American Political Science Review (2019) 113, 1, 173–194
doi:10.1017/S0003055418000618

From Isolation to Radicalization: Anti-Muslim Hostility and Support for ISIS in the West

TAMAR MITTS *Columbia University*

What explains online radicalization and support for ISIS in the West? Over the past few years, thousands of individuals have radicalized by consuming extremist content online, many of whom eventually traveled overseas to join the Islamic State. This study examines whether anti-Muslim hostility might drive pro-ISIS radicalization in Western Europe. Using new geo-referenced data on the online behavior of thousands of Islamic State sympathizers in France, the United Kingdom, Germany, and Belgium, I study whether the intensity of anti-Muslim hostility at the local level is linked to pro-ISIS radicalization on Twitter. The results show that local-level measures of anti-Muslim animosity correlate significantly and substantively with indicators of online radicalization, including posting tweets sympathizing with ISIS, describing life in ISIS-controlled territories, and discussing foreign fighters. High-frequency data surrounding events that stir support for ISIS—terrorist attacks, propaganda releases, and anti-Muslim protests—show the same pattern.

INTRODUCTION

Between 2011 and 2016, about 30,000 foreign fighters traveled to Syria and Iraq to join the Islamic State (Benmelech and Klor 2018). Fighters came to ISIS from all over the world, many from Western countries like France, Britain, Belgium, Germany, and the United States. A large number of Western recruits were radicalized online by consuming extremist content on the Internet and social media (Carter, Maher, and Neumann 2014; Vidino and Hughes 2015). Online radicalization was not limited to certain social groups or those with national grievances; rather, recruits came from different backgrounds, age groups, education, and income levels (Greenberg 2016). Why did so many Westerners come to support groups like the Islamic State? How could one organization attract so many individuals to a conflict not their own?

This study brings together research on violent extremism and radicalization, along with the literature on immigration in the West, to examine how anti-Muslim sentiment is linked to radicalization and support for the Islamic State in Western European countries. I argue that hostility toward Muslims in the West can lead individuals to seek comfort and acceptance elsewhere, making radical messages promulgated by foreign rebels seem attractive. A large body of research on immigration to the West studies factors that facilitate or inhibit immigrant integration, with a particular focus on economic outcomes (Dancygier and Laitin 2014). This literature emphasizes the powerful role that natives' attitudes play in this context, and points to cultural, economic, and psychological factors that determine natives' acceptance, or lack of acceptance, of immigrants in social and economic settings (Hainmueller and Hopkins 2014).

A recent strand of this important body of work has focused on discrimination against Muslim immigrants in particular, empirically documenting the central role of anti-Muslim discrimination in facilitating Muslims' lack of integration. In France, for example, Adida, Laitin, and Valfort (2016) found that Muslims and non-Muslims are often caught in a vicious cycle in which the latter discriminate against the former, falsely equating "Muslim" and "Jihadist," and Muslims, in turn, tend to distrust non-Muslims and withdraw from French society, thus perpetuating their nonintegration. But this body of research has yet to examine other outcomes of discrimination. Focusing primarily on social and economic integration, it has not systematically considered how native attitudes toward immigrants might increase the likelihood of jihadi radicalization.

One of the most distinctive aspects of the Islamic State's recruitment strategies is its extensive use of social media. The organization has not only been distributing provocative content to general audiences on the Internet, it has also been using social networks on *Twitter*, *Facebook*, and related platforms to attract new members from all over the world. Twitter has been particularly popular, as it enabled fast and large-scale public dissemination of content. Studies documenting the usage of Twitter by Western foreign fighters have noted that it played a central role in their radicalization process by intensifying their mental and emotional connection to war events on the ground (Carter, Maher, and Neumann 2014). Potential recruits found it appealing to connect to the organization through Twitter, as the platform enabled the anonymous consumption of radical and extremist ideas, without being exposed to the risk of physically interacting with a recruiter (Berger 2015). In fact, the organization's online radicalization operation has been so

Tamar Mitts, Assistant Professor, School of International and Public Affairs, Columbia University, tm2630@columbia.edu.

I thank Susanne Baltes, Eli Berman, Abhit Bhandari, Christopher Blattman, Jasper Cooper, Lindsay Dolan, Page Fortna, Grant Gordon, Guy Grossman, Kolby Hanson, Macartan Humphreys, Sarah Khan, Sung Eun Kim, Egor Lazarev, Summer Lindsey, Joshua Mitts, Suresh Naidu, Jacob Shapiro, Kunaal Sharma, Jack Snyder, Camille Strauss-Kahn, Lauren Young, and the Journal's editor and anonymous reviewers for their helpful feedback. I also benefited from comments from seminar participants at Brown, Chicago, Columbia, Dartmouth, Emory, Michigan, Northwestern, NYU, Stanford, and Vanderbilt. I thank Geffen Huberman for her research assistance. Replication files are available at the American Political Science Review Dataverse: https://doi.org/10.7910/DVN/5QUCW7.

