Social Media Narratives Across Platforms in Conflict: Evidence from Syria

Erin Walk, University of Pennsylvania
Elizabeth Parker-Magyar, Harvard University
Kiran Garimella, Rutgers University
Ahmet Akbiyik, Harvard Kennedy School
Fotini Christia, Massachusetts Institute of Technology

How do representations of violent conflict differ across social media platforms? We constructed and analyzed comparable datasets of public messages and images from elite- and meso-level Syrian actors posting on three popular social media platforms. Our findings show that complementary if divergent discussions of violence remain central even amid a period of relative de-escalation. Narratives on Twitter reaching an international audience contextualize violence within the conflict's master cleavages, while on Telegram, they address a more local audience and emphasize the violence's day-to-day impacts. A site with stricter surveillance, Facebook features more loyalist narratives. Paired with a sample of users' responses to an open-ended questionnaire, our results show that users across platforms diversify their presentation of violence to reach domestic and international audiences and to accommodate technical affordances, with consequences for how both audiences and researchers understand ongoing conflict.

More than 10 years after the outset of protests calling for the removal of Syrian President Bashar Al Assad, Syria's ongoing war has killed hundreds of thousands and forced more than half of the country's population to flee their homes. Today, foreign parties have intervened on all sides, and the war has long since taken on ethnic and sectarian dimensions. Social media has played a key role from the outset of the Syrian uprising, as a principal avenue for Syrians to internally organize, share information, and make sense of the conflict. Over time, the social media platforms Syrians use have proliferated and shifted. How and why do Syrians' representations of the ongoing violence differ as they communicate across platforms?

Like social media users more broadly, individuals in conflict use social media for a broad variety of purposes. They can use social media to signal their support for protest movements (Steinert-Threlkeld 2017), to recruit (Mitts 2019), to spread hate speech (Siegel and Badaan 2020), or to access information amid heightened insecurities (Schon 2021). Research shows that individuals in conflict face a variety of considerations when they post online, such as their international and domestic audiences (Esberg and Siegel 2023; Zeitzoff 2017), whether their conversations are public or private (Mitts 2022), or the extent to which they face repression or surveillance (Gohdes 2020).

We further this research by underscoring how the different ways those experiencing conflict communicate publicly

Erin Walk (erin.walk@asc.upenn.edu) is a postdoctoral fellow at the Annenberg School for Communication at the University of Pennsylvania, Philadelphia, PA 19104. Elizabeth Parker-Magyar (eparkermagyar@fas.harvard.edu) is the Moulay Hicham Alaoui Postdoctoral Fellow at the Weatherhead Center for International Affairs at Harvard University, Cambridge, MA 02138. Kiran Garimella (kg766@rutgers.edu) is an assistant professor of library and information science at Rutgers University, New Brunswick, NJ 08901. Ahmet Akbiyik (ahmetakbiyik@g.harvard.edu) is a PhD candidate in Political Economy and Government at the Harvard Kennedy School, Cambridge, MA 02130. Fotini Christia (cfotini@mit.edu) is Ford International Professor of the Social Sciences and director of the Institute for Data, Systems, and Society (IDSS) at the Massachusetts Institute of Technology, Cambridge, MA 02142.

This project received IRB approval from MIT COUHES, #2104000360. Kiran Garimella thanks the Michael Hammer Fellowship for financial support. Replication files are available in the *JOP* Dataverse (https://dataverse.harvard.edu/dataverse/jop). The empirical analysis has been successfully replicated by the *JOP* replication analyst. An online appendix with supplementary material is available at https://doi.org/10.1086/732997.

Published online March 17, 2025.

 ooo

across multiple mediums or platforms can shape the aggregate representation of a conflict's violence. Online and offline, the ways elite and meso-level actors construct narratives around ongoing violence (Brubaker and Laitin 1998; Horowitz 1992), or contextualize ongoing events by embedding them in broader patterns of meaning, drive individuals' subsequent mobilization and beliefs about escalation or termination of conflict (Corstange and York 2018; Fabbe, Hazlett, and Sınmazdemir 2019). These online representations have also become fundamental avenues for foreign audiences to understand ongoing conflict (Lynch, Freelon, and Aday 2014).

To study these narratives, we construct comparable datasets of elite- and meso-level Syrian actors with social media accounts posting publicly on three widely used social media platforms in Syria in the war's later years, when fighting continues even as it is unlikely to shift the war's trajectory toward regime persistence and foreign intervention remains salient. We study Twitter, the staple social medium in academic research; Telegram, which many Syrian citizens use as a principal source of information about the conflict; and Facebook, the most widely adopted social media platform globally and in Syria. In total, we analyzed 7.2 million messages and 1.5 million images posted across the three platforms between a turning point in violence in October 2017 and December 1, 2020. We supplement our findings by analyzing messages from users who maintain accounts and post on all three platforms and among users who share an overarching affiliation in the conflict. We use the responses to a structured, open-ended questionnaire we sent to a stratified random sample of users in our dataset to draw out the mechanisms driving our findings.

Our results show that discussions of violence in its many forms remain central even during a period of relative de-escalation while pointing to the distinctive ways that individuals use these respective platforms in the Syrian context. We highlight that, in addition to these platforms' technical affordances (Bucher and Helmond 2018), the domestic or international audiences that users reach drive the aggregate narratives that social media users in conflict present. In the Syrian conflict, Twitter comprises a channel to foreign audiences such that narratives are more likely to concentrate on political dynamics, macronarratives, and outside interveners. Telegram is a platform that captures a more local story for Syrians experiencing the conflict, where content centers on the war's quotidian violence and its direct effects. Facebook fits somewhere in between as a broader and more social platform for conversation among Syrians that more heavily features loyalist narratives.

Our article makes three primary contributions. First, it presents evolving narratives of an ongoing conflict during its later stages. Amid attention to how narratives shape protest mobilization and civil war onset, there have been fewer examinations of narratives during periods where violence has de-escalated, especially online. Our results suggest that, during this phase, users of different platforms who follow actors even with similar overarching affiliations will often consume broadly different narratives around a conflict. Second, our findings emphasize that users' depiction of violence differ not only according to their social context, status in exile, or exposure to international scrutiny (Esberg and Siegel 2023; Zeitzoff 2017) but also according to the social media platforms they adopt. Third, researchers increasingly use available social media data to measure concepts and outcomes of interest. By highlighting how actors adjust behavior to accommodate technical affordances and to instrumentally reach different audiences, we underscore how researchers' selection of platforms of study can impact their inferences when using social media data to study concepts of interest.

THEORY: NARRATIVES AROUND VIOLENCE ON SOCIAL MEDIA

Violent conflicts feature a secondary struggle over their meaning (Horowitz 1992 cited in Corstange and York 2018). Because the processes through which conflict participants ascribe meaning to ongoing violence can legitimize or undermine future violence (Brubaker and Laitin 1998; Corstange and York 2018; Fearon and Laitin 2000), conflict actors often compete to frame violence in favorable terms in order to maximize support. Beyond strategic outcomes, the collective adoption of narratives around violence underpins the construction of collective identities in conflict (Pearlman 2016).

Given social media's increasing centrality to conflict processes (Zeitzoff 2017), the information, opinions, and perspectives that elite- and meso-level actors share on these platforms provide insight into the meaning they ascribe to ongoing events within a conflict and contribute to subsequent mobilization, identity construction, and attitudes toward termination. Amid scrutiny of how the availability of social media reshapes conflict and how behavior online does or does not diverge from behavior offline, less is known about how narratives and behavior differ across social media platforms during conflict. A broad literature in communications emphasizes the affordances that structure users' range of actions and interactions on a platform. The concept refers to what material technologies such as media technologies allow people to do (Bucher and Helmond 2018, 2). We distinguish between what we term the technical affordances of a platform—for instance how the platform maps onto social network structure and whether the platform blocks actors or truncates speech—and the audience affordances that a platform allows users to reach.

