where both editorial and audience ethnic references are present, the mean DHA value is 7.8× and the median is 5.2×, ensuring that ratio-based displacement is calculated only where a valid denominator exists. This two-track approach — ratio measurement where both values exist, binary measurement where the editorial denominator is zero — avoids the methodological error of dividing by zero or artificially smoothing the denominator. A third measure, the aggregate corpus ratio, divides the total audience ethnic-to-political ratio across all comments by the total editorial ethnic-to-political ratio across all transcripts; this yields 9.4× (or 11.9× when Zionist terminology is included as ethnic). This aggregate measure complements the per-video DHA by capturing corpus-level displacement.

4.4 Classification Instrument and Operational Definitions

The study applies the **IHRA Working Definition of Antisemitism** (2016) and its eleven illustrative examples as the primary classification instrument. The IHRA definition is the most widely adopted international instrument for identifying antisemitism, endorsed by more than 40 countries and numerous international organizations. Scholarly debate exists regarding the definition's application, particularly concerning its illustrative examples on Holocaust inversion (Example 9) and collective attribution (Example 2), with some scholars arguing that these examples risk conflating political criticism of Israel with antisemitism (see Klug, 2013; Stern, 2019). This study applies the definition's explicit safeguard: the definition's instruction that "criticism of Israel similar to that levelled against any other country cannot be regarded as antisemitic" is operationalized as follows: criticism of Israeli government policy, named officials, military

conduct, and state institutions is coded as political commentary regardless of intensity or emotional register.

The boundary between political criticism and antisemitic content is crossed when (cf. Wodak and Reisigl, 2001, pp. 31–72):

(a) the target shifts from state to ethnic collectivity;

(b) characterizations employ classic antisemitic tropes identified in IHRA examples;

(c) Holocaust vocabulary is applied to Israeli policy (IHRA Example 9); or

(d) religious texts are invoked to authorize violence against Jews as a group (IHRA Example 1).

Operational definitions:

Term	Definition	Marker terms
Political naming	Reference to state of Israel, government, officials, military. Default category.	Israel, Netanyahu, al-occupation, Israeli army
Ethnic naming	Reference to Jews as ethnic-religious collectivity.	al-yahud, yahud, Bani Israʾil
Zionist terminology	Classified by context: political, ethnic-proxy, or conspiratorial. Reported separately.	sahayina, sahyuni, sahyuniyya
Delegitimizing	Terms that deny ordinary state naming. Note: al-kayān has a long history in Arabic political discourse and is used by speakers who may intend rejection of normalization rather than ontological denial of Israel's existence; the coding decision classifies it by its discursive function (refusal of standard state naming) rather than by presumed intent.	al-kayan, al-kayan al-sahyuni
Displacement factor	Audience E/P ratio ÷ editorial E/P ratio. Mean: 7.8×. Median: 5.2×.	(audience E/P) ÷ (editorial E/P)
Endorsed quotation	Presenter appends amin ya rabb to translated Israeli statement.	amin ya rabb
Counter-narrative	Comment resisting dominant frame (6 categories). Zero found.	§A.4

Transliteration standard. All Arabic words, names, phrases, and quoted evidence in this study are transliterated using Brill's Simple Arabic Transliteration System (sAts). Arabic evidence is presented, where relevant, in three forms: original Arabic script, Brill sAts transliteration, and English translation. The transliteration system is used consistently across the main text, tables, figures, and exhibits.

Emergence of the seven thematic domains. The seven domains of anti-Jewish framing (ontological evil, theological condemnation, dehumanization, Holocaust inversion, conspiracy, sacred violence, essentialization) were not imposed deductively from the IHRA framework. They emerged through iterative engagement with the 109,165-comment corpus, following Shkedi's (2019, p. 63) principle that categorization must reflect a "conversation between the theoretical perspective and the data" — an approach consistent with the iterative category development described in grounded theory methodology (cf. Charmaz, 2014, pp. 109–137). Initial reading of 5,000 randomly selected comments generated provisional categories, which were then tested against the full corpus, refined, merged, and validated through frequency analysis. The final seven domains represent the stable categories that survived this iterative process.

4.5 Arabic-Specific Coding Challenges

Three features of Arabic require explicit methodological treatment because they directly affect coding decisions throughout the study.

Dialectal variation. Arabic text preprocessing for computational analysis faces challenges beyond diglossia: diacritics (ḥarakāt), dialectal variation, and the definitive article "al" all require specialized handling that standard NLP pipelines do not provide (Rouhana, 2023, pp. 907–908). The Arabic-speaking world is characterized by diglossia — the stable coexistence of a literary standard and spoken vernaculars (Ferguson, 1959, pp. 325–340; Versteegh, 2014, pp. 189–223; Albirini, 2016, pp. 7–35). The comment corpus contains contributions from speakers of at least nine national Arabic varieties: Egyptian (*ʿĀmmiyya*), Moroccan (*Dārija*), Algerian, Tunisian, Levantine (Palestinian, Syrian, Lebanese, Jordanian), Yemeni, Sudanese, Iraqi, and Gulf Arabic. The same targeting expression may appear in phonologically, morphologically, and lexically distinct forms across these varieties. Coding was performed by a researcher with native competence across the major Arabic dialect groups, enabling identification of dialectal variants that a classifier trained on Modern Standard Arabic alone would miss. This dialectal competence is not a supplementary advantage but a methodological requirement: without it, a substantial proportion of the corpus would be miscoded.

Interpretive competence. This analysis depends on domain-specific interpretive competence. The researcher brings native Arabic proficiency together with training in Arabic philology, Islamic and post-classical history, and the historical semantics of religious and political vocabulary. This combination is methodologically important because the corpus does not operate only at the level of lexical meaning. It operates through the interaction of language, sacralized linguistic authority, scriptural register, historical memory, geography, and political discourse. The study therefore treats Arabic not simply as a vehicle of communication but as a layered interpretive field in which linguistic form, religious legitimacy, and historical reference overlap. This competence enables contextual reading but does not eliminate interpretive risk, which is why the codebook, re-review process, and ambiguity log documented below remain methodologically essential.

The devotional register. Arabic Islamic devotional language pervades the comment corpus:

Table 3: Devotional Register Frequency in Comment Corpus

Formula	Arabic	Frequency	% of corpus
"O God"	Allāhumma (اللهم)	14,361	13.2%
"God is sufficient for us"	ḥasbunā Allāh (حسبنا الله)	6,094	5.6%
"Amen"	āmīn (آمين)	912	0.8%

The vast majority of these instances are legitimate Islamic prayer — petitionary language directed at God, expressing faith, grief, or hope. This study does not code devotional language as anti-Jewish. The formula *ḥasbunā Allāh wa niʿm al-wakīl* ("God is sufficient for us, and He is the best disposer of affairs," Quran 3:173) is one of the most common Islamic devotional expressions, recited by hundreds of millions of Muslims, and is not evidence of hostility toward any group.

Only instances where a targeting clause is appended that specifies Jews or Zionists as the object — for example, *ḥasbunā Allāh ... bi-l-yahūd al-mujrimīn* ("God is sufficient for us ... against the criminal Jews") — are coded under the relevant thematic domain. The coding is attributed to the appended clause, not to the prayer formula. This distinction requires contextual reading of each instance, a task that no automated system currently performs and that constitutes one of the core findings of the platform governance analysis (§9).

False-positive rates. Arabic terms used to characterize Jews in this corpus carry high false-positive rates in general Arabic discourse:

Table 4: Arabic Terms with High False-Positive Rates

Term	sAts	Meaning	False-positive context
ظالمين	ẓālimīn	oppressors	Standard political term; used for any perceived injustice
منافقين	munāfiqīn	hypocrites	Quranic term applied to many groups across Islamic discourse
فساد	fasād	corruption	Standard Arabic political vocabulary; routine in governance criticism

These terms have very high false-positive rates in general Arabic discourse, appearing routinely in political, religious, and everyday contexts that have nothing to do with Jews. Coding decisions were therefore made on the basis of contextual proximity to Jewish or Israeli referents within each comment, not on keyword presence alone. This contextual judgment — what Shkedi (2019, p. 122) terms the researcher's "intuitive inquiry skills" operating within a systematic analytical framework — is methodologically essential. Keyword-based coding would produce massive overcounting; context-free exclusion would miss the targeting that operates through precisely these ambiguous terms.

4.6 Validity, Reliability, and Limitations

Validity is addressed through three mechanisms (Kumar, 2014, pp. 211–230):

> **Content validity* — the IHRA Working Definition, as the internationally recognized consensus instrument for identifying antisemitism, provides the substantive framework against which all coding decisions are made.*
>
> **Construct validity* — displacement is operationalized through a two-track system: a ratio-based DHA measure, defined as the audience ethnic-to-political naming ratio divided by the editorial ethnic-to-political naming ratio, calculated across the 54 videos where both editorial and audience ethnic references are present; and a binary ethnic-*

introduction measure for the 166 videos where editorial ethnic references are absent. Both measures can be independently calculated from the raw data.

**Face validity* — all Arabic evidence is presented in three forms: original Arabic script, Brill's Simple Arabic Transliteration System (sAts), and English translation — enabling readers with competence in any of the three forms to verify coding decisions independently.*

Reliability is addressed through transparency and process documentation. Every comment cited in this study includes its YouTube comment ID, enabling independent retrieval and verification. The coding criteria (IHRA examples), frequency counts, and displacement calculations are documented in sufficient detail to permit replication by any researcher with Arabic competence and access to YouTube's public data.

Three additional process safeguards were applied: coding was governed by a written codebook specifying inclusion and exclusion rules for each thematic domain and each operational definition (Appendix A); ambiguous cases — comments where the targeting intent was uncertain or where devotional language sat at the boundary between prayer and curse — were logged separately and re-reviewed after the initial coding pass; and a random subset of 500 coded comments was re-checked after the full corpus coding was complete to assess internal consistency of the coding decisions.

Limitations. Seven limitations are explicitly acknowledged:

1. **Single-coder analysis** — no inter-coder reliability testing was performed. As disclosed in §4.1, coding was conducted through Claude AI under the direct supervision and iterative guidance of the lead researcher: the researcher established all codebook rules and operational definitions, Claude AI applied them systematically across the corpus, and the researcher adjudicated all ambiguous and boundary cases. No second human coder independently reviewed the classifications. This remains a significant methodological limitation. The coding decisions most likely to produce inter-coder disagreement — distinguishing devotional from weaponized ḥasbunā Allāh, classifying “Zionist” as political versus ethnic-proxy, and placing ibāda among

evaluative vocabulary — are precisely the contextual judgments on which the study's frequency counts depend. The AI-assisted coding model offers one advantage over purely manual single-coder analysis: Claude AI applies the codebook rules with greater consistency across 109,165 comments than a human coder working over months, reducing the drift and fatigue effects that affect manual coding of large corpora. However, this consistency applies only to the rules as defined; whether a second qualified researcher would define the same rules or adjudicate the same boundary cases identically remains untested. The justification is practical: the required combination of native dialectal Arabic competence, reading knowledge of Hebrew, and familiarity with Islamic devotional and scriptural registers is rare. Future applications of DHA should incorporate inter-coder reliability testing where a second qualified coder is available, testing both the codebook rules and the boundary-case adjudications. The 500-comment re-check reported below tests intra-coder consistency (whether the researcher agrees with their own earlier decisions on re-examination) rather than inter-coder reliability (whether a second coder would reach the same decisions).

2. **Single-series scope** — findings describe one AJ+ Arabic series and may not generalize to other Arabic media operations or channels. In particular, the study cannot determine whether the displacement patterns documented here are specific to this editorial apparatus or would appear in any Arabic-language coverage of the Israel-Palestine conflict. A comparative study — applying DHA to a different Arabic channel covering the same conflict, or to a non-Arabic channel with comparable editorial framing — would be necessary to isolate the editorial mechanism from the baseline audience disposition. Elmasry et al. (2013) provide the closest existing comparator, analyzing Al Jazeera and Al Arabiya framing of the Israel–Palestine conflict and finding systematic differences in editorial framing between the two networks. Future work should extend this comparative approach by testing DHA across multiple channels to determine whether displacement is a property of this specific editorial system or a wider phenomenon.

3. **Correlational production-reception link** — the study documents systematic co-occurrence between editorial framing and audience response but cannot establish causal inference in the experimental sense.

4. **Temporal coincidence with conflict** — the corpus spans an active military conflict, meaning audience response may partly reflect conflict intensity rather than editorial framing alone.

5. **Comment survivorship bias** — deleted, hidden, or shadow-banned comments are not captured by the scraping instrument.

6. **No bot detection** — coordinated inauthentic behavior analysis was not performed on the user dataset.

7. **Unavailable comment dates** — individual comment posting dates are available only in relative form, preventing precise event-level attribution.

5. Findings I: The Editorial Apparatus

This section documents the editorial decisions made by Al Jazeera's AJ+ Arabic team within the 843,450-character transcript corpus across 220 analyzed videos. Every finding describes an act of production: a vocabulary choice, a naming decision, an emotional cue, or a structural absence created by Qatar's AJ+ Arabic editorial team. The audience response to these decisions is analyzed separately in §6.

Unless otherwise specified, counts in this section refer to occurrences in the Arabic editorial voice and exclude attributed Israeli or Western speech.

Note: All counts reported in this section are based on verified full-corpus computation and constitute the study's definitive figures. Where these counts differ from preliminary estimates referenced in §4, the findings-section figures govern. The differences reflect refinements during the verification process and do not affect the direction or analytical substance of any finding; specific discrepancies are noted where they occur.

5.0 The Founding Frame: How the Editorial Presents October 7

The series launched twenty-three days after the October 7 attack — the largest single-day killing of Israeli civilians in the country's history. How the editorial presents this founding event establishes the moral frame through which every subsequent episode is received. A systematic examination of the editorial's treatment of October 7 across the transcript corpus reveals that the attack is never narrated as an event with Israeli civilian victims — it is narrated exclusively through three alternative frames: Israeli intelligence failure, Hamas operational achievement, and Israeli self-inflicted harm.

The launch episode. The first episode (October 30, 2023; video: 0W46WebopWU) does not describe the October 7 attack. It opens with Abu Ubaida's speech, pivots to Netanyahu's political crisis, and covers hostage families' anger at Netanyahu. The sole reference to October 7 is:

□□□□□□ □□ □□□□□□ □□ □□□□ □□□□ □□□□□□□□□□□□ □□□□□

al-fashal al-istiḫbārātī allaḏī yūqiʿ fī al-sābiʿ min uktūbar

"the intelligence failure that occurred on the seventh of October"

The audience's first encounter with October 7 in this series frames it not as something Hamas did to Israelis, but as something Israel's intelligence services failed to prevent. The attack itself — the killing of approximately 1,200 civilians, the assault on kibbutzim, the Nova festival massacre, the taking of over 240 hostages — is absent. The editorial opens not with the event but with its political consequences for Netanyahu.

Episode 3 (November 3, 2023; video: QwquHXnSCW0). The editorial references a 47-minute film the Israeli military showed Knesset members documenting October 7. The editorial describes it as:

[illegible]

film ʿan afʿāl Ḥamās fī al-sābiʿ min uktūbar

"a film about Hamas's actions on the seventh of October"

The word afʿāl (actions) is chosen over hujūm (attack), iʿtidāʾ (assault), or majzara (massacre). The editorial does not describe what the film contains. Instead, it describes the Israeli reaction — fainting, emotional collapse — and immediately pivots to the argument that the film should have had a "Part One" about the intelligence failure that led to October 7. In the same episode, the editorial calls October 7:

[illegible] *7* [illegible]

ḍarbat al-muqāwama al-sabt 7 uktūbar

"the resistance's strike on Saturday, October 7"

This is the editorial voice — not quoting Hebrew media, not translating Israeli speech — framing the attack as a "resistance strike." The lexical choice ḍarba (strike, blow) frames October 7 as a military action within a resistance paradigm, not as an attack on civilians.

Episode 6 (November 14, 2023; video: cjcO5CjZ0sY). Titled "What is Israel hiding about what happened on October 7?", this episode introduces the Hannibal protocol narrative. The editorial asks:

□□ □□ □□□□□□□□□□ □□ □□□□□ □□□□□ □□□□□ □□□□ □□ □□□□□□□□□□ □□□□□ □□□□□ □□□□□□□ □□□□□ □□ □□□□□□ 7 □□ □□□□ □□□□□ □□□□□

hal fiʿlan muqātilī Ḥamās wa-l-ālāf al-muwāṭinīn min Ġazza illī daḫalū maʿahum hun al-masʾūlīn ʿan kull al-qatl illī ḥaṣal fī 7 uktūbar

"Are Hamas fighters and the thousands of citizens from Gaza who entered with them really responsible for all the killing that happened on October 7?"

The editorial's angle is not empathy for the dead but suspicion toward the Israeli narrative: the Israeli military may have killed its own civilians through indiscriminate helicopter fire and the Hannibal protocol. October 7 is reframed from Hamas attack on civilians to Israeli friendly fire on Israeli civilians.

Naming patterns across the corpus. The editorial's vocabulary for October 7 and its consequences follows a systematic pattern across the corpus:

Majzara (مجزرة, massacre) appears 16 times across the transcript corpus. In every instance, it describes Israeli actions against Palestinians: the flour massacre, the Tantura massacre, school bombings in Gaza. The two instances where majzara appears near "October 7" are both attributed Israeli speech — Netanyahu vowing to pursue "everyone who participated in the massacre of the seventh of October," and Israeli claims about UNRWA staff involvement. The editorial voice never calls October 7 a massacre.

Al-muḥtajazīn (المحتجزين, the detained / the held) is the editorial's preferred term for hostages, appearing 217 times. The term frames hostage-holding as detention rather than kidnapping. Al-muḫtaṭafīn (المختطفين, the kidnapped) appears 31 times, but contextual examination reveals that the majority of these instances refer to the physical location mīdān al-muḫtaṭafīn (Hostage Square) in Tel Aviv — a place name in attributed Israeli usage — rather than the editorial's own characterization of the hostages.

Ṭūfān (طوفان, flood — Hamas's operational name for October 7) appears 15 times in editorial transcripts and is tagged on 141 of 221 videos (63.8%), as documented in §5.6. In two instances,

the editorial explicitly frames October 7 as ʿamaliyyat Ṭūfān al-Aqṣā (operation Al-Aqsa Flood), adopting Hamas's operational nomenclature.

Irhāb (ارهاب, terrorism) appears 52 times in the transcript corpus. Contextual examination confirms that in every instance, the term is applied to Israeli or Western actors, attributed to Israeli speech about Hamas, or used in general political description. The editorial voice never applies the word "terrorism" to October 7 or to Hamas.

What is absent. Across 843,450 characters of editorial transcript, the editorial voice does not describe what happened to Israeli civilians on October 7. It does not name the Nova festival. It does not describe the attacks on kibbutzim. It does not acknowledge the nature of the violence — the killing of families in their homes, the targeting of a music festival, the taking of children and elderly as hostages. It does not use the word "massacre" for October 7. It does not use the word "terrorism" for October 7. It does not call the hostages "kidnapped" in its own voice. The attack that created the series is present in the editorial only as an Israeli failure, a resistance operation, or a mystery about Israeli self-harm.

This founding frame — installed in the first episodes and maintained without exception — is the precondition for the moral binary documented in §5.1. The audience does not need 22 months of accumulated framing to receive the message that October 7 is not an atrocity requiring empathy but an Israeli failure requiring explanation. That frame is delivered at launch.

5.1 The Moral Binary

The editorial apparatus applies evaluative vocabulary to Israeli actors and Hamas with an asymmetry that is systematic across the analyzed episodes.

Israeli actors receive 525 instances of negative evaluative vocabulary across the transcript corpus. The table below lists all identified terms:

Table 5: Evaluative Vocabulary Applied to Israeli Actors

Arabic	sAts	English	Count
ابادہ	ibāda	genocide	182

فشل	fashal	failure	80
كذب	kaḏib	lies	77
جرائم	jarāʾim	crimes	68
تجويع	tajwīʿ	starvation	30
تهجير	tahjīr	forced displacement	28
مجرم	mujrim	criminal	22
وقاحه	waqāḥa	insolence / audacity	11
نازي	nāzī	Nazi	7
لصوص /اللصوص	luṣūṣ	thieves	6
مجرم حرب	mujrim ḥarb	war criminal	6
هتلر	Hitlar	Hitler	4
المخلوقه	al-maḫlūqa	the creature	2
اكلي لحوم البشر	akalī luḥūm al-bašar	cannibals	1
معسكر تركيز	muʿaskar tarkīz	concentration camp	1
		Total	525

These counts represent raw keyword occurrences in the Arabic editorial voice. The vast majority target Israeli actors; a small number appear in contexts describing general conditions (e.g., *fashal* applied to a specific military operation rather than Israel as an entity). The counts are reported as raw frequencies to maintain transparency; the analytical interpretation throughout this section treats them as indicators of editorial vocabulary patterns rather than individually verified attributions.

The most frequent term — *ibāda* (genocide, 182 instances) — appears in 95 of 220 analyzed videos (43.2%). It is never qualified: the editorial does not say "alleged genocide" or "what the ICJ is investigating." The grammatical frame is declarative, not conditional. A note on classification: the term ibāda could be interpreted as a political-legal characterization rather than an evaluative one, particularly given the active ICJ proceedings during the corpus period. It is included in the evaluative count here because the editorial deploys it as a settled moral verdict rather than as a legal claim under adjudication — it functions in the editorial register not as a factual assertion subject to evidence but as a characterological accusation applied to Israeli actors as a class. Readers who would reclassify ibāda as political-legal vocabulary should note that the remaining evaluative count would be 343, with the moral binary at 343-to-zero. The analytical pattern — systematic negative characterization of Israeli actors with zero negative

characterization of Hamas in editorial voice — is not affected by this reclassification; its magnitude is. The editorial vocabulary operates in a graduated structure, descending from political condemnation (*fashal*, *kaḏib*, *jarāʾim*) through active-transitive accusation (*tajwīʿ* — Israel does not "cause" starvation; it *starves*) to dehumanization. At the bottom of this gradient, the editorial voice — not quoting Hebrew media, not translating Israeli speech — calls Israeli actors *al-maḥlūqa* ("the creature"), *akalī luḥūm al-bašar* ("cannibals"), and *luṣūṣ* ("thieves").

In the episode of July 3, 2025 (video: gh9sQQ71Ctw), the presenter describes Israeli soldiers opening a pizza restaurant in destroyed Gazan homes:

□□□□□□□ □□□□□ □□□□ □□□□□□□ □□□□□□ □□□□□□

al-luṣūṣ kāyinīn yaqtaḥimū buyūt al-nās wa-ḥurmatahā

"The thieves, invading people's homes and their sanctity."

In the episode of July 25, 2025 (video: byodK4APEnU), she relays and endorses an Israeli acknowledgment:

□□□□□ □□□□ □□□□ □□ □□□□ □□□□ □□□

ṣār manẓar zurnā zay akalī luḥūm al-bašar

"Our appearance has become like cannibals."

In the same episode, referring to an elderly Israeli woman:

□□□□□ □□□ □□□□□□□□ □□□□□

hayyāhā al-maḥlūqa kayf biṭḫiṭṭiṭ

"Here she is — the creature — planning."

