

# Countering Violent Extremism: Do Community Engagement Efforts Reduce Extremist Rhetoric on Social Media?\*

Tamar Mitts<sup>†</sup>

April 6, 2017

## Abstract

Over the past few years, efforts at countering violent extremism (CVE) have increased around the world. In the United States, much of the focus has been on community engagement – programs aiming to reduce radicalization by empowering local communities to identify warning signs of extremism before individuals engage in violence. The emphasis on community engagement is rooted in the idea that local knowledge held by families, neighbors, and friends is crucial for countering radicalization. Understanding whether engaging communities is effective is of paramount importance, especially with the rising accessibility of extremist materials on the Internet and social media. However, to date, there has been little systematic study of the effectiveness of community engagement programs in reducing radicalization in the United States. This paper uses new geo-located data on the online behavior of Islamic State supporters and their followers on Twitter, along with information on community engagement activities held by the Department of Homeland Security’s Office for Civil Rights and Civil Liberties during the Obama Administration from 2014 to 2016, to examine whether community engagement events are associated with reductions in pro-ISIS content on Twitter in these localities. The findings show that community engagement activities are followed by a decrease in online pro-ISIS rhetoric, especially in areas that have held a large number of these events.

---

\*This research was approved by the Columbia University Institutional Review Board under protocol IRB-AAAQ7657. I thank Shaked Doron for her research assistance.

<sup>†</sup>Assistant Professor (August 1, 2017), Gerald R. Ford School of Public Policy, University of Michigan. Ph.D. (expected May 2017), Columbia University, Department of Political Science. Email: tm2630@columbia.edu

# 1 Introduction

Countering radicalization and support for violent extremism is becoming a central policy area around the world. The rise of Islamic State (ISIS) and its ability to recruit individuals via online propaganda and social networks has intensified efforts to find solutions to extremist violence (Vidino and Hughes, 2015a). Since 2011, tens of thousands have radicalized in support for ISIS, some seeking to become foreign fighters and others attempting to plot terror acts in their home countries (Schmitt and Sengupta, 2015). What can be done to mitigate this wave of violence and extremism? Are there certain policies that governments can implement to counteract the narrative that Islamic State and other groups promote among their followers? While there are several possible responses to extremism,<sup>1</sup> this paper focuses on one specific strategy: engaging communities in thwarting radicalization in the United States. Community engagement has been a central counter-extremism policy pursued by governments around the world, particularly since 9/11 (Briggs, 2010; Challgren et al., 2016; Romaniuk, 2015).

Engaging communities is rooted in earlier models of community policing developed in the 1990s.<sup>2</sup> Unlike professional policing that focuses on the response and prosecution of crime, community policing emphasizes crime prevention by addressing specific needs of local communities and by involving citizens in police activities (Cordner, 2014). For example, instead of relying solely on the police to actively intervene and put a stop to drug dealing, community members assist by monitoring and reporting the activity of drug dealers in their neighborhood. In a similar manner, engaging communities in countering extremism is based on the idea that extremist violence can be prevented by involving local communities in efforts to detect early signs of extremism.

This logic relies on several assumptions. First, that radicalization is a process that begins with a cognitive stage in which an individual embraces an extremist ideology, which culminates in behavioral manifestations of radicalization like committing violence (Neumann, 2013; Sedgwick, 2010). Some scholars do not agree with this assumption, arguing that not all extremist behaviors are preceded by the adoption of radical ideologies (Borum, 2011). Many, however, agree that

---

<sup>1</sup>For example, countering extremist propaganda online (Fernandez, 2015) or initiating educational programs among vulnerable populations (Aldrich, 2014)

<sup>2</sup>In the United States, community policing was formally enacted in the United States in the 1994 Violent Crime Control and Law Enforcement Act. see <https://www.congress.gov/bill/103rd-congress/house-bill/3355>

radicalization follows a sort of continuum from mere ideology to actual violence (Horgan, 2008; Wilner and Dubouloz, 2010; Crossett and Spitaletta, 2010).

Second, models of community engagement rest on the premise that the cognitive phase of radicalization can be detected by people who are close to a radicalizing individual. As members of local communities tend to have close personal connections, they are most likely to notice changes in behavior and thus serve as “early warning systems” of violent extremism (Briggs, 2010). Third, proponents of community engagement assume that individuals who experience cognitive radicalization can be swayed away from the radicalization path by members of their community. By addressing local communities’ concerns and engaging in activities that reinforce a sense of belonging in the broader society, governments believe that they can counter the “us versus them” narrative that extremist groups promulgate in their recruitment propaganda (Executive Office of the President of the United States, 2011).

While many agree that stopping violent extremism is of paramount importance, not everyone believes that government-sponsored community engagement is the right way to go. Civil rights and Muslim advocacy groups strongly criticize efforts to engage Muslim communities in countering radicalization. The Council on American-Islamic Relations (CAIR), for example, has argued that government-sponsored community engagement is likely to be ineffective, as community figures working with the government are viewed as not credible by radicalizing individuals. In addition, CAIR stressed that community engagement efforts unjustly focus on Muslim communities, even though far-right extremist violence has led to far greater casualties than violence by Islamic extremists (Council on American-Islamic Relations, 2016). The American Civil Liberties Union (ACLU) has stressed that involving Muslim communities in countering extremism pressures community members to monitor each other, which can harm community cohesion and can easily lead to government overreach (American Civil Liberties Union, 2016*a*). Indeed, focusing only on the Muslim community stigmatizes Muslims and reinforces Islamophobic stereotypes, which can be counterproductive from a counter-radicalization standpoint (Patel and German, 2015).

For these reasons, it is of crucial importance to study the effectiveness of community engagement to counter extremism. However, empirically evaluating these efforts is a challenging task. One reason is that community engagement events are not randomized; focusing on the behavior of a very small minority in society, they almost always selectively target specific communities. Moreover, an

‘effective’ counter-radicalization program requires obtaining some sort of measurable evidence of a “decrease” in radicalization. This is challenging, as observing the absence of radical sympathies does not necessarily imply a decline but may simply reflect the absence of a tendency toward extremist ideologies from the outset. Third, many counter-extremism efforts are not reported to the public,<sup>3</sup> which makes it challenging to systematically study them.

In this paper, I take advantage of publicly reported community engagement events held by the Department of Homeland Security’s Office of Civil Rights and Civil Liberties (CRCL) from 2014 to 2016, and combine them with high-frequency, geo-located panel data on tens of thousands of individuals who follow Islamic State accounts on Twitter, to examine whether community engagement activities are systematically associated with changes in pro-ISIS rhetoric at the local level. The use of high-frequency Twitter data helps overcome the challenge of ‘non-evidence’ in evaluating counter-extremism efforts, as it serves as a continuous measure of online rhetoric of individuals at risk of radicalization. Analyzing over a hundred community engagement events in a Difference-in-Differences design, I show that community engagement activities are followed by a significant decrease in online pro-ISIS chatter, especially in localities in which CRCL has held a large number of events. The next section describes in detail the context of the study, and the role of community engagement events in the Obama Administration’s strategy to counter violent extremism in the United States.

## 2 Countering violent extremism in the United States

In August 2011, the Obama Administration initiated a counter radicalization strategy, “Empowering Local Partners to Prevent Violent Extremism in the United States,” to prevent extremist violence in its territories (Obama, 2011). The plan focused on three main areas. First, it sought to increase and strengthen the government’s engagement with local communities whose members may be targeted by violent groups. This effort was based on the notion that community members with personal

---

<sup>3</sup>In fact, in February 2016, ACLU filed a lawsuit under the Freedom of Information Act against the Department of Homeland Security, Department of Justice, Federal Bureau of Investigation, Office of the Director of National Intelligence, Department of State, Department of Health and Human Services, and the Department of Education for not releasing records of their countering violent extremism activities. See: [https://www.aclu.org/sites/default/files/field\\_document/cve\\_foia\\_complaint\\_2.10.16.pdf](https://www.aclu.org/sites/default/files/field_document/cve_foia_complaint_2.10.16.pdf)

connections to radicalizing individuals — for example, teachers, friends, or family members — are best positioned to detect changes in behavior that might convey early signs of extremism. Building relationships and trust with local communities was seen as important to accessing crucial information on individuals at early stages of radicalization and to provide an opportunity for communities to give feedback on the government’s CVE efforts. When setting out its strategy, the Obama Administration stated:

*“Engagement is essential for supporting community-based efforts to prevent violent extremism because it allows government and communities to share information, concerns, and potential solutions. Our aims in engaging with communities to discuss violent extremism are to (1) share sound, meaningful, and timely information about the threat of radicalization to violence with a wide range of community groups and organizations, particularly those involved in public safety issues; (2) respond to community concerns about government policies and actions; and (3) better understand how we can effectively support community-based solutions.”*<sup>4</sup>

In addition to community engagement, the initiatives focused on increasing training for government and law enforcement on preventing radicalization and extremism, and seeking ways to counter the propaganda spread by violent groups on the Internet and social media.

The significance that the American government placed on finding solutions to violent extremism increased with the rise of Islamic State, its vast online propaganda machine, and its efforts to recruit foreign fighters around the world (Vidino and Hughes, 2015*a*). Even though individuals had radicalized in America prior to the rise of ISIS, the pace at which the group attracted supporters was unprecedented compared to prior conflicts. From 2014 to 2016, over a hundred individuals have been charged in the U.S. with criminal behavior related to Islamic State (Vidino and Hughes, 2015*b*). Activities that led to charges included providing material support to ISIS and its affiliates, traveling or planning to travel to Syria to become foreign fighters, or plotting violent attacks in the territories of the United States (Greenberg, 2016). While those who displayed ‘behavioral radicalization’ in the United States are only a tiny minority, security agencies estimate that the number of people who sympathize with ISIS’s ideology — those who display ‘cognitive radicalization’ — is much larger, possibly in the thousands (Vanden Brook, 2015).

---

<sup>4</sup>Obama (2011), p. 5

In order to prevent radicalization and violence, the United States CVE strategy seeks to support locally-based activities that can sway individuals who are at the early stage of radicalization from the path of extremism (Bjelopera, 2012; Challgren et al., 2016). In late 2014, the government launched a “Three City Pilot” program in three cities in the United States: Boston, Los Angeles and Minneapolis–St. Paul in order to create local solution to ISIS-inspired radicalization. Recommendations from the program included, among other things, increasing local communities’ understanding of extremism with training, enhancing collaboration between communities and law enforcement, and building networks between public and private groups to counter extremism (Challgren et al., 2016; Vidino and Hughes, 2015*a*).

At the federal level, several agencies have been tasked with implementing the government’s CVE strategy.<sup>5</sup> In this paper, I focus on the activities of the Department of Homeland Security. The Office of Civil Rights and Civil Liberties (CRCL) in the Department of Homeland Security is responsible for the Department’s community engagement efforts. Its goals include “promoting the respect of civil rights and civil liberties in policy creation and implementation” and “communicating with individuals and communities whose civil rights and civil liberties may be affected by Department activities” (U.S. Department of Homeland Security, 2016). In the past few years, CRCL has been holding various community engagement events across the United States to facilitate relationships with local communities and to enhance counter extremism efforts. These activities include (Office for Civil Rights and Civil Liberties, 2016):

- **Community roundtables.** Events that bring together government officials from the federal, state, and local level and leaders from American Arab, Muslim, South Asian, Middle Eastern, Somali, Sikh, Latino, Jewish, an Asian/Asian Pacific Islander communities, in order to strengthen relationships and engagement.
- **Consultation with communities on CVE.** Events in which CRCL representatives meet with local communities to share information, discuss community concerns related to extremism, and receive community input on the effects of the Department’s policies on the ground.

---

<sup>5</sup>Specifically, the Federal Government tasked the Department of Homeland Security, the National Counter Terrorism Center, and Federal Bureau of Investigation, and the Department of Justice implement the governments countering violent extremism strategy (Executive Office of the President of the United States, 2011).

- **Community awareness briefing.** Meetings in which community members and law enforcement officials are presented with information on the process of radicalization and foreign fighter recruitment by violent groups, in order to increase awareness and knowledge of the phenomenon.
- **Community resilience exercise.** Events in which law enforcement and local communities participate in a half-day exercise designed to build trust and communication, and to empower communities against violent extremism. The training includes discussion of a hypothetical scenario of violent activity, and evaluates the way in which it affects law enforcement and community members. The exercise concludes with developing a local plan to prevent extremism.

Many community engagement meetings held by CRCL do not exclusively focus on extremism. Instead, they cover a wide range of issues related to the Department’s activities. The collaboration between communities and the government is meant to facilitate a “shared sense of belonging” among communities and government officials, which arguably helps undermine extremist propaganda (Executive Office of the President of the United States, 2011).

Figure 1: Community roundtables in Atlanta, GA and Phoenix, AZ



*Photo credit: Islamic Speakers Bureau, Atlanta, Islamic Community Center of Phoenix.<sup>6</sup>*

### 3 Criticism of the U.S. CVE program

That said, many civil rights and Muslim advocacy organizations across the United States, such as the American Civil Liberties Union (ACLU) and the Council on American-Islamic Relations (CAIR), have strongly opposed the American CVE strategy on several grounds. First, they argue that the program is ineffective. Community engagement events sponsored by the government are not likely to be viewed as credible by individuals attracted to extremist propaganda. As the ideology promoted by ISIS and other groups intentionally calls for fighting the American government, initiatives stemming from the government are likely to be viewed with suspicion (Council on American-Islamic Relations, 2015).

In addition, these groups argue that United State's CVE program lacks clear leadership, receives attention only after terrorist attacks, and tends to depend on the whim of local authorities (Council on American-Islamic Relations, 2016; Patel and German, 2015). Since Muslim communities are already targeted with hate crimes and Islamophobia, especially after terrorist attacks, efforts of CVE programs can further a sense of alienation and hostility that can feed into grievances capitalized upon by groups like Islamic State. Finally, critics of the United State's CVE strategy argue that the program is based on a false model of radicalization, which assumes that violent behavior can be predicted by the expression of certain beliefs. However, numerous empirical studies have shown that there is no single path to radicalization (American Civil Liberties Union, 2016*b*; Patel and German, 2015).

The third argument against the American counter-extremism strategy claims that it is unjust. Even though the majority of casualties since 9/11 have been caused by far-right terrorism, counter-radicalization efforts almost always target Muslims (Council on American-Islamic Relations, 2016; Patel and German, 2015). Critics argue that the exclusive focus on Muslim communities stigmatizes Muslims and Islam, and tends to encourage anti-Muslim sentiment. In addition, community engagement events are not as benign as they might seem, because government agencies use them for spying and intelligence gathering. This sort of infiltration of community spaces solely on the basis of religion is unfair, according to civil rights activists (American Civil Liberties Union, 2016*b*).

---

<sup>6</sup><http://archive.constantcontact.com/fs036/1101532851599/archive/1108633466286.html>,  
<http://iccpaz.com/dhs-hosts-crcl-roundtable-at-the-iccp/>

<http://iccpaz.com/dhs-hosts-crcl-roundtable-at-the-iccp/>



Furthermore, community engagement frequently tasks members with monitoring each other's behavior. Such monitoring can create a climate of fear, discourage free expression of political opinions, and can be used by the government to suppress dissent (American Civil Liberties Union, 2016*b*; Patel and German, 2015). Finally, critics of CVE argue that the strategy imposes heavy costs on Muslim communities by harming community cohesion, increasing suspicion among members, and by framing community relations with the government only on the basis of security issues (American Civil Liberties Union, 2016*b*; Patel and German, 2015).

These are all very important considerations that should be taken into account when evaluating counter-radicalization policies. To date, however, there has not been a systematic empirical analysis of the link between community engagement events and observable measures of pro-ISIS radicalization. This study brings together new geo-located data on the online behavior of individuals who follow Islamic State accounts on Twitter, who might be at risk of radicalization,<sup>7</sup> along with data on community engagement events held during the Obama Administration by the Department of Homeland Security's Office of Civil Rights and Civil Liberties. I use these new sources of data to examine whether there is a link between community engagement activities and pro-ISIS rhetoric on Twitter.

This is an important step forward in understanding the potential effects of community engagement, as there has been very little systematic empirical evaluations of these programs, especially in the United States. I should note, however, that while this study sheds light on the possible impact of community engagement on the behavior of ISIS followers online, it does not allow concluding that changes in pro-ISIS chatter reflect a decline in radicalization. It is equally possible that these events reduce pro-ISIS chatter by suppressing political expression. Below, I describe the data collection, research design, and results, and discuss several tests to examine the alternative explanation that community engagement events might be discouraging online expression.

---

<sup>7</sup>Online measures of pro-ISIS rhetoric can plausibly proxy underlying support for extremism. Among more a hundred individuals charged in the United States with providing material support for ISIS or plotting a violent attack on the organization's behalf, about 63% used social media when they were radicalizing, and among those, 86% expressed their support for ISIS in publicly viewable posts. This study uses large amounts of publicly viewable posts by ISIS supporters in the United States to measure support for violent extremism. See Appendix Table A3 for more details.

## 4 Data

In this section, I describe the data collection for this study. First, I describe how I collected information on the timing and location of community engagement events held by the Department of Homeland Security’s Office of Civil Rights and Civil Liberties in various locations across the United States. Second, I describe how I collected data on the online rhetoric of individuals located in the United States who follow Islamic State accounts on Twitter. Finally, I discuss how I matched ISIS followers to community engagement events taking place in their areas based on granular information on their geographic location.

### 4.1 Community engagement events

The Office of Civil Rights and Civil Liberties (CRCL) in the Department of Homeland Security has been holding community engagement events since 2003 (Bjelopera, 2012). The first event took place in Dearborn, Michigan, and community activities soon expanded to other cities in the United States (Schlanger, 2011). In the end of 2010, CRCL began publishing monthly newsletters in which it provided information on its community engagement activities:


*“This is the first of CRCL’s new monthly newsletters. Our goal is to inform members of the public about the Office’s activities, including how to make complaints; ongoing and upcoming projects; opportunities to offer comments and feedback ... Public engagement with diverse American communities plays a key role in the DHS mission to protect America while preserving our freedoms ... We are hard at work expanding our engagement program, building a strong stakeholder network of community-based organizations across the country – this newsletter is a part of that effort.”* (Schlanger, 2011)

I collected information on all events held by CRCL using these monthly reports, which began providing systematic information on events in 2011. I gathered data on the dates of these events, the cities in which they took place, and the type of engagement activity carried out in each event. These included community roundtables, community awareness briefings, and community resilience exercises, among others. Figure 2 displays the number of cities in which CRCL held community engagement activities since 2011. Figure 3 shows the number of community engagement events by month. Table A1 in Appendix A provides detailed information on each event.<sup>8</sup>

---


<sup>8</sup>In this paper, I focus on events that took place from 2014 to 2016. Thus, Table A1 describes these events in detail.

Figure 2: Number of cities in which CRCL held community engagement events, by year



*Note:* The figure presents the number of cities in which The Office for Civil Rights and Civil Liberties (CRCL) in the Department of Homeland Security held community engagement events each year since 2011.

Figure 3: Number of CRCL community engagement events, by month



*Note:* The figure presents the number of community engagement events held each month by The Office for Civil Rights and Civil Liberties (CRCL) in the Department of Homeland Security since 2011.

## 4.2 Islamic State supporters on Twitter

To evaluate the possible impact of these community engagement events on the behavior of individuals attracted to ISIS's ideology, I used original Twitter data on Islamic State supporters in the United States, which comes from a larger database on ISIS-affiliated accounts across the world (Mitts, 2017). Below, I provide an overview of the data collection procedure, which included (i)

identifying accounts of Islamic State supporters on Twitter and downloading information on their posting history, (ii) coding the extent to which their posts reflected extremist ideology, and (iii) predicting their geographic location using network data. Appendix B provides more details on the data collection method.

#### **4.2.1 Identifying Islamic State accounts on Twitter**

First, I identified about 15,000 accounts of Islamic State activists — accounts that actively disseminated ISIS propaganda online — that were flagged for suspension from Twitter by the group Controlling Section (@CtrlSec). Controlling Section has been monitoring, since 2015, Twitter accounts identified with ISIS and publicly flagging them for suspension. I downloaded every available piece of information on these accounts before they were suspended from Twitter, including user-level data such as profile picture, description, and self-described location, as well as complete historical tweet timelines. In addition to the core list of ISIS activists, I collected user-level data and tweeting history for all the followers of these accounts, which amount to about 1.6 million users. The followers group includes individuals who follow one or more ISIS activist accounts.

#### **4.2.2 Measuring online expression of extremist ideology**

Using the historical tweet timelines for these accounts, I measured the extent to which each tweet represented pro-ISIS content. Specifically, I used supervised machine learning to classify tweets in four different languages (English, Arabic, French, and German) into one or more of the categories listed below.<sup>9</sup>

1. *Travel to Syria or foreign fighters* - tweets describing interest or intent to travel to Syria, and/or discussion of foreign fighters
2. *Sympathy with ISIS* - expressions of support or sympathy with the Islamic State, its ideology and its activities in territories under its control
3. *Life in ISIS territories* - tweets describing the life of ISIS activists in the territories controlled by the Islamic State

---

<sup>9</sup>I also coded whether tweets represented discourse on Islam in general, but in this paper I focus on topics that can capture pro-ISIS rhetoric.

