Countering Violent Extremism and Radical Rhetoric

Online Appendix

Contents

1 Community Engagement Events

I collected information on community engagement events from newsletter reports published by the Office of Civil Rights and Civil Liberties (CRCL) in the Department of Homeland Security. CRCL has been holding community engagement events since 2003 (Bjelopera, 2014). The first event took place in Dearborn, Michigan, and community engagement activities soon expanded to other cities in the United States. 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, 2010)

I collected information on all events held by CRCL using these monthly reports. 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. Figure A1 shows an example of a newsletter report from August 2015. In this study, I focus on events that took place between 2014 and 2016 due to the availability of Twitter data from this time period. Table A1 provides information on the timing and location of each of the 78 community engagement events analyzed in the paper.

Figure A1: Community Engagement Activities in CRCL's August 2015 Newsletter

CRCL on the Road, August

August 17 – Denver, Colorado
CRCL convened meetings with the U.S. Attorney's Office and diverse community stakeholders.

August 25 – Denver, Colorado
CRCL convened its quarterly community engagement roundtables with diverse ethnic and community-based organizations.

August 27 – Los Angeles, California
CRCL convened its quarterly community engagement roundtables with diverse ethnic and community-based organizations.

August 24 – Philadelphia, Pennsylvania
CRCL convened meetings with diverse community stakeholders.

August 26 – Atlanta, Georgia
CRCL convened its quarterly community engagement roundtables with diverse ethnic and community-based organizations.

August 27 – Boston, Massachusetts
CRCL participated in the BRIDGES roundtable with local federal partners and diverse community leaders.

Table A1: Community Engagement Events Analyzed in the Study

Date	City	State	Date	City	State
2014-01-09	Atlanta	GA	2015-03-31	Atlanta	GA
2014-01-16	Los Angeles	CA	2015-04-09	Chicago	IL
2014-02-11	Minneapolis	MN	2015-04-27	Houston	TX
2014-02-27	Houston	TX	2015-04-30	Seattle	WA
2014-03-05	Phoenix	AZ	2015-06-06	Phoenix	AZ
2014-03-13	Denver	CO	2015-06-23	Chicago	IL
2014-03-27	Chicago	IL	2015-06-25	Atlanta	GA
2014-04-10	Denver	CO	2015-07-23	Chicago	IL
2014-04-30	Los Angeles	CA	2015-08-25	Denver	CO
2014-04-30	New York	NY	2015-08-26	Atlanta	GA
2014-05-08	Tampa	FL	2015-08-27	Boston	MA
2014-05-21	Minneapolis	MN	2015-08-27	Los Angeles	CA
2014-06-12	Atlanta	GA	2015-09-17	Tampa	FL
2014-06-12	Houston	TX	2015-10-08	Boston	MA
2014-06-24	Chicago	IL	2015-10-29	Chicago	IL
2014-07-30	Denver	CO	2015-10-29	Houston	TX
2014-08-14	Seattle	WA	2015-11-18	Columbus	OH
2014-08-22	Orlando	FL	2015-11-23	Denver	CO
2014-08-27	Los Angeles	CA	2015-12-16	Atlanta	GA
2014-08-28	New York	NY	2015-12-17	Denver	CO
2014-10-21	Chicago	IL	2015-12-17	Los Angeles	CA
2014-10-30	Boston	MA	2015-12-17	Tampa	FL
2014-11-04	Houston	TX	2016-01-13	Detroit	MI
2014-11-07	Minneapolis	MN	2016-01-25	Denver	CO
2014-11-13	Los Angeles	CA	2016-01-28	Boston	MA
2014-11-20	Atlanta	GA	2016-02-11	Chicago	IL
2014-12-04	Atlanta	GA	2016-02-18	Minneapolis	MN
2014-12-15	Houston	TX	2016-02-23	New York	NY
2014-12-18	Tampa	FL	2016-03-10	Tampa	FL
2015-01-21	Seattle	WA	2016-03-23	Detroit	MI
2015-01-22	Boston	MA	2016-03-28	Houston	TX
2015-01-28	Chicago	IL	2016-03-29	Dallas	TX
2015-01-28	Detroit	MI	2016-03-30	Columbus	OH
2015-02-12	Tampa	FL	2016-03-30	Phoenix	AZ
2015-02-19	Denver	CO	2016-04-07	Atlanta	GA
2015-02-24	Columbus	OH	2016-04-09	Los Angeles	CA
2015-02-25	Phoenix	AZ	2016-05-04	Los Angeles	CA
2015-03-16	New York	NY	2016-05-17	Portland	OR
2015-03-24	Minneapolis	MN	2016-05-18	Seattle	WA

2 Islamic State Sympathizers 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 around the world. Below, I provide an overview of the data collection procedure, which included identifying accounts of Islamic State supporters on Twitter, downloading information on their posting history, and coding the extent to which their posts reflected pro-ISIS rhetoric.

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. In this study, I use a subset of these data on users located in the United States, which amounts to 30,358 accounts.

Measuring Online Expressions of Support for ISIS

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.[1]

1. *Travel to Syria or foreign fighters* - tweets describing interest or intent to travel to Syria, and/or discussion of foreign fighters

[1]When developing my training set, I coded content into additional categories, including anti-West sentiment, references to Islam (expressions of faith, Islamic quotes, and prayers and/or requests for prayers), as well as Islamophobia (content describing discrimination against Muslims). I did not use these categories in my analysis, which focused on pro-ISIS rhetoric.

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

4. *Syrian war*- tweets describing events in the Syrian civil war and/or discussion/analysis of those events

For each of the four languages, 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 Figure Eight 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. See Figure A2 for an example of instructions for the classification task in the Figure Eight platform.

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 Arabic, I similarly removed punctuation and numbers, and used the R package `arabicStemR` to stem Arabic text (Nielsen, 2017).[2] 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-idf) transformation to down-weight the frequency of very common phrases in the corpus.

Since Twitter textual data are very noisy, and pro-ISIS content was rare, many tweets in the database were coded as unrelated to any of the above categories. 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.

[2]See `https://CRAN.R-project.org/package=arabicStemR` for more details.

Figure A2: Tweet content classification task instructions for Figure Eight 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 Figure Eight 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.

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 λ by cross-validation to maximize the area under the ROC curve. Tables A3 – A6 (taken from Mitts (2019)) show model performance statistics from 10-fold cross validation for each topic and language.