Received: April 19, 2017; revised: January 5, 2018; accepted: August 13, 2018. First published online: November 26, 2018.

vast and extensive that many security agencies found it challenging to keep track of every aspect of these activities (McCaul et al. 2015).

In this study, I take advantage of the presence of this widespread radicalization in the West, and the availability of large amounts of public Twitter data, to examine whether anti-Muslim hostility is linked to support for ISIS in Europe.[1] Using an original method described in the body of the article, I collected granular data on the social media activity of about 15,000 accounts of ISIS activists, as well as the full social network of their followers across the world ($N \approx 1.6$ million). I monitored the online behavior of ISIS activists and their followers in real time, capturing their activity prior to account suspension, and recorded textual and image content, which I use for analysis.

Using computer science methods to predict the physical geographic location of Twitter users, I matched user-level data to local-level administrative data from the four European countries with the highest share of Western foreign fighters: France, the United Kingdom, Germany, and Belgium (Barrett et al. 2015). I collected data on levels of unemployment, the share of immigrants and asylum seekers in each locality, and local-level vote share for far-right, anti-Muslim parties in recent elections across Europe. As voting for far-right parties strongly correlates with anti-Muslim sentiment,[2] I use vote share for these parties as a local-level measure of anti-Muslim hostility, examining whether it predicts support for ISIS on social media.

I developed several measures of online radicalization and support for ISIS on Twitter. Using supervised machine learning, I classified millions of tweets in English, Arabic, French, and German along various dimensions of ISIS support. These include expressing sympathy with ISIS, tweeting about the life of fighters in ISIS-controlled territories, and expressing an interest in traveling to Syria or becoming foreign fighters. In addition, I classified tweets as containing anti-West rhetoric to examine how Western ISIS sympathizers might refer to their own countries. I kept track of which users were flagged as ISIS activists by several hacktivist groups, and also noted when they were suspended from Twitter.

The results show that local-level vote share for far-right, anti-Muslim parties in France, the United Kingdom, Germany, and Belgium correlates significantly with online radicalization. In substantive terms, an increase of one percentage point in the local-level vote share for far-right parties is associated with a 6% and 5% increase, respectively, in the probability of a user being flagged as ISIS-affiliated and being among the top 1% posters of radical content. A one percentage-point increase in the right-wing vote share is associated with an average increase of up to 10,000 pro-ISIS tweets across the entire sample, including tweets sympathizing with ISIS, discussing life in ISIS territories, and expressing interest in foreign fighters and travel to Syria.

[1] Focusing on Twitter sheds light on the public behavior of Islamic State sympathizers on social media. I leave for future research the study of television, other websites or encrypted social media.

[2] See more information in the body of the article.

As the relationship between pro-ISIS radicalization and support for far-right parties is complex and may also run in the other direction or be driven by omitted variables, I run several additional tests. First, I take advantage of the high-frequency nature of Twitter data and examine whether events that likely spur sympathy with ISIS among potential recruits, such as terrorist attacks, propaganda dissemination events, and anti-Muslim protests, are immediately followed by increased posting of pro-ISIS content, especially in areas with high far-right support. Second, I examine whether the results might be driven by the local presence of minority populations. In analyses with data available only in the United Kingdom, I include covariates for the proportion of Muslims, Arabs, Pakistanis, Bangladeshis, and foreign-born in each local area. After controlling for these covariates—many of which are negatively or not correlated with radicalization measures—I find that vote share for the far-right remains strongly positively associated with posting pro-ISIS content on Twitter.

RADICALIZATION

Why do individuals living in Western countries begin to support groups like the Islamic State? What attracts people to ISIS's extremist ideology? A large literature has sought to explain the causes of radicalization and violent extremism, especially in the context of militant jihad. Most agree that radicalization involves a change in ideology or beliefs that support indiscriminate violence against civilians for political reasons, or a group that represents this ideology and actions (Borum 2011; Sedgwick 2010; Wilner and Dubouloz 2010). Scholars view radicalization as a process that occurs over time, in which a person becomes increasingly committed to extreme and violent worldviews.[3]

While models of radicalization vary, most involve the following stages. First, an individual begins to find extremist ideology appealing by interacting with others who have radicalized, or by exploring extremist content on the Internet—a phenomenon that has been more frequent in recent years (Walter 2017). In the second stage, the person becomes increasingly committed to the ideology and begins to vocally express radical sympathies or take actions to show affiliation with the cause (Borum 2011; Wiktorowicz 2005). Finally, an individual might take violent actions, though some argue that radicalization need not culminate in violence (Neumann 2013; della Porta 2018).