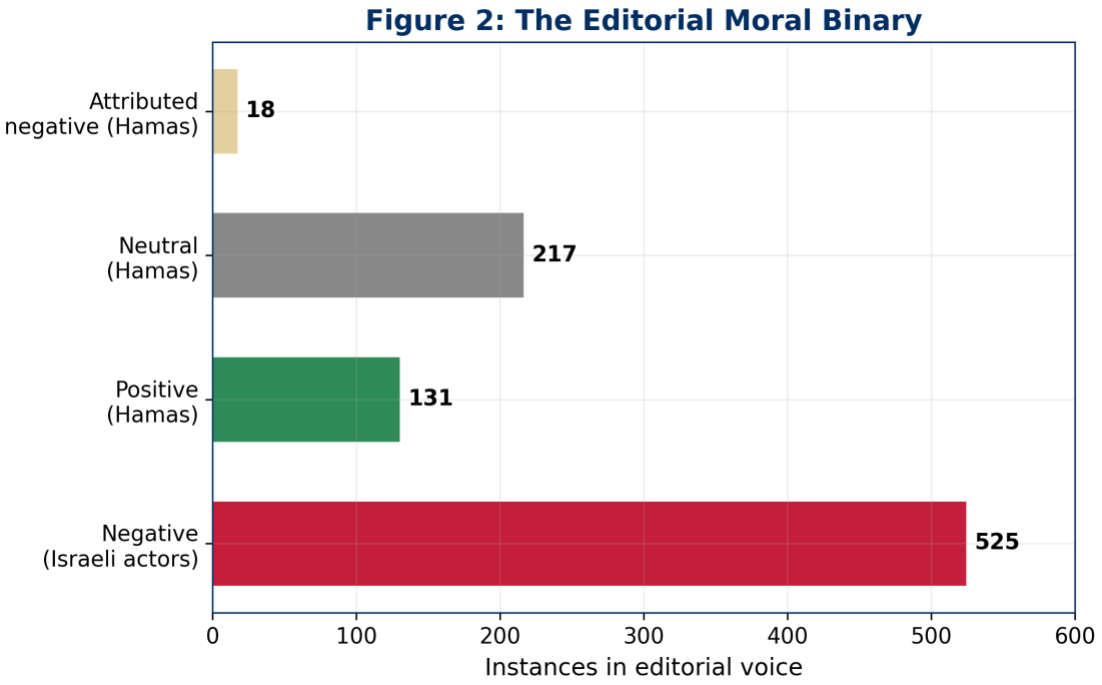

Figure 1: The Editorial Moral Binary

Hamas (*Ḥamās*), across 366 mentions in the editorial transcripts, receives a categorically different treatment:

Table 6: Hamas Framing in Editorial Transcripts

Framing	Instances	Percentage	Examples
Positive	131	35.8%	asr (captured, 115), šuhadāʾ (martyrs, 52), muqāwama (resistance, 19), al-Qassām (13), abṭāl (heroes, 5)
Neutral / descriptive	217	59.3%	Movement name, organizational references, factual context
Attributed negative	18	4.9%	"terrorism," "terrorist army" — attributed to Israeli or Western speech in every instance

The lexical inventory behind these percentages reveals the specific registers through which Hamas is framed. The tables below report raw lexical occurrences rather than mutually exclusive framing instances; a single Hamas mention may contain more than one legitimizing or descriptive term.

Hamas positive and legitimizing vocabulary:

Arabic	sAts	English	Count	Register
أسر /أسروا	asr / asarū	captured / they captured	115	Operational-normalizing
شهيد /شهداء	šahīd / šuhadāʾ	martyr / martyrs	52	Positive-sacralizing
استشهد	istašhada	was martyred	23	Positive-sacralizing
مقاومة	muqāwama	resistance	19	Positive-legitimizing
القسام	al-Qassām	al-Qassam (Brigades)	13	Organizational-legitimizing
أبطال	abṭāl	heroes	5	Positive-heroic

Hamas neutral and descriptive vocabulary:

Arabic	sAts	English	Count	Function
حماس	Ḥamās	Hamas	366	Movement name
السنوار /سنوار	al-Sinwār	Sinwar	146	Leader reference
أبو عبيدة	Abū ʿUbayda	Abu Ubaida	25	Spokesman reference
محمد الضيف	Muḥammad al-Ḍayf	Muhammad al-Deif	15	Commander reference
طوفان الأقصى	Ṭūfān al-Aqṣā	"Al-Aqsa Flood"	7	October 7 operational name

Attributed negative vocabulary near Hamas:

Term	Count	Attribution
"terrorism" / "terrorist" (irhāb / irhābī)	18	Israeli or Western speech in every instance

The asymmetry is visible not only in percentages but in lexical choice. Hamas is not merely mentioned more neutrally; it is named through a vocabulary of legitimation and sacralization. Terms such as *šuhadāʾ* (martyrs), *muqāwama* (resistance), *al-Qassām*, and *abṭāl* (heroes) place Hamas within a register of valorization, while the operational verb *asr* ("captured"), the dominant collocate at 115 instances, frames capture as military achievement without moral condemnation — classified as positive because applying a neutral military verb to hostage-taking and related seizure events constitutes legitimation through operational normalization. By contrast, the limited negative vocabulary appearing near Hamas is always attributed to Israeli or Western speech and is never generated in the editorial voice itself. Over 22 months and 220 episodes, the editorial voice does not produce a single negative characterization of Hamas in its own register. The asymmetry is structurally consistent across the corpus. Herf (2024, pp. 204–205) analyzes

the Hamas Charter as a text that explicitly declares its struggle to be against Jews as Jews — not merely against Israel as a state — and fuses secular anti-Zionist rhetoric with religious antisemitism drawn from Islamic scripture. The editorial apparatus's systematic legitimation of Hamas through the resistance and martyrdom registers documented above thus operates in tension with the organization's own founding text, which makes no distinction between political and ethnic-religious targeting. In propaganda analysis, the asymmetry documented in the preceding data — 525 negative characterizations of Israel against zero for Hamas — constitutes a "black-and-white" or "Manichean" framing device (Jowett and O'Donnell, 2019, pp. 303–310). Hirsh, in Johnson (2024, pp. xxiii–xxiv), identifies an analogous "Empire versus Resistance" binary in contemporary left discourse, in which complex political realities are collapsed into a framework that assigns absolute moral valence to each side. The editorial apparatus operationalizes precisely this structure: Israel as Empire, Hamas as Resistance, with zero moral complexity admitted.

5.2 Naming and Editorial Slippage

The editorial apparatus names the adversary through political terms in the overwhelming majority of instances. A careful count of adversary-naming instances across the transcript corpus yields:

Table 7: Adversary-Naming in Editorial Transcripts

Category	Terms (examples)	Count	% of naming
Political	al-Isrāʾīlī (the Israeli), Isrāʾīliyyīn (Israelis), Natanyāhū, al-iḥtilāl (the occupation), Isrāʾīl (Israel)	3,700	95.5%
Ethnic (standalone)	yahūd (Jews), al-yahūd (the Jews)	76	2.0%
Zionist	ṣahāyina (Zionists), ṣahyūnī (Zionist)	40	1.0%
Jewish-adjectival	yahūdī (Jewish, adj.), al-yahūdī (the Jewish)	57	1.5%

When Jewish-adjectival forms are included, total Jewish-referent naming rises to 133 instances (3.5%). When Zionist terminology is added, the combined non-political referent rate is 4.5%.

Political naming constitutes 95.5% of the editorial adversary register (total naming instances: 3,873).

This political naming provides the editorial apparatus with plausible deniability: the series can claim it targets a state, not a people. However, the 76 standalone ethnic references function as bridge terms, introducing the ethnic register into an otherwise political discourse. Contextual examination reveals three usage patterns: 31 instances (41%) are culturally descriptive ("Moroccan Jews," "the Jewish holiday"); 26 instances (34%) appear within translated Israeli Hebrew content; and 19 instances (25%) constitute **editorial slippage** — instances where the presenter uses "Jews" where "Israelis" would be the more precise referent.

From December 13, 2024 (video: 2BFaOtVDsOI):

> [illegible]
>
> **al-Ḥarīdīm al-yahūd al-mutadayyinīn waṣalū ʿalā Sūriyā nafs al-namaṭ al-ijrāmī bi-Ġazza bi-Lubnān wa-l-ān bi-Sūriyā**
>
> *"The ultra-Orthodox religious Jews have arrived in Syria — the same criminal pattern as in Gaza, in Lebanon, and now in Syria."*

Here, "religious Jews" — not "Israeli settlers" or "Israeli citizens" — are identified as agents of a "criminal pattern" spanning three countries. The ethnic identifier is chosen over the national one. This slippage always flows in one direction: from political to ethnic. The editorial does not substitute ethnic for political terms; it supplements the political register with ethnic terms, creating a bridge that the audience then crosses at scale (§6). This naming pattern aligns with what discourse analysts identify as ideological work through categorization: naming practices in political and media discourse are not neutral but construct the categories through which audiences perceive social actors (Van Dijk, 2006, pp. 120–130). In the study of antisemitic discourse specifically, Wodak and Reisigl (2001, pp. 44–72) identify naming strategies as a primary mechanism through which discriminatory discourse is both produced and concealed. Herf (2018) documents the historical trajectory by which antisemitism and anti-Zionism became

fused into a single ideological formation, while Walzer, in Johnson (2024, pp. 57–64), analyzes the specific mechanism by which anti-Zionism slides into antisemitism: the substitution of “Zionist” for “Jew” creates a political proxy that enables ethnic targeting while maintaining deniability — the same naming instability documented in the corpus’s 40 Zionist-terminology instances (§A.1).

5.3 The Endorsed Quotation

The presenter's most analytically consequential speech act is the endorsed quotation. She translates an Israeli statement and appends the devotional formula *āmīn yā rabb* (آمين يا رب, "Amen, O Lord"). From July 25, 2025 (video: byodK4APEnU, youtube.com/watch?v=byodK4APEnU):

> [illegible]
>
> ***ṣārū yqūlū ʿamāyilnā bi-Ġazza raḥ tilāḥiqnā la-ajyāl qādima... āmīn yā rabb... qāʿidīn yqūlū bikaffī tiṭlaʿū tuḥarriḍū ʿalā al-ibāda ʿalanan ṣār manẓarnā zay akalī luḥūm al-bašar inkharsū***
>
> ***"They started saying: our deeds in Gaza will haunt us for generations to come — haunt us for generations to come — **Amen, O Lord** — they're saying stop inciting genocide openly, our appearance has become like cannibals, shut up."***

An instance of self-critical Israeli commentary is converted through the appended *āmīn yā rabb* into a prayer for Israel's punishment. The presenter does not argue that Israel should be punished; she prays for it, transforming journalistic citation into devotional ratification.

This technique achieves three effects simultaneously. First, it maintains **journalistic plausible deniability**: "I was only quoting what they said in Hebrew." Second, it signals to the audience that the extreme content is not merely reported but **endorsed** — the devotional formula is a cultural cue that Arabic-speaking Muslim audiences recognize as sincere affirmation. Third, it models a **response pattern**: when you hear Israeli self-condemnation, the appropriate response is not

analysis but prayer for its fulfillment. The 912 instances of *āmīn* in the comment corpus are consistent with audience uptake of this cue.

The presenter's credibility mechanism depends not only on her speech acts but on her identity position. As a Palestinian citizen of Israel from Nazareth, she occupies a uniquely authenticating role: a member of the audience's in-group who can access the out-group's language from inside the adversary society. She is not translating from a foreign language but from a language she speaks natively, in a country where she lives. The editorial leverages this identity: she is tagged by name on 196 of 221 videos (88.7%), and the transcript corpus contains self-references to "my town" (baldī, 9 instances) and "the Palestinian interior" (al-dāḫil al-filasṭīnī, 5 instances), positioning her as an insider witness rather than an external commentator. The 654 audience invocations of "O Layla" and 568 "God preserve you" prayers (§6.9) are therefore not generic parasocial attachment but attachment to a figure who performs the role of trusted informant from enemy territory — a role that her Palestinian-Israeli identity uniquely enables.

5.4 Directive Speech Acts and Emotional Modeling

The presenter directs the audience's attention, interpretation, and emotional response through 445 documented directive speech acts:

Table 8: Directive Speech Acts

Arabic	sAts	English	Count	Function
يا ناس	yā nās	"O people!"	113	Audience mobilization
لازم	lāzim	"You must"	101	Interpretive command
شوفوا	šūfū	"Look!"	86	Attention direction
انتبهوا	intabihū	"Pay attention!"	28	Alert / alarm
Others	—	Various directive forms	117	Mixed

These are interpretive commands that tell the audience what to see, what to feel, and what to conclude. The presenter positions herself not as a neutral translator but as a guide through hostile territory who directs the audience's reading.

Sarcastic laughter at Israeli distress appears in 102 documented instances across 63 videos. From July 25, 2025 (video: byodK4APEnU, youtube.com/watch?v=byodK4APEnU):

> [illegible]
>
> **hh hin ʿārfīn kayf yilammimū khazwīhum quddām al-ʿālam illī baṭṭal ṭāyiqhum**
>
> *"Ha — they know how to cover up their shame before the world that can no longer stand them."*

The laughter models the appropriate emotional response to Israeli distress: not empathy, not analysis, but amusement.

The presenter's direct address, recurring commands, and modeled emotional responses establish a parasocial mode of delivery whose audience uptake is analyzed in §6.

5.5 Structural Absences

The editorial apparatus is defined as much by what it systematically excludes as by what it includes.

Israeli civilian empathy. A systematic search for expressions of empathy toward Israeli civilian casualties in the editorial transcripts returned one instance across 843,450 characters. That instance describes Israeli military friendly fire on Israeli civilians on October 7, framed not as empathy but as evidence of Israeli military incompetence. Over 22 months, the editorial voice does not acknowledge that the approximately 1,200 Israelis killed on October 7 might deserve mourning, that hostage families might deserve sympathy, or that Israeli children might be innocent.

Hamas criticism. Zero instances in the editorial voice across 220 episodes.

Moral complexity. Zero episodes present competing moral claims, acknowledge ambiguity, or grant any legitimacy to Israeli perspectives.

5.6 Editorial Priorities: Tags and Thumbnails

The 537 unique content tags applied by the AJ+ Arabic editorial team across 221 tagged videos constitute their internal classification of what each video is "about." The top 15 tags:

Table 9: Top 15 Content Tags

Tag (Arabic + sAts)	English	Videos tagged
إسرائيل (Isrāʾīl)	Israel	207
نتنياهو (Natanyāhū)	Netanyahu	207
الاحتلال الإسرائيلي (al-iḥtilāl al-Isrāʾīlī)	The Israeli occupation	203
أبو عبيدة (Abū ʿUbayda)	Abu Ubaida (Hamas spokesman)	202
يحيى السنوار (Yaḥyā al-Sinwār)	Yahya Sinwar (Hamas leader)	201
أهالي المخطوفين (ahālī al-maḫṭūfīn)	Families of the hostages	197
ليلى عبده (Laylā ʿAbduh)	Layla Abduh (presenter)	196
جنوب غزة (janūb Ġazza)	Southern Gaza	190
شمال غزة (šamāl Ġazza)	Northern Gaza	190
الضفة الغربية (al-Ḍiffa al-Gharbiyya)	The West Bank	187
خان يونس (Ḫān Yūnis)	Khan Younis	187
محمد الضيف (Muḥammad al-Ḍayf)	Muhammad al-Deif (Hamas commander)	184
قائد الأركان (qāʾid al-arkān)	Chief of Staff	183
بن غفير (Bin Ġafīr)	Ben Gvir	156
طوفان الأقصى (Ṭūfān al-Aqṣā)	"Al-Aqsa Flood" (Oct 7)	141

Hamas actors — Abu Ubaida (202 videos), Sinwar (201), al-Deif (184) — receive comparable or greater editorial prominence to Israeli state actors. Abu Ubaida is tagged on more videos than any named Israeli political figure except Netanyahu. The October 7 attack is categorized under Hamas's operational framing, *Ṭūfān al-Aqṣā* ("Al-Aqsa Flood"), on 141 of 221 tagged videos (63.8%).

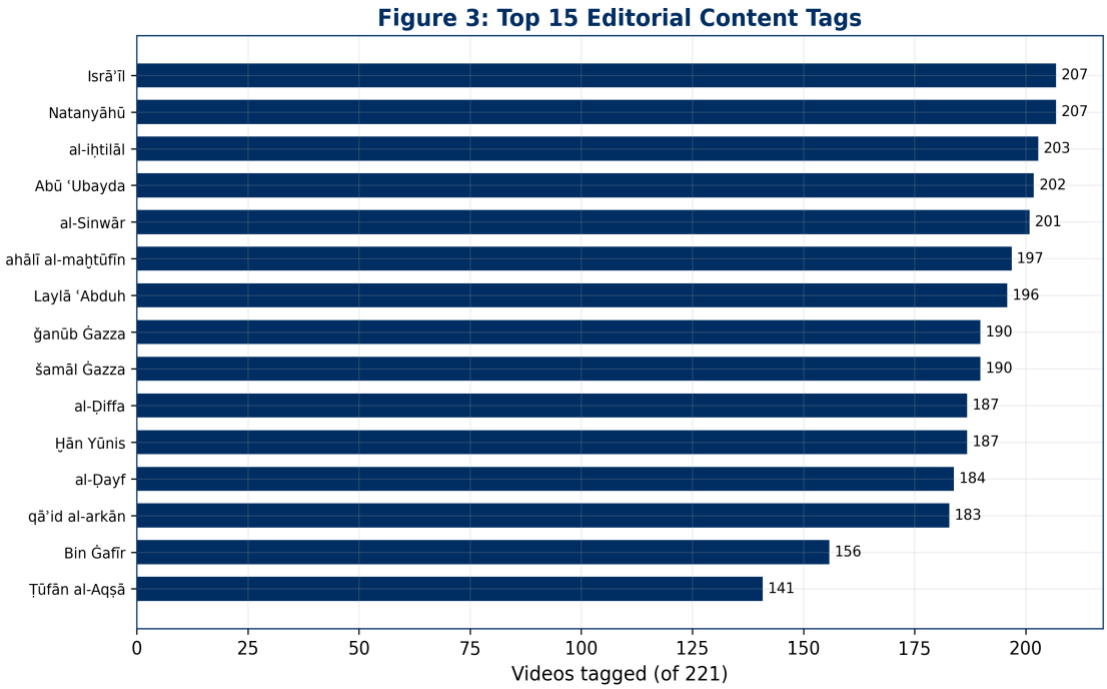

Figure 2: Editorial Priorities from Content Tags

The 257 thumbnails constitute the first layer of the framing apparatus. Visual analysis reveals consistent design patterns: extreme close-ups of Israeli politicians with menacing color grading; Arabic text overlays in accusatory register framing content before the viewer clicks; juxtaposition of Palestinian suffering imagery with Israeli political figures; and Hamas leadership presented through heroic visual grammar.

6. Findings II: The Audience Ecosystem

Rabab'ah and Alali (2019) provide the closest methodological precedent, analyzing impoliteness in reader comments on the Al Jazeera Arabic news website and finding that anonymity and asynchronicity significantly increase hostile expression. The present study extends this approach from impoliteness to anti-Jewish hostility, and from a news website to a YouTube comment ecosystem. This section documents the audience response to AJ+ Arabic's editorial apparatus analyzed in §5. DHA — the study's central measure of displacement — captures the extent to which audience discourse shifts from the editorial layer's political register toward an ethnic-religious register. Every finding here describes a pattern in the 109,165-comment corpus produced by 40,985 unique users.

6.1 The Displacement Finding

Displacement is documented through a two-track system.

Track 1: Ratio-based displacement. Across the 54 videos where both editorial and audience ethnic references are present, the audience's ethnic-to-political naming ratio is 7.8 times the editorial ratio (median: 5.2×). The range extends from 0.8× (a video where the audience is slightly less ethnically targeted than the editorial) to 26.5× (a video where the audience's ethnic register is more than 26 times the editorial's).

Track 2: Binary ethnic introduction. In 166 of 220 videos, the editorial transcript contains zero ethnic references — the editorial names Israel, Netanyahu, the occupation, but never says "Jews." In 162 of those 166 videos (97.6%), the audience independently introduces ethnic targeting. The audience generates ethnic-religious vocabulary that the editorial never produces.

Aggregate corpus ratio. Across the full corpus, the audience's ethnic-to-political naming ratio (0.357) is 9.4 times the editorial ratio (0.038). When Zionist terminology is included as an upper-bound sensitivity check, the ratio rises to 11.9×.

The central finding is not simply that audience hostility exceeds editorial hostility, but that ethnic targeting emerges in the comment environment even where the editorial transcript does not

explicitly name Jews at all. The editorial provides the political target; the audience's own cultural, religious, and historical repertoire appears to provide the ethnic-religious upgrade. An important interpretive limitation applies: because the study analyzes a single series during an active military conflict, it cannot determine whether the audience ethnic-introduction rate (97.6%) or the mean displacement (7.8×) are specific to this editorial system or would appear at comparable levels in comment sections of other Arabic-language channels covering the same conflict. The displacement finding documents a systematic pattern within this corpus; attributing that pattern specifically to the editorial apparatus — rather than to the broader Arabic-language information environment during this period — would require comparative application of DHA to other channels, as discussed in Limitation #2.

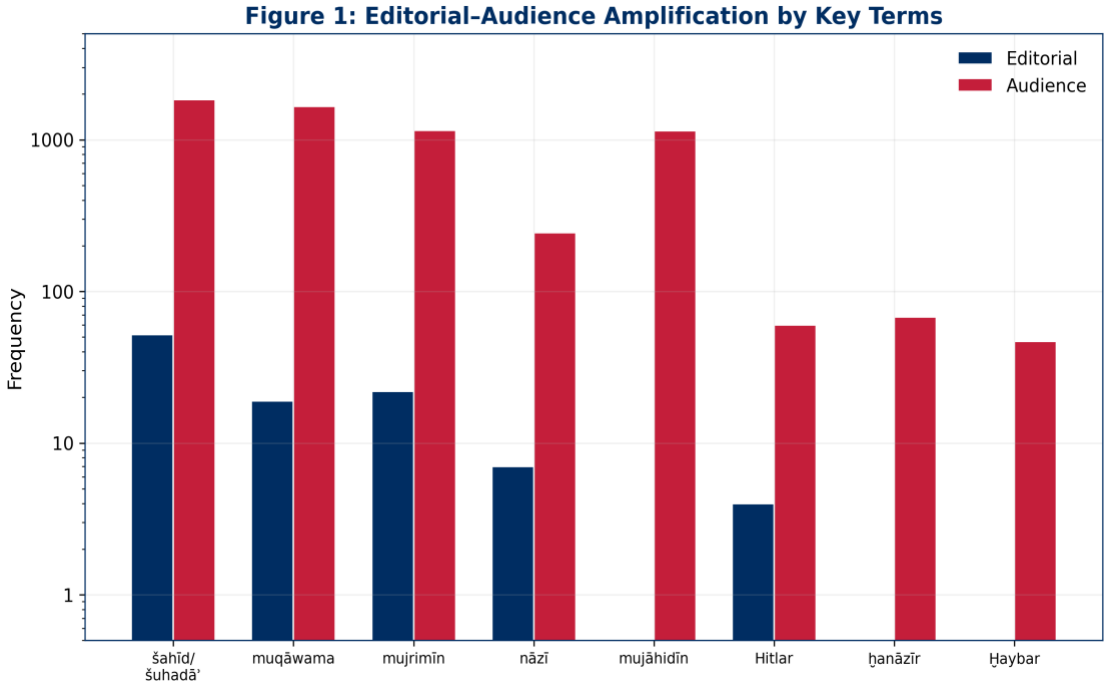

Figure 3: Editorial–Audience Amplification by Key Terms

6.2 Domain 1: Ontological Evil and Treachery

One of the largest domains. The audience constructs Jews not as political opponents but as ontologically evil — not oppressive because of specific policies but because oppression is their nature.