4. *Syrian war*- tweets describing events in the Syrian civil war and/or discussion/analysis of those events
5. *Anti-West* - anti-West rhetoric, criticizing Western countries' foreign policy and military operations in the Middle East

I asked human coders from two crowdsourcing platforms, Amazon Mechanical Turk and Crowdflower, to manually label a training set of randomly selected Twitter posts in Arabic, English, French and German. Each tweet was labeled by three coders, and label(s) were retained for a given tweet only if at least two out of the three coders assigned the same label(s) to the tweet.<sup>10</sup> Using the labeled training set, I predicted the content of all unlabeled tweets using the elastic-net generalized linear model (Friedman, Hastie and Tibshirani, 2010), where the regularization parameter  $\lambda$  was selected by cross-validation. The algorithm employed information on the words in each labeled tweet to 'learn' the categorization rules to classify unlabeled tweets. Appendix B provides information on the coding procedure and model performance.

### 4.2.3 Predicting ISIS supporters' geographic locations

In order to estimate where ISIS activists and followers are located, I employed a spatial label propagation algorithm developed by Jurgens (2013).<sup>11</sup> This algorithm predicts users' geographic location using geo-location information available in the network, along with information on network structure and the strength of ties between users. Since a very small share of users enabled geo-tagging of their tweets or provided location information in their accounts, I predicted the geographic locations for all users to avoid relying on the small selected subset of users with reported locations. Appendix B provides more details on the method, along with information on its prediction accuracy and stability.

As the goal of this study is to match ISIS supporters to community engagement events held by the Department of Homeland Security, I used the predicted geographic coordinates of users to determine which were located in the United States. Thus, in this paper, I analyze the online rhetoric of about 47,000 ISIS-affiliated accounts predicted to be located in America, examining whether they

---

<sup>10</sup>The coders were proficient in the languages of the tweets that they labeled.

<sup>11</sup>The use of location prediction in social network research is a growing field (for example, see Backstrom, Sun and Marlow (2010); McGee, Caverlee and Cheng (2013); Jurgens et al. (2015)). To the best of my knowledge, this study is one of the first applications of these methods in the study of Islamic State online networks.

changed their pro-ISIS rhetoric in the aftermath of community engagement events. Specifically, and as described in more detail in Section 5, I use Difference-in-Differences estimations for each community engagement event to compare changes in pro-ISIS rhetoric by individuals located in the area of the event to changes in such rhetoric by all other individuals. Figure 4 displays the predicted locations of these accounts (blue dots), along with the locations of CRCL community engagement events from 2014 to 2016 (orange dots).

Table 1 provides summary statistics on the tweeting patterns of ISIS sympathizers in the United States. The top panel shows the number of tweets that each user posted on each topic from 2014 to 2016. For example, ISIS followers in America posted an average of about 9 tweets discussing travel to Syria or foreign fighters, with a maximum of 230 tweets, and an average of 5.5 tweets expressing sympathy with ISIS, with a maximum of 145 tweets.<sup>12</sup> The bottom panel provides additional information on these accounts, such as whether they were flagged for suspension by the group Controlling Section (about 0.2%), whether they were already suspended (2.6%), and the number of ISIS-affiliated accounts that each user followed (the mean being about 3 accounts, with a maximum of 1,793).

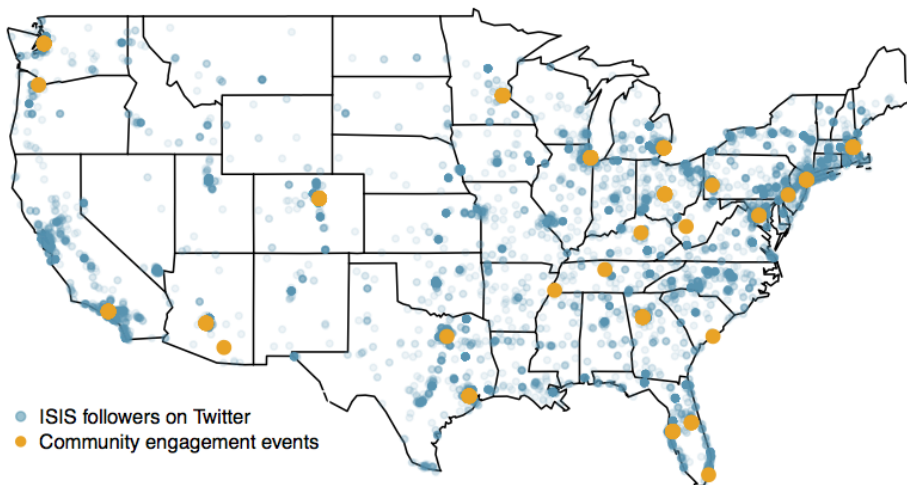
Table 1: Summary statistics for ISIS supporters in the United States

Statistic	N	Mean	St. Dev.	Min	Max
Travel to Syria or foreign fighters (#)	35,248	8.824	19.092	0	230
ISIS sympathy (#)	35,248	5.597	12.358	0	145
Life in ISIS territories (#)	35,248	8.946	19.719	0	262
Syrian war (#)	35,248	4.153	9.467	0	118
Anti-West (#)	35,248	4.721	10.513	0	151
All topics (#)	35,248	21.666	46.075	0	561
Flagged as an ISIS activist (0/1)	47,296	0.002	0.041	0	1
Suspended by Twitter (0/1)	47,287	0.026	0.159	0	1
ISIS accounts following (#)	47,287	2.970	17.431	0	1,793

*Note:* The table reports summary statistics for accounts of ISIS activists and followers located in the United States. The number of tweets for each topic reflect tweets that coded 1 if their predicted value of belonging to the topic (i.e., sympathy with ISIS, life in ISIS territories, travel to Syria or foreign fighters, Syrian war, or anti-West) was above the mean of the predicted values for that topic, and 0 if not.

<sup>12</sup>In Table 1, a tweet was coded as belonging to a topic if its predicted value from the classification model described in section 4.2.2 was above the mean of the predicted values for this category.

Figure 4: ISIS supporters and CRCL Community Engagement Activities in the United States



*Note:* The figure plots the predicted locations of accounts of ISIS activists and their followers in the United States (blue dots). In addition, it shows the locations in which CRCL held community engagement events from 2014 to 2016 (orange dots).

### 4.3 Creating datasets for each community engagement event

In order to facilitate analysis of the relationship between community engagement activities and pro-ISIS rhetoric on Twitter, I created separate datasets of Twitter posts produced by ISIS sympathizers in the United States, around each of 112 engagement events held by the Department of Homeland Security from 2014 to 2016. For each event, I identified the tweets posted by ISIS sympathizers in the 7, 14, 21, and 30 days before and after the event. I created two binary indicators to (i) differentiate between posts appearing before and after each event, as well as (ii) distinguish between tweets posted by individuals located in or out of the area of the event. Specifically, for each community engagement event, I created the variable *Post*, which is coded 1 when a tweet appeared after the event and 0 otherwise, and a variable *In event area*, which is coded 1 when a tweet was posted by an individual located in the area of the event and 0 otherwise. Finally, to quantify the extent to which each post expressed pro-ISIS rhetoric, I used the predicted values generated for each tweet by the classification model described in section 4.2.2, for each of the five content categories. Table A1 in Appendix A provides summary statistics for each of these 112 datasets. Each row represents a different dataset for a different community engagement event, and the columns show the distribution

of the variables described above in each of the 112 datasets.

## 5 Research design and results

Since this study analyzes the relationship between pro-ISIS rhetoric and over a hundred community engagement events, it is impractical to present regression results for each event separately. To uncover patterns underlying all events held by the Department of Homeland Security from 2014 to 2016, I employ two types of analysis. First, I conduct a pooled analysis where I examine all community engagement events simultaneously. Second, I carry out meta analysis of the results of individual events, as described in detail below. Meta analysis is a useful tool for the purpose of this study, as it allows systematically evaluating the relationship between community engagement events and online pro-ISIS rhetoric when considering many events simultaneously. In addition, it enables examining how the relationship varies as a function of event characteristics. In the first part of this section, I describe the Difference-in-Differences model I used to analyze community engagement events when all events are pooled together. In the second part, I describe the meta analysis method I employed to evaluate the overall impact of 112 community engagement events on the rhetoric of ISIS sympathizers in the United States.

### 5.1 Identification strategy

The key identifying assumption in this Difference-in-Differences design is that in the absence of a community engagement event, individuals located in the event area and individuals who do not would follow parallel trends in their online expression of pro-ISIS rhetoric. While it is possible that community engagement events target specific areas that might be more prone to have individuals “at risk” of radicalization (Obama, 2011), the *over-time changes* in pro-ISIS online posting should not be significantly different between the groups before the occurrence of community engagement events.

To empirically test this identification assumption, I visually examine whether the two groups display parallel trends before CRCL community engagement events. Figure 5 plots pre- and post-time trends in pro-ISIS rhetoric — a standardized variable capturing online posts on all topics: (i) travel to Syria or foreign fighters, (ii) ISIS sympathy, (iii) life in ISIS territories, (iv) Syrian war, and



(v) anti-West — for the group of individuals located in event areas (black) and the group of those who do not (gray). The x-axis is the number of days between the date of a community engagement event and the date in which ISIS followers posted on Twitter. In order to observe time trends for all events simultaneously, the figure normalizes, for all CRCL events, the difference in days between community engagement events and the timing of Twitter posts. I calculated the average pro-ISIS content by each of the two groups in each day, and applied nonparametric smoothing piecewise to the pre- and post- time periods for each group, using a Gaussian kernel with a bandwidth of 30 days.

Figure 5 shows that the trends in pro-ISIS rhetoric over time for individuals located in event areas and those who do not are parallel prior to the day in which CRCL held community engagement events. Only after community engagement events we observe a shift in those trends, where pro-ISIS content by individuals in event areas decreases, but the rhetoric of those outside of event areas does not change. Interestingly, we also observe that the average pro-ISIS rhetoric is overall higher for individuals located in event areas. This suggests that CRCL may intentionally target locations that might have greater numbers of individuals at risk of radicalization.


## 5.2 Difference-in-Differences model

To examine the relationship between community engagement events and pro-ISIS rhetoric on Twitter, I estimate the following least squares regression model:

$$y_{i,j,k} = \beta_1 Post_{i,j,k} + \beta_2 In\ event\ area_{i,j,k} + \beta_3 (Post_{i,j,k} \times In\ event\ area_{i,j,k}) + \alpha_k + \varepsilon_j \quad (1)$$

where  $y_{i,j,k}$  is the predicted value of a given topic (i.e., travel to Syria or foreign fighters, sympathy with ISIS, life in ISIS territories, Syrian war, anti-West) for tweet  $i$  posted by user  $j$  surrounding event  $k$ ;  $Post$  is an indicator coded 0 where a tweet appears before event  $k$  and 1 afterwards;  $In\ event\ area$  is an indicator coded 1 when a tweet was posted by an individual who is predicted to be located in the area of event  $k$ , and 0 otherwise; and  $\alpha_k$  is an event fixed effect. In this specification,  $\beta_3$  is the Difference-in-Differences coefficient of interest, reflecting how the change in pro-ISIS rhetoric after community engagement events is different for individuals located in event areas, compared to the change in rhetoric of users outside event areas. Standard errors are clustered at the user level

Figure 5: Pro-ISIS rhetoric: Parallel trends



*Note:* The Figure presents pre- and post-time trends in pro-ISIS rhetoric — a standardized variable capturing online posts on all topics: (i) travel to Syria or foreign fighters, (ii) ISIS sympathy, (iii) life in ISIS territories, (iv) Syrian war, and (v) anti-West — for the group of individuals located in event areas (black) and the group of those who do not (gray). The x-axis is the number of days between the community engagement event date and the day in which ISIS followers posted on Twitter. I calculated the average pro-ISIS content by each of the two groups in each day, and applied nonparametric smoothing piecewise to the pre- and post- time periods for each group, using a Gaussian kernel with a bandwidth of 30 days.

to account for serial correlation in tweet content posted by the same user. All outcome variables are standardized.

### 5.3 Pooled analysis results

Table 2 reports the results estimated from a pooled Difference-in-Differences analysis of the relationship between 112 community engagement events and pro-ISIS rhetoric on social media, captured 7, 14, 21, and 30 day before and after each event. The results encompass data from 32,694,069 (7-day results), 64,848,887 (14-day results), 100,860,573 (21-day results), and 147,372,103 (30-day results) geo-located tweets generated from 2014 to 2016 by individuals who follow Islamic State

accounts on Twitter. The coefficient on the interaction term,  $Post \times In\ event\ area$ , represents the Difference-in-Differences coefficient of interest.

Overall, the findings show that community engagement events are systematically associated with reductions in pro-ISIS rhetoric on Twitter. Each of the six panels in Table 2 reports the results for a different content category. It can be seen that the Difference-in-Differences coefficient,  $Post \times In\ event\ area$ , is negative and statistically significant at the 1% to 5% level for almost all categories. When analyzing data from 30 days before and after the events and considering all topics together (Table 2, Panel 1, Column 4), the results show that community engagement events are linked to a reduction of about 4 percentage points in pro-ISIS rhetoric among individuals located in event areas. This is a relative decrease of 18% over the baseline coefficient of 0.233 standard deviations, which is the pre-post decrease in pro-ISIS content among individuals located in areas where CRCL held community engagement events.

The findings also hold when considering the content categories separately. For example, posting on the topic ‘Travel to Syria or foreign fighters’ decreases by about 3 percentage points after community engagement events in areas where they take place, which is a relative decrease of 16% over the baseline coefficient of 0.189 standard deviations. Similarly, discourse on the topic ‘ISIS sympathy’ decreases by about 2 percentage points (a relative decrease of 26% over the baseline) and posting on ‘Life in ISIS territories’ is reduced by about 3 percentage points (a relative decrease of 13% over the baseline). Tweets discussing the Syrian civil war decrease by about 2 percentage points after community engagement events in areas where they take place, which is a relative decrease of 22% over the baseline of 0.1 standard deviations. Interestingly, the results for anti-West rhetoric, which includes a lot of anti-America tweets,<sup>13</sup> do not significantly change after community engagement events among individuals located in event areas. When considering shorter time windows, such as 14 and 21 days before and after community engagement events, the results hold as well, but the coefficient estimates have slightly smaller magnitudes.

---

<sup>13</sup>For examples, see Table B13 in Appendix B.

Table 2: Pooled Diff-in-Diff Analysis: Community engagement and pro-ISIS rhetoric on Twitter

	(1) 7 Days	(2) 14 Days	(3) 21 Days	(4) 30 Days
<b>1. All Topics</b>				
Post	-0.007*** (0.001)	0.008*** (0.001)	0.028*** (0.001)	0.049*** (0.001)
In event area	0.204*** (0.011)	0.215*** (0.009)	0.228*** (0.009)	0.233*** (0.008)
<b>Post × In event area</b>	<b>-0.024*</b> (0.014)	<b>-0.039***</b> (0.011)	<b>-0.033***</b> (0.008)	<b>-0.042***</b> (0.007)
$R^2$	0.01	0.009	0.008	0.006
<b>2. Travel to Syria or foreign fighters</b>				
Post	-0.005*** (0.001)	0.005*** (0.001)	0.017*** (0.000)	0.033*** (0.001)
In event area	0.167*** (0.007)	0.171*** (0.006)	0.181*** (0.006)	0.189*** (0.005)
<b>Post × In event area</b>	<b>-0.014</b> (0.01)	<b>-0.023***</b> (0.007)	<b>-0.026***</b> (0.006)	<b>-0.032***</b> (0.005)
$R^2$	0.005	0.004	0.004	0.003
<b>3. ISIS sympathy</b>				
Post	-0.002*** (0.001)	0.004*** (0.000)	0.012*** (0.000)	0.021*** (0.000)
In event area	0.052*** (0.007)	0.058*** (0.005)	0.065*** (0.005)	0.067*** (0.005)
<b>Post × In event area</b>	<b>-0.011</b> (0.01)	<b>-0.016**</b> (0.007)	<b>-0.013**</b> (0.005)	<b>-0.018***</b> (0.005)
$R^2$	0.002	0.002	0.002	0.001
<b>4. Life in ISIS territories</b>				
Post	-0.006*** (0.001)	0.005*** (0.001)	0.024*** (0.001)	0.043*** (0.001)
In event area	0.245*** (0.011)	0.255*** (0.008)	0.268*** (0.008)	0.271*** (0.007)
<b>Post × In event area</b>	<b>-0.012</b> (0.013)	<b>-0.031***</b> (0.01)	<b>-0.029***</b> (0.007)	<b>-0.034***</b> (0.007)
$R^2$	0.008	0.007	0.006	0.005
<b>5. Syrian war</b>				
Post	-0.007*** (0.001)	0.003*** (0.001)	0.016*** (0.000)	0.028*** (0.001)
In event area	0.08*** (0.008)	0.088*** (0.006)	0.096*** (0.006)	0.1*** (0.005)
<b>Post × In event area</b>	<b>-0.016</b> (0.011)	<b>-0.022***</b> (0.008)	<b>-0.017***</b> (0.006)	<b>-0.022***</b> (0.006)
$R^2$	0.004	0.003	0.003	0.002
<b>6. Anti-West</b>				
Post	-0.004*** (0.001)	0.003*** (0.001)	0.013*** (0.000)	0.022*** (0.001)
In event area	0.007 (0.008)	0.02*** (0.006)	0.026*** (0.005)	0.028*** (0.005)
<b>Post × In event area</b>	<b>0.023</b> (0.016)	<b>0.001</b> (0.01)	<b>0.000</b> (0.007)	<b>-0.005</b> (0.006)
$R^2$	0.003	0.002	0.002	0.002

*Note:* The table reports coefficients estimated from a pooled Difference-in-Differences analysis of the relationship between 112 community engagement events and pro-ISIS rhetoric on social media, captured 7, 14, 21, and 30 day before and after each event. The results encompass data from 32,694,069 (7-day results), 64,848,887 (14-day results), 100,860,573 (21-day results), and 147,372,103 (30-day results) geo-located tweets generated from 2014 to 2016 by individuals who follow Islamic State accounts on Twitter. All outcome variables are standardized. The analysis includes event fixed effects and clustered standard errors (reported in parentheses) at the user level. \* p<0.10, \*\* p<0.05, \*\*\* p<0.01.

## 5.4 Meta analysis

Another approach to examine whether community engagement activities influence pro-ISIS rhetoric online is to carry out meta analysis of the Difference-in-Differences estimations of individual events.<sup>14</sup> Meta-analysis is “the statistical synthesis of the data from separate but similar, i.e. comparable studies, leading to a quantitative summary of the pooled results” (Last et al., 2001). I treat the results from each community engagement event as a separate ‘study,’ in order to quantitatively estimate their combined effect. As the underlying analysis for individual events uses regression models, I follow the recommendations of existing research on meta-analysis of regression slopes, and use standardized (“beta”) coefficients when carrying out the meta analysis (Becker and Wu, 2007; Peterson and Brown, 2005).

### 5.4.1 Model

For each content category (i.e., travel to Syria or foreign fighters, sympathy with ISIS, life in ISIS territories, Syrian war, anti-West), I estimated 112 Difference-in-Differences regressions, corresponding to different community engagement events held by the Department of Homeland Security’s Office for Civil Rights and Civil Liberties in various locations in the United States. The assumption behind meta analysis models is that each individual estimate corresponds to a true latent coefficient, measured with some error:

$$y_i = \theta_i + \varepsilon_i \quad (2)$$

In the equation above,  $y_i$  represents the  $\beta_3$  coefficient from Equation (1) for study  $i$ ;  $\theta_i$  represents the (unknown) true coefficient; and  $\varepsilon_i$  is a sampling error, assumed to be distributed normally with mean 0 and variance  $v_i$  (this is without loss of generality by the Central Limit Theorem). If  $N_i$  is the sample size of the  $i$ th study, the sampling variance is calculated as follows:

$$v_i = \frac{(1 - y_i^2)^2}{N_i - 1} \quad (3)$$

---

<sup>14</sup>Specifically, I estimate the following Difference-in-Differences model for each event:  $y_{i,j} = \beta_1 Post_{i,j} + \beta_2 In\ event\ area_{i,j} + \beta_3 (Post_{i,j} \times In\ event\ area_{i,j}) + \varepsilon_j$  This is the same model as Equation in (1), but it does not include event fixed effects, as each event is estimated separately. In addition, in order to facilitate the meta analysis, all variables in these models (independent and dependent variables) are standardized.

