

RESEARCH NOTE—The Enemy of My Enemy Is Not My Friend: Arabic Twitter Sentiment Toward ISIS and the United States

Abstract (125 words): A counter-intuitive finding emerges from an analysis of Arabic Twitter posts from 2014–2015: Twitter participants who are negative toward the Islamic State of Iraq and the Levant (ISIS) are also more likely to hold negative views of the United States (US). This surprising correlation is due to the interpretations of two sets of users. One set of users views the US and ISIS negatively as independent interventionist powers in the region. The other set of users negatively links the US with ISIS, often asserting a secretive conspiracy between the two. The intense negativity toward the United States in the Middle East seems conducive to views that, in one way or another, cause citizens to link the United States and ISIS in a conspiratorial manner.

1. Introduction

The United States' (US) war against the Islamic State of Iraq and the Levant (ISIS) has been waged both on the ground in Iraq and Syria and online in social media (Greenberg 2016). Social media plays an increasingly important role in a range of conflicts (Zeitsoff 2017), changing how these conflicts are fought and studied. In the context of ISIS in particular, researchers have made use of ISIS's social media presence to study the causes of radicalization into support for ISIS (Magdy et al. 2015; Mitts 2018), the nature of the networks of individuals who support and oppose ISIS (Bodine-Baron et al. 2016), the measurement of ISIS's success (Siegel and Tucker 2018; Barceló and Labzina 2018), and how to counter violent extremism online (Mitts 2020). This paper also uses social media data but examines a novel question: How do publics in and adjacent to a conflict update attitudes toward involved entities, particularly ones toward whom they have held negative attitudes? Only 0.4 to 6.4% across the Arab world agree with ISIS's goals (Tessler, Robbins, and Jamal 2016), and only 15 to 34% of Arabs hold favorable views toward the US (Vice 2017; Jamal et al. 2015; Lynch 2007). Though intensely disliked, these two entities have also been fighting a war against each other. In such a scenario, how do citizens pick a side, if they pick one at all? Do their attitudes shift because of a stronger dislike for one of them (i.e. an "enemy of my enemy is my friend" dynamic), or do other factors guide their attitudes?

We address these questions by examining Arabic Twitter posts about the US and ISIS from 2014 to 2015, building on previous work on politics in social media (Barceló and Labzina 2018; Jamal et al. 2015; Zeitsoff et al. 2015) and anti-Americanism (Sokolov et al. 2018; Corstange 2016; Katzenstein and Keohane 2007). Our results are not consistent with an "enemy of my enemy is my friend" interpretation in this context; instead, Arab Twitter publics simultaneously dislike the US and ISIS. We find that those who are critical of ISIS are more negative toward the US than

both those whose content is neutral toward ISIS as well as Arabic-speaking Twitter in general. We also find that expressed dislike of both entities increases as the US intervenes against ISIS. We provide evidence that those on Twitter who express negativity toward the US and ISIS can be divided into two distinct groups. One group negatively views the US and ISIS but sees them as acting independently; the other sees them as linked, often in a conspiratorial manner; and what ties these two groups is a strong negative sentiment toward intervening forces in the region.

2. Possible Outcomes and Explanations of Attitudes toward ISIS and the US

This project was initiated as an inductive exercise. For this reason, in this section we outline the outcomes of interest that were most likely at the outset of our data gathering exercise and then describe the theories that most likely explain those outcomes.

Outcome 1: The Enemy of My Enemy Is My Friend

A first possible outcome would be to find that negative sentiment toward ISIS correlates with positive sentiment toward the US. This outcome would be consistent with an “enemy of my enemy is my friend” dynamic, wherein one traditional enemy is supported because of their offer to aid in fighting another enemy. Such a dynamic has been used as a framework to study topics central to international relations, including: states' decisions to intervene in civil wars (Findley and Teo 2006), interstate disputes (Joyce et al. 2013), the likelihood of terrorists attacking nationals of countries which are allied with the terrorists' home country (Plümper and Neumayer 2009), and flows of international trade (Polachek et al. 1999). Studying this dynamic is important because an “enemy of my enemy is my friend” logic also reflects a reality of dyadic behavior—that it is dependent on attitudes and allegiances toward third parties—that often goes unacknowledged (Cramer and Desmarais 2016).

Similarly, an “enemy of my enemy is my friend” dynamic may govern the behavior of individuals and public opinion, a factor that is often important in mediating a state's response in the arena of foreign policy (Foyle 1997). In the realm of individual beliefs and attitudes, this first possible outcome of our empirical work would be most in line with a view that public opinion is a function of perceived self-interest on the part of citizens. Scholars have modeled opinion formation in this way in a variety of contexts. In international trade, it is often posited that public opinion on trade policy is primarily a function of who stands to gain or be hurt by said policies (Mayda and Rodrik 2005). Some have found that changes in public opinion as it relates to foreign policy are similarly “reasonable,” i.e. they reflect shifts in the real interests of the nation (Shapiro and Page 1988; Isernia et al. 2002). In the context of this study, it could be that the US's engagement against a common foe (ISIS) causes those who express negative sentiment toward ISIS to embrace at least some aspects of US involvement in the fight against ISIS.