In this study, I focus on the second stage, examining the online behavior of individuals who have already expressed interest in the Islamic State by choosing to follow ISIS accounts. My study does not consider what makes an individual begin to find extremist ideology appealing (the first stage), nor do I examine what tips an individual toward violence (the third stage). My focus is on expressions of support for the Islamic State and the ideology that it promotes, among those who may already

[3] Following Walter (2017), I define an extreme ideology as one that significantly deviates from the ideology held by the majority of the population which the group seeks to represent and/or control.

be embarking down the path of radicalization, as evidence that the radicalization process is progressing.

Researchers seeking to explain radicalization and the rising global appeal of extremist groups have emphasized the role of ideology.[4] Walter (2017) showed that in recent civil wars, rebel groups upholding extreme ideologies have been more successful than moderate groups in attracting supporters, a phenomenon she attributes to the strategic use of ideology by groups in competitive environments. Hegghammer (2010) argued that the rise in the number of Muslim foreign fighters since the 1980s is driven by the emergence of a new ideological movement, "populist pan-Islamism," that paints the world as a place that threatens the existence of Muslims worldwide. Empirically, Malet (2013) found that armed groups portraying conflicts as posing an alarming danger to both local rebels and foreign supporters have historically attracted most foreign fighters.

The micro-level implication of these theories is that individuals are more likely to support extremist groups if they are receptive to radical ideologies. But what makes someone receptive to extremism in the first place? Why do people located thousands of miles away from the location of civil wars pay attention to violent propaganda promoted in these wars? Over four decades of scholarship on terrorism has ruled out the notion that personality types explain one's propensity for extremism (Borum 2011). Instead, scholars point to structural factors, such as social, economic, or political grievances (Bass 2014; Lyons-Padilla et al. 2015), the powerful role of social networks (Dalgaard-Nielsen 2010; Wiktorowicz 2005; Mousseau 2011), and thrill and identity-seeking (Bass 2014; Bayman and Shapiro 2014; Nussio 2017) as explanations for radicalization.[5]

This article focuses on a slightly different explanation, arguing that experiences of social isolation can exacerbate a process of radicalization. I contend that in Western countries, local and personal exposure to anti-Muslim hostility can increase individuals' attraction to extremist jihadi ideologies. Prior research on jihadi radicalization in the West has shown, using case studies, that experiences of discrimination led individuals to radicalize (Borum 2011; Wiktorowicz 2005; Wilner and Dubouloz 2010). While not focusing on radicalization as an outcome, related work on the impact of anti-Muslim discrimination has shown that it tends to inhibit integration and assimilation (Adida, Laitin, and Valfort 2016; Bryan 2005; Gould and Klor 2016). Indeed, recent evidence from the United States suggests that failed integration of Muslim immigrants can increase support for violent extremism (Lyons-Padilla et al. 2015).

I argue that groups like the Islamic State seek to attract isolated individuals in the West, by providing an alternative 'virtual community' on the Internet and social media. A large number of people who radicalized and joined ISIS from Western countries began embracing the organization's ideology when searching for belonging and identity (Shane, Apuzzo, and Schmitt 2015; Vidino and Hughes 2015). Hoda Muthana, for example, an American student from Alabama, was radicalized on social media after opening a secret Twitter account without her parents' knowledge. After interacting with ISIS supporters on Twitter, she adopted radical interpretations of Islam and eventually traveled to Syria to join the organization (Hall 2015). Ali Shukri Amin, an American teenager from Virginia, found solace from his troubled life in the virtual communities of ISIS activists on Twitter. In the end, Amin disconnected from his family and friends, spread ISIS propaganda to thousands of followers online, and recruited one of his friends to travel to Syria to become a foreign fighter (Robinson 2015; Shane, Apuzzo, and Schmitt 2015).

Indeed, evidence from the United States shows that among over a hundred individuals charged with providing material support for ISIS or plotting a violent attack on the organization's behalf, about 62% used social media when they were radicalizing, and among those, 86% expressed their support for ISIS in publicly viewable posts.[6] The Internet and social media seem to play a central role in exposing Western individuals to paths of radicalization. These findings are consistent with research on the social media usage of European foreign fighters, which shows that online social networks played a dominant role in fighters' radicalization process (Carter, Maher, and Neumann 2014). However, our knowledge of the online behavior of ISIS supporters and its relation to real-world events is currently very limited. In this article, I examine whether online indicators of pro-ISIS radicalization are stronger for individuals experiencing greater levels of anti-Muslim hostility.