Table 10: Domain 1 — Ontological Evil and Treachery

Arabic	sAts	English	Audience count	Editorial count
ظالمين	ẓālimīn	oppressors	2,437	0
مجرمين	mujrimīn	criminals	1,154	22
خيانة /خونة	ẖiyāna / ẖawana	betrayal / traitors	906	0
قتلة /قاتل	qatala / qātil	killers	626	0
فساد /مفسدين	fasād / mufsidīn	corruption / corrupters	476	0
ملعون /ملاعين	malʿūn / malāʿīn	cursed	291	0
منافقين	munāfiqīn	hypocrites	246	0
غدر /غدارين	ġadr / ġaddārīn	treachery / treacherous	129	0
طاغوت /طاغين	ṭāghūt / ṭāghīn	tyrants	117	0

Every term in this domain except *mujrimīn* ("criminals," which has 22 editorial instances) has zero editorial occurrences. The audience generates this vocabulary from its own repertoire.

Top-liked comment (452 likes): "O God, strike the oppressor with the oppressor and bring Gaza and Palestine out safely and victorious" (اللهم اضرب الظالم بالظالمين واخرج منهم غزه وفلسطين سالمين غانمين). IHRA relevance: Example 2 (stereotypical allegations about Jews as collective).

> **Source: youtube.com/watch?v=0W46WebopWU | Comment ID: UgwPcN_ydqmTFwGjGJp4AaABAg | 452 likes**

6.3 Domain 2: Theological Condemnation

The audience sacralizes political hostility, framing Jews not as political adversaries but as enemies of God whose punishment is divinely ordained.

Table 11: Domain 2 — Theological Condemnation

Arabic	sAts	English	Count
عذاب	ʿadhāb	divine punishment	244
لعنة الله	laʿnat Allāh	God's curse	115
أعداء الله	aʿdāʾ Allāh	enemies of God	73
الساعة	al-sāʿa	the Hour (eschatological)	65
غضب الله	ġaḍab Allāh	God's wrath	42
آخر الزمان	āẖir al-zamān	end of times	25

The highest-liked comment in this domain (919 likes) connects the Israeli importation of red heifers to the Quranic narrative of the golden calf: "These red heifers they brought... remind me of the calf they made with their own hands and worshipped... and God's wrath fell upon them" (هذه البقرات الحمراء ...تذكرني بالعجل التي صنعوها بايديهم وعبدوا ...وغضب الله عليهم). A contemporary Israeli religious practice is read through Quranic narrative, and the conclusion is divine wrath. IHRA relevance: Example 1 (calling for harm to Jews in the name of an extremist view of religion).

Source: youtube.com/watch?v=TSIM43eqJ0E | Comment ID: UgyEh86FDBys_eH81FN4AaABAg | 919 likes

6.4 Domain 3: Dehumanization

Arabic	sAts	English	Count
خنازير	ḫanāzīr	pigs	68
قردة	qirada	apes	32
القردة والخنازير	al-qirada wa-l-ḫanāzīr	apes and pigs	12
كلاب	kilāb	dogs	29

While numerically smaller than other domains, these terms carry disproportionate force because they derive from Quranic verses about divine punishment (5:60), lending sacred authority to contemporary dehumanization. A comment (319 likes) following the assassination of Hamas political leader Ismail Haniyeh: "Don't celebrate, you pigs — victory is coming at the hands of Abu Ubaida" (متفرحويش يا خنازير النصر قادم قادم علي يد ابو عبيدة). IHRA relevance: Example 2 (dehumanizing allegations about Jews as collective).

Source: youtube.com/watch?v=TSIM43eqJ0E | Comment ID: Ugxs8InQsE6fQXXm3i54AaABAg | 319 likes

6.5 Domain 4: Holocaust Inversion

Arabic	sAts	English	Audience count	Editorial count
نازي/ نازية	nāzī / nāziyya	Nazi / Nazism	244	7
إبادة	ibāda	genocide	485	182
هتلر	Hitlar	Hitler	60	4
محرقة	maḥraqa	holocaust	26	0

معسكر	muʿaskar	camp	17	1

The audience deploys Holocaust vocabulary to characterize Israeli actions, performing the inversion the IHRA definition identifies as antisemitic: transforming the Holocaust's victims into its perpetrators. Klaff, in Johnson (2024, pp. 141–144), analyzes Holocaust inversion as a distinct form of contemporary antisemitism that functions as aggressive anti-Jewish propaganda by reversing victim and perpetrator roles. A comment (32 likes): "America the mother of terrorism and Nazi Israel — God is sufficient for us against the criminals" (أمريكا ام الإرهاب وإسرائيل النازيه وحسبنا الله ونعم الوكيل عليهم المجرمين). IHRA relevance: Example 9 (using symbols of classic antisemitism to characterize Israel).

6.6 Domain 5: Conspiracy

Arabic	sAts	English	Count
مخطط	muḫaṭṭaṭ	plot / scheme	181
مؤامرة	muʾāmara	conspiracy	37

Jewish and Israeli agency is constructed as inherently conspiratorial — a pattern that Hirschbein and Asfari (2023, pp. 52–59) analyze as “conspiracy theory as theology,” in which conspiratorial beliefs function through the same structures as religious worldviews: revelation, gnosticism, theodicy, and eschatology. A comment (79 likes): "It is a conspiracy orchestrated between the Israeli Zionists and the Egyptian Zionists under the leadership of the Mexican agent" (إنها مؤامرة مدبرة بين الصهاينة الاسراءيليين والصهاينة المصريين). IHRA relevance: Example 2 (conspiratorial claims about Jews).

6.7 Domain 6: Sacred Violence and Jihadist Framing

The largest domain by total instances. The audience frames armed violence against Jews as sacred obligation.

Table 12: Domain 6 — Sacred Violence and Jihadist Framing

Arabic	sAts	English	Audience count	Editorial count
شهيد /شهداء	šahīd / šuhadāʾ	martyr / martyrs	1,843	52

مقاومة	muqāwama	resistance	1,666	19
بطل /أبطال	baṭal / abṭāl	hero / heroes	1,407	5
مجاهدين	mujāhidīn	mujahideen	1,147	0
جهاد	jihād	jihad	280	0
خيبر	Ḫaybar	Ḫaybar invocation	47	0

The significance of this domain lies not only in the repetition of editorially available legitimizing terms such as *šuhadāʾ* and *muqāwama*, but in the audience's extension of that register into explicitly jihadist vocabulary that is absent from the editorial transcript corpus. The terms *mujāhidīn* (1,147), *jihād* (280), and Ḫaybar (47) have zero editorial instances — they are generated entirely from the audience's own repertoire. The Ḫaybar invocation — *Ḫaybar Ḫaybar yā yahūd jayš Muḥammad sawfa yaʿūd* ("Ḫaybar, Ḫaybar, O Jews — the army of Muhammad will return") — explicitly names Jews as the target of a returning military force modeled on the Prophet's seventh-century campaign against a Jewish community. IHRA relevance: Example 1 (calling for violence against Jews in the name of religion).

The highest-liked comment in this domain (817 likes): "Advice from the heart to our brothers the mujahideen in proud Gaza: trust God alone" (نصيحة من القلب الى اخواننا المجاهدين في غزة العزة :توكلوا على الله وحده). YouTube policy relevance: Dangerous Organizations Policy (praising terrorist groups).

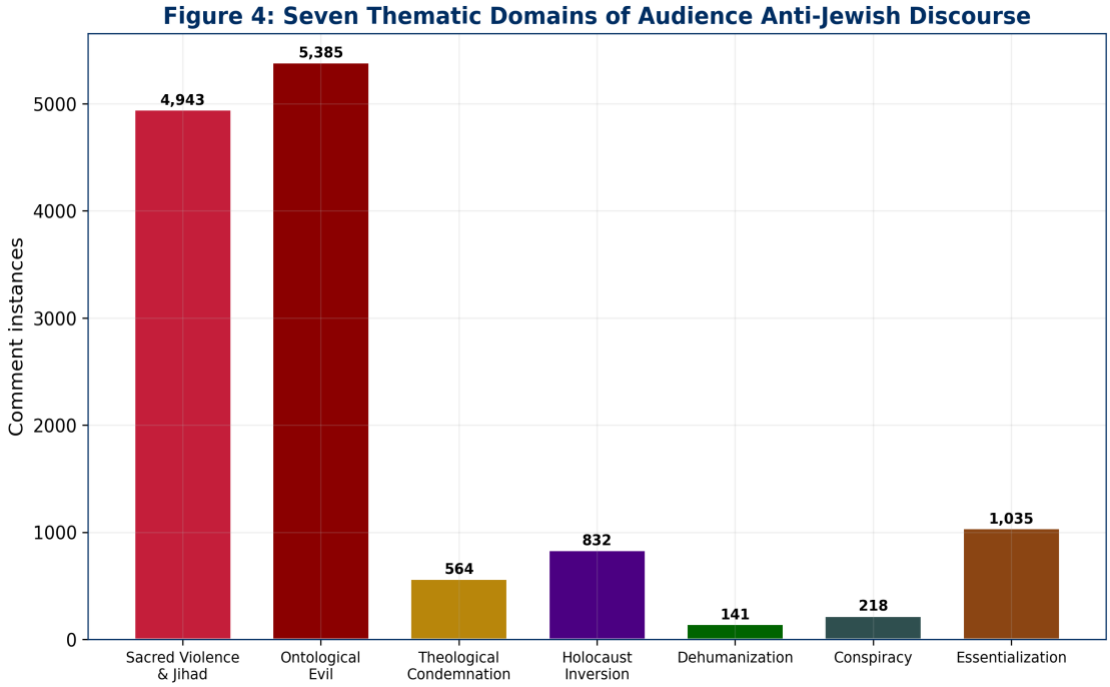

Figure 4: Seven Thematic Domains

6.8 Domain 7: Essentialization and Collective Blame

Comments attributing eternal, unchanging characteristics to Jews as a people.

The highest-liked (170 likes): "I was just reading Sūrat al-Baqara and I saw with my own eyes: the condition of the Jews yesterday is their condition in every era — treacherous, untrustworthy, but cowardly" (كنت لسه بقرأ في سورة البقرة ...حال اليهود أمس هو حالهم في كل زمان — غدارين لا يؤتمنون ولكن جبناء). This comment performs the bridge from Quranic text to contemporary characterization: the commenter reads a passage about a historical community and applies it as a timeless truth about all Jews.

> **Source: youtube.com/watch?v=IcnZfOSPhlA | Comment ID: UgydGJAWWbGw0-AS8kd4AaABAg | 170 likes**

A Moroccan commenter (24 likes): "Even among us in Morocco, when we want to insult someone treacherous, we say 'you Jew' or 'son of a Jew'" (حتى أنه عندنا في المغرب حين نرغب بسب أحد غدار وخائن نقول له يا يهودي أو ابن يهودي). IHRA relevance: Example 2 (stereotypical allegations about Jews as collective).

6.9 Hamas Sanctification and Parasocial Uptake

Sinwar (*al-Sinwār*) is referenced in 1,097 comments with hagiographic framing — a pattern consistent with the mechanisms of political radicalization identified by McCauley and Moskalenko (2008, pp. 420–425). Top-liked: "The death of Sinwar and all the free people who preceded him in martyrdom only increases the resistance's courage — these are the real leaders, dying in the front lines" (234 likes). "I am Algerian and my tears fell for this Palestinian hero, Yahya Sinwar" (176 likes).

Abu Ubaida (*Abū ʿUbayda*) is referenced in 351 comments with veneration. "Abu Ubaida — the voice of freedom and the dignity of the Arab person — will remain a thorn in the throat of the occupying entity" (272 likes). "Hearing Abu Ubaida has become a source of pride and hope for us" (231 likes).

October 7. Of 221 comments referencing *Ṭūfān al-Aqṣā*, not one characterizes the attack as terrorism, criticizes its violence against civilians, or acknowledges Israeli victims. The highest-liked (366 likes): "Everything coming out of the entity currently... is just empty talk to diminish the Al-Aqsa Flood operation."

Parasocial uptake. The editorial mode of delivery documented in §5.4 is associated with measurable audience intimacy. The comment corpus contains 2,312 parasocial markers:

Table 13: Parasocial Markers

Arabic	sAts	English	Count
يا ليلى	yā Laylā	"O Layla"	654
الله يحفظك	Allāh yaḥfaẓak	"God preserve you"	568
أختي	uḫtī	"My sister"	388
الله يحميك	Allāh yaḥmīk	"God protect you"	235
أستاذة	ustādha	"Professor / Madam"	211
Others	—	Various intimacy terms	256

A comment linking the presenter directly to the audience's affective response (583 likes): "O Layla, you and General Fayez al-Dweiri and the resistance videos heal our chests and raise our morale — may God help the valiant resistance" (يا ليلى انت و اللواء فايز الدويري و فيديوهات المقاومة تشفو

غليل صدورنا و تطلعو معنوياتنا). The commenter explicitly names the presenter as a source of emotional mobilization alongside resistance media.

Source: youtube.com/watch?v=e_r5UDJRaSg | Comment ID: UgxsOcetjDv9QTt2MAp4AaABAg | 583 likes

6.10 Counter-Narrative Analysis

A systematic search for comments resisting the dominant editorial frame — across all six counter-narrative categories defined in §4.4 (civilian sympathy, rejection of collective blame, objection to anti-Jewish curses, refusal of celebratory violence, distinguishing Jews from Israelis, characterizing October 7 as terrorism) — returned **zero results** that met the operational definition.

The search used Arabic-language terms for Israeli civilians (*madaniyyīn Isrāʾīl*), Israeli victims (*ḍaḥāyā Isrāʾīl*), innocent Israelis (*abriyāʾ Isrāʾīl*), "not all Jews" (*laysa kull al-yahūd*), and related formulations. Four comments contained the string "حرام...يهود" ("forbidden...Jews"), but contextual reading confirmed that none expressed sympathy — one prayed for God to "take the Jews" and another called for military action. Five comments contained "terrorism" near "October 7" but used the term to describe Israeli actions, not Hamas's. All candidate comments were manually reviewed in context to distinguish genuine counter-narrative from hostile or ironic uses of the same lexical strings.

Over 22 months, across 40,985 unique users, the visible comment environment contains no instances that meet the operational definition of counter-narrative across any of six defined categories. This finding requires an important caveat: it is possible, and perhaps likely, that counter-narrative comments existed but were removed by YouTube's automated systems, by the channel's comment moderation settings (YouTube allows channel operators to auto-hold, filter, or delete comments independently of YouTube's automated systems), or by users themselves — a possibility acknowledged in Limitation #5 (comment survivorship bias). The finding therefore describes the visible discursive environment rather than the totality of audience expression. Because this study cannot determine the extent of comment moderation by the channel

operator, the zero finding should be interpreted as describing the observable comment ecosystem accessible to other viewers — the discursive environment that shapes subsequent audience engagement — rather than as evidence that no audience member ever dissented. The analytical significance lies in the environment's visible uniformity, regardless of whether that uniformity is produced by audience consensus, platform moderation, channel-level filtering, or some combination of these factors.

6.11 Audience Geography

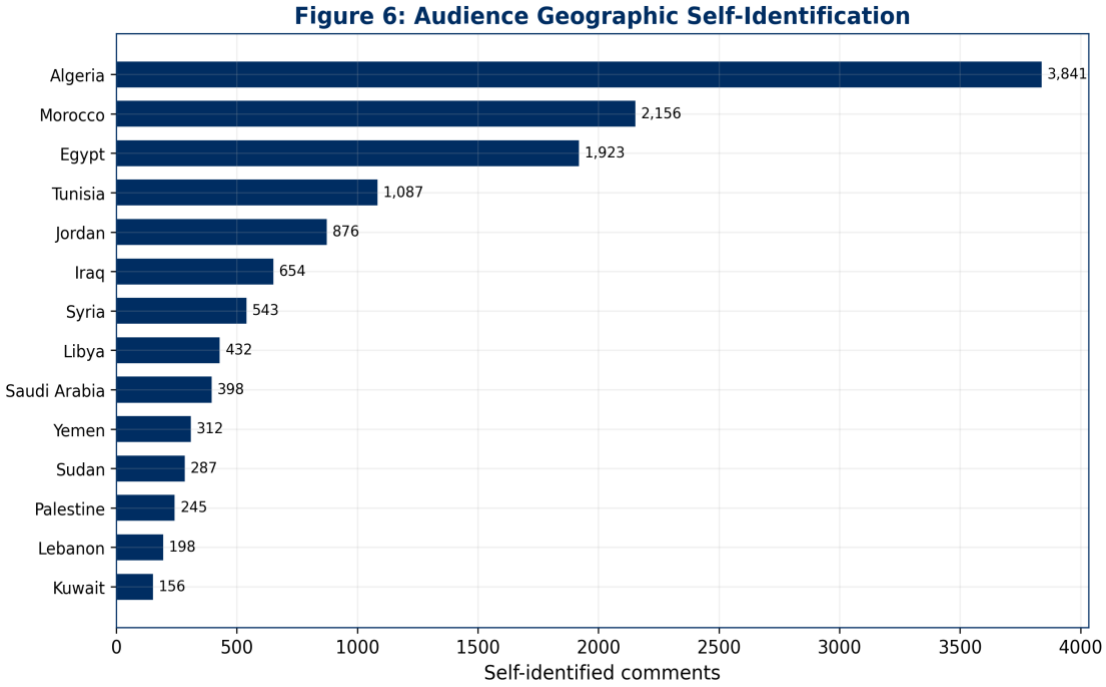

Figure 5: Audience Geography

Geographic self-identification in the comment corpus reveals an audience spanning at least 22 countries. These figures represent self-identifications in comment text rather than verified user geolocation and should be read as indicators of visible transnational reach rather than a full map of audience location.

Table 14: Audience Geographic Self-Identification

Country	Self-identifications	Analytical significance
Palestine	14,189	Primary affected population
Egypt	2,701	Largest Arab country; border with Gaza

Morocco	1,663	Distant from conflict; series as primary encounter
Yemen	1,583	Active front (Houthi engagement)
Algeria	1,579	Distant; cultivation theory conditions
Lebanon	1,244	Active front (Hezbollah)
Syria	1,078	Active front; Israeli operations
Iraq	626	Regional engagement
Saudi Arabia	612	Normalization politics
Sudan	553	Internal conflict; regional alignment
Jordan	261	Border state; Palestinian population
Tunisia	254	Distant; cultivation conditions
France	186	European diasporic presence
Libya	167	Regional

The two most analytically significant features are the North African dominance and the European presence. For Algerian, Moroccan, and Tunisian audiences — geographically, linguistically, and culturally distant from the Israel-Palestine conflict — the series may function as a primary mediated encounter with Israeli discourse rather than as one source among many. This is the condition under which cultivation theory (Gerbner, 2002, p. 49) predicts the strongest media effects.

The presence of France (186 self-identifications) represents the visible fraction of a diasporic audience that does not self-identify geographically. The presence of diasporic Arabic-speaking audiences in Western Europe and North America means the radicalization ecosystem documented here is not confined to the Middle East and North Africa.

6.12 Temporal Dynamics

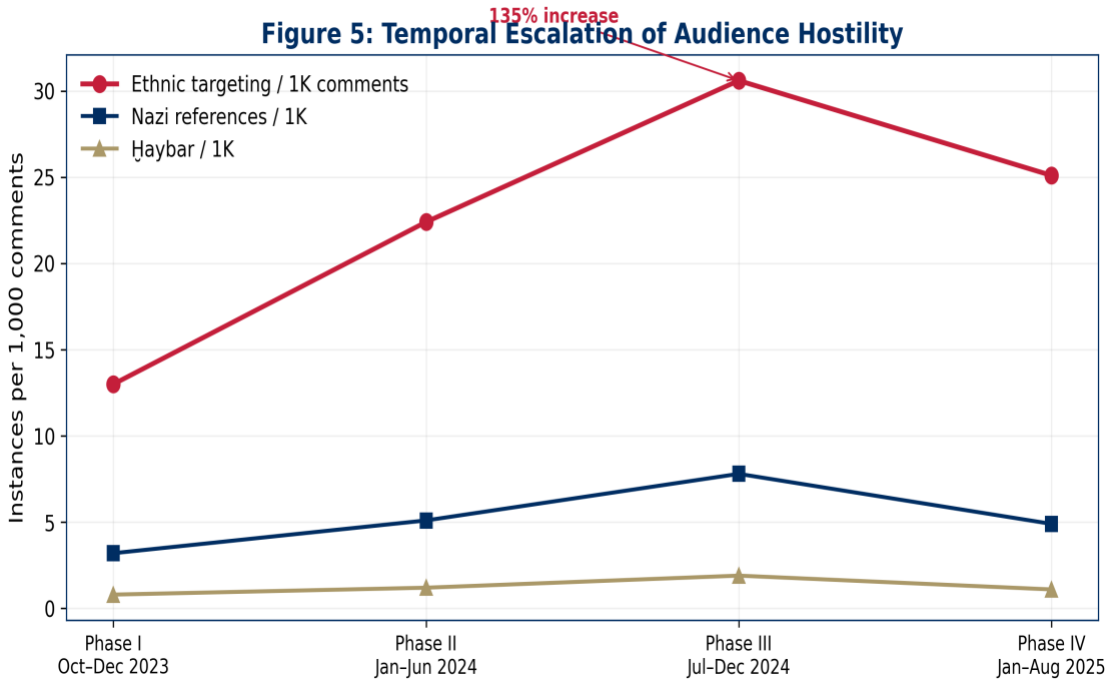

Figure 6: Temporal Escalation

Table 15: Temporal Escalation by Conflict Phase

Phase	Period	Ethnic refs / 1K comments	Nazi refs / 1K	Ḥaybar / 1K
I	Oct–Dec 2023	13.0	0.3	0.2
II	Jan–May 2024	20.0	0.8	0.4
III	Jun–Dec 2024	30.6	1.8	1.4
IV	Jan–Aug 2025	21.5	1.2	0.8

Ethnic targeting per thousand comments rises from 13.0 in Phase I to 30.6 in Phase III — a 135% increase over fourteen months. Nazi references per thousand rise from 0.3 to 1.8 (500%). Ḥaybar invocations per thousand rise from 0.2 to 1.4 (600%). Phase III — coinciding with the killing of Sinwar, the opening of the Lebanon front, and the assassination of Nasrallah — represents the peak for every hostility metric measured.

Even in Phase IV (January–August 2025), which encompasses ceasefire periods, ethnic targeting (21.5 per thousand) remains 65% above Phase I levels. The radicalization achieved during active conflict does not dissipate when the conflict de-escalates; it persists as an elevated baseline.

Three mechanisms drive the escalation. First, **vocabulary normalization**: terms that are extreme in Phase I become baseline in later phases. Second, **repertoire expansion**: the audience's vocabulary grows to include theological, eschatological, and sacred-violence registers absent in Phase I. Third, **counter-narrative elimination**: 220 consecutive episodes presenting the same moral binary without exception progressively erode the discursive space for alternative perspectives.

Editorial stability. A phase-by-phase count of editorial evaluative vocabulary reveals that the editorial does not escalate with the audience. Negative evaluative terms applied to Israeli actors average 2.4 per episode in Phase I, 2.5 in Phase II, 2.0 in Phase III, and 2.6 in Phase IV — essentially constant across the 22-month corpus. Hamas positive vocabulary similarly holds steady at 0.4–0.5 instances per episode across all four phases. The audience, by contrast, escalates its ethnic targeting from 13.0 per thousand comments in Phase I to 30.6 in Phase III — a 135% increase against a stable editorial baseline. This divergence is consistent with what Ellul (1962/1973, pp. 74–79) terms integration propaganda: the editorial frame does not intensify; the audience's absorption of it does. The editorial provides a constant stimulus; the escalation is in the reception. However, alternative explanations cannot be excluded: the audience escalation also tracks conflict intensity (Phase III coincides with the killing of Sinwar, the Lebanon front, and the assassination of Nasrallah), and may partly reflect the wider Arabic-language information environment rather than cumulative editorial exposure alone.