Outcome 2: The Enemy of My Enemy Is Not My Friend

A second possible outcome is that negative sentiment toward ISIS is correlated with negative sentiment toward the US. Such an outcome would be consistent with at least two separate interpretations, the first of which we label (1) *independent negativity*. Under this interpretation, members of the Arabic Twitter universe have a firm, long-term dislike of the US; they believe that US actions are hostile to Arabs. They also perceive ISIS as hostile to their values and perhaps directly threatening to themselves. But, in spite of the US's engagement against ISIS, they do not connect their negativity toward the US with their negative opinions of ISIS; these two entities are viewed as independent of each other. The second we label (2) *linked negativity*. Under this interpretation, citizens' negativity toward the two entities is connected, with two possible sub-interpretations. First, under the (A) *non-conspiratorial* interpretation, citizens may believe that US intervention

has inadvertently led to the growth of ISIS. They view US actions as integral to allowing for the rise of ISIS but do not purport a direct connection between the two. Alternatively, citizens may view the US as responsible for ISIS is under a (B) *conspiratorial* interpretation. Under this interpretation, the US and ISIS are in cahoots, with either the US purposefully creating ISIS or engaging in efforts to support and strengthen its power in the Middle East.

Some of these interpretations, especially (2B), may appear outlandish. However, research on attitude formation has demonstrated that accuracy is not always individuals' goal when updating opinions and beliefs. One of the most prominent sets of literature on this topic is research on motivated reasoning, which notes that people often form or update opinions as a way of reaching a preferred conclusion rather than out of a desire to be accurate (Taber and Lodge 2006). Much of this research has focused on the US, and in that context has found that motivated reasoning takes a partisan hue. For instance, Hellwig et al. (2008) find that citizens evaluate their government's ability to negotiate in international trade disputes based on their partisan identity; similarly, Gaines et al. (2007) find that motivated reasoning functions in citizens' evaluations of foreign policy, in their case the war in Iraq, because citizens interpretations of facts varied by partisan identity. Additionally, this research has found that motivated reasoning is strongest for the issues that individuals care about the most. Studies of motivated reasoning have often focused on topics that are controversial or about which individuals are likely to hold strong opinions, for instance stem cell research (Nyhan and Reifler 2010) or gun control and affirmative action (Taber and Lodge 2006). In studies where the public has less information or is less opinionated, such as in opinions on foreign aid (Hurst et al. 2017), scholars have found individuals less susceptible.

Our context, though outside the US, could see similar factors at work because of the Arab Middle East's experience with intervention by the US and other foreign powers, making the operation of motivational reasoning likely. As noted, anti-US sentiment is very strong in the Middle East, and the US is seen consistently as one of the most salient and important countries to the foreign policy of Middle Eastern countries (Lynch 2007; Vice 2017). Negative attitudes toward the US are integrally tied to US intervention in the region (Jamal et al. 2015)—even if that intervention was against another country. As Furia and Lucas (2006) note, Arab attitudes toward foreign countries are primarily “based on the military and non-military policy behaviors of those countries in regard to *regionally salient issues*,” particularly the Iraq war and Palestine. Additionally, experience with real foreign plots in the past adds to the potential to interpret the present in conspiratorial terms. As just one example, secret negotiations between foreign powers (the so-called Sykes-Picot agreement) led to French and British mandates of control over some Middle East states post WWII. Any belief in what seem like outlandish theories about foreign powers cannot be fully understood without the recognition of historical experiences like this, many of which are still salient to Arabs across the Middle East to this day (Gray 2010; Silverstein 2000).

Summary

We focus on two possible outcomes of our empirical exploration: one where negativity toward ISIS is correlated with positivity toward the US (enemy of my enemy is my friend), and one where ISIS negativity is correlated with a similar negativity toward the US. A confluence of negativity could be happenstance—i.e. individuals' negative views toward ISIS and the US could remain independent of each other, despite US intervention against ISIS. However, an interpretation based on motivational reasoning and on the Arab Middle East's historical experience would predict that individuals draw a connection between ISIS and the US, potentially in a conspiratorial manner. In

the sections that follow, we establish that negativity toward ISIS coincides with negativity toward the US in Arabic Twitter, examine alternative explanations for the correlation, and then examine the connections Arabic Twitter users make between ISIS and the US.

3. Data and Analysis

Our data consist of all public Arabic Twitter posts 2014–2015, which we gain access to and analyze with textual analysis models on the platform of Crimson Hexagon (CH), a social media analytics company. Throughout the rest of the paper, we will refer to several textual analysis models we used to estimate category proportions or classify Twitter users. Although some aspects of the analysis were carried out on our own, all model estimates come from proprietary classifiers provided on CH's platform. All of the main models are *supervised*, meaning they required human input in categorizing texts. Categorizing human expression using automated methods can be a tricky business, especially in languages like Arabic where multiple dialects exist in our texts. We take as many steps as we can to ensure that our results are accurate. For each model requiring human coding, we followed a rigorous training procedure with two Arabic-speaking coders who classified hundreds of Twitter posts for this project. Information on each model are provided in the applicable sections of the paper, and additional information and example training texts can be found in the online appendix. Additionally, for tasks where classification is so tricky as to make automated methods untenable—such as in our detailed content analysis in section 6, where we attempt to parse conspiratorial and non-conspiratorial linkages of the US and ISIS—we eschew automated methods for a more qualitative approach.

Our data and results should be interpreted with an eye toward the limitations of social media data. For instance, the conclusions we draw are limited by who uses Twitter in the Middle East. We do not analyze data from other mass media sources, which may be similar or different to our

social media data and also may influence (or be influenced by) what is expressed on social media. The Middle East media environment is a complex set of interacting sources, and social media data is just a sliver of that. However, concerns about the representability of our data are tempered by the fact that for our period of interest—the years 2014–2015—Twitter use was relatively widespread throughout the Middle East. Outside of the gulf, results from the 2016 Arab Barometer (the closest to our date range of interest) indicate that sizable portions of the population in Lebanon (27%), Palestine (25%), Egypt and Algeria (both 19%), Morocco and Tunisia (both 17%), and Jordan (15%) reported having a Twitter account.¹ In the Gulf states, and especially in Saudi Arabia, the UAE, and Kuwait, Twitter usage has been much higher from an even earlier date (Dubai School of Government 2012). This indicates that our sample of Twitter users covers a broad swath of the Arab Middle East, while over-representing the gulf states and, as with much of social media, skewing younger than the general population.