In the meta analysis, I assume that the sample of Islamic State supporters in the United States represents the same underlying population, and thus I estimate fixed effects meta analytic models (Patall and Cooper, 2008; Viechtbauer et al., 2010). This assumption is reasonable, since the same population is being estimated over and over again for each event. The only change between events is the categorization of tweets as being ‘in the event area’ or not, or being posted before or after the event.<sup>15</sup> The fixed-effects model estimates the underlying true average effect using weighted least squares:

$$\bar{\theta}_w = \frac{\sum_{i=1}^k w_i \theta_i}{\sum_{i=1}^k w_i} \quad (4)$$

In the equation above,  $\bar{\theta}_w$  represents the weighted average of the true latent coefficients estimated for each community engagement event ( $\theta_i$ ), where the weight is inverse-proportional to the sampling error:  $w_i = \frac{1}{v_i}$ . In other words, the model gives more weight to studies with smaller sampling variance.

## 5.5 Meta analysis results

Table 3 reports the meta analysis results for 112 community engagement events held between 2014 and 2016. The coefficients represent the weighted average of the coefficients estimated for each event (i.e.,  $\bar{\theta}_w$  from equation (4)), measured in standard deviation units, where the time window surrounding each event is set to 30 days before and 30 days after the event. As in the pooled analysis, we find that community engagement events are systematically associated with reductions in pro-ISIS rhetoric on Twitter.

Since the coefficients reported in Table 3 are measured in standard deviation units, which are somewhat hard to interpret substantively, I report in Table 4 the percent change reflected in each Difference-in-Differences coefficient. The percent change is calculated as the coefficient  $\beta_3$  (*Post*  $\times$  *In event area*) divided by the coefficient  $\beta_2$  (*In event area*), and reflects the change in pro-ISIS content after community engagement events for individuals located in event areas, compared to

---

<sup>15</sup>This may raise the concern that the same tweet  $i$  might be coded differently in different estimations. For example, in the estimation of event A, tweet  $i$  might be coded 1 for the variable *In event area*, but in the regression of event B is it coded as 0. At worst, this misclassification will bias the results of event B towards zero (by making the outcomes for the treatment and control groups more similar to each other, on average), but it is important to preserve the same population for the meta analysis.

pro-ISIS content generated in these areas before the events. Table 4 also reports the results when the impact of the events are measured with different time windows: 7, 14, and 21 days.

As in the pooled analysis, the results in Table 4 show that the relationship between community engagement events and pro-ISIS rhetoric on Twitter is overall negative, and becomes more strongly negative as the window around the event expands. For example, when comparing the content of tweets in the 7 days before and after community engagement events, discourse on foreign fighters decreases by about 5%, but the difference is only marginally significant with a  $p$ -value of 0.09. The percent change for other topics is also negative, but not statistically significant. For the 14 and 21 day estimations, community engagement events are associated with a significant decrease of 5-8% in discourse on foreign fighters, 10-13% decrease in tweets expressing sympathy with ISIS, 4-5% decrease in tweets describing life in ISIS territories, and 8-10% decrease in discussion of the Syrian war. In these estimations, the anti-West topic has a positive percent change, but it is not statistically significant.

Finally, the results are strongest in terms of magnitude and significance when considering estimations using the 30 day window. The rightmost column in Table 4 shows that 30 days after community engagement events discussion on foreign fighters decreased by more than 11%, tweets sympathizing with ISIS decreased by almost 19%; discussion of life in ISIS territories decreased by 8%, and tweets describing the Syrian war decreased by 14%. Unlike the pooled analysis, the meta analysis results for the 30-day window show that anti-West content significantly decreased by almost 17% after community engagement events in areas where they were held.

Taken together, results from the pooled and meta analyses show, systematically across over a hundred community engagement events and tens of thousands of individuals, that engagement activities are followed by a decrease in pro-ISIS rhetoric. At least during the Obama Administration, these events were meant to share information, give feedback, and build trust between communities and the government. The finding that the impact of these events is strongest after 30 days might mean that it takes communities time to identify, counsel, and help people who show signs of radicalization in their areas. If one assumes that reduced pro-ISIS rhetoric on Twitter reflects a decline in radicalization, then these results suggest that engaging communities in countering extremism might be effective. To the best of my knowledge, this is one the first systematic examinations of the possible impacts of community engagement activities in the United States. Much of the CVE

strategy implemented in America so far has not been based on rigorous empirical research (Vidino and Hughes, 2015a). This study reveals important patterns relating to these initiatives that prior work has not been able to observe.

Table 3: Meta Analysis: Community engagement and pro-ISIS rhetoric on Twitter

	Estimate	Std. Err.	P-value
<b>1. All topics</b>			
Post	2.36***	0.01	0.00
In event area	0.78***	0.01	0.00
Post $\times$ In event area	-0.09***	0.01	0.00
Intercept	0.00	0.01	0.71
<b>2. Travel to Syria or foreign fighters</b>			
Post	1.56***	0.01	0.00
In event area	0.65***	0.01	0.00
Post $\times$ In event area	-0.07***	0.01	0.00
Intercept	0.00	0.01	0.81
<b>3. ISIS sympathy</b>			
Post	1.01***	0.01	0.00
In event area	0.22***	0.01	0.00
Post $\times$ In event area	-0.04***	0.01	0.00
Intercept	0.00	0.01	0.86
<b>4. Life in ISIS territories</b>			
Post	2.08***	0.01	0.00
In event area	0.97***	0.01	0.00
Post $\times$ In event area	-0.08***	0.01	0.00
Intercept	0.00	0.01	0.74
<b>5. Syrian war</b>			
Post	1.31***	0.01	0.00
In event area	0.35***	0.01	0.00
Post $\times$ In event area	-0.05***	0.01	0.00
Intercept	0.00	0.01	0.83
<b>6. Anti-West</b>			
Post	1.07***	0.01	0.00
In event area	0.10***	0.01	0.00
Post $\times$ In event area	-0.02**	0.01	0.04
Intercept	0.00	0.01	0.86

*Note:* The table shows coefficients estimated from a meta analysis of the relationship between 112 community engagement events and pro-ISIS rhetoric on social media, captured 30 day before and after each event. The results encompass data from 147,141,409 geo-located tweets generated from 2014 to 2016 by individuals who follow Islamic State accounts on Twitter. All variables are standardized. Coefficients are multiplied by 100 for presentation purposes.



Table 4: Meta Analysis: Community engagement and pro-ISIS rhetoric on Twitter (different time windows)


	7 Days		14 Days		21 Days		30 Days	
	% Change	P-value	% Change	P-value	% Change	P-value	% Change	P-value
Travel to Syria/foreign fighters	-4.88*	0.09	-5.26**	0.01	-7.76***	0.00	-11.33***	0.00
ISIS sympathy	-10.99	0.23	-10.72*	0.08	-12.97***	0.00	-18.51***	0.00
Life in ISIS territories	-2.08	0.26	-3.96***	0.00	-5.35***	0.00	-8.01***	0.00
Syrian war	-8.97	0.18	-9.73**	0.01	-8.48***	0.00	-14.26***	0.00
Anti-West	32.93*	0.07	2.83	0.83	0.77	0.93	-16.97**	0.04
All topics	-4.42*	0.07	-5.91***	0.00	-7.13***	0.00	-11.86***	0.00

*Note:* The % Change reflects the change in pro-ISIS content after community engagement events for individuals located in event areas, compared to pro-ISIS content in these areas before the events. In technical terms, it represents the meta analysis results for the Diff-in-Diff coefficients ( $Post \times In\ event\ area$ ) across 112 community engagement events, divided by the coefficient  $In\ event\ area$  (where  $Post = 0$ ).

## 5.6 Heterogeneity

In this section, I expand the prior analysis by looking at how the results vary with event-level characteristics, such as the timing of the event, the number of engagement activities held in a given location before the event, as well as the type of event. To examine how the estimated coefficients of community engagement events vary over time, I regressed the estimates of these events on the dates in which they took place. Figure 6 presents a meta-analytic scatterplot, in which the observed estimates for community engagement events (measured as a standardized index combining all content categories) are plotted against the date of the event. The resulting regression line has a negative and statistically significant slope, indicating that pro-ISIS rhetoric decreased more strongly after community engagement events taking place in the later part of 2015 and 2016.

Figure 6: The effect of community engagement events on pro-ISIS rhetoric, by date of event




*Note:* The figure presents the meta-analytic scatterplot of the observed effects estimated for individual community engagement events, where the dependent variable is a standardized index of all topics, calculated 30 days away from the event. The x-axis plots the date in which each event was taking place. The point sizes are proportional to the inverse of the standard errors, which means that events with larger samples have larger points. The predicted average effects are included (with corresponding 95% confidence intervals), calculated from the meta-analysis model.

Next, I examine how the Difference-in-Differences coefficients vary with the number of events

held in each area. If engaging communities in countering extremism is effective, then more events might lead to a greater reduction in pro-ISIS tweets by ISIS sympathizers in these localities. Figure 7 shows a meta analytic scatterplot of the estimated coefficients for each event, plotted against the number of community engagement activities taking place in each location prior to the event. Here, too, we find a negative and statistically significant relationship, which might suggest that more community engagement activities held in an area might be more effective for countering online support for extremism.

Figure 7: The effect of community engagement events on pro-ISIS rhetoric, by number of events in a city



*Note:* The figure presents the meta-analytic scatterplot of the observed effects estimated for individual community engagement events, where the dependent variable is a standardized index of all topics, calculated 30 days away from the event. The x-axis plots the number of community engagement events held in each city at the time of each event. The point sizes are proportional to the inverse of the standard errors, which means that events with larger samples have larger points. The predicted average effects are included (with corresponding 95% confidence intervals), calculated from the meta-analysis model.

Finally, in Table 5, I summarize the heterogeneous results we find when considering different event types. The table reports, for each type of activity (community roundtable, community resilience exercise, and community awareness briefing), the estimated pooled Difference-in-Differences

coefficients in the first row, and in the second row, the estimated difference from the pooled Difference-in-Differences result for each event type. The results show that community roundtables — events in which government representatives and members of various communities meet to strengthen relationships and engagement — are associated with a stronger negative coefficient of community engagement on pro-ISIS rhetoric when compared to all other event types. A similar finding is reported for community awareness briefings — events in which community members and law enforcement officials are presented with information on the process of radicalization and foreign fighter recruitment. The results for community resilience exercises also have a negative relationships but the results are not statistically significant.

Table 5: Community engagement and pro-ISIS rhetoric, by event type

	Estimate	Std. Err.	P-value
Estimated $\bar{\theta}$	-0.07***	0.02	0.00
Community roundtable	-0.03*	0.02	0.07
Estimated $\bar{\theta}$	-0.09***	0.01	0.00
Community awareness briefing (CAB)	-0.10*	0.06	0.08
Estimated $\bar{\theta}$	-0.09***	0.01	0.00
Community resilience exercise (CREX)	-0.02	0.14	0.89

*Note:* The table shows coefficients estimated from a meta analysis of the relationship between 112 community engagement events and pro-ISIS rhetoric on social media, captured 30 day before and after each event. The table reports, for each type of activity (community roundtable, community resilience exercise, and community awareness briefing), the estimated pooled Difference-in-Differences coefficients in the first row, and in the second row, the estimated difference from the pooled Difference-in-Differences result for each event type. All variables are standardized. Coefficients are multiplied by 100 for presentation purposes.

## 6 Do community engagement events suppress expression?

One primary objection to the results described above is that the findings do not reflect a decline in radicalization, but the suppression of expression. As described in section 3, community engagement activities might discourage individuals from expressing their opinion, views, and beliefs, as they facilitate a climate of fear by tasking community members with monitoring each other’s behavior (American Civil Liberties Union, 2016b). Thus, the observed decline in pro-ISIS rhetoric after community engagement activities might be driven by individuals’ abstention from expressing their opinions on Twitter.

Relatedly, the lower number of pro-ISIS tweets might be caused by the migration of ISIS sympathizers to private social media platforms. Feeling more strongly monitored by community members, individuals who are interested in Islamic State’s ideology might choose to abandon the public Twitter platform altogether. Finally, community engagement might increase surveillance and government intervention in the lives of ISIS sympathizers, which could result in a reduction in their public expression of pro-ISIS sentiment. This might happen, for example, when government agencies request Twitter to suspend accounts of individuals who are accused of supporting extremism. These behaviors will result in an overall reduction in pro-ISIS content, but this decline would not necessarily reflect de-radicalization.

To examine these possibilities, I carry out several additional estimations. First, I evaluate whether community engagement events suppress overall expression on Twitter. If they do, we should observe ISIS supporters located in event areas reduce the number of tweets — regardless of their content — after community engagement events. I counted the number of Twitter posts that ISIS sympathizers posted in each locality in the 7, 14, 21, and 30 days before and after each event. I aggregated the tweet-level dataset to a locality-time level dataset, where the locality is a geographic unit (city, town, etc.), and the time is a binary indicator coded 1 for tweet-sums appearing after an event and 0 otherwise.<sup>16</sup> I then estimate the following least squares regression, for each window length:

$$y_{ikt} = \beta_1 \text{Post}_{ikt} + \beta_2 \text{In event area}_{ikt} + \beta_3 (\text{Post}_{ikt} \times \text{In event area}_{ikt}) + \alpha_i + \varepsilon_k \quad (5)$$

In the equation above,  $y_{ikt}$  represents the number of tweets in location  $k$  at time  $t$  posted before and after event  $i$ ;  $\text{Post}_{ikt}$  is an indicator coded 0 where a tweet-sum is calculated for tweets appearing before event  $i$  and 1 afterwards;  $\text{In event area}_{ikt}$  is an indicator coded 1 for sums of tweets posted by individuals predicted to be located in the area of event  $i$ , and 0 otherwise; and  $\alpha_i$  is an event fixed effect. As before,  $\beta_3$  is the Difference-in-Differences coefficient of interest, reflecting how the change in the number of tweets after the event is different for locations in which community engagement

---

<sup>16</sup> Locality geographical data was taken from the United States’ census Topologically Integrated Geographic Encoding and Referencing (TIGER) database. Shape files for geographical units were taken from “Cartographic Boundary Shapefiles - Metropolitan and Micropolitan Statistical Areas and Related Statistical Areas.” ([https://www.census.gov/geo/maps-data/data/cbf/cbf\\_msa.html](https://www.census.gov/geo/maps-data/data/cbf/cbf_msa.html))

events took place, compared to the change the number of tweets in locations where events did not take place. Standard errors are clustered at the locality level.

Table 6 shows the results. It can be seen that community engagement events are not systematically associated with changes in the number of tweets posted by ISIS sympathizers in areas where they take place. The interaction term  $Post \times In\ event\ area$  is null in all window sizes, and its sign is not consistent. These results suggest that community engagement events are not affecting the number of tweets posted by ISIS supporters on Twitter.

Table 6: Community engagement events and the number of tweets

	7 days	14 days	21 days	30 days
Post	-14.49*** (3.36)	-9.56*** (2.44)	1.66 (1.40)	104.08*** (23.11)
In event area	254.26*** (97.51)	503.64** (198.11)	715.11*** (266.85)	900.49*** (347.68)
Post $\times$ In event area	-1.97 (53.54)	27.69 (120.02)	127.95 (124.84)	194.44 (154.67)
Constant	134.78*** (28.78)	229.19*** (49.37)	326.47*** (71.99)	936.82*** (225.13)
Event fixed effects	✓	✓	✓	✓
$R^2$	0.016	0.015	0.015	0.027
Observations	29,970	36,026	39,696	18,488

*Note:* Standard errors in parentheses, clustered at the locality level.

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Next, I evaluate whether community engagement events might increase surveillance of ISIS followers on Twitter. While this is difficult to measure, it is possible to use data on the suspension rate of ISIS supporters to proxy for increased surveillance. Using user-level data,<sup>17</sup> I created a week-by-week panel data for ISIS followers (i.e., accounts that follow core ISIS accounts), in which I measured (i) whether they were suspended from Twitter and (ii) whether they were flagged for suspension by Controlling Section (@ctrlsec). Replicating the meta analysis in equations (2) through (4), I estimate the relationship between community engagement events and these outcomes. Due to data availability limitations, I am only able to study thirty community engagement events taking place between January and June of 2016.

<sup>17</sup>These data come from continuous observations user-level data in the ISIS activists/followers database collected by the author, which are refreshed over time as users change their profile data.

Table 7 presents the meta analysis results. The coefficients represent the weighted average of the coefficients estimated for each event, measured in standard deviation units, where the time window surrounding each event is set to 30 days before and 30 days after the event. As before, each event was first estimated separately in a Difference-in-Differences regression, where the variable *In event area* differentiated between individuals who were located in the area of a community engagement event and those who did not, and the variable *Post* differentiated between user-level data observed before and after the event. The coefficient on the interaction term,  $Post \times In\ event\ area$ , represents the Difference-in-Differences coefficient combining all events.

The results show that community engagement events are not associated with a greater suspension rate of individuals located in the event area. In panel A in Table 7, the coefficient on  $Post \times In\ event\ area$  is positive, but not statistically significant. Interestingly, the result is different for the flagging outcome described in Panel B in Table 7. Here, we can see that community engagement events are associated with lower flagging rate of accounts of individuals located in event areas. As the flagging of accounts for suspension by Controlling Section (@ctrlsec) is strongly linked to the content that these accounts disseminate (see Table B14 in Appendix B), this suggests that the lower flagging rate is driven by the lower number of pro-ISIS tweets posted in event areas after community engagement events. This findings hold across all window lengths, as can be seen in Table 8.

Taken together, the results of this analysis do not support the argument that community engagement events suppress ISIS sympathizers' overall expression on Twitter. Nonetheless, while ISIS supporters located in event areas did not reduce the number of tweets that they posted, they seemed to have have changed their *content*: after community engagement activities, ISIS sympathizers expressed less pro-ISIS rhetoric. This result might be interpreted as a sign of de-radicalization, but it can also be driven by ISIS supporters limiting their expression of specific (e.g., pro-ISIS) topics on Twitter. Similarly, the finding that community engagement events are not systematically associated with greater suspension or flagging of ISIS accounts might be the result of users' greater awareness to monitoring. Thus, while the findings show a systematic decrease in pro-ISIS rhetoric after community engagement events, they do not allow determining the reason behind this decline.

Table 7: Meta analysis: Community engagement and account-level changes

	Estimate	Std. Err.	P-value
<b>A. Suspended from Twitter</b>			
Post	2.12***	0.04	0.00
In event area	0.25***	0.04	0.00
Post $\times$ In event area	0.03	0.04	0.50
Intercept	0.00	0.04	1.00
<b>B. Flagged as an ISIS activist</b>			
Post	12.68***	0.05	0.00
In event area	0.64***	0.05	0.00
Post $\times$ In event area	-0.16***	0.05	0.00
Intercept	-1.23***	0.05	0.00

*Note:* The table shows coefficients estimated from a meta analysis of the relationship between 30 community engagement events and accounts suspension and flagging, captured 30 day before and after each event. All variables are standardized. Coefficients are multiplied by 100 for presentation purposes.

Table 8: Meta analysis: Community engagement and account-level changes (different time windows)

	Change (%)	P-value	Change (%)	P-value	Change (%)	P-value
Suspended by Twitter	8.20	0.70	6.34	0.70	10.55	0.50
Flagged as an ISIS activist	-33.67***	0.00	-35.12***	0.00	-24.66***	0.00

*Note:* The % Change reflects the change in suspension and flagging rates after community engagement events for individuals located in event areas, compared to suspension and flagging rates in these areas before the events. In technical terms, it represents the meta analysis results for the Diff-in-Diff coefficients ( $Post \times In\ event\ area$ ) across 112 community engagement events, divided by the coefficient  $In\ event\ area$  (where  $Post = 0$ ).



## 7 Conclusion

Over the past few years, efforts to counter radicalization and violent extremism have increased across the world. In the United States during the Obama Administration, a large portion of CVE activities focused on building trust and engagement with local communities. These initiatives were premised on the idea that community members are best positioned to help radicalizing individuals because of their local-level, context-specific knowledge and expertise. While many initiatives to engage communities have taken place in recent years, there has been little systematic empirical research on how they might affect online extremist behaviors by individuals at risk of radicalization. In this study, I sought to shed light on the impact of community engagement activities held by the Department of Homeland Security’s Office of Civil Rights and Civil Liberties (CRCL), by combining granular data on community engagement events with information on the online behavior of Islamic State sympathizers in the United States.

Results from over 100 community engagement activities show that these events were systematically and significantly associated with a reduction in pro-ISIS rhetoric on Twitter among individuals located in event areas. Specifically, the data show that CRCL events were followed by a decrease in discourse on foreign fighters or travel to Syria, reduction in tweets expressing sympathy with ISIS, and a decrease in the number of tweets discussing the Syrian civil war and life in ISIS-controlled territories. The patterns in this study were robust to a large number of community engagement activities and tens of thousands of individuals. However, the results were inconclusive with respect to whether the reduction in pro-ISIS rhetoric was caused by de-radicalization or by the suppression of political expression in these areas. Further research is needed to shed light on this important question.