Creating a Pro-ISIS Index

The article uses a pro-ISIS index that captures pro-ISIS sentiment across the four content categories (sympathy with ISIS, life in ISIS-controlled territories, ISIS's actions in the Syrian civil war, and foreign fighters). To create the index, I summed the predicted content scores from the classification models for each topic, and used min-max normalization to rescale the variable to range between zero and one. Table A2 below provides tweet-level summary statistics for the index and its component scores.

Table A2: Summary statistics of pro-ISIS index (tweet-level data)

	Number of tweets	Mean	St. Dev.	Min	Max
Sympathy with ISIS	15,140,867	0.133	0.285	0.000	1.000
Life in ISIS territories	15,140,867	0.215	0.345	0.000	1.000
Travel to Syria or foreign fighters	15,140,867	0.055	0.175	0.000	1.000
Syrian war	15,140,867	0.229	0.233	0.000	1.000
Pro-ISIS index	15,140,867	0.632	0.559	0.020	4.000
Pro-ISIS index (normalized)	15,140,867	0.154	0.140	0.000	1.000

Since the analysis presented in the article is at the user level, I aggregated the pro-ISIS index for each user in the one week before and 1-4 weeks after community engagement events. In the "pre" period, the aggregated variable reflects the average of the index for each user in the week before community engagement events. In the "post" periods, it is the average of the index in each time window – first week, two weeks, three weeks, or four weeks after community engagement events. Summary statistics for the aggregated variable, for each time period, is shown in Table A9.

Table A3: Model performance (English)**

	anti-west	is-sympathy	is-life	syria-travel-ff	syrian-war
Accuracy	0.9899	0.9868	0.9784	0.9960	0.9802
Sensitivity	0.9855	0.9781	0.9628	0.9921	0.9699
Specificity	0.9941	0.9955	0.9943	1.0000	0.9907
Pos Pred Value	0.9939	0.9954	0.9940	1.0000	0.9906
Neg Pred Value	0.9862	0.9787	0.9635	0.9920	0.9702
Precision	0.9939	0.9954	0.9940	1.0000	0.9906
Recall	0.9855	0.9781	0.9628	0.9921	0.9699
F1	0.9897	0.9867	0.9781	0.9960	0.9801
Prevalence	0.4936	0.4962	0.5019	0.5020	0.5019
Detection Rate	0.4865	0.4853	0.4831	0.4979	0.4867
Detection Prevalence	0.4895	0.4876	0.4860	0.4979	0.4914
Balanced Accuracy	0.9898	0.9868	0.9785	0.9960	0.9803

Note: Model performance metrics are taken from Mitts (2019).

Table A4: Model performance (Arabic)**

	anti-west	is-sympathy	is-life	syria-travel-ff	syrian-war
Accuracy	0.9866	0.9828	0.9928	0.9948	0.9816
Sensitivity	0.9843	0.9825	0.9855	0.9965	0.9635
Specificity	0.9889	0.9831	1.0000	0.9931	1.0000
Pos Pred Value	0.9887	0.9828	1.0000	0.9929	1.0000
Neg Pred Value	0.9846	0.9830	0.9858	0.9967	0.9643
Precision	0.9887	0.9828	1.0000	0.9929	1.0000
Recall	0.9843	0.9825	0.9855	0.9965	0.9635
F1	0.9865	0.9826	0.9927	0.9947	0.9814
Prevalence	0.4972	0.4942	0.4984	0.4925	0.5029
Detection Rate	0.4894	0.4856	0.4912	0.4908	0.4845
Detection Prevalence	0.4950	0.4941	0.4912	0.4943	0.4845
Balanced Accuracy	0.9866	0.9828	0.9928	0.9948	0.9818

Note: Model performance metrics are taken from Mitts (2019).

Table A5: Model performance (French)**

	anti-west	is-sympathy	is-life	syria-travel-ff	syrian-war
Accuracy	0.9955	0.9948	0.9927	0.9968	0.9940
Sensitivity	0.9909	0.9933	0.9887	0.9938	0.9885
Specificity	1.0000	0.9963	0.9969	1.0000	0.9992
Pos Pred Value	1.0000	0.9963	0.9971	1.0000	0.9993
Neg Pred Value	0.9913	0.9933	0.9884	0.9936	0.9892
Precision	1.0000	0.9963	0.9971	1.0000	0.9993
Recall	0.9909	0.9933	0.9887	0.9938	0.9885
F1	0.9954	0.9948	0.9928	0.9969	0.9938
Prevalence	0.4998	0.5054	0.4993	0.5065	0.4911
Detection Rate	0.4953	0.5020	0.4935	0.5034	0.4855
Detection Prevalence	0.4953	0.5039	0.4950	0.5034	0.4858
Balanced Accuracy	0.9954	0.9948	0.9928	0.9969	0.9939

Note: Model performance metrics are taken from Mitts (2019).

Table A6: Model performance (German)**

	anti-west	is-sympathy	is-life	syria-travel-ff	syrian-war
Accuracy	0.9793	0.9648	0.9710	0.9772	0.9777
Sensitivity	0.9696	0.9564	0.9693	0.9879	0.9772
Specificity	0.9896	0.9717	0.9727	0.9662	0.9775
Pos Pred Value	0.9894	0.9693	0.9711	0.9679	0.9793
Neg Pred Value	0.9688	0.9609	0.9705	0.9869	0.9775
Precision	0.9894	0.9693	0.9711	0.9679	0.9793
Recall	0.9696	0.9564	0.9693	0.9879	0.9772
F1	0.9793	0.9627	0.9701	0.9778	0.9780
Prevalence	0.5057	0.4756	0.4896	0.5150	0.4974
Detection Rate	0.4902	0.4549	0.4746	0.5088	0.4860
Detection Prevalence	0.4953	0.4694	0.4886	0.5254	0.4969
Balanced Accuracy	0.9796	0.9641	0.9710	0.9771	0.9774

Note: Model performance metrics are taken from Mitts (2019).

3 Predicting Geographic Location of ISIS Sympathizers

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 $u \rightarrow [n_i, ..., n_m]$. 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 A3 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 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

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 $u \in (U - domain(L))$ *(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**
            $E'(u) = \arg\min_{x \in L} \sum_{y \in L} distance(x, y)$ *(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)

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.