Some may also be concerned about the presence of bots in our data. A first concern relating to Twitter bots is the presence of state actor bots, i.e. state-led propaganda efforts involving automated Twitter accounts. There is evidence of state manipulation using such accounts (Leber and Abrahams 2019), and it could be the case that state actors are behind the patterns we see. Second, inauthentic content from “ISIS bots” is also a reasonable concern—ISIS was notorious for its use of bots to spread its extremist propaganda. However, the bad publicity Twitter received over ISIS's use of its platform led to large-scale action targeting ISIS bots in 2014 and 2015 (Berger and Morgan 2015; Berger and Perez 2016). Additionally, our results are the opposite of the result one would expect if our findings were driven by ISIS bots tweeting negative about the US. Lastly, our

¹ As a point of comparison, in the US, 23% of all Internet users reported using Twitter in 2015 (Duggan 2015).

examination of our findings across different frequency levels, i.e. different levels of Twitter usage, can also help alleviate concerns about bots in our data. Bots are unusually prolific Twitter users—for instance, regarding ISIS bots, Berger and Morgan (2015, p. 28) find that the typical ISIS account Tweets 7.3 times per day. In Section 5, however, we find that our results hold for low-frequency users.² We cannot completely rule out the influence of bots, either ISIS bots or those directed by state entities, on our results, but the fact that our results hold across frequency levels can somewhat mitigate those concerns.

4. Aggregate Analyses

There are three main models we present in this section, all of which are based on the supervised algorithm described by Hopkins and King (2010). Each examines data at the aggregate level, where our quantity of interest is the percentage of posts (Tweets) that fall into certain categories. Of these analyses, one examines posts about the US (*US Analysis*), a second examines posts about ISIS (*ISIS Analysis*), and a third examines posts containing references to both entities (*Combined Analysis*) from January 1, 2014 to December 31, 2015. These analyses use hundreds of user-trained posts that coders categorize into negative, neutral, and positive categories. The model then uses the training texts to categorize the universe of posts for each analysis. For the *US Analysis*, we differentiate between political and social content in the neutral and negative categories. For the *Combined Analysis*, we omit the positive category (for lack of posts) and add negative categories based on the entities with which the US is associated, differentiating posts that associate the US with ISIS from those that associate it with another disliked entity (e.g. the Syrian Regime, Iran, or Shias) and from those that do not associate the US with any entity.

² See section 6.2 of the online appendix for a more detailed look at this.

The results from these aggregate analyses, which are shown in Table 1, demonstrate a high degree of negativity toward both the US and ISIS. In the *US Analysis* (column 1), over two-thirds of the traffic is negative, mostly political in nature; the rest is neutral, with negligible positive traffic. In the *ISIS Analysis* (column 2), over half of the traffic is neutral, and of the traffic that expresses an opinion (i.e. excluding the neutral traffic), eighty percent is negative. Looking at posts that mention both entities in the *Combined Analysis* (column 3), we find that about 40% of the traffic expresses a negative opinion about at least one of the entities. Of this negative traffic, over half associates the US with ISIS or another disliked entity, while just under half does not associate the US with anyone.

Table 1 – Results of the Aggregate Analyses

	US Analysis (N = 88,092,712)		ISIS Analysis (N = 103,228,106)		Combined Analysis (N = 5,832,491)	
Negative	1a. Neg. Pol.	60%	1. Neg.	34%	1a. Neg., US & ISIS	12%
	1b. Neg. Soc.	11%			1b. Neg., US & Syria/Iran/Shia	11%
					1c. Neg., Other	18%
Neutral	2a. Neu. Pol.	28%	2. Neu.	58%	2. Neutral	58%
	2b. Neu. Soc.	1%				
Positive	3. Pos.	<1%	3. Pos.	8%	3. N/A	

Notes: Abbreviations used for categories are Neg. = Negative, Pol. = Political, and Soc. = Social. Columns do not sum to 100% due to rounding.

5. Group-Based Analysis

To examine the correlation between users’ views towards the US and ISIS, we use the textual analysis models from the Aggregate Analyses on subsets of users defined by post frequency and levels of negativity toward ISIS, an analysis that we refer to as the *Group-Based Analysis*. For this analysis, we have access to a very large sample— a randomly selected set of 1,078,832 posts (262,906 unique users) from the ISIS Analysis—rather than the population. Using this sample, we

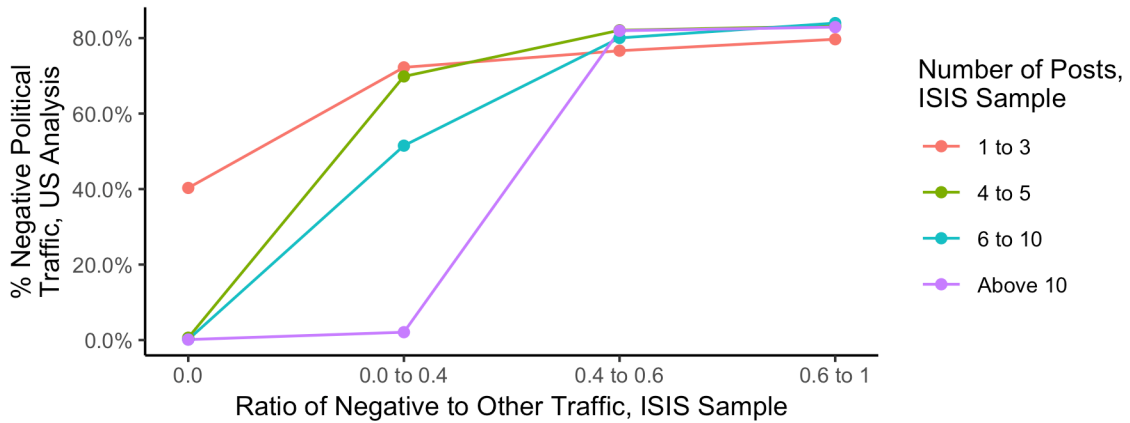


Figure 1 – Percent negative political traffic in the Group-Based Analysis.