We place particular emphasis on how the diverse audiences that actors can reach on social media shape their presentation of the conflict. Building on evidence that changes in exilees' domestic and international networks underpin shifts in their dissent (Esberg and Siegel 2023), we primarily focus on the different ways that individuals frame violence when reaching or mobilizing domestic and international audiences. Divisions across domestic and international audiences should be especially salient in highly polarized conflict environments where foreign intervention and support from diaspora actors can have disproportionate impact on conflict outcomes (Walter 1997). With some forms of social media constituting a primary avenue to reach international audiences (Lynch et al. 2014), different platforms can then become grounds for narrative contestation at the international level (Zeitzoff 2017) as users attempt to influence international public opinion about the nature and meaning of ongoing violence. These audience considerations are especially likely to be salient in conflicts where international intervention shapes outcomes.

Social media also remains a site for connection among more proximate, domestic communities. Even in protracted conflicts where daily violence is unlikely to shift conflict trajectories, this violence nevertheless continues to generate uncertainty and anxiety for those whose future is tied to its outcome (Pearlman 2016), and violence shifts form rather than diminishes entirely (Kalyvas 2006). Those experiencing conflict highly value information, which is needed to maintain personal safety or make life-altering decisions like when to flee a frontline or when to return home (Schon 2021). Because these individuals value practical information around the impact of that violence and its immediate local trajectory, we should expect individuals will share more discussion of local violence and its impacts when addressing communities most directly impacted by conflict.

The diverse audiences that users are able to reach across platforms should shape the information and analysis they share around events, shaping aggregate narratives across platforms. On platforms where individuals experiencing conflict communicate with one another, actors are more likely to share details around the immediate, local, and quotidian impacts of violence. Conversely, where actors are also communicating with external audiences, they are more likely to frame events within a conflict's master cleavage, aggregate victimization, and its broader trajectories while neglecting highly local, low-level, or quotidian violence. As a result, aggregate narratives about the causes and consequences of violence can differ not only according to individuals' high-level loyalties or affiliation but also according to the context in which they describe events. In this case, narratives should differ across platforms even among users who are in the same parties in a conflict. These narratives can have substantive impact because they shape individuals' continued support for conflict processes like mobilization and termination (Brubaker and Laitin 1998; Corstange and York 2018; Horowitz 1992).

Technical affordances also impact the messages individuals share. This study examines communication from elite- and meso-level actors across three social media platforms, each with salient technical affordances: Twitter, Facebook, and Telegram. A staple medium in academic research, Twitter's hashtag and retweet technical affordances and norm of mostly public usage can allow content to reach broad audiences (Bonilla and Rosa 2015). Both Twitter and Facebook regularly remove content for violating its guidelines. Facebook also more stringently requires individuals to verify their offline identity, heightening security risks in a context where all actors use online information to surveil and target (Gohdes 2020). While Telegram gained international and academic attention for its initial reticence to filter extremist content, it is also well known for its private messaging rooms, perceived security, and ease of access, which can foster strong community-focused groups (Urman, Ho, and Katz 2021).

These technical affordances amount to important differences in the communities that users reach across these platforms. While Twitter comprises a channel to a foreign audience, Telegram is valued as a site for communication among socially proximate communities. Facebook, with its broad user base but norm of connections among offline social contacts, fits somewhere in between. As we discuss below, the technical affordances and divergent audiences ultimately shape how users present violence on each of them.

CASE: CIVIL WAR IN SYRIA

We study how diverse audiences and technical affordances shape the way elite- and meso-level actors report, describe, and contextualize violence during a period of relative de-escalation in Syria's civil war. Syria's war evolved from a 2011 protest movement that centered on removing Syrian President Bashar Al-Assad from power, with regime opponents tying their movement to the Arab Spring protests that elsewhere toppled the Egyptian and Tunisian presidents. Protests evolved by 2012 into a largely symmetric armed conflict that has killed more than 400,000, forced millions to flee, and internally displaced millions more.

The war has been shaped by foreign intervention since its earliest months. Russia, Iran, Turkey, the United States and its coalition partners, Gulf nations, the Lebanese militia Hezbollah, and Israel have all intervened. The period of our study from October 2017 through 2020 coincides with a de-escalation in the conflict's violence, beginning with the US-backed International Coalition's defeat of the Islamic State

in its self-proclaimed capital of al-Raqqa. Four principal actors controlled swathes of territory in this period: the Syrian government and its allies, the US-backed Syrian Democratic Forces (SDF), Turkey and affiliated armed groups, and the formerly Al-Qa'eda–affiliated Hayat Tahrir Al Sham.

The collective ebb of violence during this period led Syria to fall from international headlines. Throughout our study time frame, foreign media consistently reported that the war was drawing to a close and that Syrian President Bashar Al-Assad had mostly won (Hubbard 2020). But on the ground, symmetric and insurgent violence continued, and some regions witnessed the worst episodes of displacement since the conflict began. Even in more stable areas and before the worsening of the COVID-19 pandemic, the country's currency collapsed and Syrians struggled to access basic goods. While redoubled intervention shifting the overall course of the conflict did not occur, no settlement had been reached determining postconflict governance in several regions.

Amid this fragmentation and foreign intervention, major parties in the conflict continued to articulate different narratives about the causes and consequences of violence. The pro-opposition narrative centers on Syrians' continued struggle for freedom against Bashar Al Assad's unceasing authoritarian violence. This narrative combats a loyalist or progovernment narrative of Bashar Al Assad's visionary leadership in combating foreign conspiracies and sectarian hatred. In survey research, Corstange and York (2018) show how competing descriptions of the conflict drive support for different parties among Syrian refugees in Lebanon. Alrababa'h and Blaydes (2021) point to related regime narratives, emphasizing how Syrian regime media focuses on Assad's personality and conspiracies against the state.

Beyond the macrolevel narratives espoused by the conflict's principal parties, a separate axis de-emphasizes support for any party contesting the war, concentrating on Syrians' suffering, sacrifice, and increased ambivalence toward the conflict. In interviews with Syrians in exile, Pearlman (2016) underscores how collective narratives centered on overcoming political fear of the Syrian government steadily transformed over time into nebulous fears of an uncertain future amid protracted violence (21). Fabbe et al. (2019) use a survey experiment to show that framing violence as suffering rather than sacrifice generates support for conflict termination and reconciliation among Syrians surveyed in Turkey. Wedeen (2019, 3) studies cultural production in government-held areas and points to the continued, demobilizing apprehensions of what she terms an ambivalent middle. As parties splintered and violence escalated, observational survey evidence underscored high levels of ambivalence, with 80% of Syrian refugees in Jordan and Lebanon reporting that no party represents them (Arab Barometer Data Analysis Tool 2022).

Narratives of violence across social media platforms in Syria

These narratives are well represented online. Our research builds on data from three of the most prominent platforms in Syria: Twitter, Facebook, and Telegram. Though exact statistics are not available on traffic to these sites from Syrians domestically and abroad, all three platforms are widely used. Alexa, a website that ranks national-level web traffic, showed that Facebook, Telegram, and Twitter were the third, tenth, and twelfth most frequently visited websites in Syria as of December 2021. Google Trends data shows similar patterns (app. fig. A.13), with searches for downloading Facebook outstripping those for Twitter and Telegram throughout the period of our study.