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 the ISIS Twitter data, 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 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

Figure A3: Spatial Label Propagation Algorithm

(a) Round 1

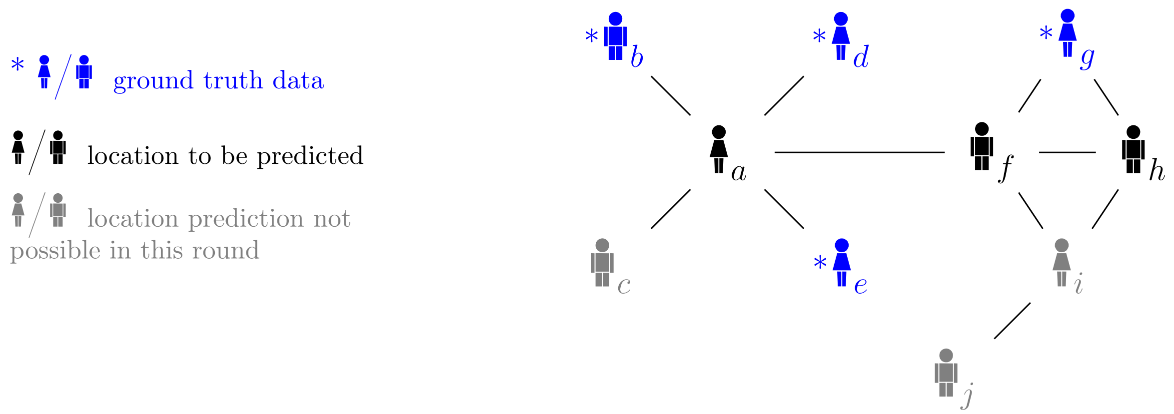

(b) Round 2

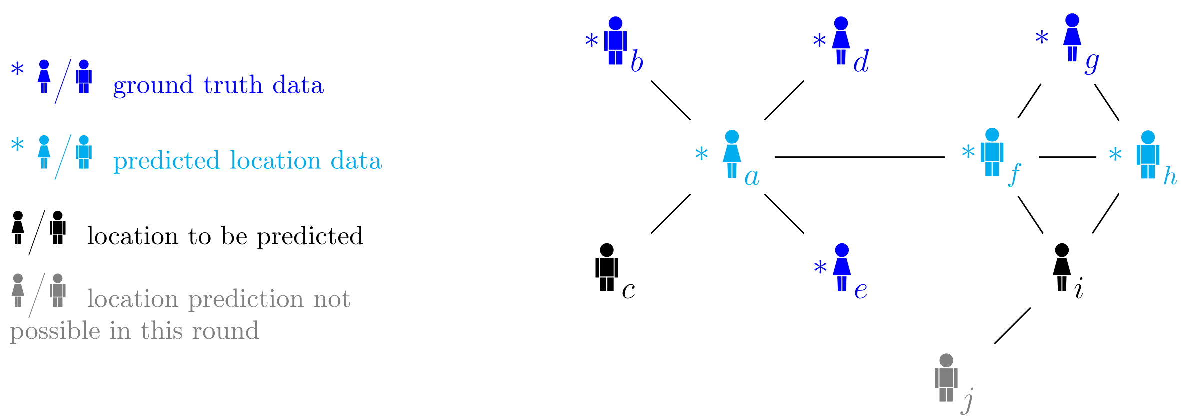

(c) Round 3

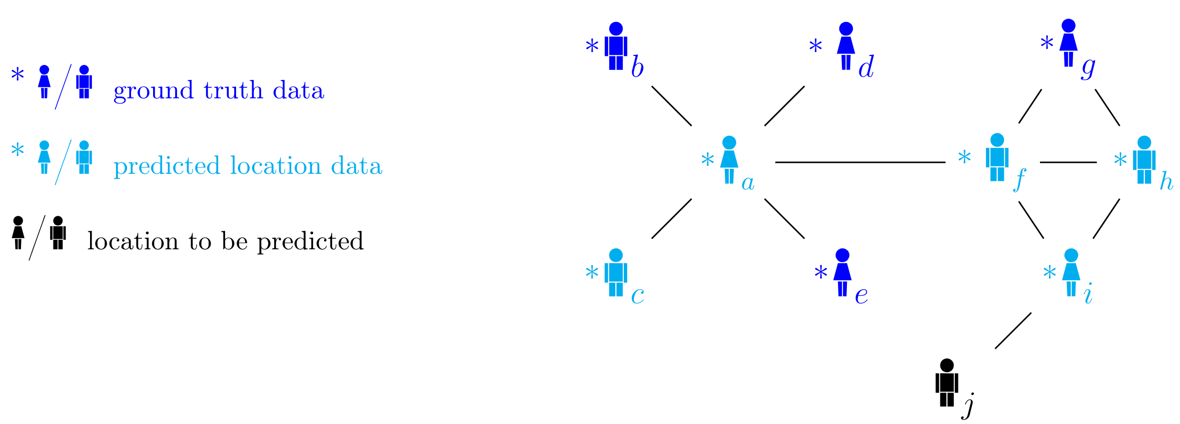

ground truth data as follows. For users with at least one geolocated tweet, I used the coordinates from an arbitrarily selected geolocated tweet. For users without any geolocated tweets but with a location field in their user profile, I looked up the location using the Google Maps and/or Bing Maps APIs (the specific API is selected arbitrarily).[3] If there was a match, I used the coordinates corresponding to this location as the user's ground truth location.

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 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. 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 θ, 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 ($N = 1,676,419$). 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

[3] Google Maps API: `https://developers.google.com/places/web-service/details`; Bing Maps API: `https://msdn.microsoft.com/en-us/library/ff701711.aspx`.

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 A4 shows the performance for the ISIS activists' accounts. Figure A5 shows the performance for the ISIS followers' accounts. 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. To account for prediction errors, I use weights that reflect the stability of the location prediction for each user in all regression models in the article.

Comparing the ISIS Sample with a Random Twitter sample

One might worry that predicting locations with the algorithm described above may not be suited for ISIS networks, as individuals in these networks are likely to be very different from ordinary citizens. While this concern is valid, and is probably true for ISIS activists that disseminate the organization's propaganda, this should not be the case for followers (who comprise over 99% of the sample). The followers are users who follow one or more ISIS activist accounts, and include a range of users, from individuals who actively support the organization, through accounts of interested citizens, to accounts that seek to counter ISIS. This means that ISIS followers are likely to be more similar to ordinary citizens than not.