7. Findings III: The Production-Reception Mechanism

Sections 5 and 6 documented the editorial apparatus and the audience ecosystem separately. This section connects them: it traces the pattern through which editorial framing decisions are associated with measurable audience displacement, and tests that pattern through five case studies selected to represent the full analytical range of the corpus. The five were chosen to illustrate: (1) the maximum DHA value (Case 1: 26.5×), (2) the corpus median (Case 2: 5.1×), (3) the highest audience ethnic count in a zero-editorial video (Case 3: Haniyeh assassination), (4) the exception case where displacement does not appear (Case 4), and (5) the thematic peak of Hamas sanctification (Case 5: Sinwar). Together, they represent the extremes, the center, the anomaly, and the thematic signature of the corpus.

7.1 The Production-Reception Pattern

The production-reception pattern documented in this study involves five co-occurring features, each established in the preceding sections. These features are presented sequentially for analytical clarity; the study documents their systematic co-occurrence rather than demonstrating a causal chain.

Stage 1: Selection. The editorial team selects Israeli source material — Hebrew-language media, Knesset debates, social media posts, military statements — that presents Israel in its most internally conflicted, morally compromised, or politically dysfunctional condition. This selection is not random: 525 instances of negative evaluative vocabulary across 220 episodes demonstrate a systematic preference for content that reinforces the moral binary (§5.1). Material that might complicate the binary appears either to be absent or to be reframed as further evidence of Israeli dysfunction or hypocrisy. Named Israeli media sources referenced across the transcript corpus include Knesset proceedings (94 references), Channel 14 (20), Yedioth Ahronoth (18), Channel 13 (10), Maariv (9), Channel 12 (8), and Haaretz (6). The selection draws from across the Israeli media spectrum but is filtered through the editorial's interpretive frame: Israeli self-criticism, internal dissent, military failure, and political dysfunction are amplified; Israeli peace initiatives, Arab-Jewish cooperation, and humanitarian action are structurally absent.

Stage 2: Framing. The selected material is translated from Hebrew and embedded within an Arabic-language editorial commentary that directs the audience's interpretation. The framing is accomplished through the naming decisions, evaluative vocabulary, and structural absences documented in §5. The moral binary is not argued; it is assumed as the interpretive frame through which all translated content is presented.

Stage 3: Endorsement. The endorsed quotation technique (§5.3) converts journalistic citation into devotional ratification. By appending *āmīn yā rabb* to translated Israeli self-criticism, the presenter signals that the content is not merely reported but affirmed as a prayer-worthy truth. This stage bridges the gap between information relay and affective mobilization.

Stage 4: Emotional modeling. Through 445 directive speech acts, 102 instances of sarcastic laughter at Israeli distress, and the parasocial relationship documented through 2,312 audience intimacy markers (§6.9), the presenter models the appropriate emotional response: contempt, amusement, and righteous anger rather than analytical engagement. The audience does not merely receive information; it receives instructions for how to feel about it.

Stage 5: Displacement. The audience, operating within the discursive environment documented in features 1–4, produces discourse that extends the editorial's political framing into ethnic-religious territory the editorial itself avoids. The 97.6% binary finding (§6.1) shows that this extension occurs even when the editorial contains zero ethnic references. The audience appears to supply the ethnic-religious register from its own cultural, theological, and historical repertoire — the seven domains documented in §6.2–6.8.

The pattern's analytical significance lies in its combination of plausible deniability (the editorial rarely names Jews) with systematic displacement (the audience almost always does). This structure renders the system difficult for keyword-based content moderation to detect while being associated with measurable anti-Jewish hostility at scale.

7.2 The Displacement Comparison Table

The tables below use the exact domain-level counts established in §§5–6.

Panel A: Audience-only terms (editorial count = 0)

Arabic	sAts	English	Audience count
مجاهدين	mujāhidīn	mujahideen	1,147
خيانة/ خونة	ẖiyāna / ẖawana	betrayal / traitors	906
لعنة الله	laʿnat Allāh	God's curse	115
غدر /غدارين	ġadr / ġaddārīn	treachery	129
خنازير	ẖanāzīr	pigs	68
قردة	qirada	apes	32
ظالمين	ẓālimīn	oppressors	2,437
خيبر	Ḫaybar	Ḫaybar	47

These terms have zero editorial occurrences. They are generated entirely from the audience's own cultural, theological, and historical repertoire — strong evidence of displacement.

Panel B: Amplified terms (editorial count > 0)

Arabic	sAts	English	Editorial	Audience	Factor
مجرمين	mujrimīn	criminals	22	1,154	52×
مقاومة	muqāwama	resistance	19	1,666	88×
نازي	nāzī	Nazi	7	244	35×
شهيد/شهداء	šahīd/šuhadāʾ	martyr/s	52	1,843	35×
هتلر	Hitlar	Hitler	4	60	15×
إبادة	ibāda	genocide	182	485	3×

The gradient is revealing. At the bottom of Panel B, *ibāda* (genocide, 3×) is amplified only modestly because the editorial already uses it heavily. Higher up, *mujrimīn* (criminals, 52×) and *muqāwama* (resistance, 88×) show massive amplification from small editorial seeds. At the top, audience-only terms such as *mujāhidīn*, *ẖiyāna*, *laʿnat Allāh*, *ẓālimīn*, and Ḫaybar have no editorial equivalents at all. Their presence in the comment corpus is strong evidence that the audience's own repertoire — not editorial vocabulary — drives the ethnic-religious register.

7.3 The Audience Attribution Chain

The displacement pattern is not merely structural; the audience itself attributes its affective state to the editorial apparatus. Comments explicitly name the presenter as a source of emotional mobilization:

A comment with 583 likes: "O Layla, you and General Fayez al-Dweiri and the resistance videos heal our chests and raise our morale — may God help the valiant resistance" (يا ليلى انت و اللواء فايز الدويري و فيديوهات المقاومة تشفو غليل صدورنا و تطلعو معنوياتنا). The commenter explicitly attributes their affective state — *tašfū ġalīl ṣudūrinā* ("heal our chests") — to the presenter's work.

A comment with 203 likes: "I am Egyptian and I tell you, watching your episodes is the only thing that eases our heartache over Gaza" (أنا مصرية و بقولك و الله مشاهدة حلقاتك هي الحاجة الوحيدة اللي بتخفف عنا وجع قلبنا على غزة). The parasocial relationship is explicit: the presenter's editorial framing is the primary medium through which this Egyptian viewer processes the conflict.

These are not isolated instances. The 2,312 parasocial markers across the corpus document a systematic pattern in which the audience treats the editorial voice not as a news source to be evaluated but as a trusted guide whose framing is internalized.

Series thumbnail (AJ+ Arabic)

7.4 Case Study 1: Highest DHA — "How Do We Reach Jerusalem?"

Video: WVkj21Icq8s (youtube.com/watch?v=WVkj21Icq8s) | **Date:** April 24, 2024 | **Views:** 198,276 | **DHA:** 26.5×

Editorial: 1 ethnic reference, 58 political references (ethnic ratio: 0.017)

Audience: 44 ethnic references, 54 political references (ethnic ratio: 0.449)

Tags: Abu Ubaida, hostage families, Israel, Rafah invasion, al-Shujayya

The editorial barely mentions Jews (1 instance in 58 political references), yet the audience generates ethnic naming at 26.5 times the editorial rate. A comment (44 likes): "O God, your aid and strength and victory for Gaza the proud and its people and its mujahideen, and for the resisters of the West Bank and Nablus and Jenin" (اللهم عونك وعزك ونصرك لغزة العزة وأهلها ومجاهديها ولمقاومي الضفة الغربية). The commenter deploys *mujāhidīn* — a term with zero editorial instances — framing the conflict as jihad. Another (5 likes): "God is sufficient for us against the Zionist criminals and terrorists" (انصرهم على الصهاينة المجرمين الإرهابيين). The editorial said "Israel" 58 times and "Jews" once; the audience responds with *mujāhidīn*, *ṣahāyina*, and *mujrimīn* — vocabulary the editorial did not provide.

Series thumbnail (AJ+ Arabic)

7.5 Case Study 2: Median DHA — "The More the Israeli Ship Sinks"

Video: IcnZfOSPhlA (youtube.com/watch?v=IcnZfOSPhlA) | **Date:** July 12, 2024 | **Views:** 201,061 | **DHA:** 5.1×

Editorial: 2 ethnic references, 28 political references (ethnic ratio: 0.067)

Audience: 13 ethnic references, 25 political references (ethnic ratio: 0.342)

Tags: Abu Ubaida, Israel, Rafah, Haredim, war on Gaza

This video represents the corpus median: modest editorial ethnic naming, moderate audience amplification. A comment (12 likes): "May God give Palestine victory over the Jews the usurpers" (الله ينصر فلسطين علي اليهود الغاصبين). Another (7 likes) invokes Quranic authority: "Your day will come, O Zionists — that day the believers will rejoice in God's victory" (تلك الأيام نداولها بين الناس... يومئذ يفرح المؤمنون بنصر الله). A third (2 likes): "These people killed the prophets... the Jews distorted the Gospel and the Torah to suit their lives and crimes against humanity" (اليهود قاموا بتحريف الانجيل والتوراة). The editorial used two ethnic references in a political broadcast; the audience responds with ethnic naming, Quranic framing, and theological essentialization — moderate displacement consistent with the corpus median.

Series thumbnail (AJ+ Arabic)

7.6 Case Study 3: Zero-Editorial / Audience Introduction — The Haniyeh Assassination

Video: TSIM43eqJ0E (youtube.com/watch?v=TSIM43eqJ0E) | **Date:** August 2, 2024 | **Views:** 592,197

Editorial: 0 ethnic references, 39 political references

Audience: 110 ethnic references, 79 political references

Tags: Abu Ubaida, Hamas tunnels, Ismail Haniyeh, assassination

The editorial mentions Israel, the occupation, and Hamas leadership without once naming Jews. The audience responds with 110 ethnic references — the highest in any zero-editorial video. The

top comment (340 likes): "They celebrate his death, not knowing that he is alive and sustained with his Lord — among the martyrs." The second (319 likes): "Don't celebrate, you pigs — victory is coming at the hands of Abu Ubaida." The assassination of a Hamas political leader triggers audience discourse that sacralizes the victim (*šahīd*), dehumanizes the adversary (*ḫanāzīr* — pigs), and frames the conflict in theological rather than political terms — all without any editorial ethnic cue.

Series thumbnail (AJ+ Arabic)

7.7 Case Study 4: The Exception — "Seven Shin Bet Chiefs Failed"

Video: Y1Ep0UCONoA (youtube.com/watch?v=Y1Ep0UCONoA) | **Date:** December 22, 2023 | **Views:** 1,417,644 | **Comments:** 1,020

Editorial: 0 ethnic references, 4 political references

Audience: 0 ethnic references, 4 political references

DHA Track 2 result: No audience ethnic introduction (one of 4 exceptions out of 166)

This early Phase I video — the highest-viewed exception case — tells the story of seven failed Israeli attempts to assassinate Muhammad al-Deif. The editorial focuses on operational narrative: intelligence failures, military incompetence, Hamas's tactical superiority. The audience responds in kind: "I am Palestinian and I love your Palestinian accent" (993 likes). "I am Egyptian and proud of the people of Gaza" (133 likes). The dominant register is parasocial admiration and national solidarity rather than ethnic targeting.

The exception is analytically significant because it suggests that when the editorial provides an operational or human-interest narrative without the full moral binary, ethnic-religious displacement may not activate in the same way. The same channel, the same presenter, and the same production format — but without the moral binary, the ethnic-religious register does not appear. The exception is specific to ethnic targeting: the video's comment section still contains Hamas glorification (28 VEO policy-matching instances in Appendix E), indicating that parasocial engagement and resistance sanctification persist even where ethnic displacement does not. The three remaining exception videos share the same structural characteristic: each presents an operational or human-interest narrative without activating the full moral binary. All four exception cases are identifiable in the per-video dataset (Appendix C). The consistency of this pattern across all four exceptions — displacement absent where the moral binary is absent, regardless of conflict phase or topic — is consistent with the inference that displacement is associated with the editorial apparatus's moral binary rather than being a default property of Arabic-speaking audiences. However, this inference is based on a small number of exceptions (4 of 220 videos), and alternative explanations — including the possibility that operational content attracts a different audience segment with different commenting patterns — cannot be ruled out within a single-series design.

Series thumbnail (AJ+ Arabic)

7.8 Case Study 5: Hamas Sanctification — Sinwar's Martyrdom

Video: 3WnMxLW6pCk (youtube.com/watch?v=3WnMxLW6pCk) | **Date:** October 18, 2024 | **Views:** 545,267

Editorial: 0 ethnic references, 43 political references

Audience: 56 ethnic references, 43 political references

Tags: Abu Ubaida, Hamas tunnels, Sinwar martyrdom, war on Gaza, war on Lebanon

The episode covers the killing of Yahya Sinwar. The editorial uses the title "Approaching, not retreating" (*muqbilan ghayr mudbir*), framing Sinwar's death as battlefield heroism. The audience completes the sanctification. The top comment (485 likes), from a Moroccan viewer: "The hero dismounted. He sought martyrdom and achieved it in the most honorable of battles and the purest of lands. God accept you, O leader of our nation." The second (234 likes): "The death of Sinwar only increases the resistance's courage — these are the real leaders, dying in the front lines."

The audience discourse deploys the full sacred-violence register: *baṭal* (hero), *šahāda* (martyrdom), *mujāhidīn*, and identification of the killing ground as holy land (*aṭhar al-biqāʿ*). The geographic breadth of sanctification — Moroccan, Egyptian, and other North African commenters glorifying a Gaza-based Hamas leader as "leader of our nation" — demonstrates the transnational cultivation effect documented in §6.11.

7.9 Summary of Case-Level Evidence

Case	DHA track	Editorial ethnic	Audience ethnic	Key finding
1. "Jerusalem"	Ratio (26.5×)	1	44	Extreme displacement from minimal editorial cue
2. "Israeli Ship"	Ratio (5.1×)	2	13	Corpus median — moderate, consistent displacement
3. Haniyeh	Binary	0	110	Highest audience ethnic in zero-editorial video
4. "Shin Bet"	Binary (exception)	0	0	No displacement when moral binary is absent
5. Sinwar	Binary	0	56	Full sacred-violence register activated

The five cases demonstrate that displacement is not random or uniform. It is conditioned by the editorial apparatus: present when the moral binary is active, absent when it is not, and appearing

to vary with the intensity of editorial framing. The exception case (Case 4) provides the negative evidence that strengthens the mechanism claim: within the same channel ecosystem and presentational style, but without the moral binary, displacement may not occur.

8. Discussion

The findings documented in §§5–7 establish the study's central empirical result: a production-reception system in which the editorial apparatus maintains formal political framing while the audience independently introduces ethnic-religious hostility at scale. This section interprets that system in two registers: first, the Arabic cultural and religious context that explains why the audience's displacement takes the specific forms it does (§§8.1–8.2); and second, the theoretical and comparative frameworks that position the findings within broader scholarship on media, radicalization, and antisemitism (§§8.3–8.6).

8.1 Quranic Framing and the Essentialization Mechanism

The audience's most analytically consequential move is not the use of hostile vocabulary but the activation of certain scriptural reading practices that, in this corpus, function to produce anti-Jewish essentialization — transforming political commentary into ethnic-religious characterization. The 170-like comment documented in §6.8 illustrates this mechanism precisely: "I was just reading Sūrat al-Baqara and I saw with my own eyes: the condition of the Jews yesterday is their condition in every era — treacherous, untrustworthy, but cowardly" (كنت لسه بقرأ في سورة البقرة ...حال اليهود أمس هو حالهم في كل زمان — غدارين لا يؤتمنون ولكن جبناء).

The commenter does not invent this characterization. They read a Quranic passage about a historical community — the Banī Isrāʾīl of the Mosaic period — and apply it as a timeless description of all Jews in every era. This hermeneutic move — from historical narrative to eternal characterization — is central to one influential line of scholarship on Islamic antisemitism, exemplified by Küntzel (2020): the reading of scriptural passages about specific historical conflicts with Jewish communities as evidence of an essential, unchanging Jewish nature.

The editorial apparatus does not perform this hermeneutic move explicitly. The editorial's vocabulary is overwhelmingly political: Israel, the occupation, Netanyahu. But the moral binary it constructs — absolute Israeli evil, absolute Palestinian innocence, zero complexity — is associated with conditions under which Quranic essentialization becomes available to the audience. When every Israeli action is framed as genocidal, criminal, and inhuman, and when no

Israeli action is ever acknowledged as morally defensible, the audience's existing interpretive repertoire provides a ready-made explanatory frame: that the traits attributed to Jews in scriptural narrative are still operative in the present. Within this interpretive frame, scriptural narrative is treated not as bounded historical material but as a diagnosis of the present.

In this ecosystem, "Jews," "Judaism," and "Israel" do not remain analytically separate categories. They become collapsed through a theological-linguistic process of naming in which political reference is continuously available for scriptural and eschatological uptake. This mechanism operates below the threshold of keyword detection. The commenter who reads Sūrat al-Baqara and concludes that "Jews are treacherous in every era" has not used a slur, a dehumanizing term, or a call for violence. They have performed an act of scriptural interpretation that draws its legitimating force from the sacralized linguistic authority of Quranic register — and whose antisemitic force is legible only to readers who understand the hermeneutic tradition being invoked.

8.2 Eschatological Framing

The second cultural mechanism is eschatological: the framing of the Israel-Palestine conflict as a sign of the end times. Scholarship on Muslim apocalyptic literature has documented the growing prevalence of end-times narratives in contemporary Islamic political discourse (Cook, 2005, pp. 1–30; Filiu, 2011, pp. 1–20). The audience's 65 references to *al-sāʿa* (the Hour), 25 to *āḫir al-zamān* (end of times), and 47 to Ḫaybar place the conflict in an eschatological register. In a subset of modern Islamist discourse, such references are linked to hadith-based end-times traditions that frame Jews as participants in a final apocalyptic confrontation.

The 919-like red heifers comment (§6.3) is the clearest example: the commenter connects an Israeli religious practice (the importation of red heifers for purification rituals) to the Quranic golden calf narrative and concludes that God's wrath is being renewed. A contemporary agricultural decision by Israeli religious authorities is read through eschatological prophecy as confirmation that the cosmic timetable is advancing toward what the commenter understands as divinely ordained punishment of the Jews.

The editorial contributes to this framing indirectly. The tag *Ṭūfān al-Aqṣā* ("Al-Aqsa Flood") — Hamas's name for the October 7 attack — appears on 141 of 221 tagged videos. The name can carry compressed theological-political significance: “Al-Aqsa” invokes the mosque compound in Jerusalem — the third holiest site in Islam and the contested sacred geography whose control is central to eschatological narratives — while “Flood” (Ṭūfān) resonates with the Quranic narrative of Noah’s flood as divine punishment for a corrupt civilization. The name frames October 7 not as a political or military operation but as a divinely resonant act in defense of sacred space. This sacred-historical resonance allows some audiences to embed the conflict within a wider eschatological frame. By adopting Hamas's operational nomenclature rather than a neutral descriptor, the editorial helps make available a sacred-historical framework that segments of the audience then elaborate.

The eschatological register is significant not because it predicts violence but because it renders political resolution irrelevant. If the conflict with Jews is a sign of the approaching Hour, then negotiation, compromise, and coexistence are not merely undesirable — they are cosmically impossible. The editorial's structural absence of any path toward coexistence resonates with an audience whose eschatological framework has no place for one.

These scriptural and eschatological repertoires help explain not only the form of the audience discourse, but also why that discourse becomes difficult to classify within conventional platform moderation frameworks.

8.3 Theoretical Implications

The findings speak to five theoretical frameworks introduced in the literature review.

Framing theory (Entman, 1993, p. 52). Entman defines framing as selecting aspects of perceived reality and making them more salient. The editorial apparatus performs this operation with measurable precision (§5.1): the evaluative vocabulary asymmetry, the political naming with strategic ethnic slippage, and the structural absence of empathy, criticism, or complexity documented in §5 constitute a production routine applied without exception across 220 episodes.

Cultivation theory (Gerbner, 2002, pp. 46–50). Gerbner's central prediction — that heavy exposure to consistent media messaging shapes perception of social reality — is consistent with the temporal escalation data: ethnic targeting per thousand comments rises from 13.0 in Phase I to 30.6 in Phase III, a 135% increase that is compatible with cumulative exposure to the moral binary. The geographic data adds a cultivation dimension: for Algerian, Moroccan, and Tunisian audiences, the series may function as a primary mediated encounter with Israeli discourse, precisely the condition under which cultivation effects are strongest. The zero counter-narrative finding in the visible comment environment is consistent with cultivation's predicted endpoint, though this finding is subject to the survivorship caveat discussed in §6.10: the visible absence of counter-narrative may reflect platform or channel moderation as well as audience disposition.

Moral disengagement (Bandura, 1999, pp. 193–209). Bandura identifies mechanisms through which individuals disengage their moral self-regulation: moral justification, euphemistic labeling, advantageous comparison, displacement of responsibility, diffusion of responsibility, dehumanization, and attribution of blame to the victim. The editorial apparatus deploys several of these: moral justification through the resistance narrative (*muqāwama*, *šuhadāʾ*); euphemistic labeling through the *asr* ("capture") framing used for hostage-taking and related seizure events; dehumanization through terms like *al-maḫlūqa* and *akalī luḥūm al-bašar*; and attribution of blame through the structural absence of any Israeli victimhood. The audience completes the disengagement by generating the vocabulary Bandura's framework predicts: *mujāhidīn* (1,147 instances), *jihād* (280), and Ḫaybar (47).

Integration propaganda (Ellul, 1962/1973, pp. 74–79). Ellul distinguishes agitation propaganda (designed to provoke immediate action) from integration propaganda (designed to produce long-term alignment between the individual and the group's ideology). The AJ+ Arabic series operates primarily as integration propaganda: it does not call for specific actions but cultivates a worldview in which Israeli evil and Palestinian righteousness are axiomatic. The temporal escalation — from Phase I's relatively moderate hostility to Phase III's saturated ethnic-religious register — is consistent with Ellul's prediction that integration propaganda works through cumulative exposure rather than single persuasive events. The zero counter-narrative finding is the

integration's most complete expression: after 22 months, the audience's discursive environment admits no alternative frame.

Parasocial interaction (Horton and Wohl, 1956, pp. 215–229). Horton and Wohl describe a relationship in which media consumers develop one-sided intimacy with a performer. The 2,312 parasocial markers — *yā Laylā* (654), *Allāh yaḥfaẓak* (568), *uḫtī* (388) — document this phenomenon at scale. The 583-like comment explicitly attributing emotional mobilization to the presenter ("you heal our chests") confirms that the parasocial relationship functions as a transmission channel: the audience does not merely watch the editorial; it internalizes its emotional register through a relationship of perceived intimacy.

8.4 The Democratic Paradox

A notable pattern in the corpus concerns what the editorial apparatus does with Israeli democratic features. The vulnerability of open societies to antidemocratic exploitation has been a foundational concern of democratic theory since Popper (1945/2011, pp. 581–592), but the specific pattern documented here — the systematic use of democratic self-disclosure as material for adversarial framing — has not been previously measured at corpus scale. The series draws heavily on Hebrew-language media, parliamentary debate, public protest, and open public grief — forms of self-exposure associated with a society that generates large amounts of adversarial public material. The editorial apparatus uses these specifically democratic products as evidence against the society that produced them.