Due to the centrality of the conflict to life in Syria, we expect posts across these platforms to feature conversation on violence, albeit with different emphasis. Syrian journalists, civil society leaders and organizations, and armed actors have gained prominence on Twitter, as the platform has also increasingly adopted automated content policies that aggressively remove content flagged as extremist or gory. The Google Trends data in appendix figure A.13 affirms Twitter's centrality among opposition communities, as searches for Twitter downloads are highest in two highly contested governorates largely outside of regime control.

Given Twitter's technical and audience affordances, the site features greater conversation and interaction between these Syrian users and foreign audiences, whose support can provide humanitarian aid and, through intervention, influence overall conflict outcomes. We therefore expect that aggregate discussions of violence on Twitter are more likely to be embedded with macronarratives of the conflict, especially those accessible to foreign audiences. These audience contexts and technical affordances should produce narrative differences that do not solely represent differences in the composition of users on these different platforms. Instead, we expect differences in narratives between Twitter and other platforms to persist even within parties in the conflict: Pro- and anti-Assad camps, and even the same users posting on multiple platforms, should engage more with the macronarratives of conflict on Twitter than on other platforms.

Syrians have also adopted Telegram as a tool for secure connection with one another, given ease of use as a one-stop platform for direct messaging and news and a perception that the service is faster in areas with limited connectivity. Unlike the other two platforms we study, Telegram has far lower levels of international adoption. As a result, we expect

Table 1. Expectations on How Audience and Technical Affordances Impact Aggregate Narratives of Conflict Across Platforms

Platform	Twitter	Facebook	Telegram
Audience	More international	Between local and international 1-Norms of offline acquaintance	More Local
Relevant technical affordance	Hashtags allowing for international reach	Requirement to verify offline persona	Fast connection
	Limitations on sharing violent images	Limitations on sharing violent images	Perceived anonymity and security
Aggregate Narratives	Discussion of violence within macrolevel cleavages	Discussion of violence between macrolevel cleavages and local impact	Microlevel discussion of violence
	More representation of opposition narratives	More representation of loyalist narratives	

that the users in our study will more often frame violence around its local dimensions and day-to-day costs. As with Twitter, we expect these differences to persist even among individuals who share overarching affiliations and among those who post across multiple platforms.

While Facebook has a broad, international user base, the platform is more likely than Twitter to feature engagement among individuals within proximate communities that know one another offline. Given this audience context, we expect overarching representation of violence on Facebook to fit somewhere between Twitter and Telegram, with narratives of violence more closely reflecting local-level concerns even as users also engage in higher levels of distinctly social or cultural conversation. The Google Trends data aligns with our data collection results below in indicating that Facebook adoption is higher in loyalist communities. We point to one salient technical affordance: Facebook's requirement that users verify their offline identity, which renders it more difficult to post anonymously in a context with high levels of surveillance. In comparison with the other platforms' overt popularity in opposition-held areas, we therefore expect Facebook to feature more overarching loyalist narratives. As with Twitter and Telegram, we anticipate that these narrative differences will persist even when exploring narratives with affiliations and across users posting on multiple platforms. Table 1 summarizes these expectations.

Research Design and Data Collection

To examine depictions of violence on these platforms, we created parallel processes to identify, collect data from, and analyze Syrian-run accounts, channels, and groups posting public messages across these platforms. The accounts belong to what we term elite- and meso-level actors: Syrian politicians, journalists, armed actors, activists, and others who produce public-facing content that reach relatively wide audiences. While this process aimed to limit selection effects, the different nature of the platforms ultimately means that our broadest samples comprise different populations. Underscoring the replicability of our process, we believe the results we present below help to reveal the selection effects inherent in the study of social media users on any singular platform.

For each platform, we followed a process in which a Syrian researcher first used keyword searches in Arabic to select prominent accounts that were then supplemented with users who interacted with the initial sample. At each stage, we screened for Syrian-run accounts focused on commenting on Syrian news, culture, or the economy, filtering out accounts not explicitly focused on Syria generally or one of seven northern governorates.[1] Our keywords' focus on northern Syria in this period means that we capture many narratives from areas inside and outside of the Syrian regime's surveillance that have been contested through this period. The list of seed words used for searches are in appendix A.2, as is information specific to the process for each platform.

We take two steps to examine how audience and technical affordances shape narrative differences while addressing possible differences in the composition of users across platforms. We first searched for each user in our initial dataset across all platforms, creating a subset of 124 users that posted on all three platforms. We then also coded each user's political leaning (opposition, government, Kurdish groups, or neutral).

The full dataset includes 2,106 active Twitter accounts, 657 active Telegram accounts, and 2,124 active Facebook accounts. Choosing messages within the date range, the total dataset included 3,586,469 messages on Telegram, 1,787,552

1. We define northern Syria as the provinces of Aleppo, Deir Al-Zour, Hama, Hassakeh, Idlib, Qasmishli, and Al-Raqqa.

Table 2. Affiliation Data According to Platform, Images Across Platforms, and Number of Parallel Accounts Across Platforms

Platform	Messages	Leaning	Perc. Accts.
Twitter	1,787,552	Kurdish	2.8%
(n = 2,106)		Loyalist	6.9%
		Neutral	10%
		Opposition	79%
Telegram	3,586,469	Kurdish	3.7%
(n = 657)		Loyalist	6.5%
		Neutral	10.6%
		Opposition	79.0%
Facebook	1,793,444	Kurdish	10.9%
(n = 2,124)		Loyalist	52.0%
		Neutral	11.4%
		Opposition	25.7%

messages on Twitter, and 1,793,444 messages on Facebook.[2] Consistent with the Google Trends data, the Facebook data contain more Loyalist-identified pages compared with Twitter and Telegram (table 2). For all of these accounts, we also downloaded the images attached to each message. We collected 1.5 million images across these platforms, the majority of them from Telegram (table 3).

While we also coded each users' location of focus and manually removed users who identified themselves as posting from outside of Syria, the models we run below do not concentrate on location-based data because these platforms make it difficult to verify the location that all but a few users post from. There is no reason to believe that the accounts we selected on any of the platforms are more or less likely to be inside or outside of the country. Many opposition news sites feature the works of internal correspondents working with editors outside of Syria, limiting the utility of a binary distinction between internally or externally run organizations. All of these platforms are accessible across Syria, though users face intermittent internet blackouts and threats from regime monitoring (Gohdes 2020). In collecting data retroactively, we are unable to analyze individual posts or data from accounts that have been deleted. By choosing a more recent date to begin our analysis, we necessarily truncate our dataset but limit problems related to missing data. There is no reason to believe differences in data removal across platforms would bias our subsequent results.

2. In terms of word tokens, Telegram has 991 tokens that appear in at least 10,000 messages, Facebook 733, and Twitter 402. We do not find that Telegram's larger number of messages exert a disproportionate impact on topic selection.

Table 3. Number of Images Collected Across Platforms in the Dataset

Platform	Images
Facebook	257,960
Telegram	901,913
Twitter	400,341
Total	1,560,214

We pay close attention to the 124 users who post across all three platforms, who comprise Syrian journalists, media groups, and a few armed actors (app. table A.3). These 124 users post on 403 total accounts, and in our sample, we analyze 1,151,269 messages on Telegram, 264,929 tweets on Twitter, and 346,605 posts on Facebook (table 4).