To test this proposition, I obtained a random sample of Twitter users from the Twitter Streaming API, and compared it to follower and activist accounts. I used various user-level fields to examine the similarity between the samples, including the length of screen names and profile descriptions, the amount of time the accounts have been active on Twitter, whether

Figure A4: 10-Fold out-of-sample stability test (ISIS activists' accounts)**

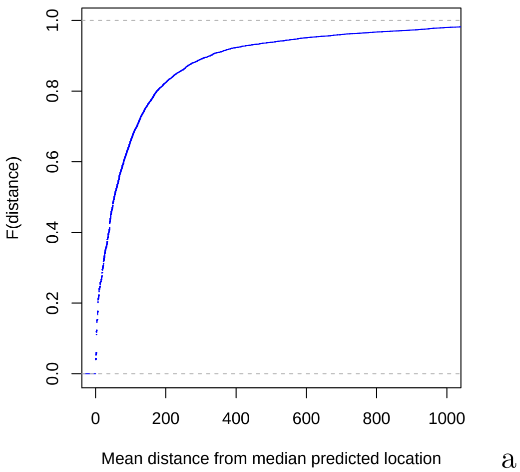

(a) Mean Stability

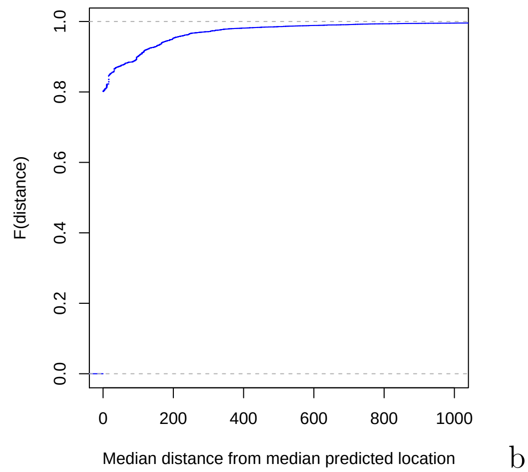

(b) Median Stability

Note: The figure plots the cumulative distribution function of the stability of location predictions of ISIS activists 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. **This figure is taken from Mitts (2019).

the accounts are geo-enabled, the number of friends, followers, and twitter posts, as well as the language used by the users.

Table A7 compares the ISIS followers sample to the random Twitter sample. In most fields, ISIS followers do not significantly differ from random Twitter users: both groups have similar length of screen names, similar network sizes, and are likely to geo-enable their accounts at a similar rate. There are four fields where the samples differ: ISIS followers are more likely to have a shorter profile description, shorter statuses, are more likely to have protected accounts, and more of them have accounts set to Arabic. Overall, however, ISIS followers are not notably different from a random Twitter sample, especially in the most important field – the size of their networks.

Figure A5: 10-Fold out-of-sample stability test (ISIS followers' accounts)**

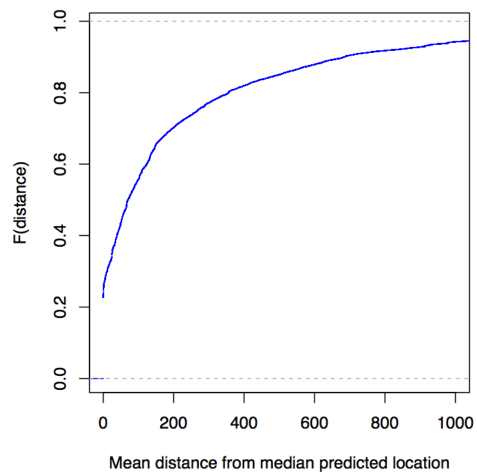

(a) Mean Stability

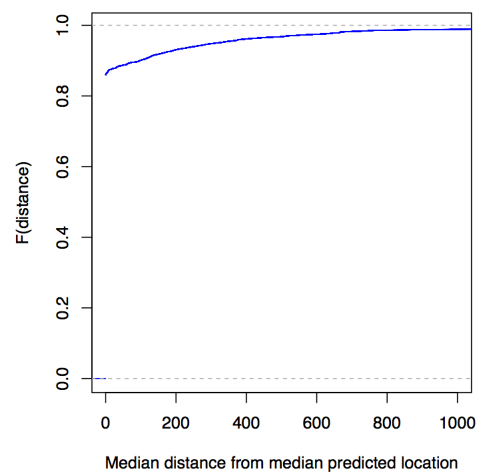

(b) Median Stability

Note: The figure plots the cumulative distribution function of the stability of location predictions of ISIS followers 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. **This figure is taken from Mitts (2019).

Table A7: Balance table: ISIS followers versus a random sample

	Random sample		**ISIS followers sample**			
	Mean	Std. Dev.	Mean	Std. Dev.	Diff.	P-value
Screen name (# characters)	10.38	2.54	10.53	2.78	-0.15	0.57
Description (# characters)	69.65	46.95	39.56	50.14	30.09	0.00***
Geo-enabled	0.34	0.48	0.26	0.44	0.08	0.12
Statuses count	38412.97	84915.98	5785.84	16758.87	32627.13	0.00***
Followers count	3677.96	12579.99	76482.71	1911304.68	-72804.75	0.23
Friends count	1769.17	7254.44	2936.38	21076.87	-1167.21	0.24
Protected	0.00	0.00	0.07	0.26	-0.07	0.00***
Account set to English	0.42	0.50	0.34	0.47	0.08	0.10
Account set to Arabic	0.11	0.31	0.44	0.50	-0.33	0.00***
Account set to French	0.07	0.26	0.07	0.26	-0.00	0.97

Table A8: Balance table: ISIS activists versus a random sample

	Random sample		ISIS activists			
	Mean	Std. Dev.	Mean	Std. Dev.	Diff.	P-value
Screen name (# characters)	10.38	2.54	10.21	2.69	0.17	0.52
Description (# characters)	69.65	46.95	49.15	52.10	20.50	0.00***
Geo-enabled	0.34	0.48	0.41	0.49	-0.07	0.15
Statuses count	38412.97	84915.98	10882.06	28366.96	27530.91	0.00***
Followers count	3677.96	12579.99	11847.67	71547.36	-8169.71	0.00***
Friends count	1769.17	7254.44	3694.59	17415.86	-1925.41	0.04**
Protected	0.00	0.00	0.09	0.29	-0.09	0.00***
Account set to English	0.42	0.50	0.37	0.48	0.05	0.34
Account set to Arabic	0.11	0.31	0.42	0.49	-0.31	0.00 ***
Account set to French	0.07	0.26	0.08	0.27	-0.01	0.77

4 Summary Statistics

In this study, I draw on two complementary Twitter datasets—one on pro-ISIS rhetoric expressed in the tweets of Islamic State sympathizers in the United States, and one on their profile metadata. The rhetoric dataset was derived from tweet text data. As explained in Section 2, I created an index variable capturing pro-ISIS sentiment expressed by each user before and after community engagement events. Table A9 shows summary statistics for this variable for different pre-post time periods.