define sixteen groups of users, classified by frequency and negativity, as seen in Table 2. The frequency data is highly skewed, as is common across Twitter as well as other social media platforms. Most users fall into the very low category (1-3 posts), with only a small proportion having more than 10 posts. Individuals’ negativity toward ISIS is estimated using a proprietary classifier, which necessarily differs from the model used in the ISIS Analysis because the unit of analysis is now users, not category proportions.³ Our choice to divide users into these sixteen groups was data-based; other grouping choices would not provide us with at least 1,000 users in each of the sixteen analysis groups. With these groups set, we then examine US-ISIS sentiment correlation by

Table 2 – User groups for the Group-Based Analysis

% Negative Traffic: Sample Volume:	Very Low (~0%)	Low (0–40%)	Medium (40–60%)	High (60–100%)
Very Low (1–3)	134,181 (51%)	4,557 (2%)	13,092 (5%)	78,933 (30%)
Low (4–5)	4,066 (2%)	4,006 (2%)	2,485 (<1%)	1,946 (<1%)
Medium (6–10)	2,237 (<1%)	3,789 (1%)	2,104 (<1%)	1,599 (<1%)
High (>10)	1,466 (<1%)	4,973 (2%)	2,685 (1%)	1,279 (<1%)

Notes: Percentages equal the proportion of the 262,906 unique users in each cell.
Because of rounding, percentages do not add to 100%.

³ More information on both models can be found in Section 2 of the Online Appendix.

estimating the category proportions from the *US Analysis* model for each of the sixteen groups. This approach is superior to alternatives: a purely aggregate-level approach could suffer from an ecological fallacy problem, while an individual-level approach would only be possible among high-frequency users.

The results of the *Group-Based Analysis* are found in Figure 1. The figure shows percent negative political US traffic on the y-axis and, on the x-axis, negativity towards ISIS (measured as a ratio). The colors differentiate the groups by number of posts in the sample. Across all levels, there is a strong positive relationship between negativity toward the US and ISIS. For the lowest-volume group (1–3 sample posts), a move from lowest to highest ISIS negativity is associated with a 39.4 percentage point increase in political negativity in the US Analysis; for those with 4–5 posts, the increase is 82.6 percentage points; for those with 6–10 posts, it is 83.8 percentage points; and, for those with more than 10 posts, it is 82.8 percentage points. On average, those who are critical of ISIS are approximately 60 percentage points more negative toward the US than those who are neutral and are even 10 percentage points more negative than the overall Twitter population. In general, the correlation is stronger for those with more than three posts—a finding which suggests that our correlation is not driven by neutral news reporting in the “low negativity” group, since news organizations generate a large amount of traffic. In the online appendix, we also check to make sure this pattern is not driven by the choice of our 16 analysis groups by examining the US-ISIS sentiment relationship at the individual level for 80 randomly-sampled high-volume users (>10 ISIS sample posts). What we find—an increase of 69.4 points in average negative political US traffic as you move from lowest to highest ISIS negativity level—closely mirrors the results in the *Group-Based Analysis*, indicating that this pattern holds regardless of our choice of analysis method.

We therefore argue that negativity toward the US is positively correlated with negativity toward ISIS. This conclusion is robust to several alternative hypotheses, the most plausible being that the correlation comes from users who are negative toward everything—imagine the rants of an “angry uncle” on Twitter. If this angry uncle hypothesis were true, we would expect to see a correlation between negativity toward ISIS and the negativity of users’ non-US and non-ISIS traffic. To test this, we apply a sentiment analysis model on all Arabic posts not including our keywords for each of our sixteen analysis groups. This automated model uses over 500,000 human-trained texts to estimate general sentiment, as opposed to a user-coded set of training texts as used for our aggregate analyses.⁴ We apply the same model to US traffic for a fair comparison, and the results from this test do not support the angry uncle hypothesis. Going from the lowest to highest ISIS negativity level is associated with an average increase of 10 percentage points in general negative traffic, compared to 17 percentage points for negative US traffic. This difference only increases in size if comparing to the user-trained US Analysis instead.

There is also strong evidence against two additional hypotheses. First, we consider that individuals’ interest in international politics may explain the US-ISIS sentiment correlation—i.e. those interested in expressing an opinion express negativity, so interest is really what matters. If true, this would imply that negativity toward ISIS is correlated with how frequently users post about the US; the data indicate that the opposite is true, with the number of posts mentioning the US going down as users, and especially high-frequency users, become more negative toward ISIS. Second, we consider the possibility that Shia Twitter users, who could be more negative toward both the US and ISIS than Sunni users, are driving the results. We test this at the country level and

⁴ More information on this model can be found in section 2 of the online appendix

find that, on average, countries with small Shia populations have the same correlation between US and ISIS negativity as those with large ones.⁵

6. What Characterizes those Negative toward ISIS and the US?

What characterizes those who are negative toward ISIS and the US? Exploring these users in-depth will help us better understand why the “enemy of my enemy is my friend” logic does not prevail for these individuals. In this section of the paper, we evaluate two possible interpretations. (1) *Independent negativity*: In spite of the US's engagement against ISIS, users do not make connections between the two entities in their negativity. Negativity toward the two is independent. (2) *Linked negativity*: Some connect their negativity toward the two entities because they hold the US responsible for ISIS. Two possible version of linked negativity are also possible. (A) *Non-conspiratorial*: Users believe that US intervention in the region has benefited ISIS, but do not believe that the US is directly intervening in favor of ISIS. (B) *Conspiratorial*: Users believe that the US and ISIS are in cahoots, with the US taking steps to purposefully help ISIS. We use data from one of our aggregate analyses and from an in-depth examination of a set of randomly-sampled users with high negativity toward ISIS to show that there is a surprisingly large number of users who believe in a conspiratorial interpretation.