Overall, this study makes several contributions to existing research. First, by providing granular, geo-located high-frequency data on the online behavior of Islamic State sympathizers in the United States, the study measures an over-time “decrease” in pro-ISIS rhetoric, which could serve as a proxy for radicalization. Online measures of pro-ISIS rhetoric can plausibly reflect underlying support for extremism: a large majority of individuals who radicalized in support for ISIS in the United States have publicly expressed their favorable views towards the organization on social media

platforms.<sup>18</sup> As most research on countering extremism has struggled with identifying measures of de-radicalization, this is an important step forward.

Second, this study contributes to current research on community engagement, which is based on sporadic empirical examples, by conducting a rigorous analysis of over 100 community engagement events. The use of multiple events allows generalizing the conclusions beyond specific examples, and enables a more nuanced analysis of the heterogeneity of the findings for different event characteristics. For example, the study found that community engagement events are more effective in areas that hold a larger number of activities, and that specific event types, such community roundtables, are associated with greater reductions in pro-ISIS rhetoric, compared to other types of engagement activities.

Third, the paper provides a model for future work seeking to study links between local events and online behavior on social media. By predicting the geographic locations of thousands of Islamic State supporters on Twitter, this study was able to incorporate an important geographic dimension to the analysis of social media data that is usually not systematically taken into account. The ability to analyze geo-located high-frequency Twitter data and match it to local activities provides an opportunity to study political behavior in new and exciting ways.

Finally, while this study has focused on community engagement, it has not examined other types of efforts to counter violent extremism. Future work might examine the effectiveness of other measures, such as countering the propaganda disseminated by groups like ISIS or integrating marginalized individuals into mainstream society. Understanding how such interventions might help stop the troubling trend of violent extremism across the world is of crucial importance.

---

<sup>18</sup>See Appendix A Table A3 for more information.

## References

- Aldrich, Daniel P. 2014. “First Steps Towards Hearts and Minds? USAID’s Countering Violent Extremism Policies in Africa.” *Terrorism and Political Violence* 26(3):523–546.
- American Civil Liberties Union. 2016a. “The Problem with “Countering Violent Extremism” Programs.”
- URL:** <https://www.aclu.org/other/problem-countering-violent-extremism-programs?redirect=problem-countering-violent-extremism-programs>
- American Civil Liberties Union. 2016b. “What Is Wrong With the Government’s “Countering Violent Extremism” Programs.”
- URL:** [https://www.aclu.org/sites/default/files/field\\_document/cve\\_briefing\\_paper\\_feb\\_2016.pdf](https://www.aclu.org/sites/default/files/field_document/cve_briefing_paper_feb_2016.pdf)
- Backstrom, Lars, Eric Sun and Cameron Marlow. 2010. Find me if you can: improving geographical prediction with social and spatial proximity. In *Proceedings of the 19th international conference on World wide web*. ACM pp. 61–70.
- Becker, Betsy Jane and Meng-Jia Wu. 2007. “The synthesis of regression slopes in meta-analysis.” *Statistical Science* pp. 414–429.
- Bjelopera, Jerome P. 2012. Countering violent extremism in the United States. Congressional Research Service, Library of Congress.
- Borum, Randy. 2011. “Radicalization into violent extremism I: A review of social science theories.” *Journal of Strategic Security* 4(4):7.
- Briggs, Rachel. 2010. “Community engagement for counterterrorism: lessons from the United Kingdom.” *International Affairs* pp. 971–981.
- Challgren, Jonathan, Ted Kenyon, Lauren Kervick, Sally Scudder, Micah Walters, Kate Whitehead, Professors Jeffrey Connor and Carol Rollie Flynn. 2016. “Countering Violent Extremism: Applying the Public Health Model.”
- Cordner, Gary. 2014. “Community policing.” *The Oxford handbook of police and policing* pp. 148–171.

- Council on American-Islamic Relations. 2015. “Brief on countering violent extremism with recommendations for U.S. Government.”.
- Council on American-Islamic Relations. 2016. “Brief on Countering Violent Extremism (CVE).”.  
**URL:** <https://www.cair.com/government-affairs/13063-brief-on-countering-violent-extremism-cve.html>
- Crossett, Chuck and J Spitaletta. 2010. “Radicalization: Relevant psychological and sociological concepts.” *The John Hopkins University* .
- Executive Office of the President of the United States. 2011. *Strategic Implementation Plan for Empowering local partners to prevent violent extremism in the United States*. Office of the President.
- Fernandez, Alberto M. 2015. “Here to stay and growing: Combating ISIS propaganda networks.” *The Brookings Project on U.S. Relations with the Islamic World U.S.-Islamic World Forum Papers 2015* .
- Friedman, Jerome, Trevor Hastie and Robert Tibshirani. 2010. “Regularization Paths for Generalized Linear Models via Coordinate Descent.” *Journal of Statistical Software* 33(1):1–22.  
**URL:** <http://www.jstatsoft.org/v33/i01/>
- Greenberg, Karen J. 2016. “Case by Case: ISIS Prosecutions in the United States.” *Center on National Security, Fordham University School of Law* .
- Horgan, John. 2008. “From profiles to pathways and roots to routes: Perspectives from psychology on radicalization into terrorism.” *The ANNALS of the American Academy of Political and Social Science* 618(1):80–94.
- Jurgens, David. 2013. “That’s What Friends Are For: Inferring Location in Online Social Media Platforms Based on Social Relationships.” *ICWSM* 13:273–282.
- Jurgens, David, Tyler Finethy, James McCorriston, Yi Tian Xu and Derek Ruths. 2015. Geolocation Prediction in Twitter Using Social Networks: A Critical Analysis and Review of Current Practice. In *Proceedings of the 9th International AAAI Conference on Weblogs and Social Media (ICWSM)*.

- King, Gary and Langche Zeng. 2001. "Logistic regression in rare events data." *Political analysis* 9(2):137–163.
- Last, John M, Robert A Spasoff, Susan S Harris and Michel C Thuriaux. 2001. *A dictionary of epidemiology*. International Epidemiological Association, Inc.
- McGee, Jeffrey, James Caverlee and Zhiyuan Cheng. 2013. Location prediction in social media based on tie strength. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*. ACM pp. 459–468.
- Mitts, Tamar. 2017. "The Radicalization and Extremist Rhetoric of Islamic State Activists Followers on Twitter."
- Neumann, Peter R. 2013. "The trouble with radicalization." *International Affairs* 89(4):873–893.
- Obama, Barack. 2011. *Empowering local partners to prevent violent extremism in the United States*. Office of the President.
- Office for Civil Rights and Civil Liberties. 2016. *Community Engagement Section*.  
**URL:** <https://www.dhs.gov/sites/default/files/publications/community-engagement-dhs-crcl-poster.pdf>
- Patall, Erika and Harris Cooper. 2008. "Conducting a Meta-Analysis." *The SAGE handbook of social research methods*. Los Angeles: SAGE .
- Patel, FAize and Michael German. 2015. "Countering Violent Extremism: Myths and Fact."  
**URL:** [https://www.brennancenter.org/sites/default/files/analysis/102915%20Final%20CVE%20Fact%20Sheet.p](https://www.brennancenter.org/sites/default/files/analysis/102915%20Final%20CVE%20Fact%20Sheet.pdf)
- Peterson, Robert A and Steven P Brown. 2005. "On the use of beta coefficients in meta-analysis." *Journal of Applied Psychology* 90(1):175.
- Ramos, Juan. 2003. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*.
- Romaniuk, Peter. 2015. *Does CVE Work? Lessons Learned from the Global Effort to Counter Violent Extremism*. Global Center on Cooperative Security.

- Schlanger, Margo. 2011. *Office for Civil Rights and Civil Liberties (CRCL) Newsletter*. Vol. 2 U.S. Department of Homeland Security.
- URL:** [https://www.dhs.gov/sites/default/files/publications/crcl-newsletter-september-2011\\_0.pdf](https://www.dhs.gov/sites/default/files/publications/crcl-newsletter-september-2011_0.pdf)
- Schmitt, Eric and Somini Sengupta. 2015. “Thousands Enter Syria to Join ISIS Despite Global Efforts.” *New York Times*.
- URL:** <http://www.nytimes.com/2015/09/27/world/middleeast/thousands-enter-syria-to-join-isis-despite-global-efforts.html>
- Sedgwick, Mark. 2010. “The concept of radicalization as a source of confusion.” *Terrorism and Political Violence* 22(4):479–494.
- U.S. Department of Homeland Security. 2016. *Office for Civil Rights and Civil Liberties*.
- URL:** <https://www.dhs.gov/office-civil-rights-and-civil-liberties>
- Vanden Brook, Tom. 2015. “ISIL activity drives up Pentagon threat level.”
- URL:** <http://www.usatoday.com/story/news/nation/2015/05/08/pentagon-security-isis/26976725/>
- Vidino, Lorenzo and Seamus Hughes. 2015a. “Countering Violent Extremism in America.”
- Vidino, Lorenzo and Seamus Hughes. 2015b. “ISIS in America: From Retweets to Raqqa (George Washington University: Program on Extremism, 2015) 7.”
- Viechtbauer, Wolfgang et al. 2010. “Conducting meta-analyses in R with the metafor package.” *J Stat Softw* 36(3):1–48.
- Wilner, Alex S and Claire-Jehanne Dubouloz. 2010. “Homegrown terrorism and transformative learning: an interdisciplinary approach to understanding radicalization.” *Global Change, Peace & Security* 22(1):33–51.

## Appendix A

Table A1 provides summary statistics for each of 112 community engagement events taking place in the United States from 2014 to 2016. Each row represents summary statistics for a different dataset, collected for each community engagement event. The table provides information on the timing and location of each event, as well as a summary of the variables in each dataset and the total number of observations. In the table, the minimum and the maximum values of all variables are 0 and 1, respectively. Thus, I report them once for each dataset.

Figure A1 presents a cumulative forest plot for each event analyzed in the meta analysis. A cumulative forest plot presents the pooled estimated coefficient as each event's estimate is added to the analysis. The figure shows a forest plot of fixed-effects meta-analysis results for the summary index of pro-ISIS rhetoric, calculated in standard deviation units. Each row represents an estimate for one event. The figure plots 95% confidence intervals for the meta-analysis model, derived from the studies' sampling variances.

Figure A2 plots influence diagnostics in the meta analysis. It allows detecting influential cases or outlying studies. Evaluating the results with and without influential cases allows testing the robustness of the meta analysis results. The figure shows that there are several influential events, but event 71 in particular has a strong influence on the results.<sup>19</sup> To examine whether the results are robust to the exclusion of this event, I re-estimated the meta analysis models without event 71. The results are reported in Table A2, which shows that the results still hold.

Table A3 provides details on the social media usage of over a hundred of individuals charged in the United States with providing material support for ISIS or plotting a violent attack on the organization's behalf. Data come from criminal complains filed against these individuals in United States courts, which describe in detail these individuals' pro-ISIS activities. I coded each case according to whether the individual used social media platforms such as Twitter or Facebook during their radicalization process. In addition, I documented whether the individual expressed publicly his or her support for the Islamic State and its ideology. Understanding whether radicalizing individual post *public* social media posts is important for this paper's data collection method, which assumes

---

<sup>19</sup>The event was a community roundtable taking place in Phoenix, Arizona on March 30, 2016.

that it is possible to observe (at least part of) one's radicalization process by scraping information on his or her online behavior. The data show that the majority of these individuals used social media when radicalizing (about 63%). Among those who used social media, the vast majority (about 86%) posted publicly their support for ISIS.



Table A1: Summary statistics for each event dataset

Date	City	State	Post		In event area		Travel/FF		ISIS sympathy		ISIS life		Syrian war		Anti-West	
			mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	n
2014-01-09	Atlanta	GA	0.80	0.40	0.0069	0.08	0.21	0.26	0.13	0.28	0.19	0.33	0.10	0.24	0.12	0.26
2014-01-16	Los Angeles	CA	0.68	0.47	0.0083	0.09	0.21	0.26	0.13	0.29	0.19	0.34	0.10	0.24	0.12	0.26
2014-02-11	Minneapolis	MN	0.52	0.50	0.0002	0.01	0.21	0.26	0.14	0.29	0.19	0.34	0.10	0.24	0.12	0.26
2014-02-27	Houston	TX	0.57	0.50	0.0003	0.02	0.22	0.26	0.14	0.29	0.20	0.34	0.10	0.24	0.13	0.27
2014-03-05	Phoenix	AZ	0.57	0.49	0.0010	0.03	0.22	0.26	0.14	0.29	0.20	0.34	0.10	0.24	0.13	0.27
2014-03-13	Denver	CO	0.56	0.50	0.0057	0.08	0.22	0.26	0.14	0.30	0.20	0.34	0.10	0.24	0.12	0.27
2014-03-27	Chicago	IL	0.51	0.50	0.0047	0.07	0.22	0.26	0.14	0.30	0.20	0.34	0.10	0.24	0.12	0.27
2014-04-10	Denver	CO	0.47	0.50	0.0063	0.08	0.22	0.26	0.14	0.30	0.20	0.34	0.10	0.24	0.12	0.27
2014-04-29	New York	NY	0.49	0.50	0.0143	0.12	0.21	0.26	0.14	0.30	0.19	0.34	0.10	0.24	0.12	0.26
2014-04-30	Los Angeles	CA	0.49	0.50	0.0100	0.10	0.21	0.26	0.14	0.30	0.19	0.34	0.10	0.25	0.12	0.26
2014-04-30	New York	NY	0.49	0.50	0.0141	0.12	0.21	0.26	0.14	0.30	0.19	0.34	0.10	0.25	0.12	0.26
2014-05-08	Tampa	FL	0.49	0.50	0.0004	0.02	0.21	0.26	0.14	0.30	0.19	0.34	0.10	0.24	0.12	0.26
2014-05-21	Minneapolis	MN	0.52	0.50	0.0003	0.02	0.22	0.26	0.14	0.30	0.19	0.33	0.10	0.24	0.12	0.27
2014-06-12	Atlanta	GA	0.56	0.50	0.0042	0.06	0.22	0.27	0.14	0.30	0.19	0.33	0.10	0.24	0.13	0.27
2014-06-12	Houston	TX	0.56	0.50	0.0005	0.08	0.22	0.27	0.14	0.30	0.19	0.33	0.10	0.24	0.13	0.27
2014-06-14	Chicago	IL	0.56	0.50	0.0004	0.02	0.22	0.27	0.14	0.30	0.19	0.33	0.10	0.24	0.13	0.27
2014-06-14	Houston	TX	0.56	0.50	0.0004	0.02	0.22	0.27	0.14	0.30	0.19	0.33	0.10	0.24	0.13	0.27
2014-06-17	New York	NY	0.55	0.50	0.0090	0.09	0.22	0.27	0.14	0.30	0.19	0.33	0.10	0.24	0.13	0.27
2014-06-20	Los Angeles	CA	0.55	0.50	0.0112	0.11	0.22	0.27	0.15	0.30	0.19	0.33	0.10	0.24	0.13	0.27
2014-06-24	Chicago	IL	0.56	0.50	0.0058	0.08	0.22	0.27	0.14	0.30	0.19	0.33	0.10	0.24	0.13	0.27
2014-07-01	Washington	DC	0.57	0.50	0.0098	0.10	0.22	0.27	0.15	0.30	0.19	0.33	0.10	0.24	0.13	0.26
2014-07-08	Washington	DC	0.56	0.50	0.0097	0.10	0.22	0.27	0.14	0.30	0.19	0.33	0.10	0.24	0.13	0.27
2014-07-29	Washington	DC	0.50	0.50	0.0092	0.10	0.22	0.27	0.14	0.30	0.19	0.33	0.10	0.24	0.12	0.26
2014-07-30	Denver	CO	0.49	0.50	0.0022	0.05	0.22	0.27	0.14	0.29	0.18	0.33	0.10	0.24	0.12	0.26
2014-08-14	Seattle	WA	0.48	0.50	0.0012	0.03	0.22	0.27	0.14	0.29	0.18	0.33	0.10	0.24	0.12	0.26
2014-08-22	Orlando	FL	0.49	0.50	0.0002	0.01	0.22	0.26	0.14	0.29	0.18	0.33	0.10	0.24	0.12	0.26
2014-08-27	Los Angeles	CA	0.49	0.50	0.0036	0.06	0.22	0.26	0.14	0.29	0.18	0.33	0.10	0.24	0.12	0.26
2014-08-28	New York	NY	0.49	0.50	0.0146	0.12	0.22	0.27	0.14	0.29	0.18	0.33	0.10	0.24	0.12	0.26
2014-09-15	Chicago	IL	0.52	0.50	0.0039	0.06	0.22	0.26	0.13	0.29	0.18	0.33	0.10	0.24	0.12	0.26
2014-09-22	Columbus	OH	0.52	0.50	0.0011	0.03	0.22	0.26	0.13	0.29	0.18	0.33	0.10	0.24	0.12	0.26
2014-10-21	Chicago	IL	0.50	0.50	0.0057	0.08	0.22	0.26	0.14	0.29	0.19	0.33	0.09	0.23	0.12	0.26
2014-10-30	Boston	MA	0.50	0.50	0.0041	0.06	0.22	0.26	0.14	0.29	0.19	0.33	0.09	0.24	0.12	0.26
2014-11-04	Houston	TX	0.51	0.50	0.0008	0.03	0.22	0.26	0.14	0.29	0.19	0.33	0.10	0.24	0.12	0.26
2014-11-07	Minneapolis	MN	0.51	0.50	0.0002	0.01	0.22	0.26	0.14	0.29	0.19	0.33	0.10	0.24	0.12	0.26
2014-11-13	Los Angeles	CA	0.51	0.50	0.0052	0.07	0.22	0.26	0.14	0.30	0.19	0.33	0.10	0.24	0.12	0.26
2014-11-13	Detroit	MI	0.51	0.50	0.0003	0.02	0.22	0.26	0.14	0.30	0.19	0.33	0.10	0.24	0.12	0.26
2014-11-20	Atlanta	GA	0.52	0.50	0.0039	0.06	0.22	0.26	0.14	0.30	0.19	0.33	0.10	0.24	0.12	0.26
2014-12-04	Atlanta	GA	0.52	0.50	0.0038	0.06	0.22	0.26	0.14	0.29	0.19	0.33	0.10	0.24	0.12	0.26
2014-12-05	Boston	MA	0.52	0.50	0.0036	0.06	0.22	0.26	0.14	0.29	0.19	0.33	0.10	0.24	0.12	0.26
2014-12-15	Houston	TX	0.58	0.49	0.0005	0.02	0.22	0.26	0.14	0.30	0.18	0.33	0.09	0.23	0.12	0.26
2014-12-18	Tampa	FL	0.59	0.49	0.0018	0.04	0.21	0.26	0.14	0.29	0.18	0.33	0.09	0.23	0.12	0.26

Summary statistics for each event dataset (cont.)