The second dataset comes from user-level metadata provided by Twitter's public APIs. This dataset includes weekly observations, sampled every seven days from January 1, 2016 to June 24, 2016, of these users' screen names, profile pictures, and the number of propaganda-disseminating accounts that they followed. I created variables that measure changes in these metadata in the week before and 1-4 weeks after community engagement events. For example, when a user changed his or her profile picture in a given week, I coded the *changed profile picture* variable as one for this user in that week. If their profile picture remained the same, I coded this variable as zero. I used the same method to measure changes in screen names. To measure changes in the following of propaganda-disseminating accounts, I counted the number of propaganda-disseminating accounts that each user followed each week. Table A10 shows summary statistics for these variables.

There are two reasons for the difference in the number of observations between the two datasets. First, in the rhetoric dataset, the pro-ISIS index is averaged for each user in each time period (pre/post), while in the user metadata, each user's profile information is observed weekly in each time period. As a result, there are up to four times as many observations in the user metadata. Second, the pro-ISIS rhetoric dataset only includes information on users who tweeted during the week preceding and the 1-4 weeks following community engagement events. Users who did not tweet over these time windows were not included in the rhetoric dataset. By contrast, user-level metadata was observed, every seven days, for every user in all time periods. While these differences affect the number of observations, the proportion of observations before and after community engagement events (*Post*), and the proportion of users inside and outside event areas (*In event area*), are almost identical in both data sources.

Table A9: Summary Statistics, Pro-ISIS Rhetoric

	N	Mean	St. Dev.	Min	Max
	7 day window				
Post	212,891	0.524	0.499	0	1
In event area	212,891	0.029	0.168	0	1
Ave. pro-ISIS rhetoric index	212,891	0.156	0.104	0.002	0.958
	14 day window				
Post	246,007	0.588	0.492	0	1
In event area	246,007	0.030	0.170	0	1
Ave. pro-ISIS rhetoric index	246,007	0.155	0.103	0.001	0.958
	21 day window				
Post	275,949	0.633	0.482	0	1
In event area	275,949	0.030	0.172	0	1
Ave. pro-ISIS rhetoric index	275,949	0.155	0.104	0.001	0.958
	30 day window				
Post	330,256	0.693	0.461	0	1
In event area	330,256	0.030	0.171	0	1
Ave. pro-ISIS rhetoric index	330,256	0.155	0.105	0.001	0.958

Table A10: Summary Statistics, User Metadata

	N	Mean	St. Dev.	Min	Max
7 day window					
Post	428,535	0.504	0.500	0	1
In event area	428,535	0.022	0.146	0	1
Changed profile picture	428,535	0.026	0.160	0	1
Changed screen name	428,535	0.001	0.035	0	1
Number of ISIS accounts following	428,485	4.658	23.624	0	1,802
14 day window					
Post	788,934	0.520	0.500	0	1
In event area	788,934	0.023	0.148	0	1
Changed profile picture	788,934	0.032	0.177	0	1
Changed screen name	788,934	0.002	0.044	0	1
Number of ISIS accounts following	788,839	4.796	24.129	0	1,802
21 day window					
Post	1,120,667	0.519	0.500	0	1
In event area	1,120,667	0.023	0.149	0	1
Changed profile picture	1,120,667	0.029	0.169	0	1
Changed screen name	1,120,667	0.002	0.041	0	1
Number of ISIS accounts following	1,120,538	4.888	24.437	0	1,802
30 day window					
Post	1,473,251	0.544	0.498	0	1
In event area	1,473,251	0.023	0.149	0	1
Changed profile picture	1,473,251	0.031	0.174	0	1
Changed screen name	1,473,251	0.002	0.044	0	1
Number of ISIS accounts following	1,473,081	4.997	24.822	0	1,802

Parallel Trends

This study uses Difference-in-Differences estimations to compare changes in pro-ISIS rhetoric by ISIS followers located in event areas to changes in such rhetoric by followers who were not. The key identifying assumption is that in the absence of a community engagement event, individuals located in the event area and individuals who are not would follow parallel trends in their online expression of pro-ISIS rhetoric.

To empirically test this assumption, I show in the article that the two groups display parallel trends before community engagement events (see Figure 5). Table A11 below presents a statistical test of the parallel trends assumption, showing that there is no difference in the time trends between the groups, except for when expanding the pre-treatment data back to 30 days before the events. To be conservative, I estimate all Difference-in-Differences models with pre-treatment data from 7 days before the events.

Table A11: Parallel Trends in the Pre-Treatment Period

	Dependent variable:			
	Pro-ISIS Content			
	(1)	(2)	(3)	(4)
Days before the event	*7 days*	*14 days*	*21 days*	*30 days*
In event area	0.006	0.006**	0.005***	0.006***
	(0.003)	(0.002)	(0.002)	(0.002)
Days before the event	−0.0003	−0.00005	0.0001	0.0001
	(0.001)	(0.0002)	(0.0001)	(0.0001)
In event area × Days before the event	0.0001	0.0002	0.0002	0.0002**
	(0.001)	(0.0003)	(0.0001)	(0.0001)
Constant	0.153***	0.154***	0.155***	0.155***
	(0.002)	(0.002)	(0.001)	(0.001)
Observations	14	28	42	60
R^2	0.549	0.362	0.383	0.415

Note: Each column represents data from different time windows. Column (1) shows data collected 7 days before the events, column (2) shows data from 14 days, and so on. The coefficient *In event area × Days before the event* reflects the difference in the time trend between individuals located in event areas and those who are not. The table shows that there is no difference in the time trends between the groups, except for when expanding the data back to 30 days before the events. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

5 Results in Tabular Form

Table A12: Pro-ISIS Rhetoric on Twitter (OLS Estimates)

	7 days	14 days	21 days	30 days
Post	0.0003	−0.0004	−0.0004	−0.0001
	(0.0004)	(0.0004)	(0.0004)	(0.0004)
In event area	0.006***	0.006***	0.006***	0.006***
	(0.001)	(0.001)	(0.001)	(0.001)
Post × In event area	−0.006***	−0.005***	−0.005***	−0.005***
	(0.002)	(0.002)	(0.002)	(0.002)
Event fixed effect	✓	✓	✓	✓
Location prediction weights	✓	✓	✓	✓
Content prediction weights	✓	✓	✓	✓
Observations	212,891	246,007	275,949	330,256
R^2	0.003	0.004	0.003	0.003

Note: *p<0.1; **p<0.05; ***p<0.01

Table A13: Changed Profile Picture (OLS Estimates)