Comparing data from the users who post across all three platforms allows us to more directly observe how audiences and technical affordances shape aggregate narratives, as it holds constant broader dynamics of selection onto different platforms. Figure 1 illustrates these differences, highlighting how the same user displays slightly different information about the conflict on three different platforms on the same day (September 10, 2020). This sample account, *Raqqa is Being Slaughtered Silently*, is a well-known media collective that has covered events in Al-Raqqa even through the period of ISIS control. The figure shows that the account posts uniquely in English on Twitter, sharing a remembrance post and a stylized sketch of a nonviolent protester killed by the Al Assad regime during 2011's revolutionary protests. On Telegram, the same account posts uniquely in Arabic only about ongoing violence. On Facebook, the account posts in Arabic about both ongoing violence and a remembrance of the revolutionary protester; however, the image of the protester is not stylized. When we quantitatively examine differences in expression among these accounts, we call data from these accounts the Parallel dataset, as opposed to the All Groups dataset.

Table 4. Number of Accounts Posting Across Multiple Platforms in the Dataset

Platforms	Parallel Accounts
Facebook–Twitter	153
Facebook–Twitter–Telegram	124
Facebook–Telegram	27
Twitter–Telegram	10

Figure 1. Users holding accounts across different platforms occasionally post different information on these platforms. On Twitter (left panel), a well-known pro-opposition outlet from Al Raqqa posted only in English a remembrance of a peaceful protester on September 10, 2020. On Telegram (center panel), the same outlet posted only news in Arabic of an ongoing clash between ISIS and government forces. On Facebook (right panel), the same outlet posted news of ongoing violence as well as a shorter remembrance of the protester, only in Arabic.

Our final models are limited to messages in Arabic, a decision we discuss in appendix A.2. Consistent with our expectations of how Syrians use these platforms, a basic analysis in appendix figure A.11 shows that Twitter users are more likely to tweet in multiple languages than users on Facebook or Telegram, even among the subset of users posting across all three platforms. Increased English language usage on Twitter echoes observations from other settings where international intervention is salient (Driscoll and Steinert-Threlkeld 2020). As we examine only Arabic language posts, our analysis is then likely conservative regarding the total difference in narratives across platforms.

Analyzing text, images, and open-ended user responses

Our primary analysis is based on unsupervised text analysis methods that we supplement with analysis of images collected alongside the messages in our dataset as well as a subset of users' responses to an open-ended questionnaire.

Text analysis: structural topic models and word embeddings. We first analyze the data in the form of Structural Topic Models (STMs; Roberts et al. 2013). Topic modeling techniques identify topics in text data, assigning each word a probability of belonging in a given topic. To prepare the text for these models, we performed standard preprocessing, first removing messages that are not in Arabic and then removing stop words and stemming the Arabic text (Nielsen 2017). Further details are available in appendix A.2.4. We include both unigrams, or single words, and bigrams, or words that are frequently co-located. Including bigrams provides additional context, as words may have different meanings when used together. We ran our models on a combined corpus with a platform variable covariate to compare and explore heterogeneity across platforms, with 25 topics to maximize coherence.

For all structural topic model analysis, expert annotators labeled each topic by looking at the 30 most salient keywords in that topic both in terms of frequency "F" and frequency and exclusivity "FREX." We present the English translations of the topic labels in the main results, with the Arabic provided in appendix A.1, figures A.8 and A.9. Below, we run the models in one combined corpus across platforms, among users posting across all platforms, and by affiliation (loyalist or opposition) across platforms. In appendix A.3.2, we present the results of models run on corpuses of just the text from each platform, showing that the combined corpus does not overrepresent topics from one single platform.

We also use word2vec to examine the contextual similarity of key conflict-related terms (Mikolov et al. 2013). Word2vec models the vocabulary of a text as a set of vectors, where each word has its own individual vector. A high cosine similarity represents a smaller distance between vectors, indicating that the two words being compared are either interchangeable, such as "cat" and "feline," or frequently used together, such as "cat" and "meow." We run our word2vec model with Continuous Bag of Words embeddings using cleaned and tokenized text with stop words included. Unlike in STMs, stop words are beneficial in word2vec models as they provide additional context and sentence structure to better understand similarities between words. Like the STMs, the word2vec model includes bigrams. To compare the words most similar to a given anchor word across platforms, we run a separate model for each platform.

Image analysis. To analyze images, we extracted image features from a deep convolutional neural network, ResNet-50 (He et al. 2016), that was pretrained on the Imagenet dataset (Deng et al. 2009). Imagenet is a large-scale hierarchical image database where images are organized according to Wordnet (Miller 1995), a lexical database of English. This process is in

line with other recent studies studying conflict and protest using images (Mitts, Phillips, and Walter 2022; Steinert-Threlkeld, Chan, and Joo 2022; Zhang and Peng 2024).

ResNet-50 converts each image to a feature vector that allows images to be compared and grouped with one another. Though the actual mechanism is slightly different, to human viewers this will manifest as grouping together images with similar colors and shapes. Pretraining the Resnet-50 model optimizes it to identify features specified in the Imagenet dataset with fewer mistakes. Thus, we can efficiently extract these same image features from our own dataset. From there, we categorized similar images based on the extracted features using k-means clustering (Dehariya, Shrivastava, and Jain 2010). We used the elbow method (Syakur et al. 2018) to determine the optimal number of clusters (app. A.4).

Annotators labeled clusters according to the images they contain. Where the model produced duplicated groupings—for instance, multiple groupings of military vehicles pictured from different perspectives—we combined similar clusters in our analysis. We then calculate the estimated topic proportion of each cluster across platforms. For additional methodological details, please see appendix 4.

Open-ended questionnaire. We contextualize our findings from this quantitative analysis with user responses to an open-ended questionnaire. The short, open-ended questionnaire focused on users' history of social media use, their goals as social media users, how their social media use has evolved over time, how platforms' policies shape their expression, and their understandings of the audience for their writing. Given our existing networks and limitations on freedom of expression in government-controlled Syria, we chose to share this questionnaire among a random sample of users we had identified as opposition affiliated. A Syrian researcher delivered a message to the account introducing the project, identifying it with the research team, asking for permission to share a set of written questions, and inviting the use of more secure messaging platforms.

Technical affordances and ethical considerations (app. A.5) shaped our ability to share the questionnaire as well as response rates with all users. In total, we sent requests to complete the questionnaire to 43 respondents, of whom 37 responded: 14 who manage accounts on multiple platforms, 6 who manage accounts only on Facebook, 10 who manage accounts only on Telegram, and 7 who manage accounts only on Twitter.

RESULTS: NARRATIVES OF VIOLENCE ACROSS PLATFORMS

We first present a central finding that violence remains prominent across all platforms despite the relative de-escalation of violence over the period. From there, we examine aggregate prevalence of content across platforms, over time, and among topics to show that Telegram features a more quotidian focus on violence and its immediate impacts, while Twitter more centrally frames ongoing violence within the conflict's master cleavages. Facebook, meanwhile, tends to feature loyalist narratives. The images these accounts share reflect a similar divergence. Throughout, we present questionnaire responses highlighting how the twin mechanisms of audience context and technical affordances shape the ways users present and describe violence across platforms. Our results show that the composition of users on these platforms drives some narrative divergence, but the persistence of these differences even among users posting across all three platforms underscores that composition alone is insufficient to explain these results.