First, we look at the *Combined Analysis* to address how often users take independent versus linked interpretations when both entities are mentioned in the same post. Table 3 summarizes the breakdown of negative posts mentioning both entities (these results corresponding to the first row and third column of Table 1). In almost half of the negative traffic, both the US and ISIS are mentioned but no connection is made between the US and other entities. However, almost a third

⁵Section 6 of the online appendix provides details on all alternative hypotheses.

of the negative traffic makes an association between the US and ISIS, while just over one quarter of it associates the US with other disliked entities. Many of these posts are conspiratorial in nature, while others merely indicate that US policy in the region has led to ISIS’s formation.

The Combined Analysis indicates that, for those on Twitter who express an opinion about both entities, independence is most common while at the same time a significant number of users connect the US and ISIS. However, there are limitations in relying solely on the Combined Analysis to examine US-ISIS sentiment correlation. The Combined Analysis only examines posts containing references to both the US and ISIS, a relatively small subset of posts and one which may be inflating our sense of how often users associate the two entities. The Combined Analysis also suffers from the same level-of-analysis issues as the Aggregate Analysis, i.e. results are at the level of the text, not the individual. Additionally, parsing out the two versions of the linked negativity interpretation—non-conspiratorial vs. conspiratorial—is a difficult task, one for which an automated textual analysis model is not well-suited. For this reason, our two coders also conduct a detailed content analysis of posts from 100 users, 25 randomly sampled from each of our high ISIS-negativity groups (i.e. the final column of Table 2).⁶ For each of these users, we examine

Table 3 – Focus of Negative Posts from Combined Analysis

US associated with ISIS	29%
US associated with Syria, Iran, and/or Shia	27%
No association between US and other entities	44%

Table 4 – Summary of Users in Detailed Content Analysis

Category:	Frequency Group:				Total
	1 to 3	4 to 5	6 to 10	Above 10	
Linked – Conspiratorial	9	9	13	12	43
Linked – Non-Conspiratorial	1	1	1	2	5
No Linkage	15	15	11	11	52

⁶ This focus on negative content is justified because the majority of conspiratorial content in both the US and ISIS Analyses falls into negative categories; see Section 5.1 of the Online Appendix.

their Twitter posts mentioning *either* the US *or* ISIS to determine how they link the two entities. For many of the users, this involved reading dozens or even hundreds of posts for those with a larger number of posts in the sample. We note, for a given user, whether (1) the user linked the US and ISIS conspiratorially in any Tweet (“Linked – Conspiratorial”); (2) the user never made a conspiratorial linkage, but did link the US and ISIS non-conspiratorially (“Linked – Non-Conspiratorial”); or (3) the user refrained from linking the two entities in any of the Tweets we read.⁷ This analysis can help us answer the following question: Among our users who are most negative toward ISIS, how many of them resort to conspiratorially linking the US and ISIS in at least some of their Tweets? Table 4 summarizes the results of this analysis, with the results broken down by the frequency with which the users Tweeted in our sample from the ISIS Analysis.⁸

From this detailed content analysis, we draw a few conclusions. First, among this group of users who are particularly negative toward ISIS, linking the US and ISIS is common. Out of 100 high-negativity users, almost half (48 users) have at least one post (and usually more) linking the two entities. Additionally, these data strongly indicate that those who link the US and ISIS tend to do so in a conspiratorial manner—only 5 of the 48 users who linked the US and ISIS avoided resorting to conspiracy theories. Lastly, these data may indicate a relationship between the frequency with which one posts and the tendency to link the US and ISIS. Out of users whose posting frequency was “Very Low” or “Low” in the sample, there are 20 combined that connect the US and ISIS; for those with “Medium” or “High” frequency there are 28 combined that do so. This difference is relatively small and therefore readers should exercise caution in drawing conclusions

⁷ Determining what constitutes a “conspiratorial” linkage is a difficult endeavor; for our purposes, we define any post that includes clearly false claims suggesting that the US created or directly supports ISIS as conspiratorial. See section 5 of the online appendix for more detail.

⁸ For a detailed breakdown along with example posts, see the online appendix materials.

from it, but it supports recent work (Guess et al. 2019) indicating a correlation of high social media engagement with endorsing conspiracy theories and other toxic messages.