	7 days	14 days	21 days	30 days
Post	−0.001***	−0.015***	−0.011***	−0.011***
	(0.0005)	(0.0004)	(0.0003)	(0.0003)
In event area	0.003	−0.005**	−0.002	0.0001
	(0.003)	(0.002)	(0.002)	(0.002)
Post × In event area	0.003	0.015***	0.009***	0.004***
	(0.003)	(0.002)	(0.002)	(0.001)
Event fixed effect	✓	✓	✓	✓
Location prediction weights	✓	✓	✓	✓
Observations	428,535	788,934	1,120,667	1,473,251
R^2	0.009	0.018	0.012	0.011

Note: *p<0.1; **p<0.05; ***p<0.01

Table A14: Changed Screen Name (OLS Estimates)

	7 days	14 days	21 days	30 days
Post	−0.0001	−0.002***	−0.001***	−0.001***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
In event area	−0.00000	−0.001*	−0.001	−0.0003
	(0.001)	(0.0004)	(0.0004)	(0.0004)
Post × In event area	−0.0001	0.001	0.001	0.0005
	(0.001)	(0.001)	(0.0005)	(0.0005)
Event fixed effect	✓	✓	✓	✓
Location prediction weights	✓	✓	✓	✓
Observations	428,535	788,934	1,120,667	1,473,251
R^2	0.0003	0.001	0.001	0.001
Note:		*p<0.1; **p<0.05; ***p<0.01		

Table A15: Propaganda Disseminating Accounts Followed (OLS Estimates)

	7 days	14 days	21 days	30 days
Post	0.038***	−0.142***	−0.273***	−0.394***
	(0.002)	(0.011)	(0.020)	(0.029)
In event area	−0.402	−0.503	−0.559	−0.757**
	(0.334)	(0.349)	(0.357)	(0.312)
Post × In event area	0.013	0.165***	0.278***	0.421***
	(0.012)	(0.026)	(0.037)	(0.086)
Event fixed effect	✓	✓	✓	✓
Location prediction weights	✓	✓	✓	✓
Observations	428,485	788,839	1,120,538	1,473,081
R^2	0.004	0.004	0.004	0.003
Note:		*p<0.1; **p<0.05; ***p<0.01		

Table A16: Engaging in Several Online Actions Simultaneously

	Rhetoric + screen name	Rhetoric + profile picture	Rhetoric + ISIS accounts following	Three or more
Post	−0.001	−0.002	−0.006	−0.001
	(0.001)	(0.002)	(0.009)	(0.001)
In event area	−0.009***	−0.101***	−0.244***	−0.012***
	(0.001)	(0.004)	(0.011)	(0.001)
Post × In event area	0.031*	0.149***	0.294***	0.021***
	(0.018)	(0.042)	(0.099)	(0.008)
Event fixed effects	✓	✓	✓	✓
Location prediction weights	✓	✓	✓	✓
Observations	51,030	51,635	54,708	116,064
R^2	0.003	0.034	0.051	0.004

Note: *p<0.1; **p<0.05; ***p<0.01

The table presents Difference-in-Differences coefficients from regressions where the dependent variables (reported in the columns) reflect different combinations of online actions. The table shows that many users in event areas took two or more actions after community engagement events, and some even engaged in three or more actions.

Table A17: Mentions of Telegram after Counter-Extremism Events (Tweet-Level Data)

	7 days	14 days	21 days	30 days
Post	0.00003	0.00002	0.00005	0.0001**
	(0.00005)	(0.00003)	(0.00003)	(0.00003)
In event area	−0.0001	−0.0001	−0.0001	−0.0001
	(0.0001)	(0.0001)	(0.0002)	(0.0002)
Post × In event area	0.00003	0.00004	0.001***	0.001***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Event fixed effect	✓	✓	✓	✓
Observations	265,649	446,125	526,568	656,934
R^2	0.00002	0.00002	0.0001	0.0001

Note: *p<0.1; **p<0.05; ***p<0.01

Table A18: Telegram Channels Launch and the Number of Tweets

	7 days	14 days	21 days	30 days
In event area	0.534	1.856***	2.220***	1.307***
	(0.421)	(0.517)	(0.521)	(0.461)
Telegram	1.779***	3.055***	4.268***	5.338***
	(0.135)	(0.197)	(0.258)	(0.305)
In event area × Telegram	−0.155	−1.156*	−1.342*	−0.308
	(0.527)	(0.643)	(0.686)	(0.675)
Observations	109,955	145,243	175,587	235,594
R^2	0.518	0.497	0.509	0.503
Note:		*p<0.1; **p<0.05; ***p<0.01		

The table reports coefficients estimated from a Difference-in-Difference analysis of the number of tweets posted by ISIS sympathizers after Telegram's channels launch in time periods following community engagement activities.

6 Robustness Tests

In this section, I examine the robustness of the results by running several additional estimations. First, I show that the results are almost identical when limiting the analysis only to individuals located in event areas or when using propensity score matching on pre-treatment data. Second, I show very similar patterns when examining tweet content with a structural topic model instead of the pro-ISIS index. Third, I show that the results are robust to several over-time and cross-sectional placebo tests. Fourth, I show that the same results hold with alternative modeling choices.

6.1 Alternative Comparisons

Since the number of locations with counter-extremism events was much smaller than the ones with no such events, one might worry that the results are driven by unobserved characteristics of locations with counter-extremism events. To address this concern, I first subsampled the data to only include cities that had at least one community engagement event, and re-ran the analysis for this subset of cities. Table A19 shows that the results are very similar when limiting the analysis only to cities with counter-extremism events. Here, the 'treatment' group consists of users who were located in a city that held a community engagement event, and the 'control' group consists of users who were located in other cities in which these events took place at other time periods. In a second test, I matched users in treatment and control areas on the basis of their pro-ISIS rhetoric in the pre-treatment period. Table A20 shows that with matching, the results remain the same, albeit the magnitude of the coefficients is slightly attenuated.