Violence on Twitter, Telegram, and Facebook across all users

Figures 2 and 3 underscore the war's continued violence and the vast diversity of forms it takes on, as they display the topics centered on violence in both the All Groups and Parallel datasets. Among both messages in the full corpus and messages from users posting across all three platforms, between 45% and 48% of all messages center on some form of violence (fig. A.10).[3]

Figure 4 shows that aggregate differences across platforms align with our understanding of their different uses. Telegram generally features a higher prevalence of violence writ large and a higher prevalence of topics centered on daily violence vis-à-vis high politics or master cleavages. The over-time trend in the top left panel shows that Telegram features consistently higher focus on all violence topics over the course of our data collection, including after a drop in violence amid a national ceasefire at the onset of the COVID-19 pandemic.

The bottom left panel of figure 4 examines the prevalence of each individual topic and shows higher topic prevalence on Telegram for "military confrontations and shelling"—67% more prevalent than on Twitter and 56% more prevalent than on Facebook—"Syrian and Turkish army movements," "air strikes in Idlib," "foreign intervention and regional negotiations," "SDF governance, insurgent violence," and "airstrikes, warnings," which features content warning civilians of the timing of potential air strikes.[4]

3. Appendix A.1 figure A.7 displays all 25 topics.

4. The effect of source on prevalence is statistically significant at the 99.9% level for all topics except "airstrikes, Idlib" between Twitter and Facebook, "contested territory, progovernment intervention" between Twitter and Facebook, and "government, revolution, Assad" between Telegram and Facebook.

1.	**Military confrontations and shelling**	**673,291 messages**	F: Idlib, regime, countryside countryside, Aleppo, Urgent, bombing FREX: anti_targeting, regime_axis, anti_tank anti_tank, Grad_rockets, gangs_axis, live_stationed
2.	**Foreign intervention, regional negotiations**	**527,112 messages**	F: Leaving, as a result, Company Company, Military, Regime, Idlib FREX: constitution_committee, Geneva, Jaifar Jaifar, syrian_constitution, agreement_towitness, Sochi
3.	**Syrian, Turkish Army movements**	**468,858 messages**	F: Army, east, countryside countryside, walking, dominated, Arabs FREX: army_controls, army_will, army_control army_control, Afrin_barrle, battle_move, controlled_army
4.	**SDF governance, insurgent violence**	**389,691 messages**	F: Debtor, fire, east east, unexploded, police, near FREX: Arrests, motorcycle, stops_motorycle stops_motorycle, arrest_person, section, city_Buseira
5.	**Airstrikes, Idlib**	**368,756 messages**	F: Idlib, Plane, War War, south, countryside, raid FREX: war_rocket, raid_targeting, rif_plane rif_plane, raid_targets, war_targets, regime_plane
6.	**Airstrikes, warnings**	**307,589 messages**	F: Airplane, conservative, meticulous meticulous, flying, minutes, revolving FREX: circling_might, flying_might, minute_army minute_army, army_minute, flying_circles, Sheikh_minutes
7.	**Eulogies, casualties**	**229,311 messages**	F: Muhammad, Hamad, Sheikh Sheikh, Tqall, Saleh, Khaled FREX: Sham_traitor, released_three, Zakaria_Eid Zakaria_Eid, Khasis_Dim, Mohammed_Hamad, traitor_Zakaria
8.	**Violence against children, displacement**	**173,847 messages**	F: A child, a child, a camp a camp, family, a woman, a warrant FREX: woman_child, child, crime_murder crime_murder, drowning, Rami_Makhlouf, Syrian_child
9.	**Contested territory, pro-government intervention**	**4,996 messages**	F: Gangs, Damascus, Enter Enter, Regime, Homs, Liberator FREX: liberated, liberated_areas, entered entered, red, execute, their homes

Figure 2. Topics related to violence in the All Groups dataset, whereas the bottom panel shows the same for the Parallel dataset. “F” indicates words that are most frequent in each topic. “FREX” indicates words that are both frequent in and exclusive to each topic. Message numbers are the total number of messages combined for that topic. All results are translated from Arabic. Arabic version in appendix A.1.

1.	**Aistrike warnings**	**178,592 messages**	F: Airplane, flying, minute minute, minutes, secrecy, war FREX: circles_south, plane_drone, drone_circles drone_circles, circles, war_circles, Humaymim_plane
2.	**Insurgent violence, anti-ISIS campaign**	**153,396 messages**	F: East, Deir, countryside countryside, Deir Al Zour, Army, Medina FREX: east, Deir, suburbs suburbs, Deir Al Zour, army, city
3.	**Military clashes, Aleppo**	**136,677 messages**	F: Aleppo, countryside, west west, city, system, Urgent FREX: Aleppo_suburbs, west_Aleppo, south_Aleppo south_Aleppo, west_Aleppo, Al Bab_city, rif_Aleppo
4.	**Airstrikes, Idlib**	**129,449 messages**	F: Idlib, bombing, debtor debtor, urban, airplane, cave FREX: result_shellings, strike_plan, shells shells, result_raid, strike_plan, strike_target
5.	**Turkish army movements**	**112,740 messages**	F: Idlib, Leaving, Urgent Urgent, Army, Military, Ridge FREX: nidaa_telegram, live_Erdogan, defense_Turkey defense_Turkey, military_ranks, good morning, nidaa_correspondent
6.	**Casualties, eulogies**	**72,289 messages**	F: Muhammad, Israel, Hamad Hamad, Al-Raqqa, Egypt, Occupation FREX: occupied_Palestine, targeted_AlRaqqa, Rami_targeted Rami_targeted, live_Palestine, Mohammed, Palestine
7.	**Explosions**	**58,739 messages**	F: Idlib, Plane, Side Side, Missiles, Countryside, Idlib FREX: shelling_rockets, Idlib_planes, war_shelling war_shelling, live_planes, Sayed_Nasrallah, helicopter_target

Figure 3. Topics related to violence in the Parallel dataset. More information in figure 2 caption.