ANTI-MUSLIM HOSTILITY AND SUPPORT FOR FAR-RIGHT PARTIES

Animosity against Muslims in the West has been rising in recent years, especially after 9/11 (Burrows 2016; Jamal 2008; Karam 2012; Naber 2008; Stack 2015). Examples include setting fire to mosques, spreading anti-Muslim graffiti, and physically attacking individuals who practice Islam. Take the case of Ms. Khola Hasan, an Islamic scholar from the U.K.'s Epping Forest region, who has been targeted by anti-Muslim violence multiple times in recent years. In an interview with *The Guardian*, she said, "I was walking down Epping High Street and a man shouted at me 'You bloody ISIS supporter.' Another time… someone stopped their car and threw an empty glass bottle at me. I was absolutely terrified" (Flaig 2016).

[4] In the global recruitment context, researchers have favored the role of ideology over other explanations, such as organizational capacity or material resources, since the latter explanations are less likely to account for the motives of foreign fighters and individuals who radicalize outside of civil wars' territories (Hegghammer 2010; Malet 2013).

[5] A broader literature on conflict participation has similarly examined the causes of mobilization into violence (Gurr 1970; Horowitz 1985; Petersen 2001; Scacco 2018; Wood 2003). Theories in this stream of work mirror many explanations posed by the radicalization literature. For a summary of the broader conflict participation literature, and how it might apply to individual-level mobilization, see Humphreys and Weinstein (2008).

[6] See full details in section S5 in the online appendix.

TABLE 1. Far-Right Support and Anti-Muslim Attitudes in Europe

	(1) Do not allow Muslims in country	(2) Disapprove immigration of different race/ethnic groups	(3) Disapprove relative marrying someone from a minority race/ethnic group	(4) Do not want a boss from a minority race/ethnic group	(5) Immigrants make crime worse
A. Far-right self placement					
Far-right self placement	0.12***	0.37***	0.99***	0.39	0.41**
	(0.03)	(0.07)	(0.27)	(0.24)	(0.18)
Constant	0.05	2.14***	2.09***	1.58***	6.78***
	(0.04)	(0.10)	(0.37)	(0.32)	(0.28)
Demographic controls	✓	✓	✓	✓	✓
R^2	0.054	0.075	0.068	0.075	0.023
Observations	3,850	3,874	3,894	3,867	3,837
B. Far-right voting					
Voted for far-right party	0.26***	0.65***	1.91***	1.49***	1.23***
	(0.05)	(0.09)	(0.34)	(0.35)	(0.24)
Constant	0.06	2.15***	2.12***	1.56***	6.77***
	(0.04)	(0.09)	(0.37)	(0.32)	(0.28)
Demographic controls	✓	✓	✓	✓	✓
R^2	0.070	0.085	0.076	0.084	0.033
Observations	3,850	3,874	3,894	3,867	3,837

Standard errors in parentheses.
*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.
Note: The table reports the correlations between voting for far-right parties in France, Belgium, Germany, and the United Kingdom and holding anti-Muslim and anti-immigrant attitudes. The *Far-right self placement* variable is an indicator coded one for individuals who identify as '10' (farthest on the right) on a 1–10 scale of left-right placement. *Voted for far-right party* is an indicator variable coded one for individuals who voted for Front National (FN) in France, United Kingdom Independence Party (UKIP) in the United Kingdom, National Democratic Party of Germany (NPD) and Alternative for Germany (AfD) in Germany, and Vlaams Belang (VB) in Belgium. The table presents estimates from ordinary least squares regressions of the outcome variables reported in columns (1) through (5) on indicators of support for far-right parties, controlling for being native-born, education, income, gender, age, and religion. Data source: European Social Survey Round 7 (2014).

Epping Forest is among the constituencies with the highest vote share for far-right parties in the United Kingdom. In the 2015 general elections, over 18% of its voters voted for far-right parties, putting the locality at the top 10% of far-right vote share in the country. A similar pattern is observed in other European localities with high far-right support. In Dartford, U.K., right-wing activists launched an "anti-halal operation," targeting Muslim restaurants selling halal food with the claim that they support terrorism by paying a zakat religious tax (Kent Online 2015). In Provins, France, where vote share for the Front National party in the 2015 Departmental Elections was above 37%, a local mosque was desecrated with anti-Muslim graffiti (Inge 2013).