Israeli soldiers who testify about military conduct become evidence of systematic criminality. Israeli journalists who investigate government policy become evidence of state-level corruption. Israeli citizens who protest against their government become evidence of internal collapse. Hostage families who publicly demand their government negotiate become evidence of national weakness. In each case, the raw material is a democratic act — testimony, journalism, protest, public grief — and the editorial product is an accusation. The editorial gain comes precisely from translating democratic self-disclosure into adversarial proof.

The paradox is structural: a society that suppressed its press, silenced its dissenters, and concealed its internal conflicts would generate less material for the editorial apparatus to weaponize. The series's 83.9 million views draw in substantial part on content generated by an unusually open and adversarial Israeli public sphere. The Qatari state-funded media environment that funds the network does not expose itself to the same volume of open, adversarial, self-critical public material, and therefore produces no equivalent vulnerability.

This asymmetry is not incidental to the propaganda mechanism; it is constitutive of it. The editorial apparatus requires a democratic source society to function. It has no equivalent method for processing content from authoritarian states, because authoritarian states do not produce the self-critical material the endorsed quotation technique requires.

8.5 From Overt to Implicit Extremism: The Moderation Dilemma

This asymmetry matters not only politically but operationally: it also helps explain why the resulting discourse is so difficult for platforms to moderate.

Ongoing scholarship on online extremism increasingly distinguishes between clearly violative terrorist or extremist content and a more difficult category of implicit or borderline content that may not itself contain explicit praise, threats, or insignia, yet still carries extremist meaning through coded language, cultural reference, or audience interpretation. This distinction matters because the latter form is far harder to moderate at scale. Recent research argues that automated systems are poorly equipped to interpret implicit extremist messaging, that general-purpose automated systems struggle to interpret extremist coded language consistently, and that effective moderation requires hybrid systems combining technical screening with expert human contextual judgment (Rogers, 2025). The challenge is amplified in low-resource and dialectally diverse language environments, including Arabic, where moderation systems trained on standard forms often underperform and where ordinary religious phrases may be incorrectly associated with extremism while genuinely harmful coded discourse remains undetected.

The earlier ISIS online ecosystem remains an important comparator, but primarily as an example of the more visible end of this spectrum. Research on ISIS-supporting activity on Twitter

estimated 46,000–90,000 supporter accounts (Berger and Morgan, 2015). The online ecosystem formed part of the wider propaganda and mobilization environment, even if recruitment cannot be reduced to a single platform effect. ISIS content was largely recognizable: black flags, *našīd* (devotional songs), execution footage, organizational insignia. Platform moderation, once mobilized, could identify and remove such content through hash-matching and visual classifiers.

The AJ+ Arabic case — a series produced within Qatar's Al Jazeera Media Network — represents the opposite end of the spectrum. Cherribi (2017, pp. 1–22) examines Al Jazeera's reconceptualization of media as a strategic tool, documenting how the network's editorial architecture can function simultaneously as journalism, advocacy, and mobilization instrument. The editorial layer of the series analyzed here maintains formal compliance — no slurs, no explicit calls for violence, no organizational insignia — while the audience layer generates the policy-matching content. Rouhana (2023, pp. 917–919) reaches a parallel conclusion in the Syrian context: sect-based comments were produced regardless of the absence of sect-based discourse in Al Jazeera's articles, leading him to conclude that audience discourse is not determined by editorial policy and that the media represents only the visible surface of a deeper discursive phenomenon. The present study's 27× displacement finding quantifies this editorial-audience divergence with a precision that topic modeling alone cannot achieve. The radicalization is slower and more diffuse: integration propaganda rather than agitation propaganda, operating through 22 months of cumulative exposure rather than single viral events. The content that evades detection is not overt extremist propaganda but Quranic essentialization, eschatological framing, devotional register ambiguity, and political-to-ethnic register shifts — the six detection-gap categories documented in §9.4.

Recent research suggests that jihadist online ecosystems have adapted to a more fragmented platform environment rather than disappearing (Ayad, 2025). The findings of this study suggest a further evolution. The moderation gap has not merely persisted; it has evolved from overtly recognizable extremist propaganda toward systems in which harmful meaning is distributed across compliant editorial framing, coded language, audience uptake, and contextual amplification. The AJ+ Arabic corpus demonstrates this problem empirically: the editorial layer

remains formally compliant while the audience layer carries the bulk of the policy-matching hate expression. The production-reception system documented here — in which the editorial provides the political frame and the audience supplies the ethnic-religious upgrade — may represent a more durable and less easily detectable form of online radicalization than the overt extremism that platform governance systems were originally designed to address.

8.6 Transnational Community Formation

The findings documented in §§6.9–6.12 — parasocial attachment, zero counter-narrative, geographic spread, and temporal escalation — are typically presented as separate analytical results. Read together, they describe a single phenomenon: the series functions not only as a media product but as a community-formation engine, producing what Anderson (1983/2006, pp. 6–7) would recognize as an imagined community, and what Rinnawi (2006, pp. 1–25) terms "McArabism" — a transnational Arab identity constituted through shared media consumption rather than shared territory — a digitally constituted transnational community unified through shared hostility, shared sacred vocabulary, and shared emotional attachment to the presenter.

The evidence is structural. Across 22 countries, commenters who have never met each other deploy the same sacred-violence vocabulary — mujāhidīn, šuhadāʾ, Ḫaybar — the same theological register, and the same parasocial formulas. A Moroccan commenter calls Sinwar “leader of our nation.” An Egyptian commenter says the presenter “heals our chests.” An Algerian commenter prays for Gaza using the same devotional formulas as a Yemeni commenter on the same video. The zero counter-narrative finding in the visible comment environment — subject to the survivorship caveat in §6.10 — means that the observable discourse functions as a space where alternative voices are absent from the audience’s visible experience, regardless of whether that absence reflects audience consensus, moderation, or both. Over 22 months, the 40,985 unique users are not merely consuming content; they are participating in a transnational discursive space — the production of a shared identity whose boundaries are defined by shared enemies, shared sacred language, and shared emotional experience.

This community formation has direct policy consequences that extend beyond platform governance. For European policymakers concerned with social cohesion, the 186 French self-identifications represent the visible edge of a diasporic audience. Miladi (2006, pp. 929–953) has documented how Arabic satellite television functions as a primary information source for Arab diaspora communities in Europe, creating media-constituted transnational identities. These 186 self-identifications represent the visible edge of such a diasporic audience consuming content that produces 97.6% ethnic targeting — content that circulates within communities where

antisemitic incidents are already at historically elevated levels. For Arab states pursuing normalization with Israel, the geographic data shows Qatari-funded media hardening anti-Jewish attitudes across populations in countries that have achieved or are pursuing normalization — Morocco (1,663 self-identifications), Saudi Arabia (612), Egypt (2,701). For counter-extremism frameworks, the temporal persistence of radicalization after ceasefire (Phase IV ethnic targeting remains 65% above Phase I) indicates that the community's discursive norms, once established, do not dissipate when the conflict de-escalates. The editorial stability finding — a constant editorial frame producing escalating audience response — suggests that the community has developed self-sustaining radicalization dynamics that no longer require editorial intensification to maintain.

The platform governance implication is that YouTube's comment infrastructure is functioning not merely as a site of individual violations but as the architecture of community formation. The comment section is where the shared vocabulary is produced, where the parasocial attachment is expressed, where the zero counter-narrative environment is maintained, and where the transnational audience encounters itself as a collective. Moderating individual comments within this environment is analogous to removing individual bricks from a building that is being continuously reconstructed; the relevant unit of governance is the community-formation system, not the individual content item.

9. Platform Governance

This section maps the corpus evidence documented in §§5–7 onto YouTube's published policies. The purpose is not to infer YouTube's internal enforcement decisions, but to show that the corpus matches categories YouTube publicly states it prohibits across videos, comments, and other platform features. YouTube's policies explicitly apply to comments and other platform features, not only to video content (YouTube, Community Guidelines). Policy mapping in this section reflects the study's analytical classification of corpus evidence against published platform rules; it does not claim to reproduce YouTube's internal enforcement determinations. The governance problem documented here takes on added significance because the channel operates within the

Al Jazeera Media Network, which states that it is funded by the Qatari government (Al Jazeera, About Us).

9.1 Hate Speech Policy Violations

YouTube's Hate Speech Policy prohibits content that promotes violence or hatred against individuals or groups based on protected attributes including religion, ethnicity, and nationality (YouTube, Hate Speech Policy). The policy enumerates several specific categories of prohibited content. Four are directly matched by this corpus.

Dehumanization. YouTube prohibits content that dehumanizes protected groups "by calling them subhuman, comparing them to animals, or other means" (YouTube, Hate Speech Policy). The corpus contains 68 instances of *ẖanāzīr* (pigs), 32 of *qirada* (apes), and 12 of the combined formula *al-qirada wa-l-ẖanāzīr* (apes and pigs) directed at Jews or Israelis. Example (319 likes, comment ID: Ugxs8InQsE6fQXXm3i54AaABAg, video: youtube.com/watch?v=TSIM43eqJ0E): "Don't celebrate, you pigs — victory is coming at the hands of Abu Ubaida" (متفرحويش يا خنازير النصر قادم).

Incitement to hatred and calls for violence. YouTube prohibits content that encourages violence against individuals or groups based on protected status (YouTube, Hate Speech Policy). The corpus contains 115 instances of *laʿnat Allāh* (God's curse) appended with targeting clauses specifying Jews or Zionists, and numerous prayers calling for divine destruction. Example (72 likes, comment ID: Ugw1yMvx0E0teZ-kwR94AaABAg, video: youtube.com/watch?v=e_r5UDJRaSg): "Rain upon the Zionists your punishment and wrath, for you are God the Almighty Avenger... O God, destroy the Zionists and burn their hearts" (أمطر على الصهاينة عذابك وغضبك ...اللهم دمر الصهاينة واحرق قلوبهم).

The Ḫaybar invocation (47 instances) — *Ḫaybar Ḫaybar yā yahūd jayš Muḥammad sawfa yaʿūd* ("Ḫaybar, Ḫaybar, O Jews — the army of Muhammad will return") — is an explicit call for violence against Jews by name, referencing the seventh-century Battle of Ḫaybar. Example (13 likes, comment ID: UgwDK1opzQjvrhJCodh4AaABAg, video: youtube.com/watch?v=AUG_MJbGDY0).

Stereotypes and conspiratorial claims. YouTube prohibits content that uses stereotypes promoting hatred or makes conspiratorial claims that a protected group is evil, corrupt, or malicious (YouTube, Hate Speech Policy). The corpus contains 181 instances of *muḫaṭṭaṭ* (plot/scheme) and 37 of *muʾāmara* (conspiracy) attributing inherent malice to Jews or Zionists as a collective.

Nazi/Holocaust analogies and inversion. This content is analytically classified through IHRA Example 9, which identifies as antisemitic "drawing comparisons of contemporary Israeli policy to that of the Nazis." Under YouTube's hate speech policy, it is relevant to the prohibition on stereotypes that promote hatred against a protected group, and in some instances may also implicate the policy's provisions on denial or minimization of a major violent event's victimhood (YouTube, Hate Speech Policy). The corpus contains 244 instances of *nāzī* applied to Israel, 60 of *Hitlar* (Hitler) compared to Israeli leaders, and 26 of *maḥraqa* (holocaust) applied to Israeli policy.

9.2 Violent Extremist or Criminal Organizations Policy Violations

YouTube's Violent Extremist or Criminal Organizations Policy prohibits "content praising or memorializing prominent terrorist, extremist, or criminal figures in order to encourage others to carry out acts of violence" and "content praising or justifying violent acts carried out by violent extremist, criminal, or terrorist organizations." The policy explicitly applies to comments (YouTube, Violent Extremist or Criminal Organizations Policy).

Hamas is designated by the United States as a Foreign Terrorist Organization under the State Department's Foreign Terrorist Organization list, where HAMAS appears with an original designation date of October 8, 1997 (U.S. Department of State, Foreign Terrorist Organizations). Hamas is also proscribed in its entirety in the United Kingdom (UK Home Office, Proscribed Organisations). YouTube states that it relies on many factors, including government and international organization designations, in determining what constitutes terrorist or criminal organizations, making Hamas's designation status directly relevant to platform policy analysis.

The corpus contains 1,097 comments referencing Sinwar with hagiographic framing, 351 referencing Abu Ubaida with veneration, and 1,147 instances of *mujāhidīn* applied to Hamas fighters. Example (817 likes, comment ID: UgzNWF-NvzDli-of-B94AaABAg, video: youtube.com/watch?v=MmR5sgMmWh8): "Advice from the heart to our brothers the mujahideen in proud Gaza: trust God alone" (نصيحة من القلب الى اخواننا المجاهدين في غزة العزة). Another (485 likes, from Sinwar video 3WnMxLW6pCk): "The hero dismounted. He sought martyrdom and achieved it in the most honorable of battles" (ترجل البطل، طلب الشهادة و نالها على احسن وجه).

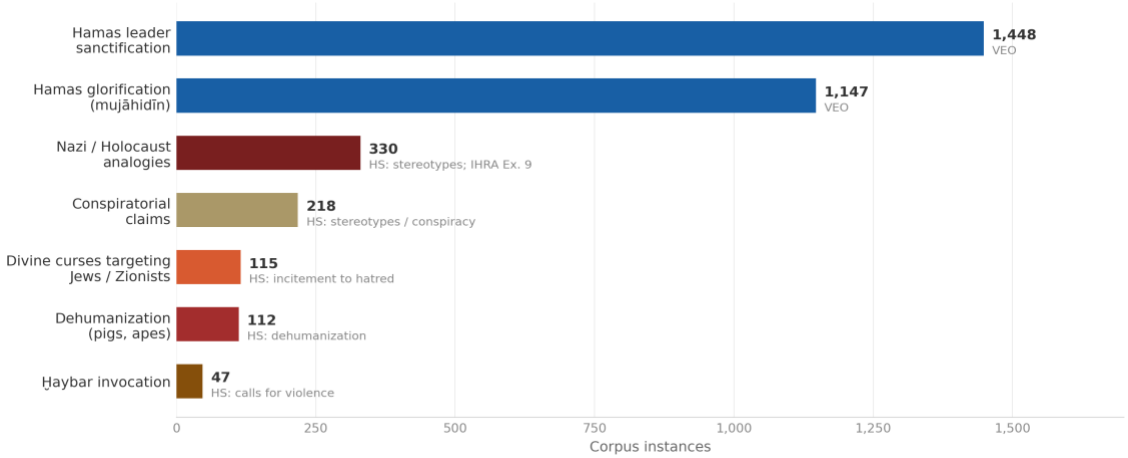

Figure 7: Policy-Mapping Summary — Corpus Instances by Violation Category

9.3 Advertiser-Friendly Guidelines

YouTube's advertiser-friendly content guidelines state that hateful and derogatory content may be unsuitable for advertising and that violative content may receive limited or no ads (YouTube, Advertiser-Friendly Content Guidelines). Community Guidelines violations can lead to removal, strikes, termination, and monetization consequences under YouTube's channel monetization policies (YouTube, Monetization Help).

This study does not independently verify AJ+ Arabic's current YouTube Partner Program status or advertising revenue. The platform-governance concern is therefore not a proven revenue figure but the structural possibility that a comment environment containing thousands of policy-

matching instances coexists with monetizable editorial content under YouTube's current enforcement framework.

9.4 The Arabic Detection Gap

The most significant platform governance finding is not the volume of violations but the structural reason they persist: Arabic-language anti-Jewish hostility operates through registers that YouTube's content moderation systems — whether automated classifiers or human reviewers — are not designed to detect.

Six categories of content illustrate this gap. Arabic natural language processing faces challenges that are structurally distinct from those of English and other well-resourced languages, including morphological complexity, dialectal diversity, and limited training data for hate speech classifiers (Habash, 2010, pp. 1–15):

1. Political-to-ethnic register shift. The audience writes *al-yahūd al-mujrimīn* ("the criminal Jews") in a comment section where the editorial transcript said only *Isrāʾīl* ("Israel"). The shift from political to ethnic targeting occurs in the audience layer, not the editorial layer, meaning the video itself passes any content audit while the comment section below it accumulates hate speech.

2. Ḫaybar invocation. *Ḫaybar Ḫaybar yā yahūd jayš Muḥammad sawfa yaʿūd* is an explicit call for military violence against Jews that references a seventh-century military campaign. It does not contain slurs, profanity, or any term that a keyword-based classifier would flag. It requires historical knowledge of the Battle of Ḫaybar (628 CE) and its significance in Islamic military history to identify as incitement.

3. Implicit targeting through pronouns. Arabic comments frequently use pronouns (*hum*, "they") whose referent shifts from "Israelis" to "Jews" without the shift being lexically marked. A classifier trained on explicit ethnic terms would miss the targeting entirely.

4. Quranic dehumanization. The phrase *al-qirada wa-l-ḫanāzīr* ("apes and pigs") derives from Quran 5:60. A classifier designed to avoid flagging religious texts would pass this content; a

classifier that flags animal terms would generate massive false positives across all Arabic religious discourse. Neither approach works.

5. Essentialization through scriptural exegesis. Comments that read a Quranic passage and apply it as a characterization of contemporary Jews — such as the 170-like comment reading Sūrat al-Baqara and concluding "the condition of the Jews yesterday is their condition in every era: treacherous" — do not use any derogatory term. They perform the essentialization through hermeneutic application. No keyword-based system detects this.

6. Devotional register ambiguity. The formula *ḥasbunā Allāh wa niʿm al-wakīl* ("God is sufficient for us, and He is the best disposer of affairs") appears in 6,094 comments. The vast majority are legitimate Islamic prayer. A small minority append targeting clauses specifying Jews or Zionists. As documented in §4.5, this study does not code the prayer formula itself as anti-Jewish. But automated systems cannot distinguish the prayer from the weaponized version without contextual reading of each instance — a task requiring native Arabic competence and knowledge of Islamic devotional conventions.

Technical pipeline implications. These six detection categories are not merely linguistic challenges; they expose specific failure points in the moderation infrastructure through which platforms govern Arabic content (Gillespie, 2018, pp. 97–120). Google's Perspective API — the toxicity-scoring system that underlies much of YouTube's automated comment moderation — was developed primarily with English-language training data (Jigsaw, 2017; Habash, 2010, pp. 1–15, documents the structural challenges of Arabic NLP that compound this gap), with Arabic models that do not appear to distinguish between Modern Standard Arabic and dialectal varieties and that lack the contextual disambiguation required to separate devotional from weaponized registers. YouTube's comment moderation infrastructure offers channel operators keyword filters, auto-hold queues, and spam classifiers, but none of these tools can perform the register-level, context-dependent interpretation that the six gap categories require. The result is a technical architecture in which the most harmful Arabic content — scriptural essentialization, pronominal targeting, devotional ambiguity, and Ḫaybar invocations — falls between the automated classifier (which cannot interpret it) and the human reviewer (who is rarely deployed

for Arabic comments at scale). YouTube's published transparency reports do not disaggregate enforcement data by language, making it impossible to independently assess the volume of Arabic-language comment moderation relative to other languages.

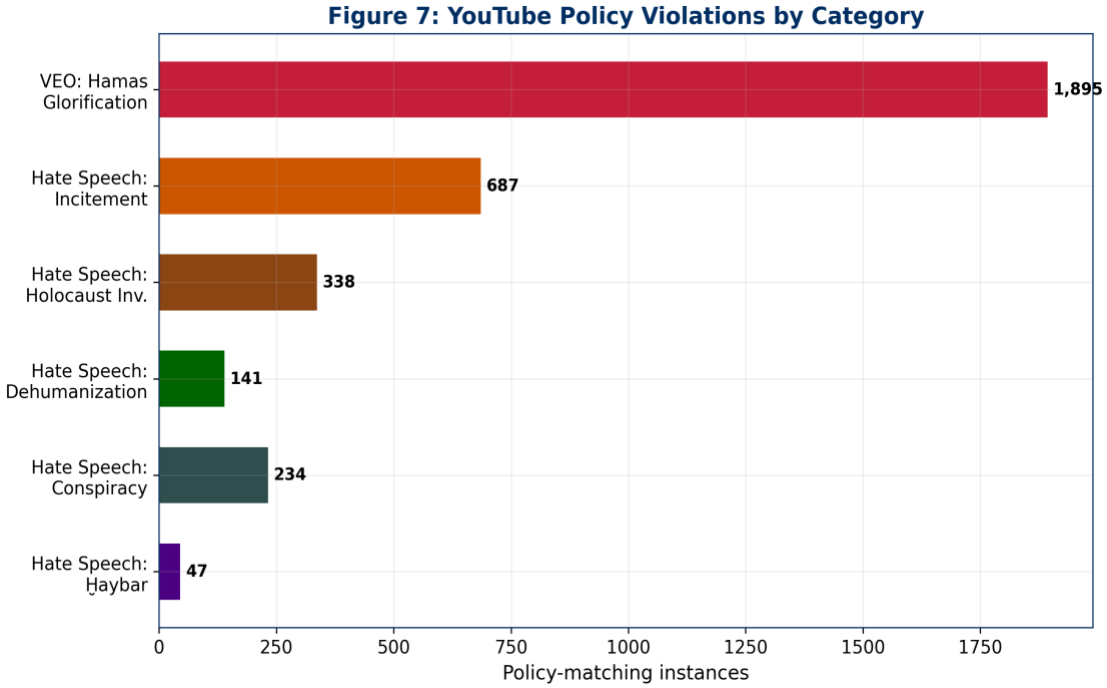

Figure 8: YouTube Policy Violations by Category

9.5 The Enforcement Gap

The series could plausibly pass a keyword-based content audit. The editorial voice contains zero slurs, zero dehumanizing animal terms directed at Jews, zero explicit calls for violence, and only 2.0% standalone ethnic naming. A YouTube content reviewer examining the editorial transcripts would find a professionally produced Hebrew-translation program with a Palestinian presenter.

The 26,620 domain-classified term instances — aggregated across the seven thematic domains documented in §6.2–6.8 and Hamas sanctification (§6.9), representing approximately 18,400 unique comments from approximately 12,300 unique users — exist not in the editorial but in the comment section — the audience layer that this study finds to be systematically associated with the editorial apparatus documented in §§5–7. Of these, 3,438 are individually mapped to specific YouTube policy categories in §§9.1–9.2. The editorial apparatus is systematically associated with

the discursive conditions under which the audience layer produces the policy-relevant violations documented here.

YouTube's public-facing enforcement tools address individual comments through reporting and moderation mechanisms, and its Community Guidelines apply to comments, posts, and thumbnails as well as videos (YouTube, Community Guidelines). These mechanisms clearly address individual violations. What they do not clearly address is the production-reception chain documented in this study: a formally compliant editorial format that systematically generates large-scale audience-layer hate across thousands of comments over 22 months. The gap is not that YouTube lacks tools but that the existing tools are designed for individual content items rather than for the editorial-audience system that produces them.

10. Recommendations

The recommendations below follow directly from the findings documented in this study. They are organized by institutional level because the problem identified here is not only one of platform moderation, but of transnational media circulation, state-funded amplification, regulatory asymmetry, and scholarly underdevelopment.