Date	City	State	Post			In event area			Travel/FF			ISIS sympathy			ISIS life			Syrian war			Anti-West		
			mean	sd		mean	sd		mean	sd		mean	sd		mean	sd		mean	sd		mean	sd	
2015-01-21	Seattle	WA	0.56	0.50		0.0027	0.05		0.20	0.26		0.13	0.29		0.17	0.32		0.08	0.22		0.11	0.25	
2015-01-22	Boston	MA	0.56	0.50		0.0015	0.04		0.20	0.26		0.13	0.29		0.17	0.32		0.08	0.22		0.11	0.25	
2015-01-28	Chicago	IL	0.53	0.50		0.0034	0.06		0.20	0.26		0.13	0.29		0.16	0.32		0.08	0.21		0.11	0.24	
2015-01-28	Detroit	MI	0.53	0.50		0.0004	0.02		0.20	0.26		0.13	0.29		0.16	0.32		0.08	0.21		0.11	0.24	
2015-02-12	Tampa	FL	0.52	0.50		0.0005	0.02		0.20	0.25		0.13	0.28		0.16	0.31		0.08	0.21		0.11	0.24	
2015-02-19	Denver	CO	0.52	0.50		0.0030	0.05		0.20	0.25		0.12	0.28		0.16	0.31		0.08	0.21		0.11	0.24	
2015-02-24	Columbus	OH	0.52	0.50		0.0021	0.05		0.20	0.25		0.12	0.28		0.16	0.31		0.08	0.21		0.11	0.24	
2015-02-25	Phoenix	AZ	0.52	0.50		0.0015	0.04		0.20	0.25		0.12	0.28		0.16	0.31		0.08	0.21		0.11	0.24	
2015-03-16	New York	NY	0.52	0.50		0.0128	0.11		0.20	0.25		0.12	0.28		0.16	0.31		0.08	0.21		0.10	0.24	
2015-03-22	Charleston	WV	0.53	0.50		0.0000	0.00		0.20	0.25		0.12	0.28		0.16	0.31		0.08	0.21		0.10	0.24	
2015-03-24	Minneapolis	MN	0.53	0.50		0.0002	0.01		0.20	0.25		0.12	0.28		0.16	0.31		0.08	0.21		0.10	0.24	
2015-03-31	Atlanta	GA	0.54	0.50		0.0022	0.05		0.20	0.25		0.12	0.28		0.16	0.31		0.08	0.21		0.10	0.24	
2015-04-09	Chicago	IL	0.54	0.50		0.0045	0.07		0.20	0.25		0.12	0.28		0.16	0.31		0.08	0.21		0.10	0.24	
2015-04-15	Houston	TX	0.55	0.50		0.0009	0.03		0.20	0.25		0.12	0.28		0.15	0.31		0.08	0.20		0.10	0.24	
2015-04-18	Minneapolis	MN	0.55	0.50		0.0001	0.01		0.20	0.25		0.12	0.28		0.15	0.31		0.08	0.20		0.10	0.24	
2015-04-27	Houston	TX	0.55	0.50		0.0009	0.03		0.19	0.25		0.12	0.28		0.15	0.30		0.08	0.20		0.10	0.23	
2015-04-28	Los Angeles	CA	0.55	0.50		0.0060	0.08		0.19	0.25		0.12	0.28		0.15	0.30		0.08	0.20		0.10	0.23	
2015-04-30	Seattle	WA	0.55	0.50		0.0031	0.06		0.19	0.25		0.12	0.28		0.15	0.30		0.08	0.20		0.10	0.23	
2015-06-06	Phoenix	AZ	0.52	0.50		0.0033	0.06		0.19	0.25		0.12	0.28		0.15	0.30		0.08	0.20		0.10	0.24	
2015-06-08	Houston	TX	0.52	0.50		0.0007	0.03		0.19	0.25		0.12	0.28		0.15	0.30		0.08	0.20		0.10	0.23	
2015-06-23	Chicago	IL	0.55	0.50		0.0029	0.05		0.20	0.25		0.12	0.28		0.15	0.31		0.08	0.20		0.10	0.24	
2015-06-25	Atlanta	GA	0.54	0.50		0.0010	0.03		0.20	0.25		0.12	0.28		0.15	0.31		0.08	0.20		0.10	0.24	
2015-07-28	Lexington	IL	0.51	0.50		0.0032	0.06		0.19	0.25		0.12	0.28		0.15	0.30		0.08	0.20		0.10	0.23	
2015-07-28	Chicago	KY	0.51	0.50		0.0000	0.00		0.19	0.25		0.12	0.27		0.15	0.30		0.08	0.20		0.10	0.23	
2015-08-17	Denver	CO	0.54	0.50		0.0019	0.04		0.19	0.25		0.12	0.27		0.15	0.30		0.08	0.20		0.10	0.23	
2015-08-24	Philadelphia	PA	0.54	0.50		0.0068	0.08		0.19	0.25		0.12	0.27		0.15	0.30		0.08	0.20		0.10	0.23	
2015-08-25	Denver	CO	0.54	0.50		0.0018	0.04		0.19	0.25		0.12	0.27		0.15	0.30		0.08	0.20		0.10	0.23	
2015-08-26	Atlanta	GA	0.54	0.50		0.0010	0.03		0.19	0.25		0.12	0.27		0.15	0.30		0.08	0.20		0.10	0.23	
2015-08-27	Boston	MA	0.54	0.50		0.0025	0.05		0.19	0.25		0.12	0.27		0.15	0.30		0.07	0.20		0.10	0.23	
2015-08-27	Los Angeles	CA	0.54	0.50		0.0052	0.07		0.19	0.25		0.12	0.27		0.15	0.30		0.07	0.20		0.10	0.23	
2015-09-13	Minneapolis	MN	0.53	0.50		0.0002	0.01		0.19	0.25		0.11	0.27		0.14	0.30		0.07	0.20		0.10	0.23	
2015-09-17	Tampa	FL	0.53	0.50		0.0003	0.02		0.19	0.25		0.11	0.27		0.14	0.30		0.07	0.20		0.10	0.23	
2015-09-24	Memphis	TN	0.54	0.50		0.0002	0.01		0.19	0.25		0.11	0.27		0.14	0.30		0.07	0.19		0.10	0.23	
2015-10-08	Boston	MA	0.55	0.50		0.0021	0.05		0.19	0.25		0.11	0.27		0.14	0.30		0.07	0.19		0.10	0.23	
2015-10-21	Phoenix	AZ	0.57	0.49		0.0011	0.03		0.19	0.25		0.12	0.27		0.14	0.30		0.07	0.19		0.10	0.23	
2015-10-29	Chicago	IL	0.58	0.49		0.0026	0.05		0.19	0.25		0.12	0.27		0.14	0.30		0.07	0.20		0.10	0.23	
2015-10-29	Houston	TX	0.58	0.49		0.0002	0.02		0.19	0.25		0.12	0.27		0.14	0.30		0.07	0.20		0.10	0.23	
2015-11-07	Boston	MA	0.59	0.49		0.0018	0.04		0.19	0.25		0.12	0.27		0.14	0.30		0.07	0.19		0.10	0.23	
2015-11-12	Minneapolis	MN	0.59	0.49		0.0005	0.02		0.19	0.25		0.12	0.27		0.14	0.30		0.07	0.19		0.10	0.23	
2015-11-18	Columbus	OH	0.58	0.49		0.0007	0.03		0.19	0.25		0.12	0.27		0.14	0.30		0.07	0.19		0.10	0.23	
2015-11-23	Denver	CO	0.61	0.49		0.0011	0.03		0.19	0.25		0.11	0.27		0.13	0.29		0.07	0.19		0.10	0.22	

Summary statistics for each event dataset (cont.)

Date	City	State	Post		In event area		Travel/FF		ISIS sympathy		ISIS life		Syrian war		Anti-West		n
			mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	
2015-12-10	Nashville	TN	0.68	0.47	0.0003	0.02	0.18	0.24	0.11	0.27	0.13	0.29	0.07	0.18	0.09	0.22	1800773
2015-12-16	Atlanta	GA	0.68	0.47	0.0008	0.03	0.18	0.24	0.11	0.27	0.13	0.28	0.07	0.18	0.09	0.22	1924451
2015-12-16	Miami	FL	0.68	0.47	0.0007	0.03	0.18	0.24	0.11	0.27	0.13	0.28	0.07	0.18	0.09	0.22	1924451
2015-12-17	Denver	CO	0.68	0.47	0.0010	0.03	0.18	0.24	0.11	0.27	0.12	0.28	0.07	0.17	0.09	0.22	1946166
2015-12-17	Los Angeles	CA	0.68	0.47	0.0026	0.05	0.18	0.24	0.11	0.27	0.12	0.28	0.07	0.17	0.09	0.22	1946166
2015-12-17	Tampa	FL	0.68	0.47	0.0005	0.02	0.18	0.24	0.11	0.27	0.12	0.28	0.07	0.17	0.09	0.22	1946166
2016-01-13	Detroit	MI	0.50	0.50	0.0003	0.02	0.18	0.24	0.11	0.26	0.12	0.28	0.06	0.17	0.09	0.21	2591121
2016-01-16	Columbus	OH	0.51	0.50	0.0004	0.02	0.18	0.24	0.11	0.26	0.12	0.28	0.06	0.17	0.09	0.21	2668235
2016-01-25	Denver	CO	0.52	0.50	0.0010	0.03	0.18	0.24	0.11	0.27	0.12	0.28	0.06	0.17	0.09	0.21	2856046
2016-01-28	Boston	MA	0.54	0.50	0.0012	0.03	0.18	0.24	0.11	0.27	0.12	0.28	0.06	0.17	0.09	0.21	2915099
2016-02-11	Chicago	IL	0.70	0.46	0.0020	0.04	0.17	0.24	0.11	0.26	0.12	0.28	0.06	0.17	0.09	0.21	4369777
2016-02-18	Charleston	SC	0.70	0.46	0.0000	0.01	0.17	0.24	0.11	0.26	0.12	0.28	0.06	0.17	0.09	0.21	4692288
2016-02-18	Minneapolis	MN	0.70	0.46	0.0002	0.01	0.17	0.24	0.11	0.26	0.12	0.28	0.06	0.17	0.09	0.21	4692288
2016-02-23	New York	NY	0.68	0.47	0.0076	0.09	0.17	0.24	0.11	0.26	0.12	0.28	0.06	0.17	0.09	0.21	4608500
2016-02-25	Atlanta	GA	0.67	0.47	0.0005	0.02	0.17	0.24	0.11	0.26	0.12	0.28	0.06	0.17	0.09	0.21	4577110
2016-02-25	Boston	MA	0.67	0.47	0.0016	0.04	0.17	0.24	0.11	0.26	0.12	0.28	0.06	0.17	0.09	0.21	4577110
2016-03-10	Tampa	FL	0.44	0.50	0.0003	0.02	0.18	0.24	0.11	0.27	0.13	0.28	0.07	0.18	0.09	0.21	4397808
2016-03-16	Seattle	WA	0.21	0.41	0.0011	0.03	0.18	0.24	0.11	0.27	0.13	0.29	0.07	0.18	0.09	0.22	4301821
2016-03-18	New York	NY	0.22	0.41	0.0091	0.10	0.18	0.24	0.11	0.27	0.13	0.29	0.07	0.18	0.09	0.22	4265043
2016-03-23	Detroit	MI	0.24	0.43	0.0008	0.03	0.18	0.24	0.11	0.27	0.13	0.29	0.07	0.18	0.09	0.22	4189790
2016-03-28	Houston	TX	0.28	0.45	0.0001	0.01	0.18	0.24	0.12	0.27	0.14	0.29	0.07	0.19	0.09	0.22	4102509
2016-03-29	Dallas	TX	0.29	0.45	0.0021	0.05	0.19	0.24	0.12	0.27	0.14	0.30	0.07	0.19	0.09	0.22	4091249
2016-03-30	Columbus	OH	0.30	0.46	0.0006	0.03	0.19	0.24	0.12	0.27	0.14	0.30	0.07	0.19	0.10	0.22	4083441
2016-03-30	Phoenix	AZ	0.30	0.46	0.0010	0.03	0.19	0.24	0.12	0.27	0.14	0.30	0.07	0.19	0.10	0.22	4083441
2016-04-07	Atlanta	GA	0.50	0.50	0.0010	0.03	0.19	0.24	0.12	0.28	0.15	0.31	0.08	0.20	0.10	0.23	4083441
2016-04-09	Los Angeles	CA	0.58	0.49	0.0050	0.07	0.19	0.25	0.13	0.28	0.15	0.31	0.08	0.20	0.10	0.23	4483780
2016-04-11	Denver	CO	0.71	0.46	0.0021	0.05	0.20	0.25	0.13	0.28	0.16	0.32	0.08	0.21	0.10	0.24	4579647
2016-04-11	Orlando	FL	0.71	0.46	0.0007	0.03	0.20	0.25	0.13	0.28	0.16	0.32	0.08	0.21	0.10	0.24	4364343
2016-04-11	Tampa	FL	0.71	0.46	0.0006	0.03	0.20	0.25	0.13	0.28	0.16	0.32	0.08	0.21	0.10	0.24	4364343
2016-05-04	Los Angeles	CA	0.72	0.45	0.0048	0.07	0.21	0.25	0.13	0.29	0.17	0.32	0.08	0.22	0.11	0.24	5605973
2016-05-17	Portland	OR	0.15	0.35	0.0006	0.02	0.20	0.25	0.13	0.29	0.17	0.32	0.08	0.22	0.11	0.24	5153951
2016-05-18	Seattle	WA	0.10	0.29	0.0022	0.05	0.20	0.25	0.13	0.29	0.17	0.32	0.08	0.22	0.11	0.24	5117649

Figure A1: Cumulative forest plot




Figure A2: Influence of individual events on meta analysis results

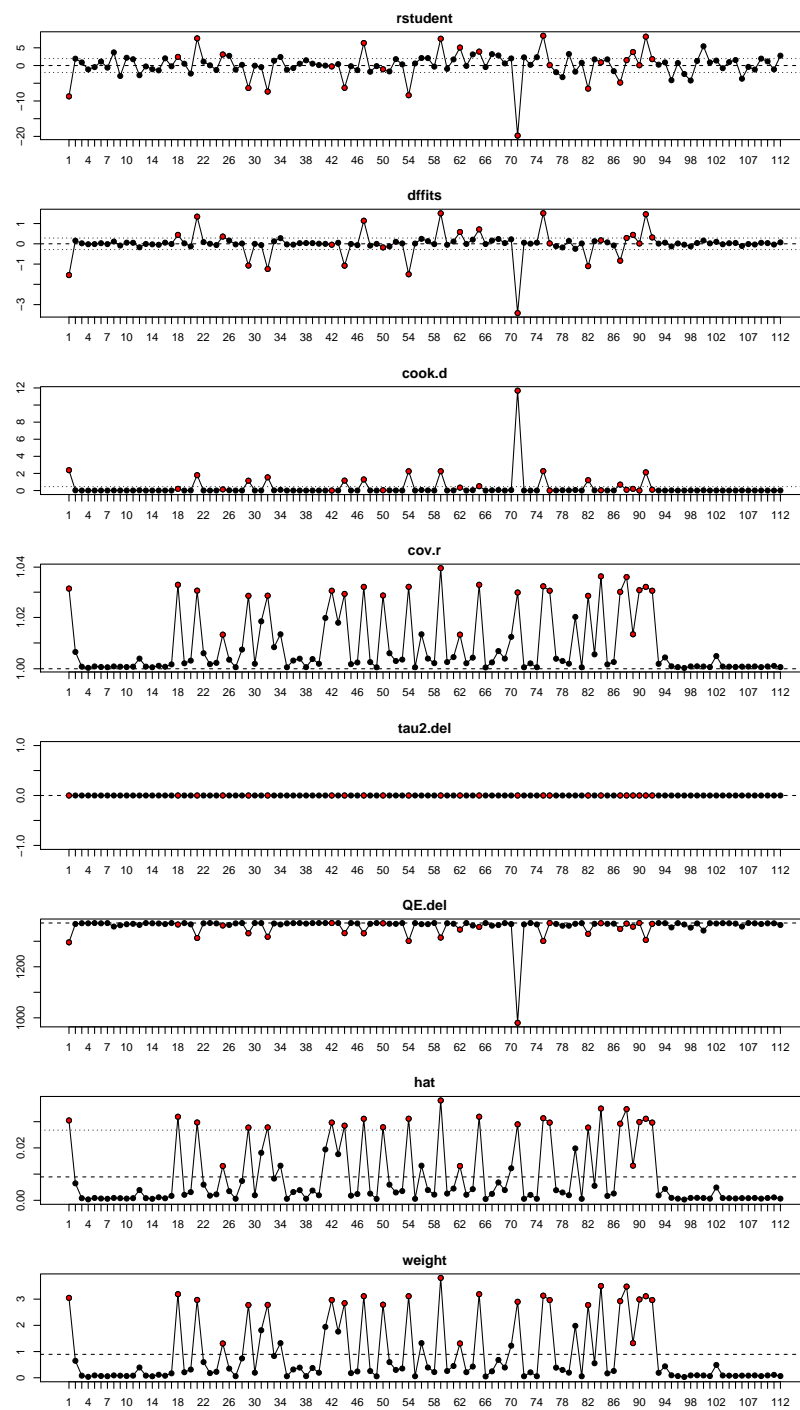


Table A2: Meta Analysis: Excluding an influential event

	Estimate	Std. Err.	P-value
<b>A. Travel to Syria or foreign fighters</b>			
Post	1.38***	0.01	0.00
In event area	0.60***	0.01	0.00
Post $\times$ In event area	-0.06***	0.01	0.00
Intercept	0.00	0.01	0.88
<b>B. ISIS sympathy</b>			
Post	0.90***	0.01	0.00
In event area	0.20***	0.01	0.00
Post $\times$ In event area	-0.03***	0.01	0.00
Intercept	0.00	0.01	0.91
<b>C. Life in ISIS territories</b>			
Post	1.84***	0.01	0.00
In event area	0.90***	0.01	0.00
Post $\times$ In event area	-0.05***	0.01	0.00
Intercept	0.00	0.01	0.82
<b>D. Syrian war</b>			
Post	1.13***	0.01	0.00
In event area	0.32***	0.01	0.00
Post $\times$ In event area	-0.03***	0.01	0.00
Intercept	0.00	0.01	0.90
<b>E. Anti-West</b>			
Post	0.93***	0.01	0.00
In event area	0.08***	0.01	0.00
Post $\times$ In event area	-0.00	0.01	0.92
Intercept	0.00	0.01	0.91
<b>F. All topics</b>			
Post	2.10***	0.01	0.00
In event area	0.71***	0.01	0.00
Post $\times$ In event area	-0.06***	0.01	0.00
Intercept	0.00	0.01	0.81

*Note:* The table shows coefficients estimated from a meta analysis of the relationship between community engagement events and pro-ISIS rhetoric on social media, captured 30 day before and after each event. The results exclude event 71 (a community roundtable in Phoenix, Arizona in March 30, 2016), which was found to be influential in the meta analysis. All variables are standardized. Coefficients are multiplied by 100 for presentation purposes.

Table A3: Social media usage by ISIS supporters in the United States

	Name	Location	Used social media	Posted public posts
1	Samy El-Goarany	New York	1	1
2	Ahmed Mohammed El Gammal	Arizona	1	1
3	Abdul Malik Abdul Kareem	Phoenix, AZ	0	0
4	Elton Francis Simpson	Phoenix, AZ	1	1
5	Nader Ehuzayel	Santa Ana, California	1	1
6	Muhanad Badawi	Santa Ana, California	1	1
7	Womain in Palestine	Palestine	1	1
8	Nicholas Michael Teausant	Acampo, CA	1	1
9	Adam Dandach	Orange County, CA	0	0
10	Enrique Marquez Jr.	Riverside, CA	0	0
11	Aws Mohammed Younis al-Jayab	Sacramento, CA	1	0
12	Mahamad Saeed Koadimati	San Diego, CA	1	0
13	Shannon Maureen Conley	Denver, CO	1	0
14	James Gonzalo Medina	Hollywood, FL	0	0
15	Harlem Suarez	Key West, FL	1	1
16	Gregory Hubbard	West Palm Beach, FL	1	0
17	Dayne Antani Christian	Lake Park, FL	0	0
18	Darren Arness Jackson	West Palm Beach, FL	0	0
19	Miguel Moran Diaz	Miami-Dade, FL	1	1
20	Robert B. Jackson	Pensacola, FL	1	1
21	Leon Nathan Davis	Augusta, GA	0	0
22	Hasan R. Edmonds	Aurora, IL	1	1
23	Jonas M. Edmonds	Aurora, IL	0	0
24	Mhammed Hamzah Khan	Bolingbrook, IL	0	0
25	Ramiz Zijad Hodzic	Saint Louis, MO	1	1
26	Sedina Unkc Hodzic	Saint Louis, MO	1	1
27	Nihad Rosic	Utica, NY	1	1
28	Mehida Medy Salkicevic	Schiller Park, IL	1	1
29	Armin Harcevic	Saint Louis, MO	1	1
30	Jasminka Ramic	Rockford, IL	1	1
31	Abdullah Ramo Pazara	Saint Louis, MO	1	0
32	Akrami I. Musleh	Brownsburg, IN	1	1
33	Alexander E. Blair	Topeka, KS	0	0
34	John T. Booker	Topeka, KS	1	1
35	Alexander Ciccolo	Adams, MA	1	1
36	David Wright	Everett, MA	0	0
37	Mohamed Elshinaway	Edgewood, MD	1	1
38	Khalil Abu Rayyan	Dearborn Heights, MI	1	1
39	Sebastian Gregerson	Detroit, MI	0	0
40	Al-Hamzah Mohammad Jawad	East Lansing, MI	0	0
41	Abdirizak Mohamed Warsame	Eagan, MN	0	0
42	Abdul Raheem Habil Ali-Skelton	Glencoe, MN	0	0
43	Mohamed Abdihamid Farah	Minneapolis, MN	0	0
44	Adnan Abdihamid Farah	Minneapolis, MN	1	1
45	Abdurahman Yasin Daud	Minneapolis, MN	0	0
46	Zacharia Yusuf Abdurahman	Minneapolis, MN	0	0
47	Hanad Mustafe Musse	Minneapolis, MN	0	0
48	Guled Ali Omar	Minneapolis, MN	0	0
49	Hamza Ahmed	Minneapolis, MN	1	1
50	"H.A.M"	Burnsville, MN	1	1
51	Abdullahi Yusuf	Inver Grove Heights, MN	1	1
52	Abdi Nur	Minneapolis, MN	1	1
53	Yusra Ismail	St. Paul, MN	0	0
54	Safya Roe Yassin	Bolivar, MO	1	1
55	Jaelyn Delshaun Young	Starkville, MS	1	1

*Note:* The table provides details on the social media usage of individuals charged in the United States with providing material support for ISIS or plotting a violent attack on the organization's behalf. Data come from criminal complaints filed against these individuals in United States courts, which describe in detail these individuals' pro-ISIS activities.