Table A19: Pro-ISIS Rhetoric on Twitter (Only Users in Event Areas)

	7 days	14 days	21 days	30 days
Post	-0.001^{*}	-0.001^{*}	-0.001^{*}	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)
In event area	0.003^{**}	0.003^{**}	0.003^{**}	0.003^{**}
	(0.001)	(0.001)	(0.001)	(0.001)
Post × In event area	-0.005^{**}	-0.004^{**}	-0.004^{***}	-0.005^{***}
	(0.002)	(0.002)	(0.002)	(0.002)
Event fixed effects	✓	✓	✓	✓
Location prediction weights	✓	✓	✓	✓
Content prediction weights	✓	✓	✓	✓
Observations	87,715	102,035	114,505	136,728
R^2	0.003	0.004	0.005	0.004
Note:				$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Table A20: Pro-ISIS Rhetoric on Twitter (Propensity Score Matching)

	7 days	14 days	21 days	30 days
Post	0.001^{***}	0.001^{***}	0.001^{***}	0.002^{***}
	(0.0003)	(0.0003)	(0.0003)	(0.0003)
In event area	0.006^{***}	0.006^{***}	0.006^{***}	0.006^{***}
	(0.001)	(0.001)	(0.001)	(0.001)
Post × In event area	-0.004^{**}	-0.003^{**}	-0.003^{**}	-0.003^{*}
	(0.002)	(0.002)	(0.002)	(0.002)
Event fixed effects	✓	✓	✓	✓
Location prediction weights	✓	✓	✓	✓
Content prediction weights	✓	✓	✓	✓
Observations	165,077	170,410	172,792	174,889
Note:				$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

6.2 Placebo Tests

Next, I present several placebo tests to examine the robustness of the paper's findings. Panel A in Figure A6 shows the results when estimating model (1) with a placebo treatment group. For each event, I re-coded a random sample of users located outside event areas as if they were located in the area of the event. I then re-estimated the model with these users as the treatment group while dropping the users who were actually located in event areas. As can be seen in the figure, there is no systematic relationship between community engagement events and the pro-ISIS rhetoric of users in the placebo treatment group. In addition, I ran a placebo test with pre-event time periods as the 'treatment.' Excluding data from the post-event period, I re-estimated model (1) when setting the placebo treatment to take place two weeks before the actual event. Panel B in Figure A6 shows that there is no systematic relationship between the placebo event dates and pro-ISIS rhetoric.

Figure A6: Placebo Tests

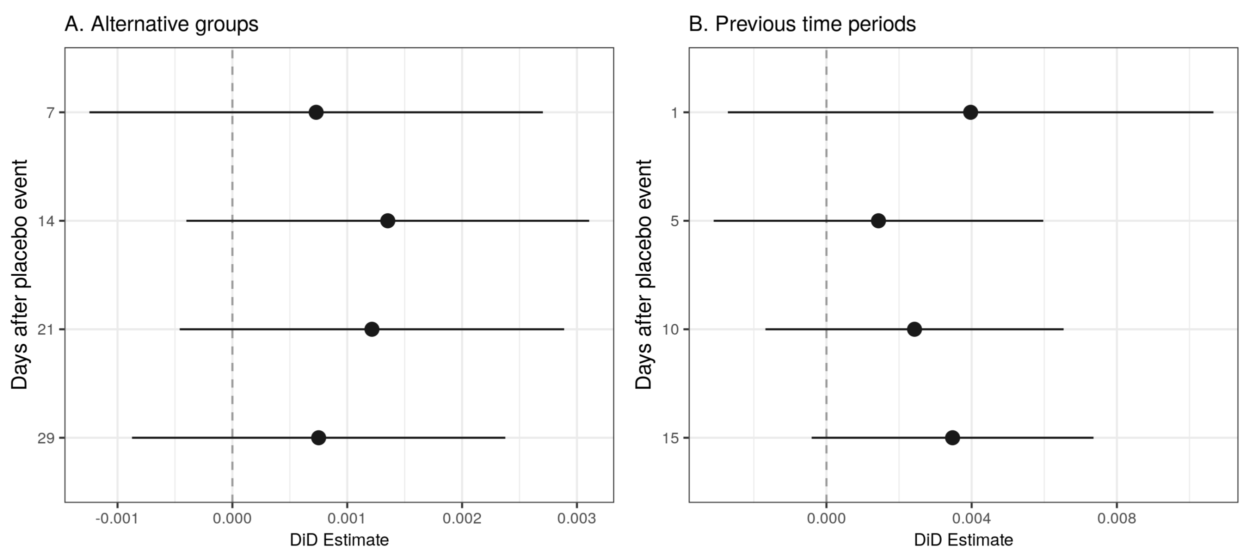

Note: Panel A presents the results with a placebo treatment group. Panel B shows a placebo test with pre-event time periods as the 'treatment,' where I excluded data from the post-event period and set the placebo treatment to take place two weeks before the actual event.

6.3 Additional Specifications

In what follows, I describe additional specifications to examine the sensitivity of the results to alternative modeling choices. Figure A7 shows the results when using an alternative proximity measure. Instead of coding users as being in event areas if their geo-location falls within the geographic boundaries of U.S. towns, I measured each user's distance in kilometers from the center of the town in which community engagement events took place. The results remain the same for users who were located approximately 5-50 kilometers from community engagement activities.

Another concern is that all estimations in the article use ordinary least squares regressions, even for binary and count outcomes. Tables A21 and A22 show that the results for changing profile pictures and screen names (binary dependent variables) hold when using logit regressions instead of OLS. Table A23 shows that the results for the number of ISIS accounts followed (a count variable) remain the same when using a Poisson regression instead of OLS.