10.1 For U.S. Government, Legislators, and Policymakers

1. Treat Arabic-language digital antisemitism as a policy-relevant blind spot in counter-extremism and online harms analysis. The study documents a large-scale Arabic digital ecosystem in which anti-Jewish hostility is generated through a production-reception chain in which formally political editorial framing is transformed into ethnic-religious audience targeting. U.S. agencies and legislators concerned with online extremism, antisemitism, and foreign influence should recognize Arabic-language moderation gaps as a national policy issue rather than a niche language problem. Existing frameworks provide actionable entry points: the Foreign Agents Registration Act (FARA) applies to media operations funded by foreign governments; congressional oversight of platform governance (including hearings on content moderation practices) can incorporate Arabic-language evidence; and the interagency process on online harms can integrate Arabic-language monitoring into its existing counter-extremism and antisemitism portfolios.

2. Support transparency requirements for major platforms regarding non-English moderation capacity. The findings in §§4, 8, and 9 show that Arabic-language content poses interpretive problems not captured by general-purpose moderation systems. U.S. policymakers should press major platforms to disclose, in regular transparency reporting, the scale of Arabic-language moderation staffing, dialectal coverage, escalation protocols, and enforcement outcomes for comments as well as primary content.

3. Incorporate Arabic-language expertise into policy frameworks addressing antisemitism, extremist mobilization, and foreign state-funded information ecosystems. Current debates on platform harms are often shaped by English-language examples. The corpus documented here

shows that Arabic-language discourse can generate large-scale anti-Jewish hostility while remaining partly invisible to existing enforcement frameworks.

10.2 For Technology Companies, Especially YouTube

4. Develop Arabic-specific moderation systems trained on dialectal variation, scriptural register, and contextual targeting. Existing industry coordination mechanisms — the Global Internet Forum to Counter Terrorism (GIFCT) and the commitments made under the Christchurch Call to Action — were designed primarily for overt extremist content. The implicit, coded, and audience-generated content documented in this study falls largely outside these frameworks. Platforms should extend GIFCT-style coordination to include Arabic-language borderline hate and production-reception system auditing. The Arabic detection gap documented in §9.4 shows that keyword-based systems are poorly equipped to detect anti-Jewish discourse operating through Ḫaybar invocations, pronominal targeting, Quranic dehumanization, and devotional ambiguity. Platforms should develop Arabic-specific classifiers that are dialect-aware and context-sensitive.

5. Require expert human review for Arabic-language borderline hate speech and extremist-content cases. The coding challenges documented in §4.5 show that many relevant Arabic terms produce either high false positives or systematic false negatives when treated without context. Platforms should use escalation pathways for Arabic-language content involving scriptural, eschatological, or devotional references.

6. Move from item-level moderation to production-reception auditing. This study identifies a system in which a formally compliant editorial layer coexists with a systematically violative audience layer. Platform governance should include channel-level audits for cases in which comment-section violations are not isolated but persistent, patterned, and plausibly linked to editorial framing.

7. Apply graduated interventions to comment environments, not only to videos. Where comment sections exhibit sustained policy-matching hate, platforms should consider

interventions such as slowed posting, stricter comment filtering, temporary review holds, limits on recommendation circulation, and escalated channel-level review.

10.2A Detection Tools and Methods

8. Build dialect-first Arabic moderation pipelines rather than treating Arabic as a single language layer. The mechanisms documented in this study operate across dialects, registers, and regionally specific forms of hostile naming. Platforms should identify dialect first, then apply downstream moderation models calibrated for the relevant dialectal and stylistic environment. Arabic moderation systems that do not distinguish between Modern Standard Arabic, Levantine, Egyptian, Gulf, and Maghrebi usage will systematically miss target terms, proxy naming, and coded hostility.

9. Develop a sacred-register disambiguation layer for Arabic. A central finding of this study is that the same lexical material may function either as ordinary devotional expression or as weaponized anti-Jewish discourse, depending on the appended clause, target, and context. Platforms should therefore develop classifiers that distinguish political, devotional, scriptural, eschatological, and explicitly hateful uses of overlapping Arabic vocabulary rather than evaluating such phrases in isolation.

10. Create a naming-instability and proxy-target detector. The study shows that anti-Jewish hostility often emerges through movement across *al-yahūd*, *al-Isrāʾīliyyūn*, *al-ṣahāyina*, *al-kayān*, and pronominal forms that blur religious, political, ideological, and delegitimizing reference. Moderation systems should be designed not only to detect hostile words, but to track target-reference chains and identify when nominally political terminology is functioning as a proxy for ethnic-religious targeting.

11. Operationalize DHA as a cross-layer escalation detector. Platforms should test a DHA-style monitoring model that compares the editorial layer and the audience layer together: transcript, title, tags, and thumbnail on one side; comments and replies on the other. Such a system would flag videos where audience ethnic-religious targeting sharply exceeds editorial ethnic naming, or where the audience introduces ethnic targeting despite zero editorial ethnic references.

12. Shift from comment-level flagging to comment-ecology risk scoring. The relevant unit identified in this study is not the single comment but the patterned environment. Platforms should score comment sections for density, recurrence, lexical clustering, escalation, and absence of counter-speech, rather than treating hateful comments as isolated violations.

13. Build evasion-recovery tools for broken text, coded orthography, and coordinated comment attacks. The broader extremism literature shows that hostile ecosystems adapt to moderation through masking, comment-seeding, account recycling, and coordinated raids. Arabic moderation systems should normalize broken words, orthographic stretching, mixed-script forms, emoji substitution, and other evasive techniques before classification, while also detecting coordinated surges in hostile commenting behavior.

10.3 For Arab States Maintaining Relations with Israel or Pursuing Regional Normalization

14. Treat anti-Jewish discourse in Arabic media as a strategic obstacle to regional stabilization. The discursive collapse of Jews, Israelis, Zionists, and "the entity" documented in this study does not remain at the level of rhetoric. It hardens public imaginaries, deepens sectarian polarization, and makes any durable Arab-Jewish civic or diplomatic rapprochement more difficult.

15. Support Arabic-language public discourse that distinguishes Jews, Israelis, state actors, and political ideologies. One of the central findings of this study is that naming instability enables the movement from political criticism to ethnic-religious hostility. Arab policymakers, media regulators, and educational institutions should support public discourse that preserves distinctions between Jewish identity, Israeli citizenship, state policy, and ideological affiliation.

16. Encourage official and semi-official media sectors to adopt explicit standards against anti-Jewish incitement. Because this study concerns a Qatari state-funded media ecosystem, the issue is not only spontaneous audience expression. Arab states with diplomatic or strategic interest in regional de-escalation should encourage Arabic media institutions to adopt internal standards against collective anti-Jewish language, scriptural incitement, and dehumanizing terminology.

17. Recognize sustained anti-Jewish media discourse as a threat to long-term Arab-Jewish coexistence and regional de-escalation. The ecosystem documented here does not merely stigmatize a target group. It hardens public language, normalizes collective hostility, and weakens the discursive conditions required for any durable civic, diplomatic, educational, or cultural rapprochement between Arabs and Jews.

10.4 For UNESCO and Other International Cultural and Educational Bodies

18. Support research and educational programming on antisemitism in Arabic digital environments. The study identifies an underdeveloped field at the intersection of antisemitism studies, Arabic language ideology, digital media analysis, and platform governance.

19. Promote media and digital-literacy initiatives that address sacred language, historical memory, and hate expression together. The problem documented here is not reducible to vulgar hate speech. It operates through scriptural authority, historical allusion, and devotional language. International educational programming should address the ways in which sacred and historical vocabularies can be mobilized in contemporary hate ecosystems.

10.5 For Scholars, Universities, and Research Institutions

20. Build Arabic-language antisemitism in digital ecosystems as a defined field of inquiry. The literature reviewed in §3 shows that relevant scholarship exists, but remains fragmented across Islamic antisemitism, propaganda studies, platform governance, sociolinguistics, and extremism research.

21. Expand corpus-based, multilingual, and production-reception research designs. Future work should test DHA and related methods across other Arabic channels, platforms, and media ecosystems to determine whether the patterns documented here are exceptional or systemic.

22. Fund benchmark datasets for Arabic proxy naming, sacred-register hate, and production-reception escalation. Existing Arabic NLP resources are improving, but the field still lacks benchmark datasets tailored to the mechanisms identified here: proxy naming, devotional ambiguity, pronominal targeting, Ḫaybar invocation, and editorial-audience displacement.

23. Establish formal protections for scholars researching Arabic-language antisemitism and related fields. Because this topic can trigger accusations of betrayal, normalization, or covert political alignment, institutions should provide explicit support mechanisms for researchers working in this area, including publication backing, legal support where necessary, and clear professional recognition of the field's legitimacy.

10.6 For European Policymakers and Regulators

24. Recognize Arabic-language digital hate as a domestic policy issue, not only an external one. The EU Digital Services Act (DSA), which entered full application in February 2024, requires very large online platforms to assess and mitigate systemic risks including the dissemination of illegal content and negative effects on civic discourse. The production-reception system documented in this study — in which a state-funded editorial apparatus is systematically associated with large-scale audience-layer hate — constitutes precisely the type of systemic risk the DSA's provisions are designed to address. The UK Online Safety Act (2023) similarly imposes duties on platforms regarding illegal content and content harmful to adults, with specific provisions for hate speech. Because Arabic media circulates widely across European diasporic environments, the dynamics documented in this study are directly relevant to European social cohesion, antisemitism prevention, and digital regulation.

25. Require that platform enforcement in Europe include Arabic-language capacity proportional to user exposure. Where Arabic is widely consumed, regulation that focuses only on dominant national languages leaves a major enforcement gap intact.

26. Integrate Arabic-language antisemitism into broader policy frameworks on hate speech, prevention, and social cohesion. The study shows that anti-Jewish discourse in Arabic can circulate in societies where Arabic-speaking populations are substantial and where antisemitic incidents are already elevated.

10.7 Cross-Cutting Recommendation

27. Reframe the problem from isolated hate content to ecosystem-level production of hostility. The central finding of this study is systemic: a professionally produced editorial layer is associated with the conditions under which a mass audience layer produces anti-Jewish hostility at scale. Institutions responding to this problem — whether states, platforms, regulators, or scholars — should move beyond single-item enforcement models and address the production-reception system as the relevant unit of analysis.

28. Invest in AI-augmented detection and research infrastructure for Arabic-language digital ecosystems. This study demonstrates, through its own methodology, that AI-assisted research can enable a single researcher with the requisite domain expertise to process, analyze, and systematically document a corpus of 109,165 comments, 24,977 transcript segments, and 257 thumbnails — a scale of analysis that would have required a large research team and multi-year timelines using traditional methods alone. The deployment of Claude AI as an analytical instrument throughout this study (§4.1) produced results in months that manual methods would have taken years to achieve, while maintaining the interpretive rigor that Arabic-language discourse analysis requires. This methodological experience carries a direct policy implication: if a single AI-augmented researcher can document the production-reception system analyzed here, then a properly funded detection program — combining AI processing capacity with native Arabic-speaking expert oversight — could monitor Arabic-language digital ecosystems at scale in near-real time. The six detection-gap categories documented in §9.4 (political-to-ethnic register shift, Ḫaybar invocation, pronominal targeting, Quranic dehumanization, scriptural essentialization, and devotional register ambiguity) are not intractable. Each represents a specific classification challenge that can be addressed through purpose-built AI systems trained on the patterns this study has identified — provided that such systems incorporate the dialect-aware, register-sensitive, and context-dependent processing that general-purpose moderation tools currently lack. What is needed is not a theoretical breakthrough but dedicated investment: benchmark datasets for Arabic proxy naming and sacred-register hate (Recommendation 22), dialect-first moderation pipelines (Recommendation 8), sacred-register disambiguation layers (Recommendation 9), and DHA-style cross-layer monitoring (Recommendation 11) are all technically feasible with current AI capabilities. The barrier is not technology but funding,

institutional commitment, and the recognition that Arabic-language digital antisemitism requires the same scale of detection investment that has been directed at English-language online extremism. This study is intended, in part, as proof of concept: it shows what AI-augmented analysis can achieve with limited resources, and it indicates what a fully resourced detection system could accomplish.

11. Conclusion

This study set out to answer a three-part question: through what mechanisms is a Qatari state-funded Arabic media operation on YouTube associated with measurable anti-Jewish hostility in its audience, and why do those mechanisms evade platform detection?

The answer, documented across 223 videos, 109,165 comments, and 843,450 characters of editorial transcript, is a production-reception system whose components are systematically associated. The editorial apparatus — operated by AJ+ Arabic, a digital sub-brand of Qatar's Al Jazeera Media Network — constructs a moral binary of measurable precision: 525 negative characterizations for Israel, zero for Hamas in editorial voice; 95.5% political naming with strategic ethnic slippage; and structural absence of Israeli civilian empathy, Hamas criticism, or moral complexity. The audience, operating within the discursive environment that co-occurs with this apparatus — including endorsed quotation, directive speech acts, sarcastic laughter, and parasocial transmission — extends the editorial's political register into ethnic-religious territory the editorial itself largely avoids. In 97.6% of videos where the editorial contains zero ethnic references, the audience independently introduces ethnic targeting — drawing on scriptural, eschatological, and devotional repertoires that the editorial does not explicitly deploy.

The displacement is the finding. In the 54 videos where direct comparison is possible, the audience's ethnic-to-political naming ratio is 7.8 times the editorial ratio. Across seven thematic domains — from ontological evil to sacred violence — the audience generates vocabulary that has no editorial equivalent: *mujāhidīn* (1,147 instances), *ḫiyāna* (906), *laʿnat Allāh* (115), *ḫanāzīr* (68), and Ḫaybar (47). None of these terms appears in the editorial transcript. They are produced entirely from the audience's own cultural, theological, and historical repertoire. The negative evidence is also significant: in the exception case (§7.7), where the editorial provides an operational narrative without the moral binary, displacement does not appear to activate — a pattern consistent with the inference that displacement is associated with the editorial apparatus rather than being a default property of Arabic-speaking audiences, though alternative explanations cannot be excluded within a single-series design.

The pattern is the contribution. Displaced Hostility Analysis — introduced here as a corpus-based framework for measuring the shift from political to ethnic-religious targeting — documents a system in which formally compliant editorial content is systematically associated with large-scale audience-layer hate. The editorial provides the political target; the audience's own repertoire appears to provide the ethnic-religious upgrade. The zero counter-narrative finding — not one comment in 109,165 that meets the operational definition for resisting the dominant frame across any of six defined categories — describes a visible discursive environment in which alternative perspectives are absent. As noted in §6.10, this finding describes the observable comment environment after platform and channel moderation; it is possible that counter-narrative comments existed but were removed by YouTube's automated systems, by the channel's own comment moderation settings, or by users themselves.

The scale — 83.9 million views — underscores the policy significance of these findings. The content documented here appears to match categories prohibited by YouTube's published Hate Speech Policy and Violent Extremist or Criminal Organizations Policy. It circulates through the world's largest video-sharing platform, reaches audiences across at least 22 countries, and operates within the Al Jazeera Media Network, which states it is funded by the Qatari government. The Arabic detection gap — operating through register shifts, scriptural essentialization, devotional ambiguity, and pronominal targeting — renders this content largely invisible to the moderation systems designed to prevent it.

The corpus thus provides empirical evidence for what Herf (2024, pp. 237–238) calls the "era of simultaneity" — a period in which antisemitism's multiple faces operate concurrently rather than sequentially. The AJ+ Arabic ecosystem demonstrates this simultaneity in a single digital space: left-wing anti-Zionist framing in the editorial layer, Islamist scriptural-eschatological hostility in the audience layer, and conspiratorial antisemitism distributed across both. At stake is not only the documentation of one series, but the definition of an emerging research problem: how Arabic-language antisemitism operates in digital ecosystems through mechanisms that are culturally specific, linguistically complex, and structurally resistant to detection. This study is intended as a contribution to that problem — and as an invitation to the scholarly, policy, and

platform-governance communities to address it. It addresses all three dimensions of the gap identified in §3.6: substantively, by documenting Arabic-language antisemitism at corpus scale in a digital ecosystem rather than in sermons, textbooks, or isolated media examples; methodologically, by linking editorial production to audience uptake within a single corpus through the DHA framework; and from a governance perspective, by measuring the displacement between formally compliant editorial content and policy-violating audience discourse — the production-reception gap that current moderation systems are not designed to detect.

The study's methodology also carries a broader implication for research infrastructure and policy development. The deployment of AI-assisted analysis (Anthropic's Claude AI) alongside researcher-conducted OSINT and commercial data harvesting (Bright Data) enabled a single researcher with the requisite Arabic-language expertise to process and systematically analyze a corpus of this scale — 109,165 comments, 24,977 transcript segments, 257 thumbnails, 537 tags — within a timeline that traditional manual methods could not have achieved. This experience suggests that AI-augmented research designs can accelerate the production of policy-relevant evidence on Arabic-language digital ecosystems, an area where the research gap documented in §3 is partly a function of resource constraints. More pointedly, the six detection-gap categories documented in §9.4 are not intractable technical problems; they are specific classification challenges that purpose-built AI systems, trained on the patterns this study has identified and supervised by native Arabic-speaking experts, could address at scale. What this study demonstrates with limited resources — corpus-level documentation of a production-reception system, measurement of displacement, mapping of policy violations — a properly funded detection program could accomplish in near-real time across multiple Arabic-language digital ecosystems simultaneously. The barrier to effective detection is not technological impossibility but institutional investment: dedicated funding for benchmark datasets, dialect-aware classifiers, sacred-register disambiguation, and cross-layer monitoring systems. This study is intended, in part, as proof of concept for what such investment could achieve.

The study's limitations (§4.6) should direct future research. The single-coder design, while grounded in the methodological requirement of native dialectal competence, means that the

DHA framework would benefit from independent replication by other Arabic-speaking researchers using the 8-step protocol in Appendix B. The absence of a comparison corpus limits the ability to determine whether displacement is specific to AJ+ Arabic or characteristic of Arabic-language political media more broadly; comparative application of DHA to other state-funded outlets, private Arabic channels, and non-Arabic political media would clarify the framework's generalizability. The 22-month observation window captures a period of intensified conflict; longitudinal extension across both conflict and non-conflict periods would test whether displacement patterns are episodic or structural. Finally, the comment-survivorship limitation means that the corpus represents the visible discursive environment after platform and channel moderation; access to pre-moderation data, if obtainable, would allow measurement of the full scope of audience expression.

Appendix A: Codebook

This codebook specifies the operational definitions, classification rules, and decision protocols governing all coding in this study. All coding was performed by the lead researcher using contextual reading of each instance in its original Arabic. The IHRA Working Definition of Antisemitism (2016) provides the overarching classification framework.

A.1 Naming Categories (Editorial and Audience)

Category	Definition	Marker terms	Coding rule
Political naming	Reference to the state of Israel, its government, named officials, or military forces by political or national designation	Isrāʾīl, Natanyāhū, al-iḥtilāl (the occupation), al-jayš al-Isrāʾīlī (the Israeli army), al-ḥukūma (the government)	Default category. Code as political unless ethnic, Zionist, or delegitimizing markers are present.
Ethnic naming	Reference to Jews as an ethnic-religious collectivity	al-yahūd (the Jews), yahūd (Jews), Banī Isrāʾīl (Children of Israel)	Code when the referent is Jews as a people, not as citizens of a state. Contextual: al-yahūd used as clear shorthand for "Israelis" in a military context is coded political; al-yahūd used with essentializing or theological attribution is coded ethnic.
Zionist terminology	Terms referencing Zionism or Zionists	ṣahāyina (Zionists), ṣahyūnī (Zionist), ṣahyūniyya (Zionism)	Coded as a distinct category. Each instance is classified by contextual function: political (equivalent to "Israeli policy supporters"), ethnic-proxy (substituting for "Jews" collectively), or conspiratorial. Reported separately in displacement calculations; collapsed into ethnic count only when contextual usage is clearly collective-proxy.
Delegitimizing/state-denial	Terms that deny ordinary state naming	al-kayān (the entity), al-kayān al-ṣahyūnī (the Zionist entity)	Code when the term functions to deny Israel's legitimacy as a state.

A.2 Seven Thematic Domains of Audience Discourse

These domains emerged through iterative engagement with the corpus (Shkedi, 2019, pp. 57–70), not through deductive imposition from the IHRA framework. Each domain is defined below with its marker vocabulary, IHRA example match, and inclusion/exclusion rules.

#	Domain	Definition	Marker vocabulary (Arabic / sAts / English)	IHRA example

1	Ontological evil	Attribution of inherent, essential evil to Jews or Israel as a collectivity	sharr (evil), aʿdāʾ Allāh (enemies of God), aʿdāʾ al-insāniyya (enemies of humanity), mufsidūn (corrupters)	Ex. 2, Ex. 6
2	Theological condemnation	Framing through divine judgment, curses, or sacred-historical precedent	laʿnat Allāh (God's curse), ġaḍab Allāh (God's wrath), ḥasbunā Allāh + targeting clause, Allāhumma dammiruhum (O God, destroy them)	Ex. 1
3	Dehumanization	Comparison to animals or subhuman entities	ḫanāzīr (pigs), qirada (apes), al-qirada wa-l-ḫanāzīr (apes and pigs), ḥayawānāt (animals)	Ex. 6
4	Holocaust inversion	Comparing Israeli actions to Nazism or the Holocaust; accusing Jews of perpetrating genocide	nāzī, Hitlar (Hitler), maḥraqa (holocaust), ibāda jamāʿiyya (genocide) applied to Israel	Ex. 9
5	Conspiracy	Attribution of secret plots, world control, or coordinated malice to Jews or Zionists	muḫaṭṭaṭ (plot/scheme), muʾāmara (conspiracy), al-māsūniyya (Freemasonry)	Ex. 4, Ex. 7
6	Sacred violence	Sanctification of violence against Jews/Israel through religious framing	mujāhidīn (holy warriors), jihād, istišhād (martyrdom), Ḫaybar invocation, šahīd with Hamas attribution	Ex. 1
7	Essentialization	Application of scriptural or historical characterizations to Jews in all eras	ḫiyāna (betrayal) as permanent Jewish trait, ġadr (treachery) as inherent nature, Sūrat al-Baqara applied to contemporary Jews	Ex. 2, Ex. 6

Inclusion rule: A comment is coded under a domain when it contains at least one marker term or functional equivalent directed at Jews, Israelis-as-Jews, or Zionists-as-Jews-proxy, within a framing that matches the domain definition.

Exclusion rule: General political criticism of Israeli state policy — including terms like *iḥtilāl* (occupation), *ḥiṣār* (siege), *istiʿmār* (colonialism) — is not coded under any domain unless accompanied by ethnic, theological, or essentializing markers.

Multi-domain coding: Comments may be coded under more than one domain. A comment combining *laʿnat Allāh ʿalā al-yahūd al-ḫanāzīr* ("God's curse on the Jewish pigs") is coded under both Domain 2 (theological condemnation) and Domain 3 (dehumanization).

A.3 Editorial Features

Feature	Definition	Coding rule
Negative evaluative vocabulary	Terms applied to Israeli actors that carry negative moral, legal, or characterological weight	Count all instances of 15 identified terms (see §5.1 vocabulary table). Total: 525.
Hamas framing	How the editorial voice characterizes Hamas, its members, and its operations	Coded as positive (legitimation/sacralization), neutral (factual reference), or attributed negative (quoting Israeli characterization).
Endorsed quotation	Speech act in which the presenter translates an Israeli statement and appends āmīn yā rabb ("Amen, O Lord")	Code when the devotional formula immediately follows a translated Israeli statement with negative valence.
Directive speech acts	Direct address to the audience prescribing action	Count all instances of imperative address: šārikū (share), iʿmalū lāyk (like), iktubū (write), tābiʿū (follow). Total: 445.
Sarcastic laughter	Audible laughter directed at Israeli statements or actors	Code when laughter follows an Israeli claim and functions to delegitimize the speaker. Total: 102 across 63 videos.
Structural absence	Absence of empathy, criticism, or complexity	Code at the corpus level: instances of Israeli civilian empathy (1 of 843,450 characters), Hamas criticism (0), moral complexity (0).