Indeed, far-right parties are one of the most prominent mobilizers of anti-Muslim sentiment in contemporary Europe. A common theme in the platforms of these parties is support for exclusionary, "nativist" populism that combines nationalism and xenophobia, seeking to ostracize groups with certain cultural, religious, or ethnic characteristics (Golder 2016). For example, France's Front National party has long blamed Muslim immigrants for many of the country's social problems, ranging from unemployment to security and national unity (Adida, Laitin, and Valfort 2016; Front National 2017). The Alternative for Germany (AfD) party, who in the 2017 German parliamentary elections gained an unprecedented share of the votes, mobilized support with an anti-Muslim xenophobic campaign (Wildman 2017).

Several scholars have suggested that far-right voting is strongly linked to anti-Muslim sentiment (Lubbers and Scheepers 2002; Norris 2005; Rydgren 2008). Using data from the European Social Survey Round 7, I tested the relationship between far-right voting and anti-Muslim attitudes in Europe. Table 1 shows that there is a strong correlation between holding anti-Muslim and anti-immigrant attitudes and self-identifying as a far-right supporter (Panel A) or voting for far-right parties (Panel B). The regressions control for demographic variables that might also explain anti-Muslim attitudes, such as being native-born, education, income, gender, age, and religious beliefs.

While the overall popularity of far-right, anti-Muslim parties in Europe has increased nationally, support for these parties still varies significantly at the local level.[7] I argue that areas with higher levels of far-right support are likely to provide a fertile ground for jihadi-inspired extremism, as

[7] See Figure S18 in the online appendix for local-level variation in the vote share for far-right parties in France, Germany, and the United Kingdom.

they present a more hostile environment for individuals at risk of radicalization. If anti-Muslim hostility has any role in attracting Westerners to the Islamic State, then we should observe more signs of radicalization among individuals located in areas where far-right parties are popular.

It is certainly possible that the relationship also runs in the opposite direction: that the presence of pro-ISIS radicals in a given locality drives support for far-right parties. However, there are strong reasons to believe that the correlation does not run only in that particular direction. Research on radicalization in the United Kingdom over the last decade has found that far-right and jihadi extremists frequently feed each other in a vicious cycle of "cumulative extremism" (Bartlett and Birdwell 2013; Eatwell 2006). In addition, the radicalization stories of Western Islamic State recruits and others who have become supportive of jihadi-inspired terrorism illustrate the powerful impact of xenophobic hostility and discrimination on people's support for violence (The Atlantic 2017; Victoroff, Adelman, and Matthews 2012; Wiktorowicz 2005).

This study systematically examines, for the first time with large scale data across thousands of locations in four countries, the local-level relationship between far-right support and pro-ISIS radicalization. In the following section, I explain how I created measures for online support for ISIS by identifying and observing in real-time the content and social media activity of individuals at risk of radicalization.

DATA

Identifying ISIS Activist and Follower Accounts on Twitter

As discussed previously, this study focuses on individuals in the second stage of radicalization: conditional on having shown some sign of interest in the Islamic State, to what extent does local-level hostility relate to greater levels of support for ISIS? I am interested in two kinds of users in this category: those which are overtly affiliated with the organization, and those who are not affiliated, but show some lesser indication of interest.

As for ISIS affiliates, the organization did maintain its own accounts on Twitter—at one time, as many as 40,000–125,000 (Berger and Morgan 2015; Isaac 2016). To systematically identify accounts associated with the organization, I tracked in real time lists published by several anti-ISIS hacking groups that since 2015 have been monitoring ISIS-affiliated accounts and publicly flagging them for suspension.[8] I designed an algorithm that between December 2015 and January 2017 continually monitored these accounts, recording information on user profiles, user locations, historical tweet timelines, and lists of friends and followers.[9] This real-time data collection enabled capturing information on accounts of about 15,000 ISIS activists before they were deleted from the Internet.

The more challenging task is to find individuals who have expressed interest in the Islamic State but are not ISIS activists. I collected information on the *followers* of all ISIS-affiliated accounts—about 1.6 million in total—to identify those who may have begun the radicalization process but have not progressed to the point of overt ISIS affiliation. The data contain user-level information, taken as "snapshots" of each user's profile at various points in time, as well as tweet-level information on over 100 million tweets posted by these users over the course of several years. Figure 1 shows the geographic distribution of ISIS activists and followers across the world.