### Social media usage by ISIS supporters in the United States

	Name	Location	Used social media	Posted public posts
56	Muhammad Oda Dakhalla	Starkville, MS	0	0
57	Justin Nojan Sullivan	Burke County, NC	1	0
58	Erick Jamal Hendricks	Charlotte, NC	1	1
59	"CW-1"	Northern District, Ohio	1	1
60	Avin Marsalis Brown	Raleigh, NC	0	0
61	Akba Johad Jordan	Raleigh, NC	0	0
62	Donald Ray Morgan	Rowan County, NC	1	1
63	Nader Saadeh	Rutherford, NJ	1	1
64	Alaa Saadeh	West New York, NJ	0	0
65	Samuel Rahamin Topaz	Fort Lee, NJ	1	1
66	"CC-1/CC-2"	Queens, NY	1	1
67	Tairod Nathan Webster Pugh	Neptune, NJ	0	0
68	Sajmir Alimehmeti	Bronx, NY	1	0
69	Abdursasul Hasanovich Juraboev	Brooklyn, NY	1	1
70	Akhror Saidakhmetov	Brooklyn, NY	1	1
71	Arbor Habibov	Brooklyn, NY	0	0
72	Dilkhayot Kasimov	Brooklyn, NY	0	0
73	Almal Zakirov	Brooklyn, NY	0	0
74	Mohimanul Bhuiya	Brooklyn, NY	0	0
75	Noelle Velentzas	Queens, NY	0	0
76	Asia Siddiqui	Queens, NY	1	1
77	Arafat M. Nagi	Lackawanna, NY	1	1
78	Ali Saleh	Fort Wayne, IN	1	1
79	Munther Omar Saleh	Queens, NY	1	1
80	Emanuel L. Luchtman	Rochester, NY	1	0
81	Mufid A. Elfgeeh	Rochester, NY	1	1
82	Farred Mumuni	Staten Island, NY	0	0
83	Terrence Joseph Mcneil	Akron, OH	1	1
84	Christopher Lee Cornell	Cincinnati, OH	1	1
85	Robert C. McCollum	Sheffield Lake, OH	1	1
86	Munir Abdulkader	West Chester, OH	1	1
87	Jalil Ibn Amer Aziz	Harrisburg, PA	1	1
88	Keonna Thomas	Philadelphia, PA	1	1
89	David Wright	Everett, MA	0	0
90	Nicholas Rovinski	Warwick, RI	1	1
91	Usama Rahim	Roslindale, MA	0	0
92	Michael Todd Wolfe	Houston, TX	0	0
93	Omar Faraj Saeed Al Hardan	Houston, TX	0	0
94	Asher Abid Khan	Spring, TX	1	0
95	"S.R.G"	Spring, TX	1	1
96	"CC-1"	Spring, TX	1	0
97	Bilal Abood	Mesquite, TX	1	1
98	Mohamad Jamal Khweis	Alexandria, VA	1	1
99	Haris Qamar	Burke, VA	1	1
100	Nicholas Young	Fairfax, VA	0	0
101	Amine El Khalifi	Fairfax, VA	1	1
102	Yusuf Abdirizak Wehelie	Fairfax, VA	0	0
103	Heather Elizabeth Coffman	Richmond, VA	1	1
104	Mohamed Bailor Jalloh	Sterling, VA	1	1
105	Ali Shukri Amin	Woodbridge, VA	1	1
106	Joseph Hassan Farrokh	Woodbridge, VA	0	0
107	Mhamoud Amin Mohamed Elhassan	Woodbridge, VA	0	0
108	Daniel Seth Franey	Montesano, WA	1	1
109	Joshua Van Haften	Madison, WI	1	1
	<b>Proportion using social media</b>		0.63	
	<b>Proportion posting public posts (among those using social media)</b>			0.86

*Note:* The table provides details on the social media usage of individuals charged in the United States with providing material support for ISIS or plotting a violent attack on the organization's behalf. Data come from criminal complaints filed against these individuals in United States courts, which describe in detail these individuals' pro-ISIS activities.



## Appendix B

### B1 Identifying ISIS activist and follower accounts on Twitter

In this project, I track lists published publicly by several anti-ISIS hacking groups to identify ISIS supporters’ accounts on Twitter. Using the Twitter APIs,<sup>20</sup> I designed an algorithm that continually monitors and records ISIS accounts identified by the hacktivist group @CtrlSec.<sup>21</sup> Immediately upon observing a new account in the @CtrlSec list, I download the complete “timeline” of tweets for the account, as well as its user profile, which includes various user-level fields, and list of the account’s friends and followers. The full list of user profile fields is given in Table B1. The database contains “snapshots” of each user’s profile at various points in time. In particular, prior to mid-May 2016, user profile snapshots were saved when the user was encountered on the @CtrlSec list or included as part of 5,000 randomly selected follower accounts for content sampling every 24 hours. Beginning in mid-May 2016, new snapshots are obtained for all non-suspended user accounts every 1-2 days, on average. The full list of data fields for each tweet is given in Table B2.

#### Downloading Twitter timelines

The dimensionality of the friends and followers is particularly challenging for historical timeline data collection. While I have identified approximately 15,000 activists thus far from the @CtrlSec postings, this has led to over 1.6 million followers and about 450,000 friends of these followers. Due to rate limits, it is impossible using the publicly available Twitter API to obtain full content timelines for 2 million accounts. Thus, I began by downloading the full historical tweet timelines of all @CtrlSec-identified “ISIS activist” accounts ( $N = 14,979$ ), as well as of all the friends of a subsample of the activists who were first observed in the database as a follower or friend, and subsequently ‘flipped’ and became flagged as activists ( $N = 193,973$ ). After completing an initial round of location prediction, I downloaded the complete historical tweet timelines of additional accounts of ISIS followers and friends predicted to be located in Europe and North America.

There are two additional sources of tweet timeline content in the dataset. The first is a so-

---

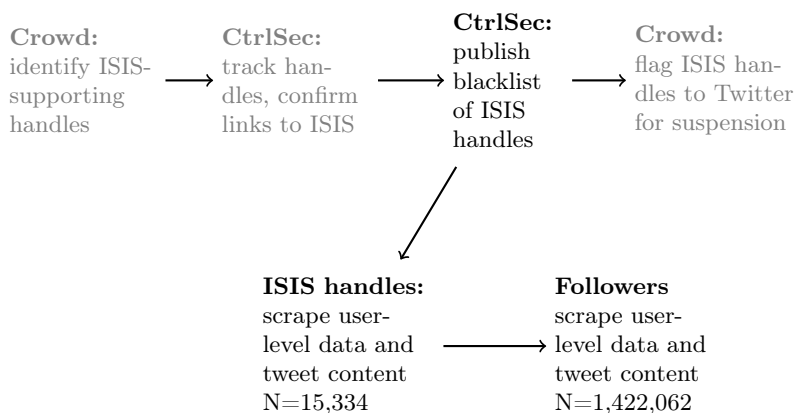
<sup>20</sup><https://dev.twitter.com/overview/documentation>

<sup>21</sup>Lists are available in these handles: @ctrlsec, @ctrlsec0, @ctrlsec1, @ctrlsec2, @ctrlsec9.

called “random sample with holes.” Since the Twitter Streaming API imposes rate limits on usage, I was only able to stream content for 5,000 users in a 24-hour period. The streaming began on December 19, 2015, and with the exception of occasional technical glitches, has been collecting data on the content posted by a random sample of 5,000 followers each day (data collection currently continues). Moreover, as noted previously, user profile information is downloaded at the same time. This ensures that user-level information (such as profile picture, number of friends, etc.), as well as account suspension status, are updated daily for this random sample.

The second source of tweet timeline data is a daily “total refresh” that began in May 2016. The Twitter API permits obtaining the a current profile snapshot for a user, which contains their most recently posted tweet, at a much faster rate limit than a full historical content download. Thus, I began to cycle through the entire database of nearly 2 million accounts on a daily basis, requesting latest profile and tweet, which leads to a complete refresh of user profiles and the latest tweet for each user in the system, as well as their suspension status, every 1-2 days on average. The total number of tweets scraped with this method was over 61 million as of August 2016.

Figure B1: Scraping ISIS accounts



*Note:* The number of Twitter users is accurate to 9/23/2016, 1:40PM ET

Table B1: List of data fields at the user level

Field Name	Description
user_id	The integer representation of the unique identifier for this User.
date_added	The datetime the user profile snapshot was added to the database.
name	The name of the user, as they've defined it. Not necessarily a person's name.
screen_name	The screen name, handle, or alias that this user identifies themselves with.
location	The user-defined location for this account's profile. Not necessarily a location nor parseable.
description	The user-defined UTF-8 string describing their account.
url	A URL provided by the user in association with their profile.
protected	When true, indicates that this user has chosen to protect their Tweets.
followers_count	The number of followers this account currently has.
friends_count	The number of users this account is following (AKA their "followings").
listed_count	The number of public lists that this user is a member of.
created_at	The UTC datetime that the user account was created on Twitter.
favourites_count	The number of tweets this user has favorited in the account's lifetime.
utc_offset	The offset from GMT/UTC in seconds.
time_zone	A string describing the Time Zone this user declares themselves within.
geo_enabled	When true, indicates that the user has enabled the possibility of geotagging their Tweets.
verified	When true, indicates that the user has a verified account.
statuses_count	The number of tweets (including retweets) issued by the user.
lang	The BCP 47 code for the user's self-declared user interface language. May or may not have anything to do with the content of their Tweets.
profile_background_image_url	A HTTP-based URL pointing to the background image the user has uploaded for their profile.
profile_image_url	A HTTP-based URL pointing to the user's avatar image.
profile_image_file	A local copy of the user's profile image.
profile_banner_url	The HTTPS-based URL pointing to the standard web representation of the user's uploaded profile banner.
profile_banner_file	A local copy of the user's profile banner.
followers	The list of the user's followers, as of the date of this "snapshot." (Only obtained for certain users such as ISIS activists.)
friends	The list of the user's followers, as of the date of this "snapshot." (Only obtained for certain users such as ISIS activists.)
suspended	A flag for whether the account was suspended.

Table B2: List of data fields at the tweet level

Field Name	Description
id	The integer representation of the unique identifier for this Tweet.
user_id	The integer representation of the unique identifier for the author of the Tweet.
date_added	The datetime that the Tweet was added to the database.
created_at	The datetime that the user account was created on Twitter.
text	The actual UTF-8 text of the status update.
source	Utility used to post the Tweet, as an HTML-formatted string. Tweets from the Twitter website have a source value of web.
truncated	Indicates whether the value of the text parameter was truncated, for example, as a result of a retweet exceeding the 140 character Tweet length. Truncated text will end in ellipsis, like this ...
in_reply_to_status_id	If the represented Tweet is a reply, this field will contain the integer representation of the original Tweet's ID.
in_reply_to_user_id	If the represented Tweet is a reply, this field will contain the integer representation of the original Tweet's author ID.
in_reply_to_screen_name	If the represented Tweet is a reply, this field will contain the screen name of the original Tweet's author.
retweet_count	Number of times this Tweet has been retweeted.
favorite_count	Indicates approximately how many times this Tweet has been "liked" by Twitter users.
lang	When present, indicates a BCP 47 language identifier corresponding to the machine-detected language of the Tweet text, or "und" if no language could be detected.
possibly_sensitive	This field is an indicator that the URL contained in the tweet may contain content or media identified as sensitive content.
coordinates	Represents the geographic location of this Tweet as reported by the user or client application.
withheld_in_countries	When present, indicates a list of uppercase two-letter country codes this content is withheld from.
quoted_status	This field only surfaces when the Tweet is a quote Tweet. This attribute contains the Tweet object of the original Tweet that was quoted.
retweeted_status	This attribute contains a representation of the original Tweet that was retweeted.

*Note:* Descriptions are copied verbatim from the Twitter REST API at <https://dev.twitter.com/overview/api>

Table B3: Number of tweets posted by all users in database, by year

Year	# tweets
2007	628
2008	3,389
2009	32,250
2010	84,913
2011	291,733
2012	826,552
2013	1,866,381
2014	3,980,438
2015	12,987,810
2016	47,107,004

*Note:* The number of tweets is accurate to 9/23/2016, 1:40PM ET.

## B2 Predicting geographic location of ISIS activists and followers

### Spatial Label Propagation algorithm

The spatial label propagation (SLP) algorithm used to predict the geographic locations of Twitter users in this paper implements the method developed by Jurgens (2013). The algorithm works as follows. First, define  $U$  to be a set of Twitter users in a social network, and for each user, let  $N$  be a mapping from the user to her friends (i.e., users to whom the user is directly connected), such that  $u \rightarrow [n_i, \dots, n_m]$ . Also, let  $L$  be a mapping of users to their known geographic locations:  $u \rightarrow (latitude, longitude)$ , and  $E$  the current mapping from users to locations.  $E$  is being updated with each iteration of the algorithm.

The algorithm works as follows. First, it initializes  $E$ , the current mapping from users to locations, with  $L$ , the ground truth data. Then, for each user who does not have location data and has friends with location data, the algorithm creates a vector,  $M$ , which stores a list of the friends' locations. Using this list of latitude and longitude coordinates, the algorithm predicts the user's location by calculating the geometric median of the locations in  $M$ . The new predicted locations from the first round are added to  $E$ , the new mapping from users to locations. The algorithm repeats itself by predicting additional users' locations in the second round, using the ground truth and predicted location data from the previous round. The algorithm stops when the stopping criterion is met (in this paper, three rounds of prediction).

Figure B2 illustrates the way in which spatial label propagation algorithms work. First, location data from users who have them are used as “ground truth” to predict the locations of users to whom they are directly connected. If a user has more than one friend with ground truth data, the geometric median is calculated to predict his or her location. The geometric median is preferred over the geometric mean, as it represent the actual location of users in the network and not a meaningless average of coordinates. In addition, it is less sensitive to outliers, which might happen when users post geo-located tweets while traveling. To give a concrete example, in Panel (a) the location of user  $a$  is predicted as the geometric median of users  $b$ ,  $d$ , and  $e$ .

In the second stage, after the first round of prediction is completed and new users have predicted location information, the algorithm carries out a second round of location predictions, which uses richer location data that is distributed across the network, incorporating both ground truth and

**Data:**  $U$ ,  $L$ , and  $N$   
Let  $E$  be the current mapping from user to location;  
Initialize  $E$  with  $L$ ;  
**while** *Convergence criteria are not met* **do**  
    Let  $E'$  be the next mapping from user to (predicted) location;  
    **for**  $u \in (U - \text{domain}(L))$  *(i.e., users who do not currently have location information)* **do**  
        Let  $M$  be a list of locations;  
        **for**  $n \in N(u)$  *(i.e., friends of user  $u$ )* **do**  
            **if**  $E(n) \neq \emptyset$  *(i.e., if the friend  $n$  has location information)* **then**  
                add  $E(n)$  to  $M$ ;  
            **end**  
        **end**  
        **if**  $M \neq \emptyset$  *(i.e., user  $u$ 's friends have location information)* **then**  
             $E'(u) = \arg \min_{x \in L} \sum_{y \in L} \text{distance}(x, y)$  *(the predicted location of user  $u$  is the geometric median of her friends' locations)*  
        **end**  
    **end**  
     $E = E'$   
**end**  
**Result:** Estimated user locations,  $E$   
**Algorithm 1:** Spatial Label Propagation (Jurgens, 2013)

predicted location data points. Panel (b) shows that in the second round, it is possible to predict the location for user  $c$  using data on the location of users  $a$ ,  $b$ , and  $e$ . In the same round, the location of user  $a$  is re-estimated, using a new data point from the predicted location of user  $f$ , in addition to the location information used in the first round, from users  $b$ ,  $d$ , and  $e$ . This process is repeated a fixed number of times or until a minimum proportion of users have predicted location data.<sup>22</sup>


I implement a slight deviation from the procedure described in Jurgens (2013). The original algorithm is designed to operate on a random sample of tweets, and not on a deep network of users who have timeline data and full lists of friends and followers. Thus, it identifies connections between individuals on the basis of “bidirectional mentions,” i.e., user A mentions user B in a tweet and vice-versa. Bidirectional mentions are used in the original algorithm as a proxy for friends on social media, as it is impractical to obtain lists of friends and followers from a random sample of tweets. However, in my database, I have actual lists of friends and followers of accounts flagged as ISIS activists. As such, while I adopt the Jurgens (2013) algorithm as-is and allow connections between

---


<sup>22</sup>I employed three iterations, which predicted locations for 1,626,350 users in the database.

Figure B2: Spatial Label Propagation Algorithm


(a) Round 1



(b) Round 2



(c) Round 3



individuals to be identified on the basis of bidirectional mentions, I also generate “artificial” tweets containing bidirectional mentions between activists and their followers and friends. This ensures that the network structure contained in my database will be faithfully reproduced in the spatial label propagation algorithm.

The SLP algorithm requires so-called “ground truth” data, i.e., users with a known location, to base the prediction of the location for users without a known location. I obtained ground truth data as follows. For users with at least one geolocated tweet, I used the coordinates from an arbitrarily selected geolocated tweet. For users without any geolocated tweets but with a location field in their user profile, I looked up the location using the Google Maps and/or Bing Maps APIs (the specific API is selected arbitrarily).<sup>23</sup> If there was a match, I used the coordinates corresponding to this location as the user’s ground truth location. To be sure, both of these methods are measured with error, but there is no reason to believe that these errors are systematically biased in any specific direction. Thus, by the law of large numbers, across the total universe of accounts with ground truth data ( $N = 287,482$ ), these errors should be inconsequential.

## Stability of location predictions

I verify the accuracy of the location prediction algorithm in the following way. The network structure in my database is relatively deep, centered around 14,979 ISIS activists for whom I have full lists of followers, as well as friends of a subset of the followers. Thus, individuals distributed across the network with ground truth data are connected to each other mainly through the ISIS activists’ accounts. This is different from flat networks studied in other SLP applications using data from random samples of tweets (Jurgens et al., 2015). As a result, cross validation using only data from accounts with ground truth information is not useful for estimating the performance of the model.

In non-network data, cross validation on the training set is useful because observations do not depend on each other. Thus,  $\hat{y}_i$ , the prediction for observation  $i$ , is simply some function of the covariates for unit  $i$  and some parameters:  $\hat{y}_i = f(x_i, \theta)$ . Taking observations out in cross validation to test the model’s prediction works well, because of the limited dependency between observations.

---

<sup>23</sup>Google Maps API: <https://developers.google.com/places/web-service/details>; Bing Maps API: <https://msdn.microsoft.com/en-us/library/ff701711.aspx>.



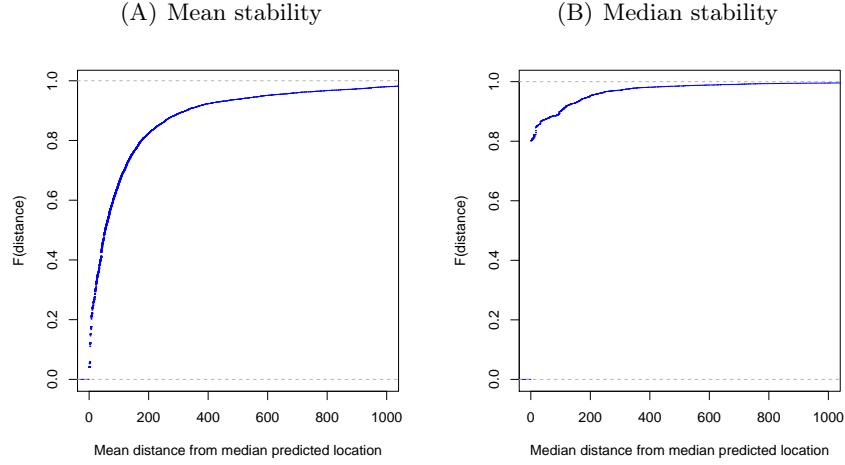
In network data, cross validation is more problematic, because observations are dependent:  $\hat{y}_i = f(\sum_j y_j, \theta)$ . Therefore, taking observations out in cross validation does not only change  $\theta$ , the parameters of the model, but also  $\sum_j y_j$ , the data used to predict  $\hat{y}_i$ . As a result, the estimations in the cross validation are likely to be biased, with greater bias for deeper networks in which the dependency between observations is higher.