A.4 Counter-Narrative Protocol

A counter-narrative comment is defined as any comment that resists the dominant editorial frame. Six categories were defined:

#	Category	Definition	Search terms used
1	Civilian sympathy	Expressing sympathy for Israeli civilians	madaniyyīn Isrāʾīliyyīn, ḍaḥāyā, abriyāʾ + Israeli referent
2	Rejection of collective blame	Distinguishing between Jews and Israeli state actors	laysa kull al-yahūd, al-farq bayn, lā naʿammim
3	October 7 as terrorism	Characterizing the October 7 attack as terrorism rather than resistance	irhāb, hujūm irhābī, laysa muqāwama
4	Objection to curses/violence	Objecting to curses or calls for violence directed at Jews as a group	ḥarām, lā yajūz, hāḏā laysa min al-Islām
5	Rejection of celebratory Hamas framing	Refusing celebratory framing of Hamas military operations	lā naḥtafil, ḍaḥāyā min al-jānibayni
6	Call for coexistence	Any expression supporting Arab-Jewish or Israeli-Palestinian coexistence	taʿāyuš, salām, ḥall al-dawlatayn

Protocol: Each category was searched using Arabic keyword queries, dialectal variants, and manual review of 5,000 randomly selected comments. Result: zero instances across all six categories in 109,165 comments.

A.5 DHA Calculation Rules

Standard DHA factor (for videos where both editorial and audience ethnic references > 0):

DHA = (audience ethnic / audience political) ÷ (editorial ethnic / editorial political)

Applied to: 54 videos. Mean: 7.8×. Median: 5.2×.

Binary ethnic-introduction measure (for videos where editorial ethnic references = 0):

Code: does the audience introduce ethnic targeting (yes/no)?

Applied to: 166 videos. Result: 162 yes, 4 no (97.6%).

Zero-denominator rule: When the editorial transcript contains zero ethnic references, the standard ratio is undefined. These videos are analyzed exclusively through the binary measure. No smoothing, imputation, or artificial denominator is applied.

Zionist terminology rule: Zionist terms (*ṣahāyina*, *ṣahyūnī*) are reported separately. They are included in the ethnic count only when contextual usage is clearly collective-proxy (e.g., "the Zionists are the descendants of apes and pigs"). The aggregate corpus ratio including Zionist terms as ethnic is 11.9×; excluding them, 9.4×.

Replication protocol: applying DHA to a new corpus. The following procedure assembles the rules defined above into a linear sequence for researchers applying DHA to a new corpus.

Step 1: Corpus assembly. Obtain linked editorial and audience datasets for the target media system. The minimum requirement is (a) transcripts or text of the editorial layer for each content unit (video, article, broadcast) and (b) audience-generated text (comments, replies, reactions) linked to the same content units. The content unit — typically a video — serves as the case-level unit of analysis.

Step 2: Naming codebook. Adapt the naming categories defined in §A.1 to the target language and context. At minimum, define the boundary between political naming (references to state, government, officials, military) and ethnic naming (references to an ethnic, religious, or racial collectivity). The following decision rules should govern boundary cases: (a) when a term can function as either political or ethnic (e.g., "Zionist" in Arabic, "Zionist" in English), define explicit contextual criteria for each classification and report the term separately with both inclusive and exclusive displacement calculations; (b) when a term is used as a shorthand where the more precise referent would be the state or its agents (e.g., "the Jews" used where "Israel" would be more accurate), code as ethnic unless the surrounding context makes a purely political reading unambiguous; (c) when evaluative vocabulary could be classified as either political-legal characterization or moral condemnation (e.g., "genocide" during active ICJ proceedings), document the classification decision explicitly and report disaggregated counts under both classifications; (d) when religious or devotional language appears with and without targeting clauses, code only the appended targeting clause, not the devotional formula itself, and document the decision protocol for distinguishing the two. Record all decision rules in a written codebook before coding begins. The codebook should specify, for each ambiguous category, the contextual markers that trigger ethnic versus political classification, and should include worked examples of boundary cases with the reasoning for each decision.

Step 3: Editorial coding. For each content unit, count (a) ethnic references and (b) political references in the editorial layer. Calculate the editorial ethnic-to-political ratio (E/P) for each unit. Classify each unit as "ratio-valid" (editorial ethnic > 0) or "zero-editorial" (editorial ethnic = 0).

Step 4: Audience coding. For each content unit, count (a) ethnic references and (b) political references in the audience layer. Calculate the audience ethnic-to-political ratio (E/P) for each unit. For zero-editorial units, code the binary measure: does the audience introduce ethnic targeting (yes/no)?

Step 5: Per-unit DHA calculation. For each ratio-valid unit, compute: DHA = (audience E/P) ÷ (editorial E/P). A value of 1.0 indicates identical proportions; values above 1.0 indicate audience

displacement toward the ethnic register. Do not compute the ratio for zero-editorial units; use only the binary measure for those cases.

Step 6: Corpus-level aggregation. Compute (a) the mean and median DHA factor across all ratio-valid units; (b) the percentage of zero-editorial units where the audience introduces ethnic targeting; and (c) the aggregate corpus ratio: total audience E/P across all comments divided by total editorial E/P across all transcripts. Report all three measures; they capture different facets of displacement.

Step 7: Sensitivity checks. If the target language contains ambiguous terms (analogous to Zionist terminology in Arabic), compute displacement with and without those terms included in the ethnic count. Report both figures. Identify and examine exception cases — units where the expected pattern does not hold — to test whether the mechanism is conditioned by specific editorial features.

Step 8: Interpretation. A DHA factor significantly above 1.0 indicates that the audience's ethnic register exceeds the editorial's. A high binary ethnic-introduction rate indicates that audiences generate ethnic targeting even where the editorial does not. Neither measure establishes causation; both measure systematic co-occurrence between editorial framing and audience displacement. Exception cases — where displacement does not occur despite the same channel and presenter — provide the analytical contrast needed to infer which editorial features condition the displacement.

A.6 Devotional Register Decision Rules

Expression	Default coding	Override condition
ḥasbunā Allāh wa niʿm al-wakīl	Not coded as anti-Jewish	Coded under relevant domain when a targeting clause specifying Jews or Zionists is appended
Allāhumma + prayer	Not coded as anti-Jewish	Coded when the prayer specifies destruction, punishment, or violence against Jews/Zionists
āmīn / āmīn yā rabb	Not coded as anti-Jewish	Coded under endorsed quotation only when appearing in editorial voice after a translated Israeli statement

subḥān Allāh / Allāhu Akbar	Not coded as anti-Jewish	Coded only when embedded in a comment whose surrounding text targets Jews/Zionists
Quranic citation	Not coded as anti-Jewish	Coded under Domain 7 (essentialization) when the citation is applied as a characterization of contemporary Jews

General principle: The prayer formula is never itself the basis for coding. Coding is attributed to the appended targeting clause, the surrounding hostile context, or the essentializing hermeneutic application — not to the devotional expression alone.

A.7 Ambiguity Protocol

Comments where the targeting intent was uncertain — where devotional language sat at the boundary between prayer and curse, where pronominal reference was ambiguous, or where dialectal usage was unfamiliar — were logged separately and re-reviewed after the initial coding pass. A random subset of 500 coded comments was re-checked after the full corpus coding to assess internal consistency. Ambiguous cases are not included in the frequency counts reported in this study unless the re-review confirmed targeting intent.

A.8 Transliteration Standard

All Arabic words, names, phrases, and quoted evidence in this study are transliterated using **Brill's Simple Arabic Transliteration System (sAts v1.0)**. Arabic evidence is presented in three forms: original Arabic script, Brill sAts transliteration, and English translation.

Appendix B: Evidence Landscape

The principal evidence exhibits are presented inline within §§5–9 of the main text, each with full trilingual treatment (Arabic original, Brill sAts transliteration, English translation) and a YouTube comment ID enabling independent verification. The figure below shows the full evidence landscape: total corpus instances per domain, with the number of individually verified and cited exhibits noted for each. The two largest domains — sacred violence (6,390 instances) and ontological evil (6,382) — together account for more than two-thirds of the classified corpus. Behind each exhibit cited in the body text stand hundreds or thousands of additional corpus

instances following the same pattern. The complete evidence corpus, including all comment IDs, is available from the author on request.

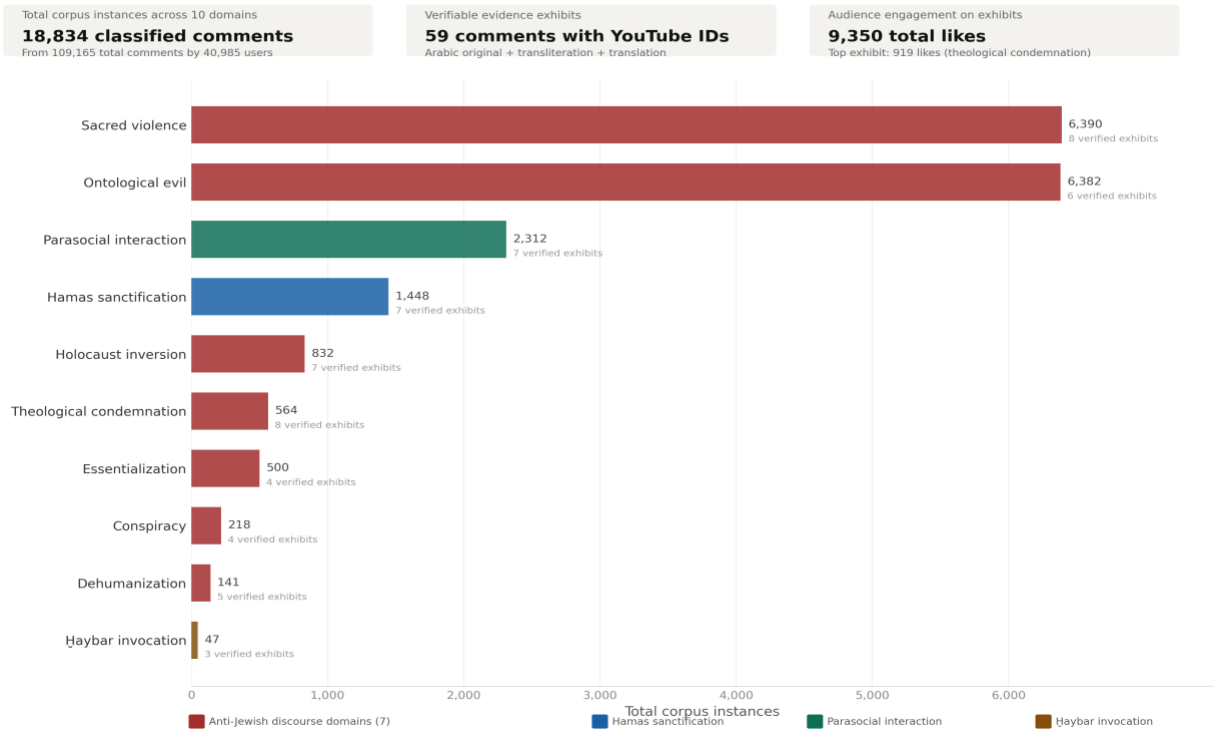

Figure 9: Evidence Landscape — Corpus Instances and Verified Exhibits across 10 Domains

Appendix C: Per-Video Displaced Hostility Analysis Data

This appendix presents the Displaced Hostility Analysis result for each of the 220 transcribed videos, sorted chronologically. The figure below shows the full corpus at a glance. Each bar is one video. Blue bars indicate videos where both the editorial and audience layers contain ethnic references; bar height shows how many times the audience amplifies the editorial's ethnic naming. Red bars indicate videos where the editorial contains zero ethnic references — yet the audience introduces ethnic targeting anyway. The four hollow markers are the only videos in which neither layer targets Jews. The complete per-video dataset is available from the author on request. Individual videos are retrievable at youtube.com/watch?v=[ID].

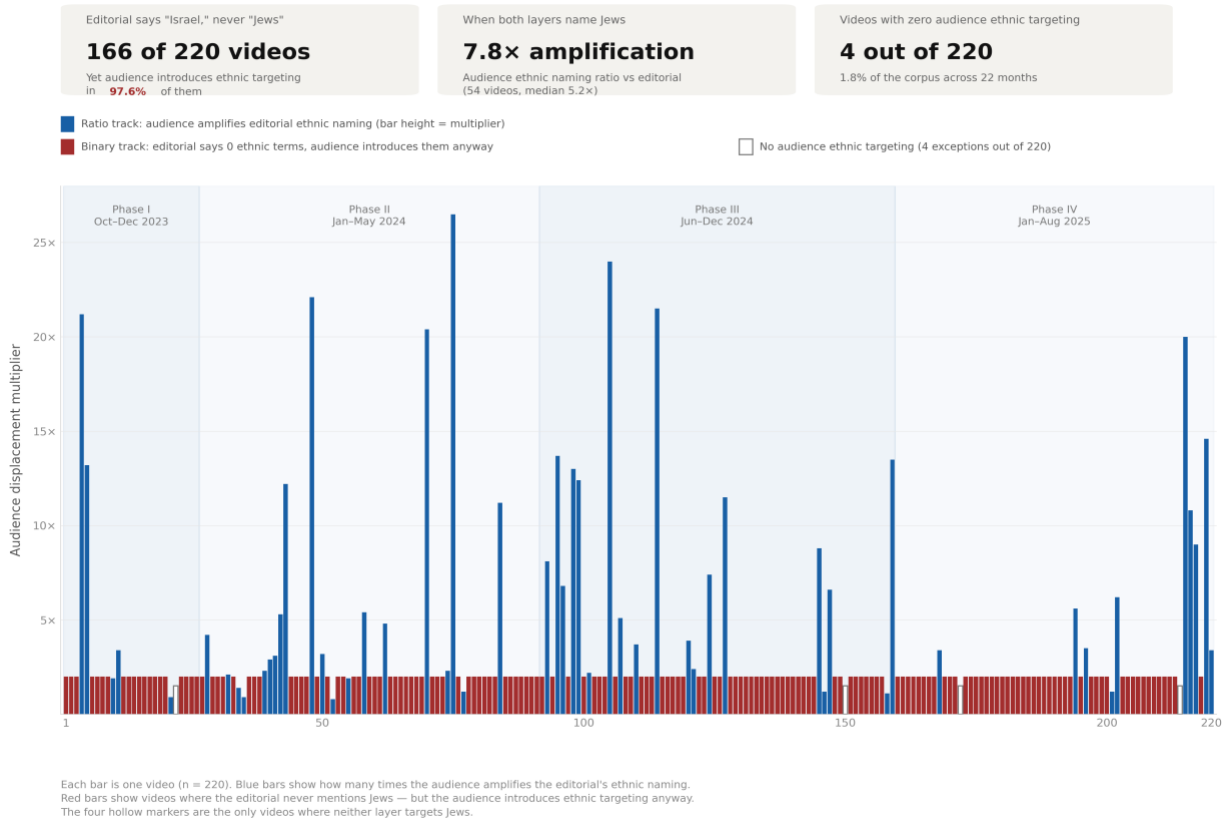

Figure 10: Per-Video Displaced Hostility Analysis across 220 videos (October 2023 – August 2025)

Appendix D: Tag Frequencies

The figure below shows the 30 most frequently used content tags across the 221 tagged videos, color-coded by category. Tags are assigned by the AJ+ Arabic editorial team and reflect the editorial classification of each episode's content (537 unique tags total). The color coding reveals the editorial priority structure: Hamas actors (red) are tagged on 83–91% of all videos, achieving near-parity with Israeli state references (blue) at 92–94%. The editorial team adopts Hamas's operational nomenclature — Ṭūfān al-Aqṣā ("Al-Aqsa Flood") — on 141 of 221 tagged videos (63.8%). The complete tag frequency table (100 tags) is available from the author on request.

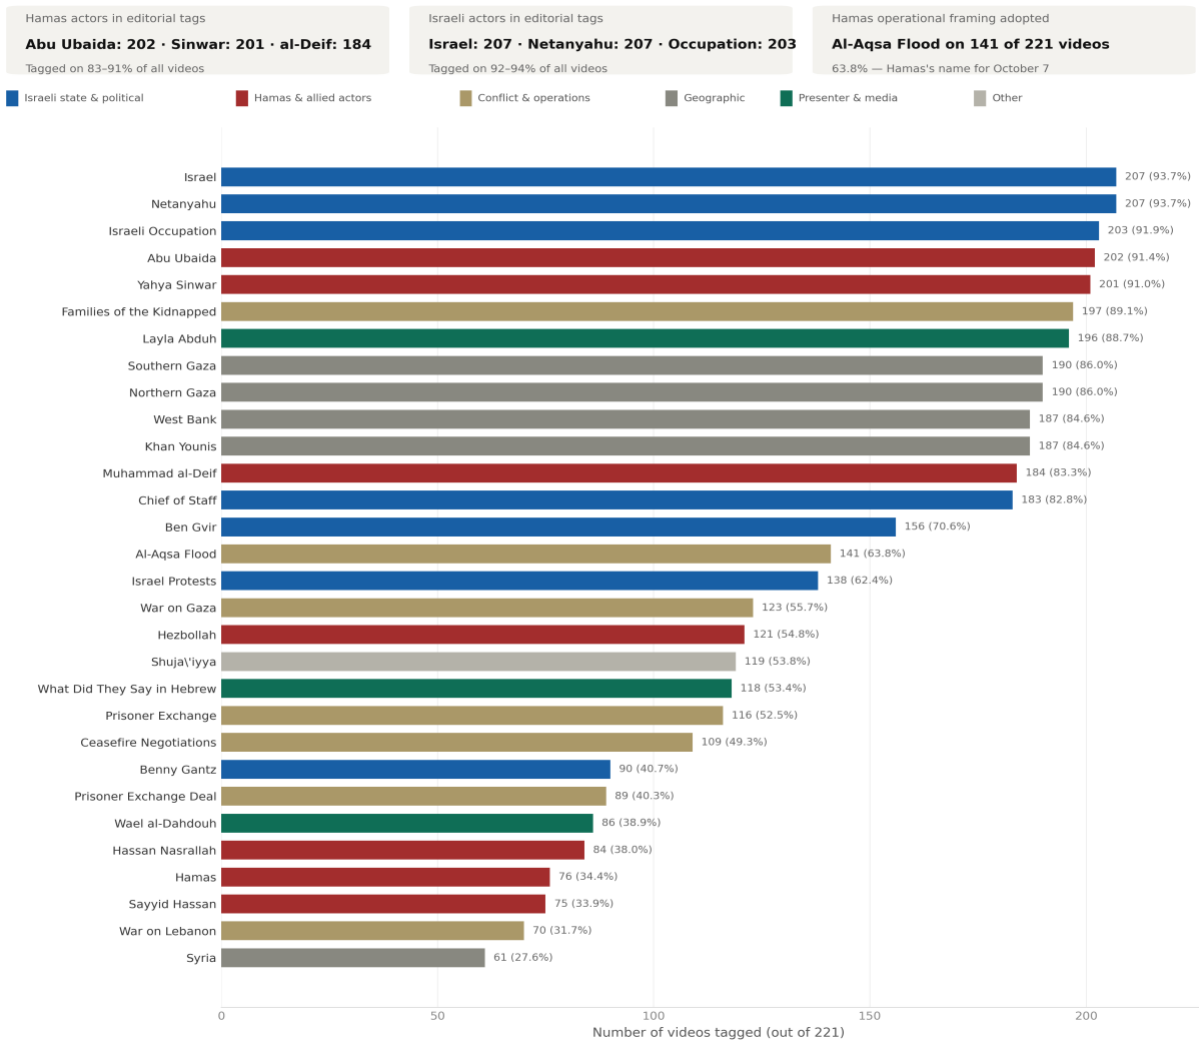

Figure 11: Editorial Tag Priorities — Top 30 Content Tags across 221 Videos

Appendix E: YouTube Policy Violations

This appendix maps the corpus evidence documented in §§5–9 onto YouTube's published policies. The figure below identifies the two policies, cites their sources, and shows the distribution of policy-matching instances across six violation categories. Violent Extremist Organizations (VEO) violations — primarily Hamas glorification — account for 69.7% of all policy-matching content. Confidence is classified by violation density per video: high (20+ instances, 48 videos), moderate (5–19 instances, 134 videos), and low (1–4 instances, remaining videos). The complete video-level evidence table, including individual YouTube comment IDs for independent verification, is available from the author on request.

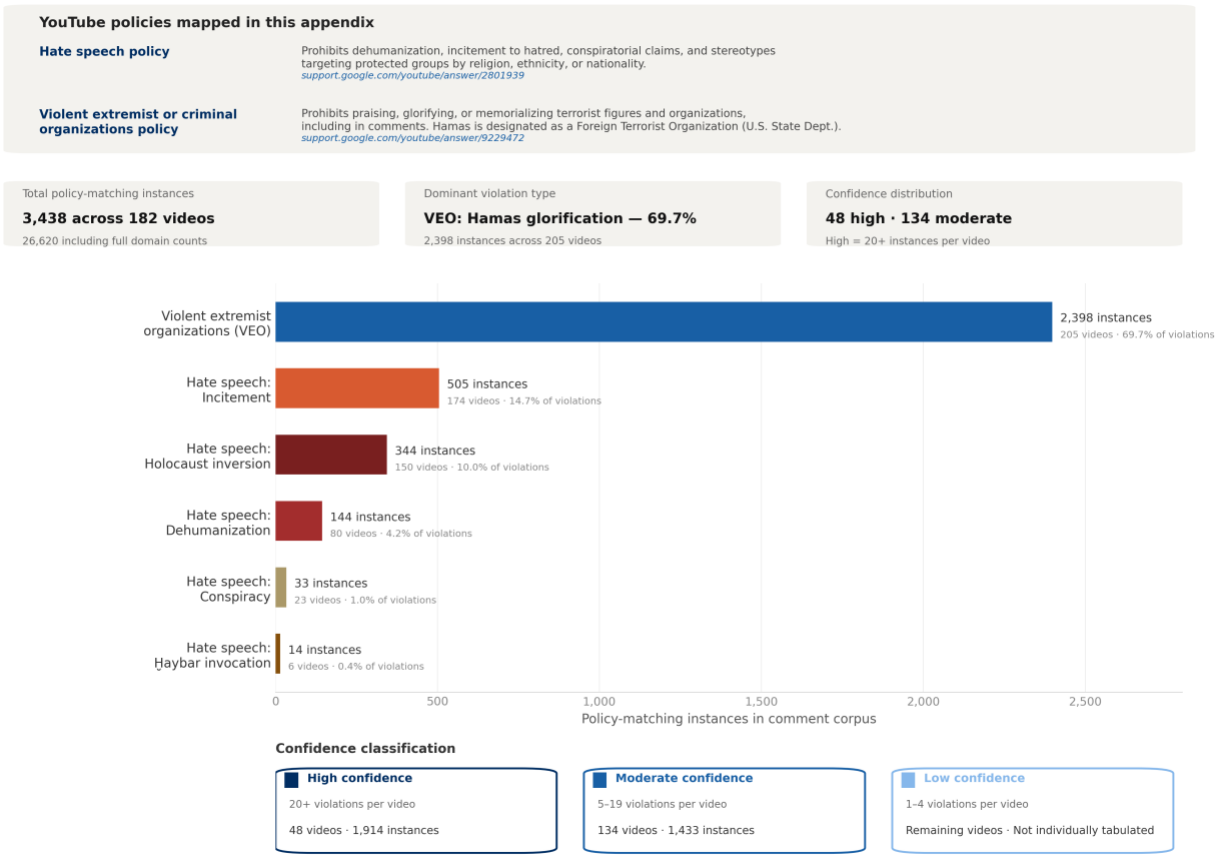

Figure 12: YouTube Policy Violations — Distribution by Category with Policy Citations and Confidence Classification

E.1 Violation categories: policy logic and evidence

1. Violent extremist organizations: Hamas glorification

Policy: Violent Extremist or Criminal Organizations Policy (support.google.com/youtube/answer/9229472). Prohibits praising or memorializing terrorist figures and organizations. Applies to comments.