Figure A7: DiD Estimates By Distance From Community Engagement Events

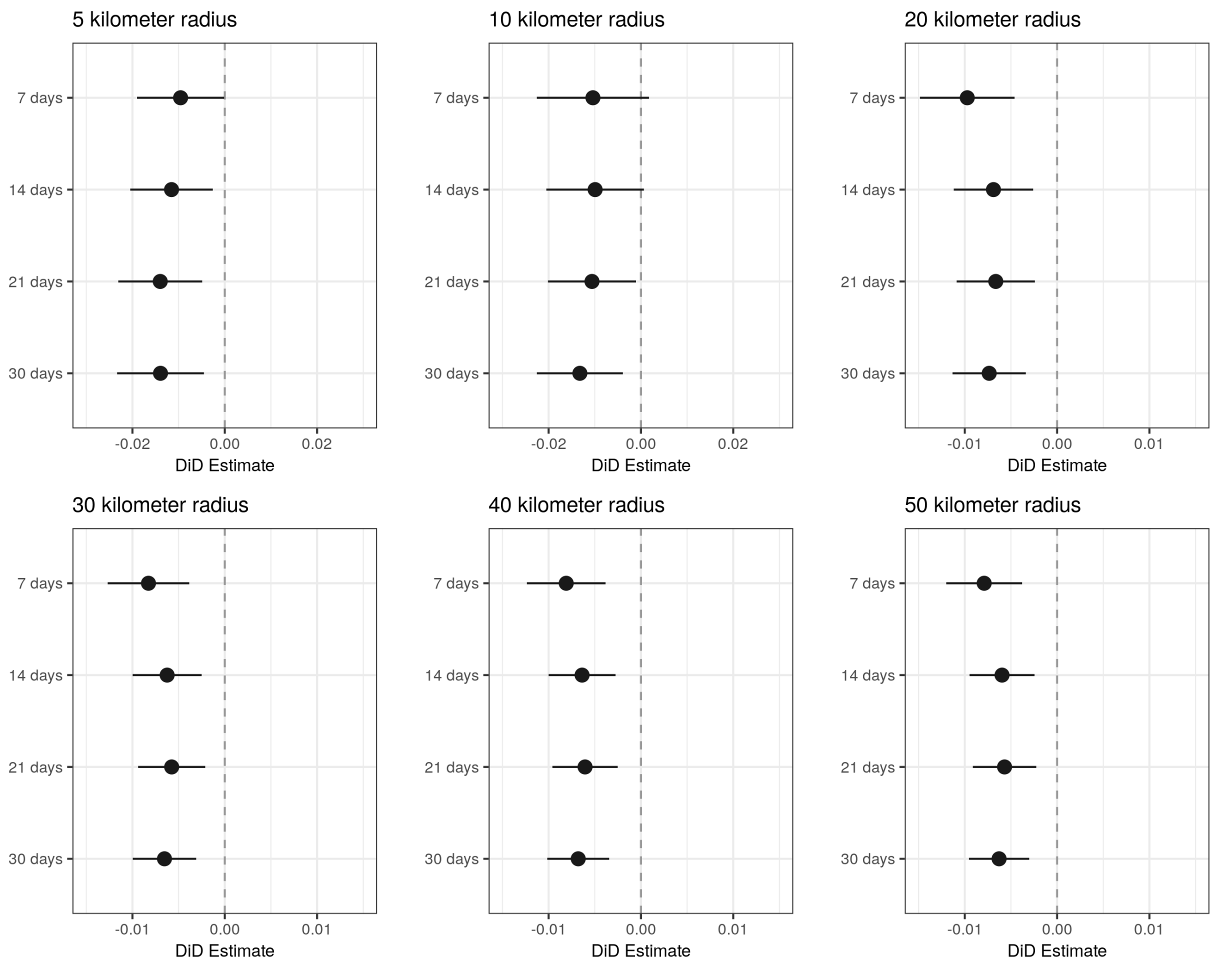

Note: The figure reports coefficients estimated from a pooled Difference-in-Differences analysis of the relationship between 78 community engagement events and pro-ISIS rhetoric on Twitter, captured a week before and up to a month after each event. Uses are considered 'treated' with the event if they are located 5-50 kilometers from the center of the town in which a community engagement event took place.

Table A21: Changed Screen Name (Logit Estimates)

	7 days	14 days	21 days	30 days
Post	−0.076	−0.829***	−0.782***	−0.654***
	(0.088)	(0.056)	(0.050)	(0.040)
In event area	−0.001	−0.351	−0.286	−0.087
	(0.415)	(0.271)	(0.227)	(0.182)
Post × In event area	−0.109	0.348	0.350	0.205
	(0.612)	(0.418)	(0.345)	(0.266)
Constant	−7.754***	−6.605***	−6.416***	−6.539***
	(0.356)	(0.168)	(0.126)	(0.114)
Event fixed effect	✓	✓	✓	✓
Location prediction stability weights	✓	✓	✓	✓
Observations	428,535	788,934	1,120,667	1,473,251
Log Likelihood	−3,991.010	−10,429.280	−13,576.420	−19,927.370
AIC	8,020.020	20,896.570	27,190.840	39,892.750

Note: *p<0.1; **p<0.05; ***p<0.01

Table A22: Changed Profile Picture (Logit Estimates)

	7 days	14 days	21 days	30 days
Post	−0.047**	−0.487***	−0.404***	−0.361***
	(0.020)	(0.013)	(0.012)	(0.010)
In event area	0.123	−0.136**	−0.035	0.036
	(0.093)	(0.067)	(0.056)	(0.048)
Post × In event area	0.133	0.486***	0.309***	0.129*
	(0.129)	(0.092)	(0.078)	(0.068)
Constant	−5.113***	−4.643***	−4.367***	−4.379***
	(0.095)	(0.059)	(0.042)	(0.037)
Event fixed effect	✓	✓	✓	✓
Location prediction stability weights	✓	✓	✓	✓
Observations	428,535	788,934	1,120,667	1,473,251
Log Likelihood	−48,607.210	−102,888.400	−138,425.100	−191,352.100
AIC	97,252.410	205,814.700	276,888.200	382,742.100

Note: *p<0.1; **p<0.05; ***p<0.01

Table A23: Propaganda Disseminating Accounts Followed (Poisson Estimates)

	7 days	14 days	21 days	30 days
Post	0.008***	−0.029***	−0.055***	−0.078***
	(0.001)	(0.001)	(0.001)	(0.001)
In event area	−0.079***	−0.099***	−0.109***	−0.149***
	(0.007)	(0.005)	(0.004)	(0.004)
Post × In event area	0.002	0.034***	0.057***	0.088***
	(0.010)	(0.007)	(0.006)	(0.005)
Constant	1.792***	1.831***	1.865***	1.886***
	(0.003)	(0.002)	(0.002)	(0.002)
Event fixed effect	✓	✓	✓	✓
Location prediction stability weights	✓	✓	✓	✓
Observations	428,485	788,839	1,120,538	1,473,081
Log Likelihood	−2,891,311.000	−5,454,973.000	−7,858,651.000	−10,519,880.000
AIC	5,782,660.000	10,909,985.000	15,717,341.000	21,039,798.000

Note: *p<0.1; **p<0.05; ***p<0.01

References

Bjelopera, Jerome P. 2014. Countering violent extremism in the United States. Congressional Research Service, Library of Congress.
URL: *https://fas.org/sgp/crs/homesec/R42553.pdf*

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/*

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.

Mitts, Tamar. 2019. "From isolation to radicalization: anti-Muslim hostility and support for ISIS in the West." *American Political Science Review* 113(1):173–194.

Nielsen, Rich. 2017. "arabicStemR: Arabic Stemmer for Text Analysis.".
URL: *https://cran.r-project.org/web/packages/arabicStemR/index.html*

Schlanger, Margo. 2010. *Office for Civil Rights and Civil Liberties (CRCL) Newsletter*. Vol. 1 U.S. Deparment of Homeland Security.
URL: *http://tinyurl.com/y227dpt3*