Both Islamic State social-media activists and their followers are in the second stage of radicalization, and thus form the sample for this study.[10] Since the article focuses on Western Europe, the sample is restricted to activists and followers who are located in France, Germany, Belgium, and the United Kingdom.[11] While ISIS activists have clear connections to the organization and are most likely to represent individuals who adopted extremist worldviews,[12] the followers' group, which represents over 99% of the data, consists of a range of users, from individuals who actively support the organization, through accounts of interested citizens, to accounts that seek to counter ISIS. While it is hard to precisely identify what drives someone to follow Islamic State accounts on Twitter, the online appendix shows that ISIS followers are largely indistinguishable from a random Twitter sample across user-level metadata.[13]

Predicting Geographic Locations

A central aspect of this study involves predicting the geographic location of Islamic State activists and followers on Twitter, in order to match them to geographic data on socio-economic variables that might correlate with online radicalization. Since a very small share of Twitter users enable geo-tagging of their tweets or provide location information in their accounts,[14] social network and computer science researchers have developed methods in recent years to triangulate a

[8] At the beginning of 2015, the group @CtrlSec, which branched out of Anonymous, asked social media users to help find ISIS accounts on Twitter (see Figure S13 in the online appendix), an effort that led to the suspension of thousands of accounts in a matter of days. Since then, the monitoring, flagging, and suspension of ISIS accounts has been continuing—for example, in early 2015 Twitter announced that it has suspended about 125,000 ISIS accounts, many of which are believed to be flagged by @CtrlSec. See: http://www.nytimes.com/2016/02/06/technology/twitter-account-suspensions-terrorism.html?_r=0; as well as: http://www.theatlantic.com/international/archive/2015/10/anonymous-activists-isis-twitter/409312/. This project leverages this information to identify ISIS activists' accounts.

[9] See section S1 in the online appendix for more information. Figures S15 and S16 in the online appendix provide visual examples of these accounts.

[10] Of course, ISIS activists are further along the radicalization spectrum.

[11] This sample includes 175,015 users, of which 854 are activists and 174,161 are followers, See more information in Table 5.

[12] See online appendix Table S26, which shows that ISIS activists are significantly more likely to show signs of radicalization.

[13] See section S2.3 in the online appendix.

[14] In this study, only 26% of the users enable geo-tagging of their tweets, and 37% provide self-described location information.