To overcome this challenge and estimate the algorithm’s performance, I designed a 10-fold out-of-sample stability test. I divided the training set into ten folds, and in each fold I randomly excluded 1/10 of the ground truth data when estimating the model. The algorithm therefore ran ten times, each time using only 90% of the training data to predict the locations of all users in the dataset ( $N = 1,626,165$ ). I assume that the out-of-sample stability of the location prediction for each user  $i$  across ten folds can proxy the algorithm’s location prediction accuracy. The logic behind this assumption is that highly unstable (stable) predictions across ten different prediction exercises likely means that the prediction is not very accurate (accurate). If a given user’s friends are distributed geographically in a manner that renders the prediction highly unstable when excluding a random portion of the friends, then it means that the geometric median of the friends’ locations is probably not a good proxy for the user’s true location. On the other hand, if leaving out friends with location data does not affect the stability of the user’s predicted location, then it means that many of the user’s friends are located in the same area, making prediction stable, and likely more accurate.

After obtaining ten different location predictions for each user in the dataset, I calculated, for each user  $i$ , the mean and median distance from the median location predicted for user  $i$ . Figure B3 shows the performance for the ISIS activists’ accounts ( $N = 14,979$ ). Figure B4 shows the performance for the ISIS followers’ accounts ( $N = 1,611,633$ ). The figures plot the cumulative distribution function of the location predictions’ stability across ten prediction estimations. In Panel (a), the stability is calculated as the mean of the predicted locations’ deviations from the median predicted location for each user across the ten folds. In Panel (b), the stability is calculated as the median of the predicted locations’ deviations from the median prediction. When using the mean stability measure, the majority of users’ predicted locations are stable around a radius of about 50 kilometers or less for activists, and 70 kilometers or less for followers. When using the median stability measure, for over 80% of the users locations are predicted with a median stability

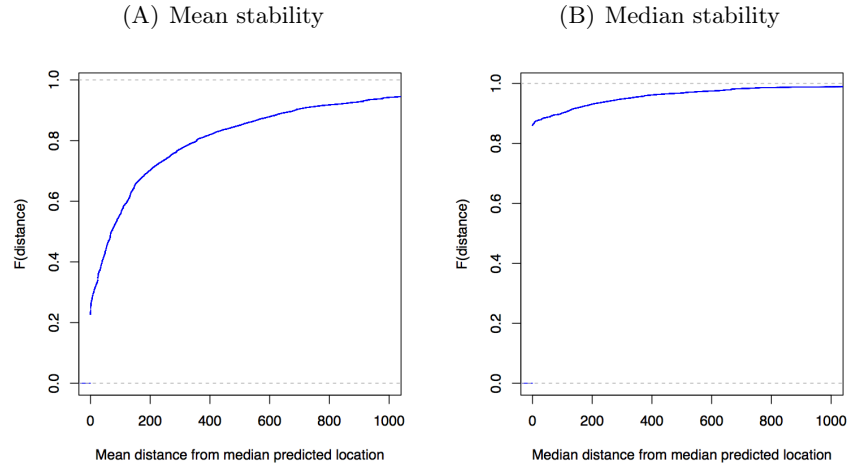
of 10 kilometers or less.

Figure B3: 10-Fold out-of-sample stability test (ISIS activists' accounts)



*Note:* The figure plots the cumulative distribution function of the stability of location predictions of ISIS activists ( $N = 14,979$ ) across ten prediction estimations when leaving out one-tenth of the training data each time. In Panel (a), the stability is calculated as the mean of the predicted locations' deviations from the median predicted location for each user across the ten folds. The  $x$  axis shows the mean distance from the median predicted location for each user. The  $y$  axis shows the probability that mean deviation is  $x$  distance or less from the user's median predicted location. In Panel (b), the stability is calculated as the median of the predicted locations' deviations from the median prediction. When using the mean stability measure, the majority of users' predicted locations are stable around a radius of about 50 kilometers or less. When using the median stability measure, for over 80% of the users locations are predicted with a median stability of 10 kilometers or less.

Figure B4: 10-Fold out-of-sample stability test (ISIS followers' accounts)



*Note:* The figure plots the cumulative distribution function of the stability of location predictions of ISIS followers ( $N = 1,611,633$ ) across ten prediction estimations when leaving out one-tenth of the training data each time. In Panel (a), the stability is calculated as the mean of the predicted locations' deviations from the median predicted location for each user across the ten folds. The  $x$  axis shows the mean distance from the median predicted location for each user. The  $y$  axis shows the probability that mean deviation is  $x$  distance or less from the user's median predicted location. In Panel (b), the stability is calculated as the median of the predicted locations' deviations from the median prediction. When using the mean stability measure, the majority of users' predicted locations are stable around a radius of about 70 kilometers or less. When using the median stability measure, for over 80% of the users locations are predicted with a median stability of 10 kilometers or less.

## B3 Classifying Twitter content

To generate the textual content outcomes in this study, I used supervised machine learning to classify tweets into several categories: (1) Anti-West, (2) Islam, (3) Sympathy with ISIS, (4) Life in ISIS territories, (5) Travel to Syria or foreign fighters, and (6) Syrian war. For each of the four languages: English, Arabic, French and German, I obtained a random sample of tweets posted by ISIS activists (i.e., the accounts that have been flagged by @CtrlSec). These tweets served as a training set for a classification model. The sizes of the training sets varied by language: English ( $N = 9,926$ ), Arabic ( $N = 10,631$ ), French ( $N = 6,158$ ), and German ( $N = 3,011$ ). Each tweet was assigned one or more of the categories by three distinct Amazon Mechanical Turk and/or Crowdfunder workers, and label(s) were retained for a given tweet if and only if there was “majority agreement,” i.e., at least two out of the three workers assigned the same label(s) to the tweet.

After obtaining the training set labels, I pre-processed the tweet text as follows. For tweets in the English, French and German languages, I removed punctuation, numbers, stop words, and applied standard word stemming algorithms for each language. For tweets in the Arabic language, I similarly removed punctuation and numbers. To pre-process Arabic tweets, I applied a standard set of Arabic text preparation techniques.<sup>24</sup>

With the pre-processed text, I generated a document-term matrix composed of unigrams and bigram tokens. That is, I obtained the frequency of individual words and two-word phrases that appeared in these tweets. I combined unigrams and bigrams in order to provide more textual structure and increase the predictive accuracy of the models. Any term included in the document-term matrix must have had appeared in at least two tweets in order to be included in the classification model. Then, I applied a term-frequency / inverse-document-frequency (tf-df) transformation to down-weight the frequency of very common phrases across the whole corpus, as is standard in automated content analysis (Ramos, 2003).

Since Twitter textual data are very noisy, and radical pro-ISIS content is rare, many tweets in the


---

<sup>24</sup>Specifically, I removed leading ‘alif lam’ with optional leading ‘waw’; leading ‘alif lam’ or double ‘lam’ at start of the text; leading ‘kaf alif lam’ with optional ‘waw’; leading ‘ba alif lam’ with optional ‘waw’; leading ‘fa alif lam’ with optional ‘waw’; leading double ‘alif’ with optional ‘lam’ and an optional leading ‘waw’; trailing ‘ha,’ ‘ya ya nun,’ ‘ya waw nun,’ ‘ha’ or ‘ha alif,’ ‘ha mim,’ ‘ha mim alif’; and single letters such as ‘waw.’ I used the code from: <http://badhessian.org/2012/08/text-normalization-and-arabic-in-r/>

database were coded as unrelated to any of the above categories. Class proportions for each language in the training set are shown in Tables B4 – B7. To facilitate statistical prediction, I followed King and Zeng (2001), randomly over-sampling pro-ISIS tweets and randomly under-sampling unrelated tweets to obtain a class proportion of 0.5 for each of the categories, for each topic, for each language.

I trained separate logit models using the labeled rebalanced training sets for each category in each language. For all specifications, I used the the elastic-net generalized linear model (Friedman, Hastie and Tibshirani, 2010), selecting the regularization parameter  $\lambda$  by cross-validation to maximize the area under the ROC curve. Figures B6 – B9 show the cross-validation curves for each language and topic. Model performance statistics for each topic and language are shown in Tables B8 – B7. The classification models for each topic and language were then employed on the full set of tweets in the database to classify each unlabeled tweet as belonging to one or more of these categories.

Figure B5: Supervised machine learning



*Note:* \* English: 9,926; Arabic: 10,631; French: 6,158; German: 3,011.

\*\* Each tweet coded by 3 coders, label retained if there was majority agreement.

\*\*\* Over-sample pro-ISIS content, under-sample unrelated tweets.

Table B4: Class proportions by topic (English)

	0	1
Anti-West	0.984577	0.015423
Islam	0.858215	0.141785
Sympathy with ISIS	0.982727	0.017273
Life in ISIS territories	0.963603	0.036397
Travel to Syria or foreign fighters	0.996607	0.003393
Syrian war	0.924532	0.075468

Table B5: Class proportions by topic (Arabic)

	0	1
Anti-West	0.998104	0.001896
Islam	0.913460	0.086540
Sympathy with ISIS	0.996777	0.003223
Life in ISIS territories	0.996777	0.003223
Travel to Syria or foreign fighters	0.999526	0.000474
Syrian war	0.981043	0.018957

Table B6: Class proportions by topic (French)

	0	1
Anti-West	0.971370	0.028630
Islam	0.890500	0.109500
Sympathy with ISIS	0.965607	0.034393
Life in ISIS territories	0.965607	0.034393
Travel to Syria or foreign fighters	0.982711	0.017289
Syrian war	0.947388	0.052612

Table B7: Class proportions by topic (German)

	0	1
Anti-West	0.959585	0.040415
Islam	0.924352	0.075648
Sympathy with ISIS	0.932124	0.067876
Life in ISIS territories	0.915026	0.084974
Travel to Syria or foreign fighters	0.947668	0.052332
Syrian war	0.915026	0.084974

Table B8: Model performance (English)

	anti-west	islamic-faith	is-sympathy	is-life	syria-travel-ff	syrian-war
Sensitivity	0.9981	0.9363	0.9992	0.9971	0.9992	0.9441
Specificity	1.0000	0.6752	0.9948	0.9948	1.0000	0.9943
Pos Pred Value	1.0000	0.7498	0.9949	0.9949	1.0000	0.9939
Neg Pred Value	0.9982	0.9106	0.9992	0.9971	0.9992	0.9478
Prevalence	0.4962	0.5097	0.5064	0.5024	0.5086	0.4949
Detection Rate	0.4953	0.4772	0.5060	0.5009	0.5082	0.4672
Detection Prevalence	0.4953	0.6364	0.5085	0.5035	0.5082	0.4701
Balanced Accuracy	0.9991	0.8057	0.9970	0.9960	0.9996	0.9692

Table B9: Model performance (Arabic)

	anti-west	islamic-faith	is-sympathy	is-life	syria-travel-ff	syrian-war
Sensitivity	0.9985	0.9627	0.9985	0.9987	0.9991	0.9583
Specificity	1.0000	0.9924	1.0000	1.0000	1.0000	1.0000
Pos Pred Value	1.0000	0.9922	1.0000	1.0000	1.0000	1.0000
Neg Pred Value	0.9985	0.9634	0.9985	0.9987	0.9990	0.9599
Prevalence	0.5039	0.5028	0.5094	0.4973	0.5115	0.5007
Detection Rate	0.5031	0.4841	0.5086	0.4967	0.5110	0.4798
Detection Prevalence	0.5031	0.4879	0.5086	0.4967	0.5110	0.4798
Balanced Accuracy	0.9992	0.9775	0.9993	0.9993	0.9995	0.9792

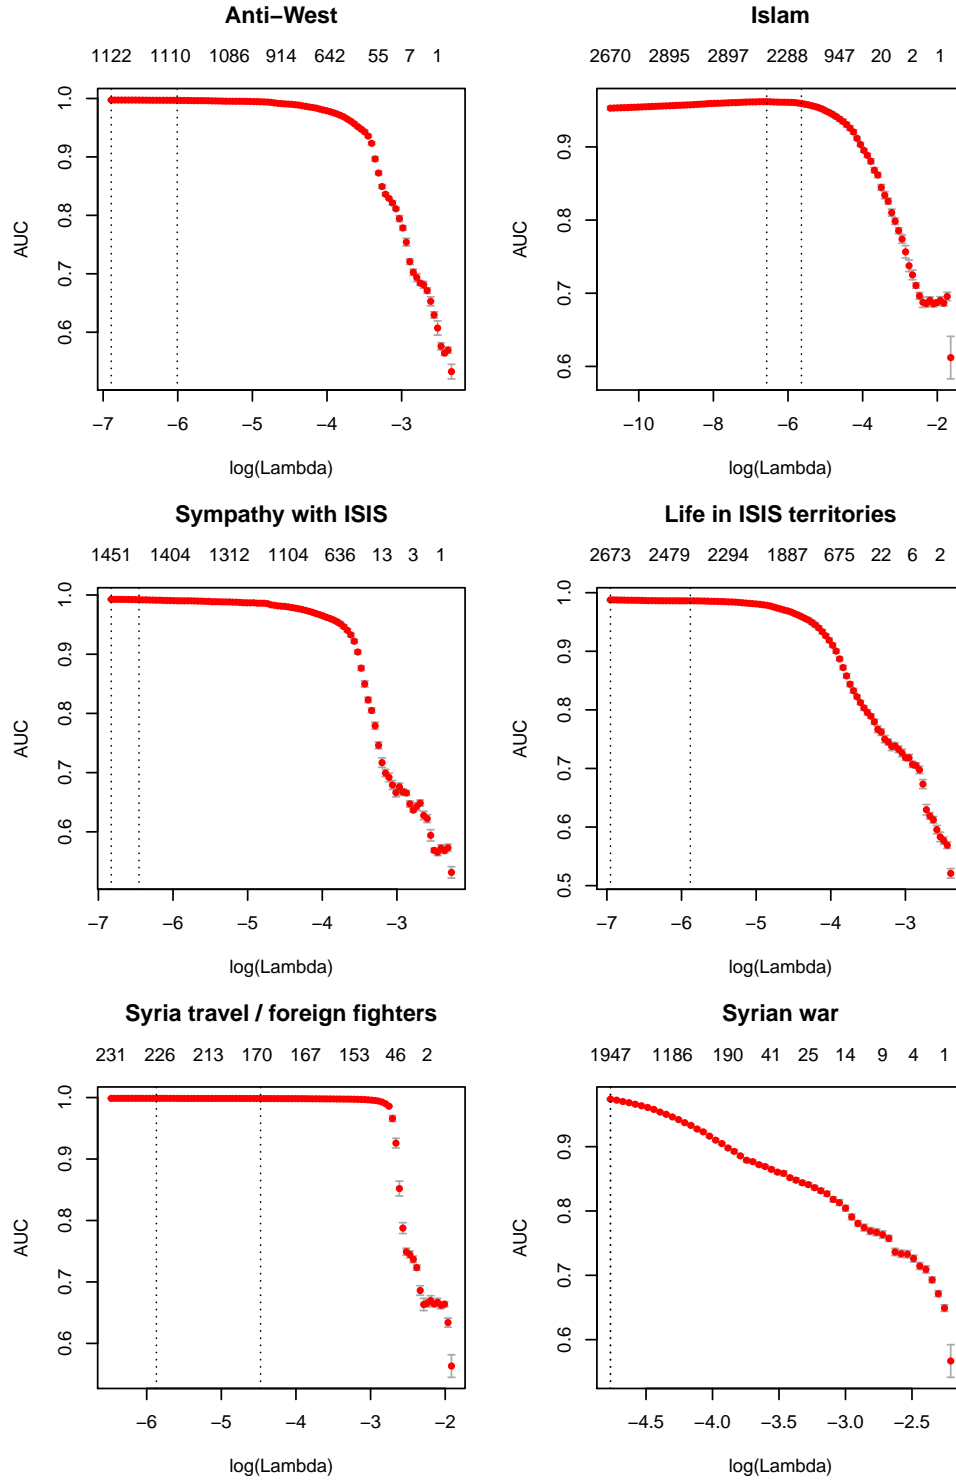
Table B10: Model performance (French)

	anti-west	islamic-faith	is-sympathy	is-life	syria-travel-ff	syrian-war
Sensitivity	0.9985	0.9910	0.9925	0.9876	0.9985	0.9926
Specificity	0.9923	0.9978	0.9952	1.0000	1.0000	1.0000
Pos Pred Value	0.9922	0.9977	0.9951	1.0000	1.0000	1.0000
Neg Pred Value	0.9985	0.9912	0.9926	0.9872	0.9985	0.9926
Prevalence	0.4951	0.4980	0.4975	0.5114	0.5127	0.5031
Detection Rate	0.4943	0.4936	0.4938	0.5051	0.5120	0.4993
Detection Prevalence	0.4982	0.4947	0.4962	0.5051	0.5120	0.4993
Balanced Accuracy	0.9954	0.9944	0.9939	0.9938	0.9993	0.9963

Table B11: Model performance (German)

	anti-west	islamic-faith	is-sympathy	is-life	syria-travel-ff	syrian-war
Sensitivity	0.9661	0.9828	0.9839	0.9770	0.9818	0.9720
Specificity	1.0000	0.9713	0.9766	0.9949	0.9766	0.9880
Pos Pred Value	1.0000	0.9729	0.9779	0.9947	0.9778	0.9869
Neg Pred Value	0.9686	0.9818	0.9829	0.9778	0.9808	0.9744
Prevalence	0.4891	0.5119	0.5135	0.4964	0.5124	0.4813
Detection Rate	0.4725	0.5031	0.5052	0.4850	0.5031	0.4679
Detection Prevalence	0.4725	0.5171	0.5166	0.4876	0.5145	0.4741
Balanced Accuracy	0.9831	0.9771	0.9802	0.9859	0.9792	0.9800

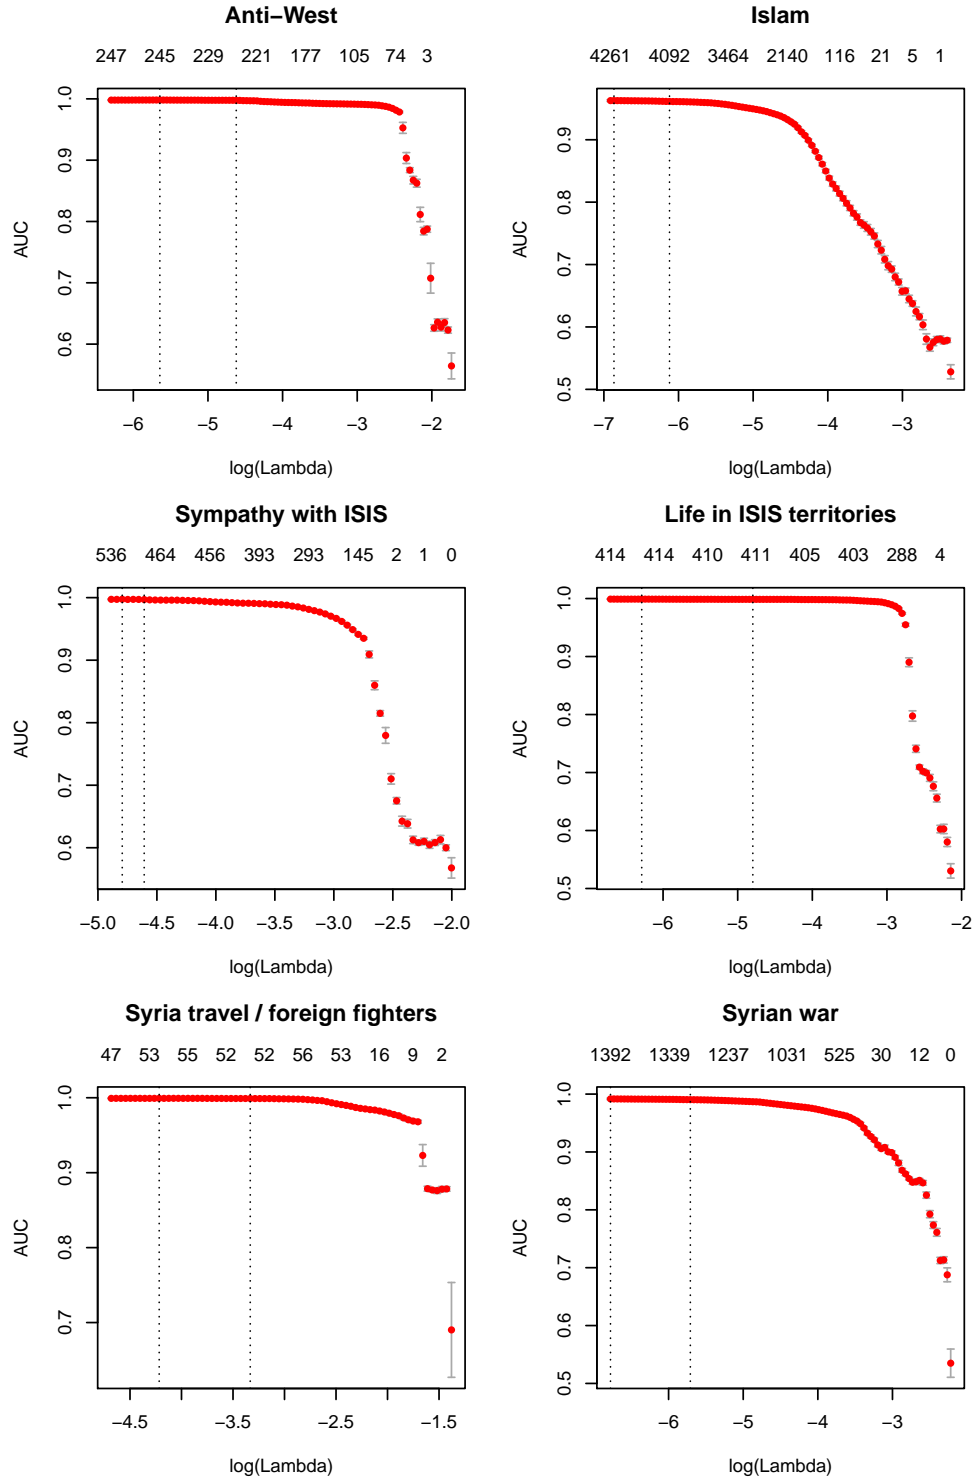
Figure B6: Cross validation for model choice (English tweets)



*Note:* The figure shows cross-validation curves for model choice in text classification of English language tweets for six topics. The cross-validation estimates for each model are shown in red dots, surrounded by error bars, plotted against the  $\lambda$  sequence. The  $y$  axis marks the Area Under the ROC Curve (AUC). Two selected  $\lambda$ s are marked by vertical dotted lines. The numbers at the top of the figures represent the number of tokens (unigrams and bigrams) used in each model.