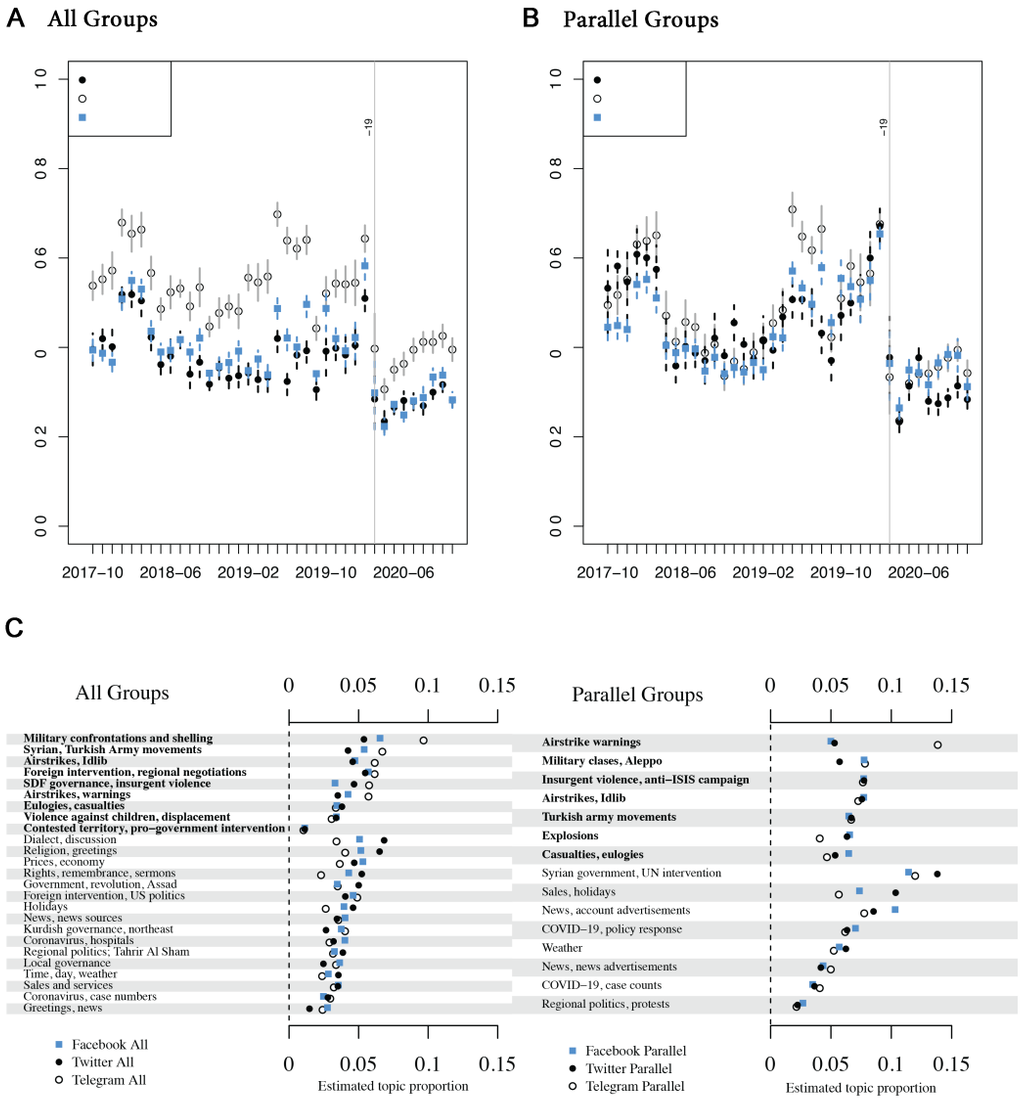

Figure 4. The top panels show the difference between Twitter, Telegram, and Facebook representation of violence-related topics over time. The grey line indicates the onset of a ceasefire and the COVID-19 pandemic in March 2020, which contributed to periods of low violence. Violence topics are those indicated in figure 3. The bottom panel shows the topic proportions for each of the topics in the All Groups and Parallel datasets, with violence topics shown in bold.

The only violence-related topics with greater prevalence in the Twitter and Facebook datasets are ones related to eulogies and casualties. Representative messages around this topic on Facebook discuss eulogies and sacrifice of fallen fighters, likely reflecting the high costs progovernment communities have faced in staffing the Syrian army (Wedeen 2019). On Twitter, prevalent topics do not relate to daily violence but instead center on political topics like "government, revolution, Assad" and "foreign intervention, US politics." Like the eulogies and casualties topic, these topics heavily focus on revolutionary remembrance. The prevalence of words centered on dialect and discussion, as well as rights, remembrance, and sermons, is also far higher on Twitter and Facebook.

Individuals who responded to the structured questionnaire pointed to two mechanisms shaping the aggregate differences in speech on these platforms: platform audience and technical affordances like speed, content monitoring, and hashtags. Describing Telegram, users emphasized that the Syrian user-base, speed, and absence of platform filtering make the platform useful for communicating key information directly to Syrian audiences, but they added that the absence of sharing features make it difficult for pages to grow. When asked about Facebook, journalists and activists noted its critical role in both their social and professional lives, emphasizing its reach to both personal and professional acquaintances. Respondents also emphasized how Facebook uses

content filtering to limit depictions of violence and prevents accounts from using pseudonyms. On Twitter, users emphasized the functional reach of hashtags to drive conversation among a diverse audience and its role in shaping global public opinion: "I use Twitter to spread the news in a more international way, as the audience is composed of elites," says one independent, opposition-aligned journalist.

Violence on Twitter, Telegram, and Facebook across parallel users

These differences are also apparent among the 124 accounts in the Parallel dataset actively posting across all three platforms (right panels of fig. 4). Given that these are the exact same users, we should expect little difference between the content they share across platforms. Though the differences in topics are indeed less pronounced, the remaining differences suggest that users' selection in different platforms or the aggregate composition of users on platforms cannot by themselves explain the results in the All Groups data.

The top right panel of figure 4 shows that the majority of topics remain violence related and that the discussion of violence on Telegram generally remains more frequent, though the difference between Telegram and Facebook disappears after the onset of the COVID-19 pandemic. While the differences in topic prevalence are smaller, we observe that Telegram channels continue to feature a steadier focus on quotidian violence in the form of airstrike warnings than Twitter and Facebook, which feature a comparatively heavier focus on remembrance of casualties and eulogies. On Twitter, discussion of the Al Assad regime and the ongoing debates around UN intervention remain more prevalent (Syrian government, UN intervention). One individual running an account that monitors Syrian air traffic to warn populations where an airstrike will occur described coming to prefer Telegram given its direct ease of use: "Given that our goal was only [to serve] the people of Idlib with a humanitarian goal, we found that Telegram is the best way to transmit the news quickly and safely." As figure 4 shows, Facebook and Twitter do feature more frequent posts about explosions, which on Facebook may capture progovernment reporting about local insurgent violence.

Questionnaire respondents who post on multiple platforms highlight that the divergent audiences on platforms helped to account for these differences. As one commented, "We work on Twitter, Facebook, and everywhere. . . . We use Twitter as a global face, and Facebook as a global and a local one" (page administrator, Civil Society Organization).

These respondents also emphasize how technical affordances directly shape their reach, with several emphasizing the role of Twitter's hashtag in allowing for messages to spread globally, even as content policies shape their expression: "We publish the same news on all our electronic platforms, but in general, if, for example, Idlib or [town] was bombed: we would use Twitter to share a certain hashtag in order to deliver this news to the whole world. Facebook allows us to publish, but like Instagram violent images are not allowed to be published. Telegram is specialized insofar that we can spread the news, with both audio and video, without being restricted" (local media activist committee page administrator in Idlib province).

Narrative differences within topics across platforms

Beyond aggregate prevalence of topics, we also see divergence within the discussion of critical topics and in association with core contested concepts like the Syrian revolution and President Bashar Al-Assad. Overarchingly, we see greater persistence of the revolutionary narratives on Twitter, reflecting a dialogue centered on remembrance of the more fundamental cleavages of the war, alongside calls for justice and accountability.

The left panel of figure 5, for instance, displays the distribution of word frequency and uniqueness within the topic related to the Syrian government, revolution, and President Bashar Al Assad in the All Groups dataset, which was most prevalent on Twitter. Figure 5 shows that Twitter, with a high prevalence and uniqueness of the words murder, Syria, rights, remembrance, crime, freedom, and martyrs, emphasizes the broad narrative of the Syrian uprising. Facebook's discussion, meanwhile, more strongly represents positive discussion of Syria's President Al Assad and ongoing governance.[5] The right panel of the figure uses word2vec (Mikolov et al. 2013) to highlight contextually similar words to Syrian President Al Assad. On Facebook, we see that words used in context with Assad are generally more positive, with references to his father, "the eternal leader." However, we also see derogatory messages, like one referencing a militia leader the Pink Panther, curse your soul, and the criminal. Twitter, meanwhile, is uniquely derogatory but nonsectarian. Telegram is also derogatory, with the criminal, but we see more contemporary and slightly sectarian references, as to Iranian militias fighting alongside Al Assad and a derogatory term for progovernment *shabiha* militias.[6] These findings again

5. Telegram features more discussion linking revolution to "war" though also concentrates on protest topics with words like "went out," often used to mean organize, as well as "protester" and "demands." A reading of representative messages from this topic on Telegram suggests this focus is on contemporary, smaller scale protest events in Syria or in neighboring countries.