Delving into the posts themselves reveals the context in which Arabic-speaking Twitter users make conspiratorial versus non-conspiratorial links between the US and ISIS. Uniting both sets of users is a focus on US intervention, as can be seen in the sample of non-conspiratorial posts displayed in Table 5.⁹ Three points about these posts deserve highlighting. First, whether it is US intervention in Iraq (posts 1, 2, and 4) or purported US support for armed groups in Syria (post 3), the negative results of US intervention are seen as creating an opportunity for ISIS to grow. Second, some of these posts highlight the existence of US sources (a former US DIA director in post 3; Michael Moore in post 4) who agree with them, seemingly as a way of legitimizing their point. Lastly, The US is one of several entities (Russia in post 1; Houthis in post 2) that these users view

Table 5 – Sample of Posts Linking the US and ISIS Non-Conspiratorially

1. The Soviet Union’s intervention in Afghanistan brought us al-Qaeda. America’s intervention in Iraq brought us ISIS. What is the Russians’ intervention in Syria going to bring us?	تدخّل الاتحاد السوفيتي في افغانستان انتج لنا القاعدة. وتدخّل امريكا في العراق انتج لنا داعش. ماذا سينتج لنا تدخّل الروس في سوريا ..!؟
2. American military orders resulted in al-Malaki handing over North Iraq to Daesh. Similar orders in Yemen left the Houthis in control	بأمر الامريكان سلمت حكومة المالكي شمال العراق لداعش بعد انسحاب الجيش وكذلك فعلت في اليمن حين أمرت الجيش بتسليم صنعاء للحوثي
3. A former director of the US Defense Intelligence Agency: The Obama administration supported al-Qaeda and the Muslim Brotherhood in Syria and made ISIS’s rise possible	مدير سابق لوكالة إستخبارات الدفاع الامريكية: إدارة أوباما دعمت القاعدة والإخوان في سوريا و مكّنت داعش من البروز
4. The American director Michael Moore displays a picture of Syrian refugees, on which was written “they didn’t create ISIS,” next to a picture of Bush and his Vice President, on which was written “they did!”.	المخرج الأمريكي مايكل مور يضع صورة لاجئين سوريين كُتب فيها: لم يصنعوا داعش. مع صورة لبوش ونائبه كُتب عليها: بل نحن!

⁹ For the sake of readability, hashtags that are not essential to the text as well as links have been removed from the posts in Tables 5 and 6.

as impinging on the sovereignty and stability of Middle Eastern states. These users do not necessarily link the US and these additional entities, but they do consider their actions comparable. For instance, in the first post, America's intervention in Iraq serves as an example of what might occur with Russian intervention in Syria.

Examining the conspiratorial posts in Table 6 indicates similarities and differences with the non-conspiratorial posts. Similar to non-conspiratorial posts linking the US and ISIS, these conspiratorial posts also focus on US intervention in the region and also mention additional entities threatening the sovereignty or stability of Middle Eastern States. However, the conspiratorial posts differ in three important ways from their non-conspiratorial counterparts. First, the additional entities mentioned not only threaten Middle Eastern states independently—they are in cahoots with the US. Whether it is Iran and the West (post 1), Israel and the US (post 2), or the US

Table 6 – Sample of Posts Linking the US and ISIS Conspiratorially

- | | |
|--|---|
| 1. ISIS is carrying out an Iranian-Western intelligence operation to strike against the Sunnis and tarnish their reputation in front of the world | داعش تقوم بعمل مخابراتي ايراني -غربي لضرب السنة وتشويه سمعتهم امام العالم |
| 2. ISIS is a Zionist and American creation and it is apparent that (even) their passports are printed in Israel or the US. | داعش صناعة صهيوي امريكية والظاهر ان جوازاتهم مطبوعة في اسرائيل او امريكا |
| 3. We will soon be publishing a dangerous video containing [details of] a major scandal about the way that America used ISIS as a justification for direct intervention in the region – watch it and focus well | سننشر بعد قليل فيديو خطير وفيه فضيحة كبرى عن طريقة استخدام أمريكا لداعش لتجعلها مسوغاً لها للتدخل المباشر في المنطقة شاهده وركزوا جيداً |
| 4. America claims to be fighting ISIS! I believe the more correct expression would be: America is supporting ISIS | أمريكا تدعي حرب #داعش !! أعتقد أن العبارة الصحيحة هكذا : أمريكا تغذي #داعش |
| 5. Don't forget the year 2014, [in which] America and Europe created ISIS in order to displace, kill, and rob money and oil from Arabs – it was the year in which we were truly eliminated, and which ended with the vile veto [reference to veto of UN resolution condemning Asad regime] | لا تنسوا عام 2014 امريكا واوروبا صنعوا داعش لتشريد العرب وقتلهم ونهب اموالهم ونفطهم كان عام القضاء الحقيقي علينا والنهاية الفيتو الحقيير. |

and Europe (post 5), these users posit a broad and variegated coalition of outside forces threatening the Middle East, Arabs, and Sunni Muslims. Second, many of these users attribute primary responsibility for ISIS's actions to the outside forces mentioned rather than to ISIS itself. In posts 1 and 5, for instance, ISIS is merely a pawn in an operation by these outside forces (including the US) to defame Sunnis and terrorize Arabs. Lastly, these users posit sinister, secretive motives behind US actions that are not as present in the non-conspiratorial posts. For users positing a conspiratorial linkage between the US and ISIS, the US's actions regarding ISIS go against its public claims (post 4), are part of secretive intelligence plots by multiple foreign entities (post 1), and are dangerous to even post about by those who choose to "tell the truth" (post 3).

This detailed content analysis has important limitations—the posts presented are just a snapshot from a 100-user sample. However, the fact that almost half of the users in this sample link the US and ISIS, most of them conspiratorially, means that even a larger sample would likely contain similar patterns. A conspiratorial view of the US's role in the region, one that sees the US as in cahoots with other intervening foreign entities and engaged in sinister plots against Arabs and Sunnis in the region, is common among those who are simultaneously negative toward ISIS.