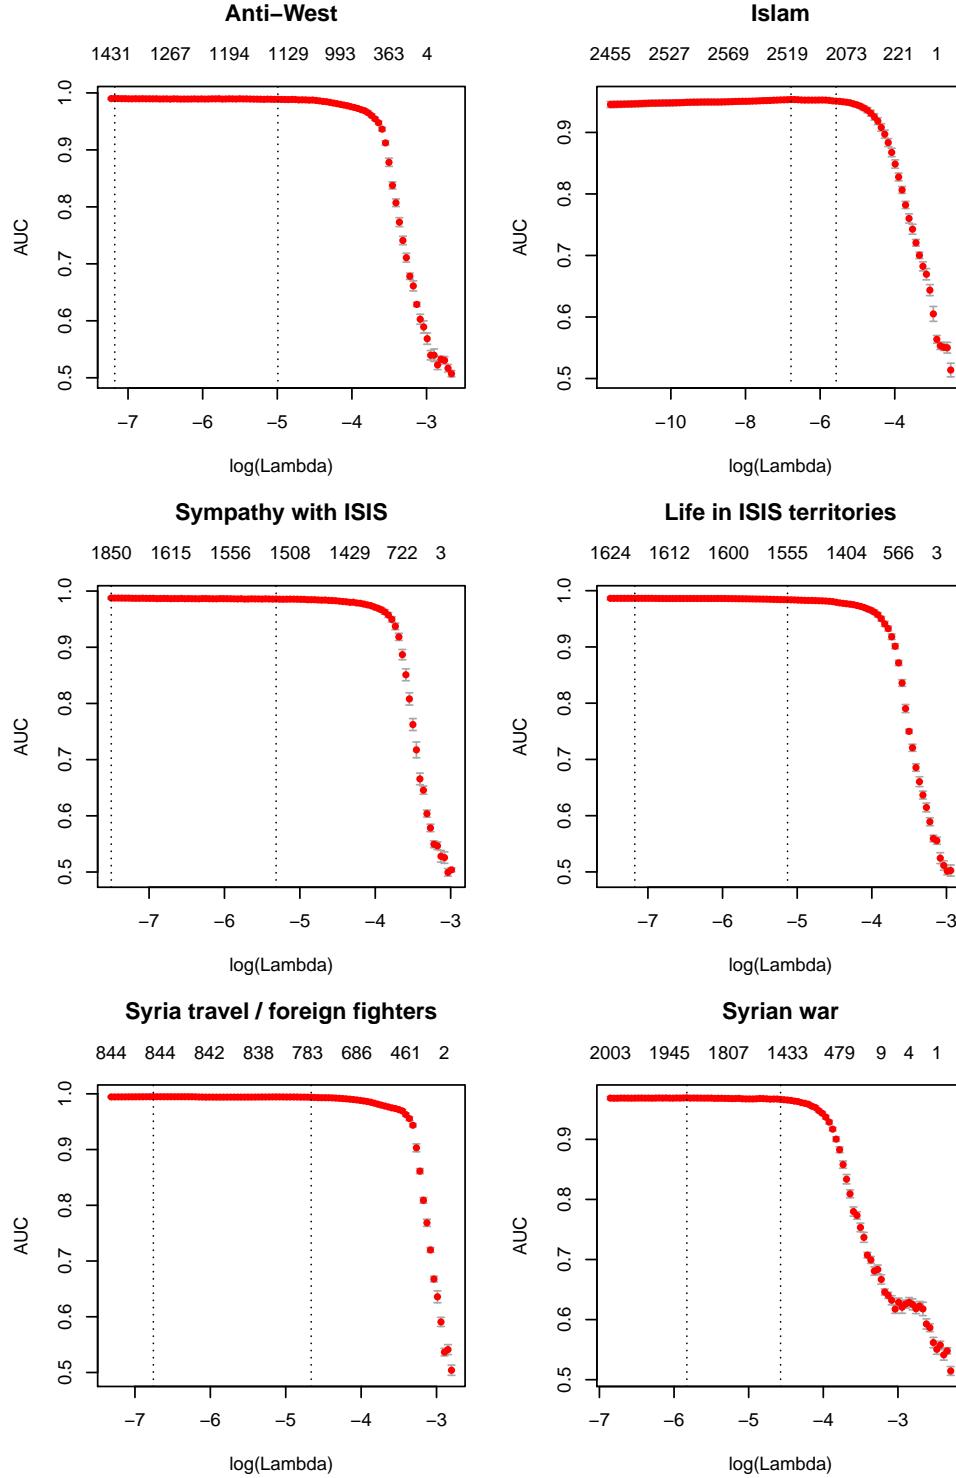


Figure B7: Cross validation for model choice (Arabic tweets)



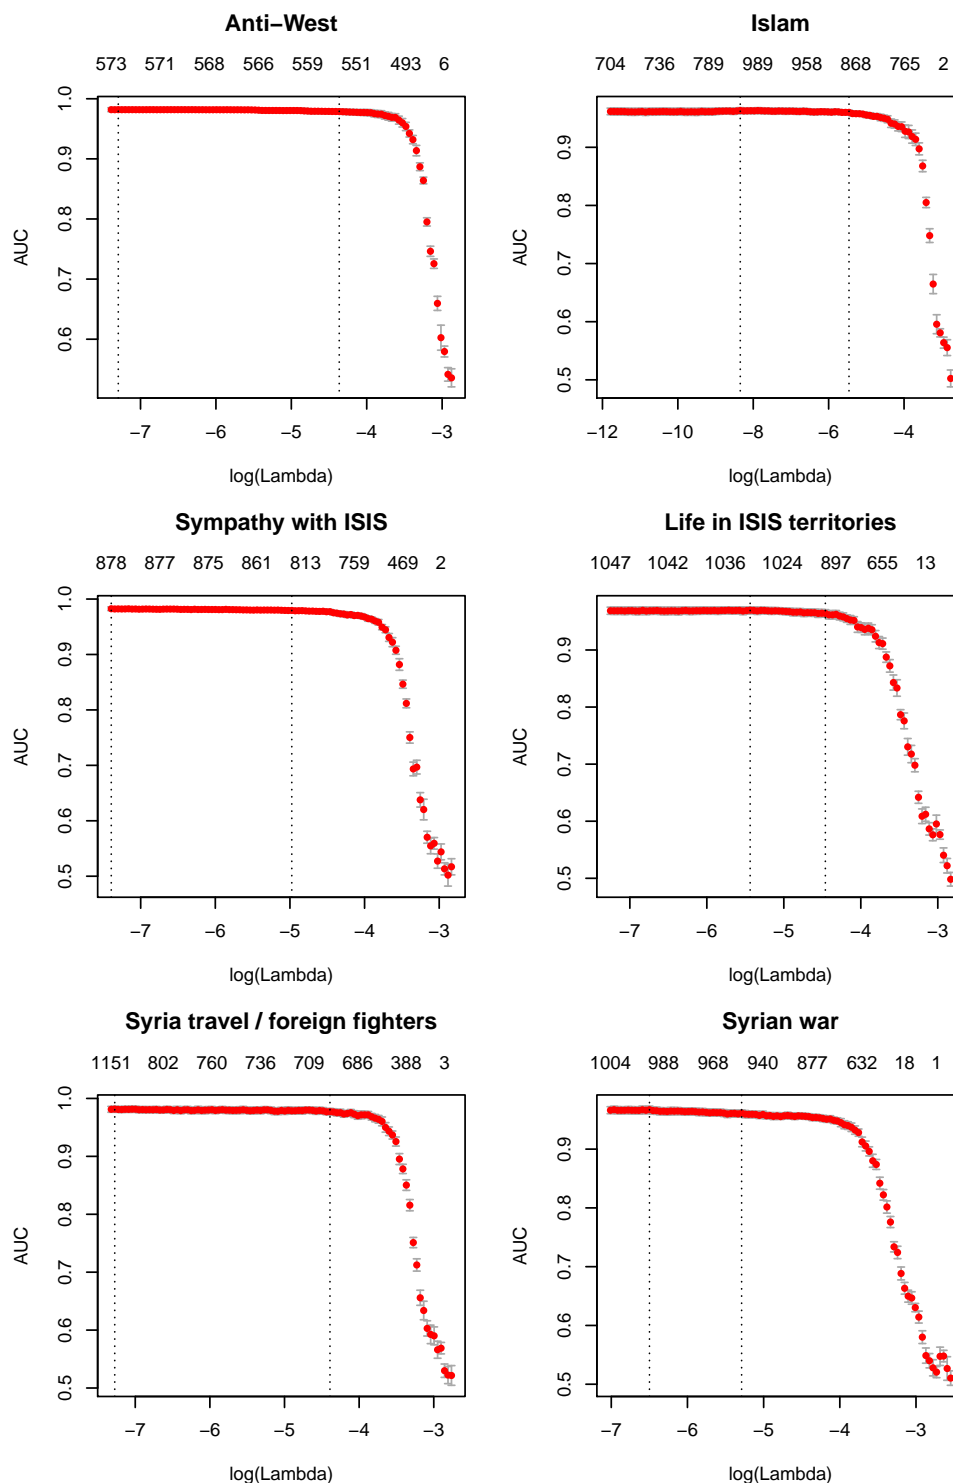
*Note:* The figure shows cross-validation curves for model choice in text classification of Arabic language tweets for six topics. The cross-validation estimates are shown in red dots, surrounded by error bars, plotted against the  $\lambda$  sequence. The  $y$  axis marks the Area Under the ROC Curve (AUC). Two selected  $\lambda$ s are marked by vertical dotted lines. The numbers at the top of the figures represent the number of tokens (unigrams and bigrams) used in each model.

Figure B8: Cross validation for model choice (French tweets)



*Note:* The figure shows cross-validation curves for model choice in text classification of French language tweets for six topics. The cross-validation estimates are shown in red dots, surrounded by error bars, plotted against the  $\lambda$  sequence. The  $y$  axis marks the Area Under the ROC Curve (AUC). Two selected  $\lambda$ s are marked by vertical dotted lines. The numbers at the top of the figures represent the number of tokens (unigrams and bigrams) used in each model.

Figure B9: Cross validation for model choice (German tweets)



*Note:* The figure shows cross-validation curves for model choice in text classification of German language tweets for six topics. The cross-validation estimates are shown in red dots, surrounded by error bars, plotted against the  $\lambda$  sequence. The  $y$  axis marks the Area Under the ROC Curve (AUC). Two selected  $\lambda$ s are marked by vertical dotted lines. The numbers at the top of the figures represent the number of tokens (unigrams and bigrams) used in each model.

Figure B10: Tweet content classification task instructions for CrowdFlower workers

## Classify Syrian Civil War Tweets (English)

### Instructions ▲

Please label each tweet by checking all labels that correctly describe its content. If a tweet does not fit any of the labels, check "None of the Above".

Category	Description
Anti-West	Anti-West rhetoric, criticizing Western countries' foreign policy and military operations in the Middle East
Islamic faith	Expressions of faith in the Islamic religion, Islamic quotes, and prayers and/or requests for prayers
IS sympathy	Expressions of support or sympathy with the Islamic State, its ideology and its activities in territories under its control
Life in IS territories	Tweets from Islamic State activists describing their life in the territories controlled by the Islamic State; includes descriptions of daily activities under Islamic State rule, fighting; things that 'market' the life in Syria to potential foreign fighters
Travel to Syria / foreign fighters	Tweets describing interest or intent to travel to Syria, and/or discussion of foreign fighters
Syrian war	Tweets describing events in the Syrian civil war and/or discussion/analysis of those events
Islamophobia	Tweets describing unfair treatment of Muslims and/or discrimination against Muslims in non-Muslim majority countries

Islam is not a religion as Christianity/Judaism nor a political belief as Capitalism/Communism but rather it is a comple...

#### Classification:

- ☐ Anti-West
- ☐ Islamic faith
- ☐ IS sympathy
- ☐ Life in IS territories
- ☐ Travel to Syria / foreign fighters
- ☐ Syrian war
- ☐ Islamophobia
- ☐ None of the Above

UK extremist's sharia law photo used in free speech ad

#### Classification:

- ☐ Anti-West
- ☐ Islamic faith
- ☐ IS sympathy
- ☐ Life in IS territories
- ☐ Travel to Syria / foreign fighters
- ☐ Syrian war
- ☐ Islamophobia
- ☐ None of the Above

*Note:* This is an example of a CrowdFlower task to classify English language tweets on various dimensions. Classified tweets are included in a training set to predict the content of unclassified tweets. The classification was carried out in English, French, Arabic, and German.

Table B12: Top 50 words across topics

	Topic	Top 50 words
1	Sympathy with ISIS	amp.will, beauti.islam, anyth.deserv, protector.htt, gambl, allah.muhammad, vis- ceg, ideolog, bros, murthad, bomb.defeat, anonym.claim, day.let, allah.martyr, idni, attaref, bomb.raid, anyth.el, learn.quran, muwahideen, child.dont, antiislam, isi, frien, insignif, abysinnia, beer, back.start, otherhezbollah, aisha, belong.pious, in- cap, yet.dont, alway.week, ghazl, ago.near, hazima, week.follow, alfurqan, amp.realli, usrussia, bizarr, upcom, backbit.hurt, palestin.sha, clemenc, concern.islam, boy.gun, shepherd, attritionlos
2	Life in ISIS territories	samir, ate, besiege, nobl, dress, today.https, bomb.etc, iraqi.children, dua.syria, rt.propaganda, saa, cctv, masharialashwaq, jaish.alislam, machin, momineen, al- loush.martyr, pkk.terrorist, islam.court, baqir, trade, charli, allegi.htt, behind.five, border.guard, assad.barrel, univ, civilian.say, outpost, eatabl, flay, dael, citi.fallen, jahiliyah, almohammad, conniv, aid.kuffar, center, hussari, qaeda, kiss, antiaq, pkk, reloc, sayyidi, judici, khaleefah, allow.sleep, like.year, ad.encrypt
3	Travel to Syria or foreign fighters	engag, dua.will, come.europ, join, elev, countrysid, martyrdom.op, sane, cemeteri, ahzamiyah, kasiki, fu, blogger, fighter, abar, aynisa, kashmiri, breakdown, foreign, shut, australian, syria, militia, kurdishwho, iraqi, martyrdom, batch.recruit, europ, braclet, age, australian.teen, bangladeshi, bangladeshi.blogger, muay, excut, fled, milit, fool.know, amp.islam, arraqqa, armi.right, get.today, graciou, pastpres, sus- tain, join.isi, amer, loot, egypt, najaf
4	Syrian war	syria, syrian, airstrik, regim, rebel, suspicion, traumat, russian, https, rt, will.face, hamid, erad, offici.tell, go.arrest, ghannam, militari, attack, againstassad, bewar, strike, aldagestani, dhawahiri, hom, onesyourself, kiss, jet, jay, southern, seiz, typo, regiment, fight, latest.twitter, amp.bomb, commit, besiege.town, civilian, ex- plo, chao.eastern, allow.enemi, mani.report, prior, build.center, heartbreak, bombard, missil, barrel, armi, enemi.get
5	Anti-West	afp.un, advic.us, bush, cant.invad, america, design, come.condemn, china.well, georg, gunmen, rppli, gwot, anyon.israel, erad, ampstop, militiasjihadist, real.terrorist, dua.ikhwan, punch, amp.tomorrow, org, aftermath.us, deplet, islamampth, usa, west- ern, attempt.stab, punish, criminalis, belief.crusad, catastroph, clark, rehman, anti- war, repercuss, alli.usa, heyena, usback, pentagon, holocaust, obama, brother.say, brother.allah, american, washington, alli.nato, democraci, gay.gambia, bashar, penni
6	Islam	allah, quran, muslim, prophet, allaah, nabi, prayer, surah, muhammad, qayyim, rather, islam, idhnillah, sin, tagh, allh, hiatus, breast, shayt, fast, qurn, habit, dawah, spous, hijab, alrahman, religion, qualifi, alnubala, nur, ayede, hellfir, qur, therebi, all, almubarak, believ, haqq, btw, bless, inshaal, aha, islaam, foremost, faith, assahab, torah, brother.allahuakbar, execut.saudia, paradis

Table B13: Examples of tweets in different topics

<i>Sympathy with ISIS</i>
Jihad is the greatest of all deeds #IslamicState Show everything from the Islamic State and other groups in Syria. It's important to hear all sides of the story. Assalam o Alaikom to All Islamic State Brothers In sha Allah we will have honor again #IslamicState
<i>Life in ISIS territories</i>
#Aljazeera reports from inside the city of #Raqqa and shows how the #IslamicState runs the daily life English Testimony from a girl in #Yarmouk Camp about the #IslamicState The glorious and mighty army of the Caliphate: Young kids ready to blow themselves up. Health services in Islamic state Wedding of an #ISIS fighter in #Raqqa: In Ribaat... Nice feeling really; Sit with the bros, drink tea, read Quran, relax & just observe the enemy! #Syria
<i>Travel to Syria or foreign fighters</i>
a lot of foreign fighters still coming in. Seems a lot responding to the call of the scholars of General March, also indicating open way in! is the door to sham open? i want your kik please akh, maybe there is sister in my country that have more money and looking for hijra too... come on join us at syam.. Dutch fighters in ar-Raqqa province #Syria
<i>Syrian war</i>
#IS fighters readying to fight an invasion of Yarmouk Camp by Assad's allies Jaysh Al-Islam and Liwa Sham Al-Rasool Massive destruction in Douma today after one of Assad's almost daily air strikes on the city. #Syria #Damascus #Syria - The evil #Assad regime lost Busra al-Harir so they tortured a 6 year old girl out of revenge... Massive explosion rocked entire of #Ramadi city. No further details yet.. #Iraq #ISIS
<i>Anti-West</i>
America has been at war 222 out of 239 years since 1776. Let that sink for a moment. If Islamic State terror is evil why would Western State war be good? US-led wars on terror have killed four million Muslims since 1990 It's sad when I am more afraid of our government then #ISIS ! At least I know #ISIS hates #America #Government =wolves Why are we shocked at ISIS brutality but not shocked by US British & European brutality?
<i>Islam</i>
Call upon Me; I will respond to you. #Quran 40:60 My identity is in who Allah says I am not in who others say I am. Allah's opinion is the only one that truly matters. To think that Allah Almighty is present with you at every given moment is the most excellent form of #faith. The beauty of Sujood is such that you whisper silently in to the ground and it's heard up in the Heavens. May allah bless you brother.....

Table B14: Correlates of flagged accounts

	Pr(flagged as ISIS account)					
	(1)	(2)	(3)	(4)	(5)	(6)
Sympathy with ISIS (# tweets)	0.002*** (0.00001)					
Travel to Syria or foreign fighters (# tweets)		0.001*** (0.00001)				
Life in ISIS territories (# tweets)			0.001*** (0.00001)			
Anti-West (# tweets)				0.002*** (0.00002)		
Syrian war (# tweets)					0.001*** (0.00002)	
Islam (# tweets)						0.001*** (0.00001)
Constant	0.008*** (0.0001)	0.008*** (0.0001)	0.010*** (0.0001)	0.008*** (0.0001)	0.011*** (0.0001)	0.006*** (0.0001)
Observations	1,052,842	1,052,842	1,052,842	1,052,842	1,052,842	1,052,842
R <sup>2</sup>	0.018	0.020	0.008	0.016	0.003	0.038

\* p<0.1; \*\* p<0.05; \*\*\* p<0.01