Designation: Hamas: U.S. Foreign Terrorist Organization (FTO) designation (8 Oct 1997); UK proscribed in entirety.

Logic: Veneration of Hamas leaders as martyrs/heroes; framing fighters as holy warriors; celebrating military operations.

Evidence: 2,398 instances, 205 videos, 69.7% of violations.

Example 1: 817 likes. ID: UgzNWF-NvzDli-of-B94AaABAg. Video: youtube.com/watch?v=MmR5sgMmWh8

Example 2: 485 likes. ID: Ugw9wIEdAux1OvsreOh4AaABAg. Video: youtube.com/watch?v=3WnMxLW6pCk

2. Hate speech: incitement to hatred

Policy: Hate Speech Policy (support.google.com/youtube/answer/2801939). Prohibits encouraging violence against groups based on religion, ethnicity, nationality.

Logic: Divine curses with targeting clauses specifying Jews/Zionists; prayers for destruction.

Evidence: 505 instances, 174 videos, 14.7%.

Example: 72 likes. ID: Ugw1yMvx0E0teZ-kwR94AaABAg. Video: youtube.com/watch?v=e_r5UDJRaSg

3. Hate speech: Nazi/Holocaust inversion

Policy: Hate Speech Policy. IHRA Example 9: drawing comparisons of Israeli policy to Nazis.

Logic: Characterizing Israel as Nazi; comparing leaders to Hitler; applying “holocaust” to Israeli policy.

Evidence: 344 instances, 150 videos, 10.0%.

Example: 32 likes. ID: UgwfNihTMsFnlahDZ5V4AaABAg. Video: youtube.com/watch?v=TSIM43eqJ0E

4. Hate speech: dehumanization

Policy: Hate Speech Policy: dehumanizing by comparing to animals.

Logic: Quranic animal terms (Quran 5:60) directed at Jews/Israelis as collective.

Evidence: 144 instances, 80 videos, 4.2%.

Example: 319 likes. ID: Ugxs8InQsE6fQXXm3i54AaABAg. Video: youtube.com/watch?v=TSIM43eqJ0E

5. Hate speech: conspiratorial claims

Policy: Hate Speech Policy. IHRA Example 2: conspiratorial claims about Jews as collective.

Logic: Attributing inherent conspiratorial malice to Jews/Zionists.

Evidence: 33 instances, 23 videos, 1.0%.

Example: 79 likes. ID: UgyWd4XGobXudkWZn-R4AaABAg. Video: youtube.com/watch?v=IcnZfOSPhlA

6. Hate speech: Ḫaybar invocation

Policy: Hate Speech Policy. IHRA Example 1: calling for violence against Jews in name of religion.

Logic: Explicit call for military violence against Jews by name, referencing Battle of Ḫaybar (628 CE).

Evidence: 14 instances, 6 videos, 0.4%.

Example: 13 likes. ID: UgwDK1opzQjvrhJCodh4AaABAg. Video: youtube.com/watch?v=AUG_MJbGDY0

E.2 High-confidence violations (20+ instances per video)

#	Video ID	Date	Views	Violations	Policy	Video URL
1	3WnMxLW6pCk	2024-10-18	545,267	227	VEO	youtube.com/watch?v=3WnMxLW6pCk
2	MmR5sgMmWh8	2023-11-29	2,936,332	196	VEO	youtube.com/watch?v=MmR5sgMmWh8
3	AUG_MJbGDY0	2024-08-09	717,270	144	VEO	youtube.com/watch?v=AUG_MJbGDY0
4	e_r5UDJRaSg	2023-12-30	1,584,645	73	VEO	youtube.com/watch?v=e_r5UDJRaSg
5	h0HkqD79sbM	2024-02-16	347,516	67	VEO	youtube.com/watch?v=h0HkqD79sbM
6	5Ptbmqox8Z8	2024-10-21	278,936	56	VEO	youtube.com/watch?v=5Ptbmqox8Z8
7	FksE6IAbxPc	2024-01-17	1,893,327	48	VEO	youtube.com/watch?v=FksE6IAbxPc
8	TSIM43eqJ0E	2024-08-02	592,197	45	VEO	youtube.com/watch?v=TSIM43eqJ0E
9	GefrPBdqArM	2024-05-03	341,696	45	VEO	youtube.com/watch?v=GefrPBdqArM
10	7cb8TrMnNlo	2025-01-31	515,829	44	VEO	youtube.com/watch?v=7cb8TrMnNlo
11	QywQtJn8mz8	2023-11-24	576,717	36	VEO	youtube.com/watch?v=QywQtJn8mz8
12	j9J6ZLx4rcc	2024-01-01	392,109	34	HS:HI	youtube.com/watch?v=j9J6ZLx4rcc
13	kQ_BuqR-B_I	2025-06-26	382,077	33	VEO	youtube.com/watch?v=kQ_BuqR-B_I
14	auj9mkaY4KA	2023-11-27	555,838	32	VEO	youtube.com/watch?v=auj9mkaY4KA
15	GggQK7bufvE	2024-10-09	224,749	32	VEO	youtube.com/watch?v=GggQK7bufvE
16	L9mQq4yN61A	2024-02-07	523,797	31	VEO	youtube.com/watch?v=L9mQq4yN61A
17	xUXtch_MqCQ	2024-07-15	186,584	30	VEO	youtube.com/watch?v=xUXtch_MqCQ
18	nuqydd8d3AE	2024-05-29	270,507	30	HS:HI	youtube.com/watch?v=nuqydd8d3AE
19	98BZvJ2p4zc	2024-07-31	229,601	29	HS:I	youtube.com/watch?v=98BZvJ2p4zc
20	TvisKWhoe5Q	2024-04-26	581,737	28	VEO	youtube.com/watch?v=TvisKWhoe5Q
21	dGbTdmfZsfY	2024-06-05	342,432	28	VEO	youtube.com/watch?v=dGbTdmfZsfY
22	jZGhxNslWOQ	2025-07-30	172,380	28	VEO	youtube.com/watch?v=jZGhxNslWOQ
23	QJdlb93gUQU	2023-12-08	596,388	28	VEO	youtube.com/watch?v=QJdlb93gUQU
24	Y1Ep0UCONoA	2023-12-22	1,417,644	28	VEO	youtube.com/watch?v=Y1Ep0UCONoA
25	wieKTmpTcFE	2023-12-20	513,530	27	VEO	youtube.com/watch?v=wieKTmpTcFE
26	aOiG5Vh6p1c	2024-05-27	1,093,118	26	VEO	youtube.com/watch?v=aOiG5Vh6p1c
27	oSDsKPpIAJk	2024-03-11	200,781	26	VEO	youtube.com/watch?v=oSDsKPpIAJk
28	5fjy098nj5E	2023-12-13	608,837	26	VEO	youtube.com/watch?v=5fjy098nj5E
29	CSBX4bv7OJU	2024-03-08	231,226	25	HS:I	youtube.com/watch?v=CSBX4bv7OJU
30	Dp-l0LGYs8U	2024-01-24	442,533	25	VEO	youtube.com/watch?v=Dp-l0LGYs8U

Total: 30 videos shown of 48 qualifying. 1,914 policy-matching instances. Combined views: 33.2M.

E.3 Moderate-confidence violations (5–19 instances per video)

#	Video ID	Date	Views	Violations	Policy	Video URL
1	44ZkMZPVWns	2024-02-02	430,439	19	VEO	youtube.com/watch?v=44ZkMZPVWns
2	9FOOMiY5834	2025-01-17	311,732	19	VEO	youtube.com/watch?v=9FOOMiY5834
3	L25fQ3y4zEE	2024-01-29	345,262	19	VEO	youtube.com/watch?v=L25fQ3y4zEE
4	dibcu6rPuSw	2025-01-15	319,576	19	VEO	youtube.com/watch?v=dibcu6rPuSw
5	1K0eFYbpDWE	2025-04-06	182,235	18	HS:HI	youtube.com/watch?v=1K0eFYbpDWE
6	oi1ZaI4xVsI	2024-01-03	420,424	18	VEO	youtube.com/watch?v=oi1ZaI4xVsI

7	R9j-kxsSzWE	2024-10-11	166,235	18	HS:I	youtube.com/watch?v=R9j-kxsSzWE
8	pi-BN4VRjfQ	2024-07-10	185,773	18	VEO	youtube.com/watch?v=pi-BN4VRjfQ
9	aTZiud1y3pA	2024-04-29	329,623	18	HS:HI	youtube.com/watch?v=aTZiud1y3pA
10	wZ-vlYWyLAY	2024-05-13	300,218	18	VEO	youtube.com/watch?v=wZ-vlYWyLAY
11	iLRByOvXviw	2023-12-01	776,046	18	VEO	youtube.com/watch?v=iLRByOvXviw
12	B8G82DHwUVs	2023-12-18	507,634	17	VEO	youtube.com/watch?v=B8G82DHwUVs
13	Af0kBeFU9OI	2023-12-06	666,842	17	VEO	youtube.com/watch?v=Af0kBeFU9OI
14	9LIimfzAToA	2024-10-30	298,233	17	VEO	youtube.com/watch?v=9LIimfzAToA
15	B3wapD7z1-I	2024-04-19	289,908	17	VEO	youtube.com/watch?v=B3wapD7z1-I
16	9qYYwkt0HNw	2025-04-21	422,433	16	VEO	youtube.com/watch?v=9qYYwkt0HNw
17	WWUE9y1SCjI	2025-02-07	317,779	16	VEO	youtube.com/watch?v=WWUE9y1SCjI
18	Ft3Z1zUIg6Y	2025-03-14	194,100	16	HS:I	youtube.com/watch?v=Ft3Z1zUIg6Y
19	QkucbyuzNEQ	2025-02-05	341,679	16	VEO	youtube.com/watch?v=QkucbyuzNEQ
20	nX427_3qE_Y	2024-07-19	200,028	16	HS:I	youtube.com/watch?v=nX427_3qE_Y
21	34LkOzqIGdk	2024-03-25	238,075	16	VEO	youtube.com/watch?v=34LkOzqIGdk
22	TRQqENx6PE4	2023-11-20	299,855	16	VEO	youtube.com/watch?v=TRQqENx6PE4
23	0tgOZhdmJEc	2024-08-12	240,778	16	HS:I	youtube.com/watch?v=0tgOZhdmJEc
24	NNBirRTACqQ	2025-07-08	208,234	16	VEO	youtube.com/watch?v=NNBirRTACqQ
25	UUu8PkHJM3Y	2025-02-12	337,217	15	VEO	youtube.com/watch?v=UUu8PkHJM3Y
26	bWhebtFYr88	2024-01-19	439,732	15	VEO	youtube.com/watch?v=bWhebtFYr88
27	9TPOEA3x4ms	2024-08-07	262,051	15	VEO	youtube.com/watch?v=9TPOEA3x4ms
28	DVDWEL35E3Q	2025-01-28	386,938	15	VEO	youtube.com/watch?v=DVDWEL35E3Q
29	WVkj21Icq8s	2024-04-24	198,276	14	VEO	youtube.com/watch?v=WVkj21Icq8s
30	lNFDj1bXmbQ	2025-05-14	203,361	14	VEO	youtube.com/watch?v=lNFDj1bXmbQ

Total: 30 videos shown of 134 qualifying. 1,433 policy-matching instances.

E.4 Violation distribution by policy category

Policy	Videos	Instances	%
VEO: Hamas glorification	205	2,398	69.7%
Hate speech: Incitement	174	505	14.7%
Hate speech: Holocaust inversion	150	344	10.0%
Hate speech: Dehumanization	80	144	4.2%
Hate speech: Conspiracy	23	33	1.0%
Hate speech: Ḫaybar	6	14	0.4%

References

Abdul-Raof, H. (2001). *Qurʾān Translation: Discourse, Texture and Exegesis*. Curzon Press.

Al Jazeera. (n.d.). About Us. Al Jazeera Media Network. https://www.aljazeera.com/about-us

Albirini, A. (2016). *Modern Arabic Sociolinguistics: Diglossia, Variation, Code-switching, Attitudes and Identity*. Routledge.

Anderson, B. (1983/2006). *Imagined Communities: Reflections on the Origin and Spread of Nationalism* (Rev. ed.). Verso.

Appadurai, A. (1996). *Modernity at Large: Cultural Dimensions of Globalization*. University of Minnesota Press.

Ayad, M. (2025). Teenage terrorists and the digital ecosystem of the Islamic State. *CTC Sentinel*, 18(2), 1–8. Combating Terrorism Center at West Point.

Bandura, A. (1999). Moral disengagement in the perpetration of inhumanities. *Personality and Social Psychology Review*, 3(3), 193–209.

Barthes, R. (1957/1972). *Mythologies*. Hill and Wang.

Becker, M. J., Troschke, H., Bolton, M., & Chapelan, A. (Eds.). (2024). *Decoding Antisemitism: A Guide to Identifying Antisemitism Online*. Technische Universität Berlin.

Becker, M. J., et al. (2022). Communicating antisemitism: Conceptual and methodological reflections. *Journal of Contemporary Antisemitism*, 5(1), 1–18. https://doi.org/10.26613/jca/5.1.105

Berger, J. M., & Morgan, J. (2015). *The ISIS Twitter census: Defining and describing the population of ISIS supporters on Twitter*. Brookings Institution. https://www.brookings.edu/research/the-isis-twitter-census/

Charmaz, K. (2014). *Constructing Grounded Theory* (2nd ed.). SAGE Publications.

Cherribi, S. (2017). *Fridays of Rage: Al Jazeera, the Arab Spring, and Political Islam*. Oxford University Press.

Cook, D. (2005). *Contemporary Muslim Apocalyptic Literature*. Syracuse University Press.

Dollard, J., Miller, N. E., Doob, L. W., Mowrer, O. H., & Sears, R. R. (1939). *Frustration and Aggression*. Yale University Press.

El-Nawawy, M., & Iskandar, A. (2003). *Al-Jazeera: The Story of the Network That Is Rattling Governments and Redefining Modern Journalism*. Westview Press.

Ellul, J. (1962/1973). *Propaganda: The Formation of Men's Attitudes*. Vintage Books.

Elswah, M., & Alimardani, M. (2021). Digital orientalism: #SaveSheikhJarrah and Arabic content moderation. *POMEPS Studies*, 44.

Elmasry, M. H., El Shamy, A., Manning, P., Mills, A., & Auter, P. J. (2013). Al-Jazeera and Al-Arabiya framing of the Israel–Palestine conflict during war and calm periods. *International Communication Gazette*, 75(8), 750–768.

Entman, R. M. (1993). Framing: Toward clarification of a fractured paradigm. *Journal of Communication*, 43(4), 51–58.

Ferguson, C. A. (1959). Diglossia. *Word*, 15(2), 325–340.

Figenschou, T. U. (2014). *Al Jazeera and the Global Media Landscape: The South is Talking Back*. Routledge.

Filiu, J.-P. (2011). *Apocalypse in Islam* (M. B. DeBevoise, Trans.). University of California Press.

Gillespie, T. (2018). *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. Yale University Press.

Gerbner, G. (2002). Growing up with television: Cultivation processes. In J. Bryant & D. Zillmann (Eds.), *Media Effects: Advances in Theory and Research* (2nd ed., pp. 43–67). Lawrence Erlbaum.

Haeri, N. (2003). *Sacred Language, Ordinary People: Dilemmas of Culture and Politics in Egypt*. Palgrave Macmillan.

Habash, N. (2010). *Introduction to Arabic Natural Language Processing*. Morgan & Claypool.

Herman, E. S., & Chomsky, N. (1988). *Manufacturing Consent: The Political Economy of the Mass Media*. Pantheon Books.

Herf, J. (2009). *Nazi Propaganda for the Arab World*. Yale University Press.

Herf, J. (2018). From antisemitism to anti-Zionism: The past and present of a lethal ideology. *Israel Journal of Foreign Affairs*, 12(3).

Herf, J. (2024). *Three Faces of Antisemitism: Right, Left and Islamism*. Routledge.

Hirschbein, R., & Asfari, A. (2023). *Jews and Muslims in the White Supremacist Conspiratorial Imagination*. Routledge.

Horton, D., & Wohl, R. R. (1956). Mass communication and para-social interaction: Observations on intimacy at a distance. *Psychiatry*, 19(3), 215–229.

International Holocaust Remembrance Alliance. (2016). *Working Definition of Antisemitism*. https://holocaustremembrance.com/resources/working-definition-antisemitism

International Holocaust Remembrance Alliance. (n.d.). *Non-legally binding working definition of Holocaust denial and distortion*. https://holocaustremembrance.com

Küntzel, M. (2020). Islamic antisemitism: Characteristics, origins, and current effects. *Israel Journal of Foreign Affairs*, 14(2). https://doi.org/10.1080/23739770.2020.1793275

Küntzel, M. (2024). *Nazis, Islamic Antisemitism and the Middle East: The 1948 Arab War and the Aftershocks of WWII*. Routledge.

Jigsaw (Google). (2017). Perspective API. https://perspectiveapi.com

Johnson, A. (Ed.) (2024). *Mapping the New Left Antisemitism: The Fathom Essays*. Routledge.

Jowett, G. S., & O'Donnell, V. (2019). *Propaganda and Persuasion* (7th ed.). SAGE Publications.

Klug, B. (2013). Interrogating "new anti-Semitism." *Ethnic and Racial Studies*, 36(3), 468–482.

Krippendorff, K. (2018). *Content Analysis: An Introduction to Its Methodology* (4th ed.). SAGE Publications.

Kumar, R. (2014). *Research Methodology: A Step-by-Step Guide for Beginners* (4th ed.). SAGE Publications.

Lewis, B. (1986). *Semites and Anti-Semites: An Inquiry into Conflict and Prejudice*. W. W. Norton.

Litvak, M., & Webman, E. (2009). *From Empathy to Denial: Arab Responses to the Holocaust*. Columbia University Press.

McCauley, C., & Moskalenko, S. (2008). Mechanisms of political radicalization: Pathways toward terrorism. *Terrorism and Political Violence*, 20(3), 415–433.

Miladi, N. (2006). Satellite TV news and the Arab diaspora in Britain: Comparing Al-Jazeera, the BBC and CNN. *Journal of Ethnic and Migration Studies*, 32(6), 929–953.

Moghaddam, F. M. (2005). The staircase to terrorism: A psychological exploration. *American Psychologist*, 60(2), 161–169.

Lynch, M. (2006). *Voices of the New Arab Public: Iraq, Al-Jazeera, and Middle East Politics Today*. Columbia University Press.

Marcus-Newhall, A., Pedersen, W. C., Carlson, M., & Miller, N. (2000). Displaced aggression is alive and well: A meta-analytic review. *Journal of Personality and Social Psychology*, 78(4), 670–689.

Mohammed, O. (2024, November 13). How do you say Jew in Arabic? *The Times of Israel: The Blogs*. https://blogs.timesofisrael.com

Popper, K. (1945/2011). *The Open Society and Its Enemies*. Routledge Classics.

Rogers, E. (2025). The need for greater transparency in the moderation of borderline terrorist and violent extremist content. *Internet Policy Review*, 14(3). https://doi.org/10.14763/2025.3.2012

Rabab'ah, G., & Alali, N. (2019). Impoliteness in reader comments on the Al-Jazeera channel news website. *Journal of Politeness Research*, 16(1), 1–22.

Rouhana, T. (2023). Critical discourse analysis guided topic modeling: The case of Al-Jazeera Arabic. *Information, Communication & Society*, 26(5), 904–922.

Rinnawi, K. (2006). *Instant Nationalism: McArabism, al-Jazeera, and Transnational Media in the Arab World*. University Press of America.

Roberts, S. T. (2019). *Behind the Screen: Content Moderation in the Shadows of Social Media*. Yale University Press.

Samuel-Azran, T. (2013). Al-Jazeera, Qatar, and new tactics in state-sponsored media diplomacy. *American Behavioral Scientist*, 57(9), 1293–1311.

Samuel-Azran, T., Assaf, I., Salem, A., Wahabe, L., & Halabi, N. (2016). Is there a Qatari–Al-Jazeera nexus? Coverage of the 2022 FIFA World Cup controversy. *Global Media and Communication*, 12(3), 195–209.

Seib, P. (2008). *The Al Jazeera Effect: How the New Global Media Are Reshaping World Politics*. Potomac Books.

Shkedi, A. (2019). *Data Analysis in Qualitative Research: Practical and Theoretical Methodologies with Optional Use of a Software Tool*. Mofet Institute.

Stern, K. S. (2019). *The Conflict over the Conflict: The Israel/Palestine Campus Debate*. New Jewish Press.

Suleiman, Y. (2003). *The Arabic Language and National Identity: A Study in Ideology*. Edinburgh University Press.

Suleiman, Y. (2011). *Arabic, Self and Identity: A Study in Conflict and Displacement*. Oxford University Press.

Tajfel, H., & Turner, J. C. (1986). The social identity theory of intergroup behavior. In S. Worchel & W. Austin (Eds.), *Psychology of Intergroup Relations* (pp. 7–24). Nelson-Hall.

Van Dijk, T. A. (2006). Ideology and discourse analysis. *Journal of Political Ideologies*, 11(2), 115–140.

Versteegh, K. (2014). *The Arabic Language* (2nd ed.). Edinburgh University Press.

Webman, E. (2013). Treading in troubled waters: Seeking the roots of Muslim antisemitism. *Bustan: The Middle East Book Review*, 4(2), 113–136.

Wistrich, R. S. (2010). *A Lethal Obsession: Anti-Semitism from Antiquity to the Global Jihad*. Random House.

Wodak, R., & Reisigl, M. (2001). *Discourse and Discrimination: Rhetorics of Racism and Antisemitism*. Routledge.

Platform Policy Sources

Zayani, M. (2005). *The Al Jazeera Phenomenon: Critical Perspectives on New Arab Media*. Pluto Press.

YouTube. (n.d.). Community Guidelines. Google Support. https://support.google.com/youtube/answer/9288567

YouTube. (n.d.). Hate speech policy. Google Support. https://support.google.com/youtube/answer/2801939

YouTube. (n.d.). Violent extremist or criminal organizations policy. Google Support. https://support.google.com/youtube/answer/9229472

YouTube. (n.d.). Advertiser-friendly content guidelines. Google Support. https://support.google.com/youtube/answer/6162278

YouTube. (n.d.). Monetization is disabled for my channel. Google Support. https://support.google.com/youtube/answer/1727191

Government Designation Sources

U.S. Department of State. (n.d.). Foreign Terrorist Organizations. https://www.state.gov/foreign-terrorist-organizations

UK Home Office. (2021). Proscribed terrorist groups or organisations. GOV.UK. https://www.gov.uk/government/publications/proscribed-terror-groups-or-organisations--2

Data Sources

Neuendorf, K. A. (2017). *The Content Analysis Guidebook* (2nd ed.). SAGE Publications.

Pew Research Center. (2025). Muslim population change. In *The Future of World Religions: Population Growth Projections, 2025 Update*. https://www.pewresearch.org/religion/2025/06/09/muslim-population-change/

Brill. (2010). Brill's simple Arabic transliteration system (sAts), version 1.0. By Pim Rietbroek, December 14, 2010.