7. Temporal Variation

As a final note, we examine how our results vary across time. The variation of our results across time in both the aggregate analyses and the *Group-Based Analysis* help illustrate the anti-interventionist nature of negative content. Figure 2 shows the percent negative traffic for the *ISIS Analysis*, the *US Analysis*, and high ISIS-negativity users from the *Group-Based Analysis*. Two insights are notable. First, in all three trend lines, negative traffic peaks as the US begins its military intervention against ISIS in mid-2014. This is the time period where the US began targeted strikes

in Syria and manned missions and ground assistance in Iraq against Islamic State targets, in response to requests from the Iraqi government. Spikes in negative traffic for the US Analysis continue in August and September, when the US and allies began large-scale airstrikes against Islamic State targets in both countries. Second, in the *ISIS Analysis*, high negative traffic also follows ISIS attacks and killings in the West (e.g. the Nov. 2015 attacks in France) and the Middle East (e.g. the killing of a Jordanian pilot in Jan. 2015).¹⁰ These results mirror findings that US intervention drives anti-Americanism in the Middle East (Jamal et al. 2015) and that ISIS attacks decrease its support (Barceló and Labzina 2018).

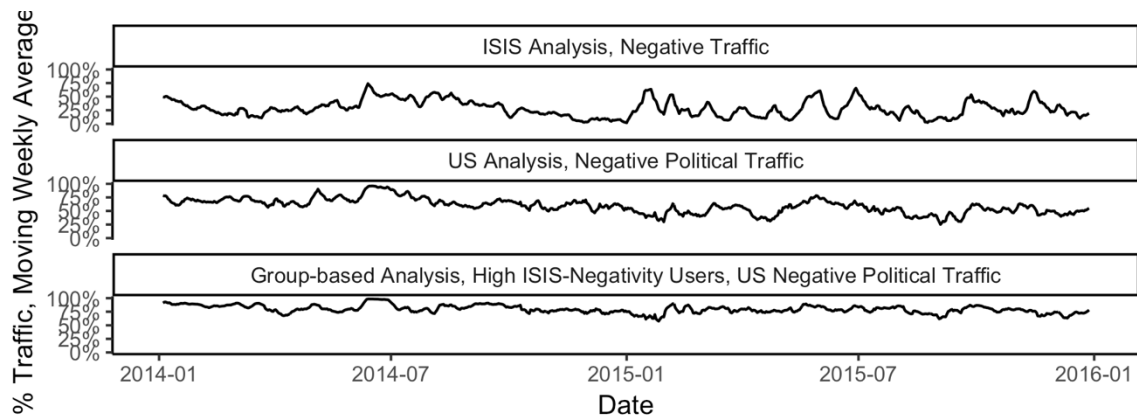


Figure 2 – Percent negative traffic in the ISIS, US, and Group-Based Analyses.

8. Discussion & Conclusion

For Arabic Twitter users posting about either ISIS or the United States during 2014–2015, we find that “the enemy of my enemy is not my friend.” On the contrary, negative views of ISIS and the United States are positively correlated, and this correlation is not spurious or due to uninteresting alternative explanations. Many Arabic Twitter users independently dislike both the United States and ISIS, but a sizeable proportion also believe in conspiratorial covert linkages between these two disliked entities.

¹⁰These temporal dynamics are further explored in section 4.1 of the online appendix.

These findings were unexpected and warrant further research on how publics in the Middle East and elsewhere interpret information they receive. As we note, we think our findings are most consistent with a story of motivated reasoning: Accounts of politics that seem implausible to well-informed outside observers often have great appeal to people with strong prior views. However, we recognize that the scope of our project limits us from exploring important questions related to this phenomenon. For instance, we cannot see the affect of outside political forces on the patterns we observe. To what extent are these patterns the result of elite cues and media framing from outside social media? How do outside actors such as these affect the acceptance and spread of conspiracy theories in the Arab Middle East? Additionally, we examine the relationship between only two entities, the US and ISIS. How do the patterns we observe generalize to attitudes toward other entities who are as involved or more in Middle East politics? As important as the US is to Middle Eastern politics, many countries of the Arab Middle East may face “enemies” more important to their citizens than the US or ISIS.

We hope these questions spur further research in this area, because it is one of great importance. Widespread acceptance of fake and inaccurate news can have a major impact on politics in the Middle East in particular, as well as elsewhere. Understanding how such views emerge and spread, and how they shape citizens’ interpretations of the actions of foreign states and non-state actors, is an important subject for continued research.

References

- Barceló, Joan and Elena Labzina. 2018. "Do Islamic State's deadly attacks disengage, deter, or mobilize supporters?" *British Journal of Political Science* 1–21.
- Berger, Jonathan M. and Jonathan Morgan. 2015. "The ISIS Twitter census: Defining and describing the population of ISIS supporters on Twitter." *The Brookings Project on US Relations with the Islamic World* 3 (20).
- Berger, Jonathan M. and Heather Perez. 2016. "The Islamic State's diminishing returns on Twitter: How suspensions are limiting the social networks of English-speaking ISIS supporters." *George Washington University Program on Extremism*.
- Bodine-Baron, Elizabeth, Todd C. Helmus, Madeline Magnuson, and Zev Winkelman. 2016. *Examining ISIS support and opposition networks on Twitter*. Santa Monica: Rand Corporation.
- Corstange, Daniel. 2016. "Anti-American behavior in the Middle East: Evidence from a field experiment in Lebanon." *Journal of Politics* 78 (1): 311–325.
- Duggan, Maeve. 2015. "Mobile messaging and social media 2015." *Pew Research Center*.
- Dubai School of Government. 2012. "Social media in the Arab world: Influencing societal and cultural change?" *Arab Social Media Report* 2 (1).
- Findley, Michael G. and Tze Kwang Teo. 2006. "Rethinking third-party interventions into civil wars: An actor-centric approach." *Journal of Politics* 68 (4): 828–837.
- Foyle, Douglas C. 1997. "Public opinion and foreign policy: Elite beliefs as a mediating variable." *International Studies Quarterly* 41 (1): 141–169.
- Furia, Peter A. and Russell E. Lucas. 2006. "Determinants of Arab public opinion on foreign relations." *International Studies Quarterly* 50 (3): 585–605.