FIGURE 1. Predicted Locations of ISIS Activists and Followers on Twitter

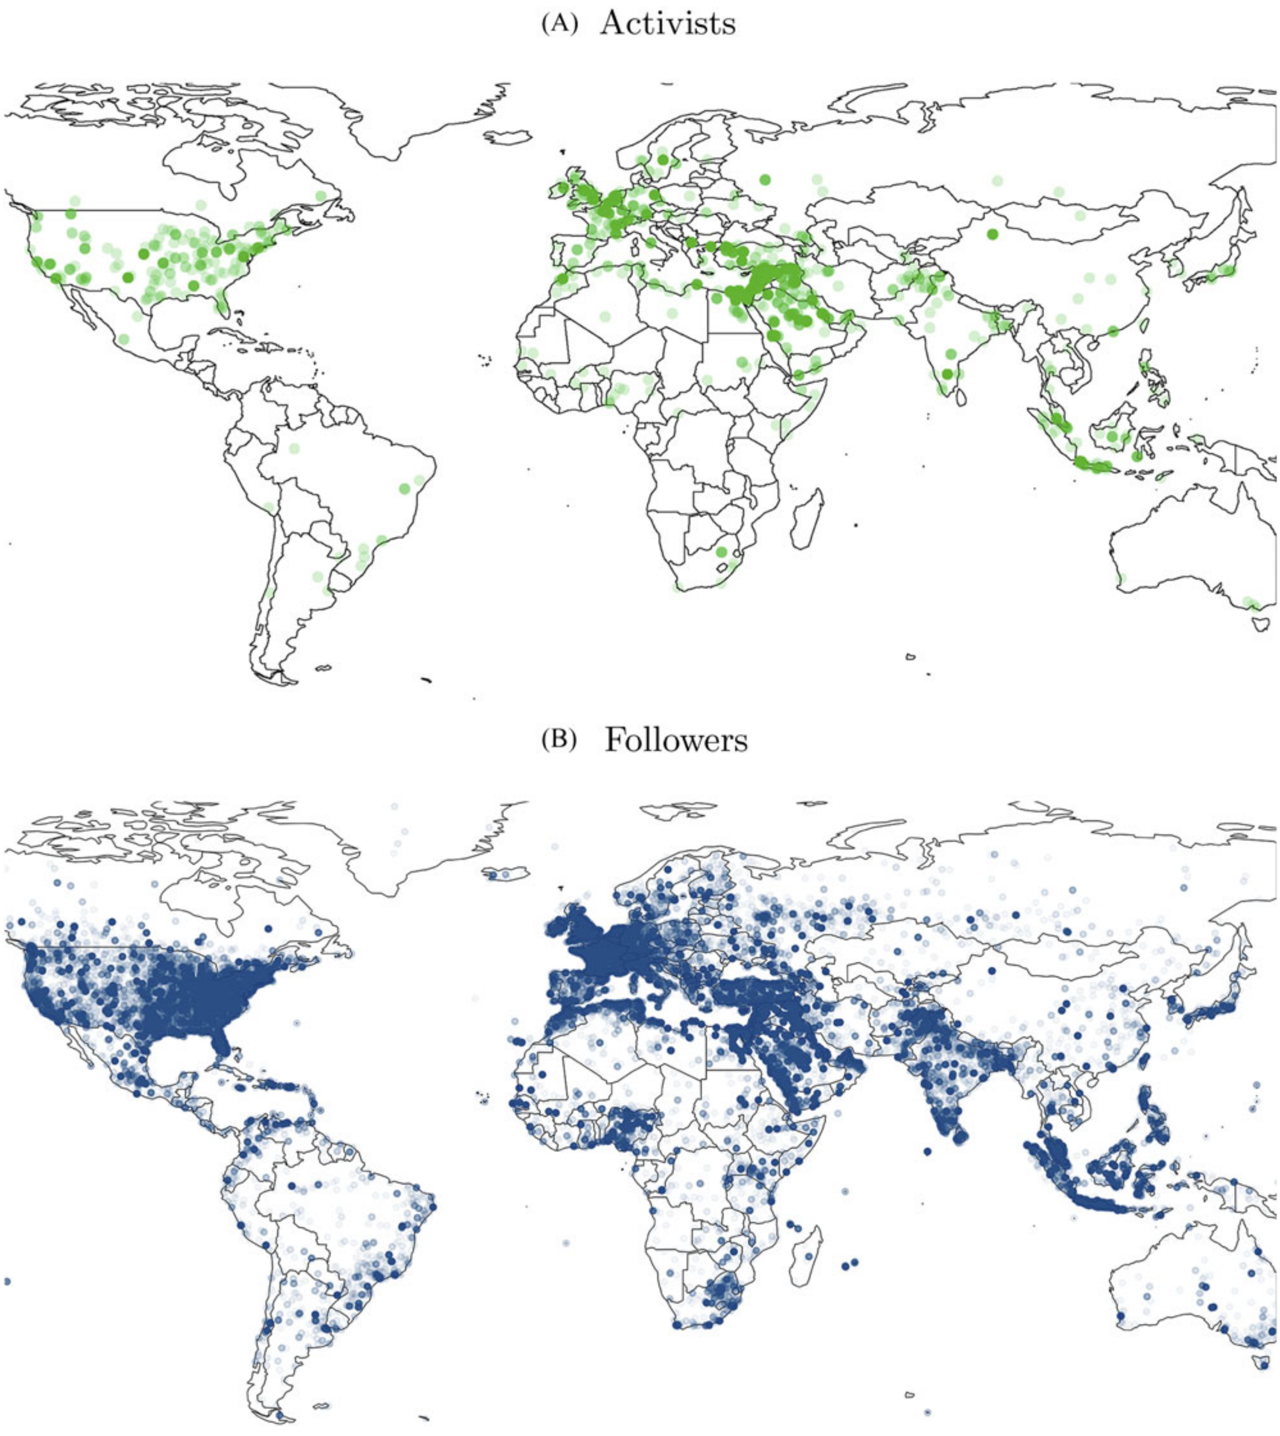

Note: Locations are predicted using a spatial label propagation algorithm (Jurgens 2013; Jurgens et al. 2015).

user's location based on locations provided by their networks of friends and followers (Backstrom, Sun, and Marlow 2010; Jurgens et al. 2015; McGee, Caverlee, and Cheng 2013). I employ a spatial label propagation algorithm developed by Jurgens (2013) to predict Twitter users' locations, which performs three rounds of prediction to maximize predictive accuracy.

Spatial label propagation algorithms rely on the finding in social network research that location information in a user's online network is a powerful predictor of a user's offline geographic location (Goldenberg and Levy 2009; McGee, Caverlee, and Cheng 2011; Takhteyev, Gruzd, and Wellman 2012). While social media platforms allow people to connect with others across the globe, recent studies have found that physical relationships in the offline world still strongly influence online social relationships. When people live their lives offline, they form relationships that subsequently transfer to the online world—e.g., co-workers or classmates who meet offline and then connect on social media platforms. As a result, a large share of individuals' online social network usually includes geographically close friends. Figure S19 in the online appendix, which is taken from Jurgens (2013), shows that across various social networks on different platforms, the majority of individuals in the network had at least one friend that was located within 4 kilometers. The online appendix provides more information on the details of this location prediction process, as well as a discussion of its out-of-sample predictive accuracy.[15]