6. In appendix A7, we use word2vec to analyze contextual associations over the related words "revolution" and "freedom" over time, showing similar results.

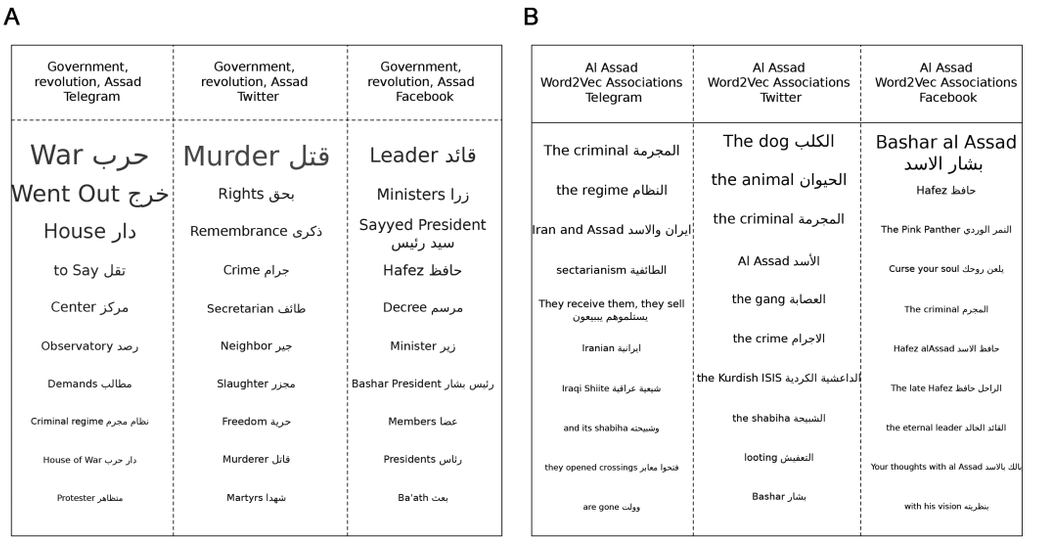

Figure 5. The left panel shows the distribution of words in the topic "government, revolution, Assad" for the All Groups dataset. The word size is in proportion to word frequency in the corpus. The color is in relation to the exclusivity of the word for the given source, though there is little visible variance since words were selected for exclusivity in the top 30 words for each source. The Arabic text in this chart currently displays the stemmed text. The right panel shows the top 10 associations with the word "Al Assad" (in Arabic) using word2vec.

reinforce Twitter's framing of the conflict amid broader pro-opposition narratives compared with more local or tactical perspectives on Telegram and more consistently loyalist narratives among our Facebook sample.

Images across platforms

The results from the image analysis mirror the findings from our analysis of the text. In total, the 100-cluster model surfaced 40 unique topics, 11 of which uniquely centered on violence (app. A.4). The top of figure 6 highlights a sample of clusters from the model: male politicians, weapons sales, bombing, and military personnel. The bottom of figure 6 shows differences in prevalence of the 11 unique violence topics in the model. Facebook features high prevalence of staged, propaganda-like images of formal military personnel, military vehicles, and warplanes and drones, consistent with the high numbers of government-held accounts in the model and emphasis on the regime's dominant propaganda. As the platform blocks images of the injured, no casualties or injured people appear on Facebook. As in text, Twitter features higher levels of focus on protest and on aggregate infrastructure destruction, while Telegram is a site for weapons sales blocked on other platforms and, like Twitter, images of the injured.

Structured questionnaire respondents placed greater emphasis on how platform content restrictions influence the images they share. Several users noted Twitter and especially Facebook's stringent content policies. One photographer noted that Telegram allows him to convey Syria's bloody reality, and facts, as is but notes that they use Twitter more frequently given a humanitarian duty to convey the suffering of Syrians and the demographic, political, and local changes in a transparent way to the whole word in the hope it will one day have an impact on changing what is occurring in Syria.

DISCUSSION

This research highlights how narratives around a single conflict diverge across three leading social media platforms. These findings are substantively important given the critical role social media has played in shaping global understanding of the Syrian conflict, one of the first conflicts in which lines between offline and online conflict engagement have become blurred (Gohdes 2020, 489). Methodologically, our research pairs collection and analysis of large amounts of social media data with a sample of users' responses to an open-ended questionnaire. We see further potential for descriptive analysis of text and image data that complements qualitative research.

We complement a body of evidence on the diversity of narratives in conflict, especially where foreign intervention is salient. Amid research that highlights the narrative contests among conflict actors broadly (Corstange and York 2018), our analysis highlights how individual actors may diversify their narratives according to the audience with whom they are communicating, even in publicly accessible speech. Users' responses to the open-ended questionnaire provide evidence for this mechanism, as respondents emphasized how platforms like Twitter reach global audiences whose intervention

Figure 6. The upper tiles show images randomly selected from different clusters, displaying how the method groups images into identifiable clusters based on certain features. The bottom image shows the estimate topic proportion of different violence-labeled clusters across platforms.

can meaningfully shape the ongoing conflict. Additional research could more systematically explore how individuals' perceptions of the audiences they reach across platforms shape their speech. Further, as our findings imply that users of different platforms are consistently exposed to different narratives around conflict, scholars could build on these insights to explore subsequent differences in their attitudes. Moreover, the timely and prevalent discussion of violence across platforms in our study shows why event datasets that include social media data have proven more complete than those that use only newspaper data (Raleigh et al. 2010; Zhang and Pan 2019). Our study shows that some platforms, like Telegram, may be more valuable, as they are more likely to report hyper-local events.

In other cases, we expect that audience and technical affordances will shape presentations of violence even as different combinations of platforms will be relevant. Given the salience of international intervention in conflict, the division between domestic and international audience contexts are likely to remain relevant even as the relevance of precise combinations of technical affordances may be more context specific.

While we created parallel procedures for account selection and compared users posting across multiple platforms,

individual users of these platforms inevitably comprise different populations of Syrian journalists, analysts, public figures, politicians, and content producers. Further studies should extend analysis of why and when users choose to adopt specific platforms (Chang et al. 2022; Hobbs and Roberts 2018), while researchers studying content on a single platform should carefully consider how the users in their sample represent their populations. As we point to two mechanisms that drive our findings, an editor at a Syria-focused news site who responded to our questionnaire reflected eloquently on the circular dynamics at play across platforms: "A social media user must determine their message, and then the audience they want to deliver it to. From there, they determine the platform they'll put it on. This platform then shapes the message. As a result, the relationship between the message, the platform, and the audience is triangular. They interact with one another, and cannot be disentangled" (editor, Syrian online news agency posting across all three platforms). Future work should take up this editor's challenge of disentangling these mechanisms.

ACKNOWLEDGMENTS

For helpful comments, the authors thank participants at AALIMS 2021, MPSA 2021, APSA 2021, Politics and Computational Social Science 2021, and Households in Conflict 2020 as well as Yale University's Comparative Politics, Washington University's Political Violence, MIT Global Diversity Lab, and Harvard's Political Violence workshops. The authors are also thankful for constructive feedback from Richard Nielsen, Tamar Mitts, Alexandra Siegel, and Zachary Steinert-Threlkeld.