- Gaines, Brian J., James H. Kuklinski, Paul J. Quirk, Buddy Peyton, and Jay Verkuilen. 2007. "Same facts, different interpretations: Partisan motivation and opinion in Iraq." *Journal of Politics* 69 (4): 957–974.
- Gray, Matthew. 2010. *Conspiracy theories in the Arab world: Sources and politics*. New York: Routledge.
- Greenberg, Karen J. 2016. "Counter-radicalization via the Internet." *The ANNALS of the American Academy of Political and Social Science* 668 (1): 165–179.
- Guess, Andrew, Jin Woo Kim, Brendan Nyhan, and Jason Reifler. 2019. "The distorting prism of social media: How online comments amplify toxicity." *Working Paper*.
- Hellwig, Timothy T., Eve M. Ringsmuth, and John R. Freeman. 2008. "The American public and the room to maneuver: Responsibility attributions and policy efficacy in an era of globalization." *International Studies Quarterly* 52 (4): 855–880.
- Hopkins, Daniel, and Gary King. 2010. "A method of automated nonparametric content analysis for social science." *American Journal of Political Science* 54 (1): 229–247.
- Hurst, Reuben, Taylor Tidwell, and Darren Hawkins. 2017. "Down the rathole? Public support for US foreign aid." *International Studies Quarterly* (61) 2: 442–454.
- Isernia, Pierangelo, Zoltán Juhász, and Hans Rattinger. 2002. "Foreign policy and the rational public in comparative perspective." *Journal of Conflict Resolution* 46 (2): 201–224.
- Jamal, Amaney A., Robert O. Keohane, David Romney, and Dustin Tingley. 2015. "Anti-Americanism and anti-interventionism in Arabic Twitter discourses." *Perspectives on Politics* 13 (01): 55–73.
- Joyce, Kyle A., Faten Ghosn, and Reşat Bayer. 2014. When and whom to join: The expansion of ongoing violent interstate conflicts. *British Journal of Political Science* 44 (1): 205–238.

- Katzenstein, Peter J., and Robert O. Keohane 2007. "Varieties of Anti-Americanism: A Framework for Analysis." In *Anti-Americanisms in World Politics*, ed. Peter J. Katzenstein and Robert O. Keohane. Ithaca, NY: Cornell University Press.
- Leber, Andrew and Alexei Abrahams. 2019. "A storm of Tweets: Social media manipulation during the Gulf crisis." *Review of Middle East Studies* 53 (2): 241–258.
- Lynch, Marc. 2007. "Anti-Americanism in the Arab World." In *Anti-Americanisms in World Politics*, ed. Peter J. Katzenstein and Robert O. Keohane. Ithaca, NY: Cornell University Press.
- Magdy, Walid, Kareem Darwish, and Ingmar Weber. 2015. "#FailedRevolutions: Using Twitter to Study the Antecedents of ISIS Support." *arXiv preprint arXiv: 1503.02401*.
- Mayda, Anna Maria and Dani Rodrik. 2005. "Why are some people (and countries) more protectionist than others?" *European Economic Review* 49 (2005): 1393–1430.
- Mitts, Tamar. 2018. "From isolation to radicalization: Anti-Muslim hostility and support for ISIS in the West." *American Political Science Review* 113 (1): 173–194.
- Mitts, Tamar. 2020. "Countering violent extremism and radical rhetoric." *International Organization*. Forthcoming.
- Nyhan, Brendan and Jason Reifler. 2010. "When corrections fail: The persistence of political misperceptions." *Political Behavior* 32: 303–330.
- Shapiro, Robert Y. and Benjamin I. Page. 1988. "Foreign policy and the rational public." *Journal of Conflict Resolution* (32) 2: 211–247.
- Siegel, Alexandra A. and Joshua A. Tucker. 2018. "The Islamic State's information warfare: Measuring the success of ISIS's online strategy." *Journal of Language and Politics* 17 (2): 258–280.

- Silverstein, Paul. 2000. "Regimes of (un)truth: Conspiracy theory and the transnationalization of the Algerian civil war." *Middle East Report* 214:6–10.
- Sokolov, Boris, Ronald F. Inglehart, Eduard Ponarin, Irina Vartanova, and William Zimmerman. 2018. "Disillusionment and anti-Americanism in Russia: From pro-American to anti-American attitudes, 1993–2009." *International Studies Quarterly* 62 (3): 534–547.
- Taber, Charles S. and Milton Lodge. 2006. "Motivated skepticism in the evaluation of political beliefs." *American Journal of Political Science* 50 (3): 755–769.
- Tessler, Mark, Michael Robbins, and Amaney Jamal. "What do ordinary citizens in the Arab world really think about the Islamic State"? Monkey Cage. July 27, 2016.
- Vice, Margaret. 2017. "In global popularity contest, U.S. and China – not Russia – vie for first." *Pew Research Center*. <https://www.pewresearch.org/fact-tank/2017/08/23/in-global-popularity-contest-u-s-and-china-not-russia-vie-for-first/>.
- Zeitsoff, Thomas, John Kelly, and Gilad Lotan. 2015. "Using social media to measure foreign policy dynamics: An empirical analysis of the Iranian–Israeli confrontation (2012–13)." *Journal of Peace Research* 52 (3): 368–83.
- Zeitsoff, Thomas. 2017. "How social media is changing conflict." *Journal of Conflict Resolution* 61 (9): 1970–91.