REFERENCES

Alrababa'h, Ala, and Lisa Blaydes. 2021. "Authoritarian Media and Diversionary Threats: Lessons from 30 Years of Syrian State Discourse." *Political Science Research and Methods* 9 (4): 693–708.

Arab Barometer Data Analysis Tool. 2022. "The Arab Barometer." https://www.arabbarometer.org/survey-data/data-analysis-tool/.

Bonilla, Yarimar, and Jonathan Rosa. 2015. "#Ferguson: Digital Protest, Hashtag Ethnography, and the Racial Politics of Social Media in the United States." *American Ethnologist* 42 (1): 4–17.

Brubaker, Rogers, and David D. Laitin. 1998. "Ethnic and Nationalist Violence." *Annual Review of Sociology* 24 (1): 423–52.

Bucher, Taina, and Anne Helmond. 2018. "The Affordances of Social Media Platforms." *SAGE Handbook of Social Media* 1:233–53.

Chang, Keng-Chi, William R. Hobbs, Margaret E. Roberts, and Zachary C. Steinert-Threlkeld. 2022. "COVID-19 Increased Censorship Circumvention and Access to Sensitive Topics in China." *Proceedings of the National Academy of Sciences* 119 (4): e2102818119.

Corstange, Daniel, and Erin A. York. 2018. "Sectarian Framing in the Syrian Civil War." *American Journal of Political Science* 62 (2): 441–55.

Dehariya, Vinod Kumar, Shailendra Kumar Shrivastava, and R.C. Jain. 2010. "Clustering of Image Data Set Using K-Means and Fuzzy K-Means Algorithms." In *2010 International Conference on Computational Intelligence and Communication Networks*. Piscataway, NJ: IEEE, 386–91.

Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. "Imagenet: A Large-Scale Hierarchical Image Database." In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. Piscataway, NJ: IEEE, 248–55.

Driscoll, Jesse, and Zachary C. Steinert-Threlkeld. 2020. "Social Media and Russian Territorial Irredentism: Some Facts and a Conjecture." *Post-Soviet Affairs* 36 (2): 101–21.

Esberg, Jane, and Alexandra A. Siegel. 2023. "How Exile Shapes Online Opposition: Evidence from Venezuela." *American Political Science Review* 117 (4): 1361–78.

Fabbe, Kristin, Chad Hazlett, and Tolga Sınmazdemir. 2019. "A Persuasive Peace: Syrian Refugees' Attitudes towards Compromise and Civil War Termination." *Journal of Peace Research* 56 (1): 103–17.

Fearon, James D., and David D. Laitin. 2000. "Violence and the Social Construction of Ethnic Identity." *International Organization* 54 (4): 845–77.

Gohdes, Anita R. 2020. "Repression Technology: Internet Accessibility and State Violence." *American Journal of Political Science* 64 (3): 488–503.

He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. "Deep Residual Learning for Image Recognition." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Piscataway, NJ: IEEE, 770–78.

Hobbs, William R., and Margaret E. Roberts. 2018. "How Sudden Censorship Can Increase Access to Information." *American Political Science Review* 112 (3): 621–36.

Horowitz, Donald L. 1992. *A Democratic South Africa?: Constitutional Engineering in a Divided Society*, vol. 46. Berkeley, CA: University of California Press.

Hubbard, Ben. 2020. "Syria's Economy Collapses Even as Civil War Winds to a Close." *New York Times*, June 15. https://www.nytimes.com/2020/06/15/world/middleeast/syria-economy-assad-makhlouf.html.

Kalyvas, Stathis N. 2006. *The Logic of Violence in Civil War*. Cambridge, UK: Cambridge University Press.

Lynch, Marc, Deen Freelon, and Sean Aday. 2014. *Syria's Socially Mediated Civil War*. Washington, DC: United States Institute of Peace.

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. "Efficient Estimation of Word Representations in Vector Space." Preprint, ArXiv, September 7. https://doi.org/10.48550/arXiv.1301.3781.

Miller, George A. 1995. "WordNet: A Lexical Database for English." *Communications of the ACM* 38 (11): 39–41.

Mitts, Tamar. 2022. "Countering Violent Extremism and Radical Rhetoric." *International Organization* 76 (1): 251–72.

Mitts, Tamar. 2019. "From Isolation to Radicalization: Anti-Muslim Hostility and Support for ISIS in the West." *American Political Science Review* 113 (1): 173–94.

Mitts, Tamar, Gregoire Phillips, and Barbara F. Walter. 2022. "Studying the Impact of ISIS Propaganda Campaigns." *Journal of Politics* 84 (2): 1220–25.

Nielsen, Richard A. 2017. *Deadly Clerics: Blocked Ambition and the Paths to Jihad*. Cambridge, UK: Cambridge University Press.

Pearlman, Wendy. 2016. "Narratives of Fear in Syria." *Perspectives on Politics* 14 (1): 21–37.

Raleigh, Clionadh, Andrew Linke, Håvard Hegre, and Joakim Karlsen. 2010. "Introducing ACLED: An Armed Conflict Location and Event Dataset: Special Data Feature." *Journal of Peace Research* 47 (5): 651–60.

Roberts, Margaret E., Brandon M. Stewart, Dustin Tingley, and Edoardo M. Airoldi. 2013. "The Structural Topic Model and Applied Social Science." Presented at Advances in Neural Information Processing Systems Workshop on Topic Models: Computation, Application, and Evaluation, Lake Tahoe, Nevada.

Schon, Justin. 2021. "How Narratives and Evidence Influence Rumor Belief in Conflict Zones: Evidence from Syria." *Perspectives on Politics* 19 (2): 539–52.

Siegel, Alexandra A., and Vivienne Badaan. 2020. "#No2Sectarianism: Experimental Approaches to Reducing Sectarian Hate Speech Online." *American Political Science Review* 114 (3): 837–55.

Steinert-Threlkeld, Zachary C. 2017. "Spontaneous Collective Action: Peripheral Mobilization During the Arab Spring." *American Political Science Review* 111 (2): 379–403.

Steinert-Threlkeld, Zachary C., Alexander M. Chan, and Jungseock Joo. 2022. "How State and Protester Violence Affect Protest Dynamics." *Journal of Politics* 84 (2): 798–813.

Syakur, M.A., B.K. Khotimah, E.M.S. Rochman, and Budi Dwi Satoto. 2018. "Integration K-Means Clustering Method and Elbow Method for Identification of the Best Customer Profile Cluster." *IOP Conference Series: Materials Science and Engineering* 336:012017.

Urman, Aleksandra, Justin Chun-ting Ho, and Stefan Katz. 2021. "Analyzing Protest Mobilization on Telegram: The Case of 2019 Anti-extradition Bill Movement in Hong Kong." *PLOS ONE* 16 (10): e0256675.

Walter, Barbara F. 1997. "The Critical Barrier to Civil War Settlement." *International Organization* 51 (3): 335–64.

Wedeen, Lisa. 2019. *Authoritarian Apprehensions: Ideology, Judgment, and Mourning in Syria*. Chicago: University of Chicago Press.

Zeitzoff, Thomas. 2017. "How Social Media Is Changing Conflict." *Journal of Conflict Resolution* 61 (9): 1970–91.

Zhang, Han, and Jennifer Pan. 2019. "CASM: A Deep-Learning Approach for Identifying Collective Action Events with Text and Image Data from Social Media." *Sociological Methodology* 49 (1): 1–57.

Zhang, Han, and Yilang Peng. 2024. "Image Clustering: An Unsupervised Approach to Categorize Visual Data in Social Science Research." *Sociological Methods & Research* 53 (3): 1534–1587.