While this method is imperfect and subject to prediction error, the rich data that it provides allow us to examine the local-level correlates of online support for Islamic State in Europe. As existing quantitative research on ISIS foreign fighter recruitment has so far remained at the country level (Benmelech and Klor 2018), this is an important step forward. In addition, while prediction errors make estimations more noisy, there is little reason to think they are plagued by systematic biases.[16] Location predictions are carried out on a very large and relatively deep network of over 1.6 million Twitter users across the world. Location prediction errors are likely to bias the results if they affect the *network* structure of individuals showing support for ISIS, e.g., by leading them to *strategically* choose friends so that their locations are systematically predicted (incorrectly) in areas with higher vote share for far-right parties. Strategic choice of friends in this way is difficult to perform systematically.[17]

Moreover, location prediction is carried out for *all* users in the database and analysis is carried out across thousands of localities in four countries. For systematic biases to be present, location predictions for ISIS supporters would have to appear systematically across countries in a pattern that correlates with far-right party vote-share locally.[18] To address the concern that Internet usage varies across rural and urban areas, regressions control for local population size.

Measuring Online Radicalization

I measure online radicalization using various user-level and tweet-level variables from the ISIS activists/followers database. First, I employ data from user-level fields to create indicators for whether a given user is flagged as an ISIS activist. Second, I use data on account suspension to code whether a user is suspended from Twitter for being associated with ISIS. Third, I use the network information in the database to count the number of ISIS accounts that each user follows. Fourth, I create textual measures for the number of pro-ISIS tweets posted by each user along several dimensions of ISIS support.

To generate the textual outcomes, I use supervised machine learning to classify tweets in English, Arabic, French, and German into one or more of these categories:[19]

1. *Sympathy with ISIS*—expressions of support or sympathy with the Islamic State, its ideology, and its activities in territories under its control.
2. *Life in ISIS territories, travel to Syria, or foreign fighters*—tweets describing the life of ISIS activists in the territories controlled by the Islamic State, posts expressing interest or intent to travel to Syria, discussion of foreign fighters, or all.
3. *Syrian war*—tweets describing events in the Syrian civil war, discussion/analysis of those events, or both.
4. *Anti-West*—anti-West rhetoric, criticizing Western countries' foreign policy and military operations in the Middle East.

Of course, anti-West sentiment often has no connection to radical ideology and can simply reflect legitimate grievances against Western countries' foreign policy. Nonetheless, I include it in a very limited way in my analysis, for two reasons. First, anti-West rhetoric has been a central part of ISIS activists' discourse on social media.[20] Second, studying anti-West sentiment sheds light on the way in which Western Islamic State sympathizers view their own countries. Since the organization's strategy has included driving a wedge between its supporters and the West,[21] this is an important topic to study. Table 2 shows examples of English language tweets for each of these topics.

[15] 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 is likely to be the case for ISIS activists, it should not be so for followers, who comprise over 99% of the sample. Section S2.3 in the online appendix shows that ISIS followers do not significantly differ from random Twitter users in many user-level fields.

[16] A test of the correlation between the prediction errors and far-right vote share shows no systematic relationship. See Table S6 in the online appendix.

[17] Another concern that may arise is that predictions will be biased for individuals who have many friends that have traveled to Syria. While it is true that an individual having a majority of friends in Syria may be erroneously predicted to be in Syria, that would not affect the results of this study, which only analyzes accounts in Western countries. Section S2.4 in the online appendix provides a detailed discussion and demonstrates via simulation the lack of bias in that sort of situation.

[18] Section S2.5 in the online appendix shows that location prediction errors are unlikely to spread users away from cities into rural areas inclined to vote for far-right parties.

[19] Life in ISIS territories, travel to Syria or foreign fighters was originally coded as two categories: one on life in ISIS territories and the other on travel to Syria. Since the two topics reflect a similar latent idea, I combined the two in this article's analysis.

[20] A qualitative examination of posts by ISIS activists showed a high number of tweets criticizing the West. See also Cunningham, Everton, and Schroeder (2017).

[21] In an essay published in the Islamic State's English-language magazine, *Dabiq*, the group stated its goal of "separating" Muslims from the West, i.e., encouraging them to immigrate to ISIS-controlled territories (Dabiq 2015).

TABLE 2. Examples of Tweets in Different Topics

Sympathy with ISIS
1. Jihad is the greatest of all deeds #IslamicState
2. Show everything from the Islamic State and other groups in Syria. It's important to hear all sides of the story
3. Assalam o Alaikom to All Islamic State Brothers
4. In sha Allah we will have honor again #IslamicState
Life in ISIS territories, travel to Syria, or foreign fighters
1. #Aljazeerah reports from inside the city of #Raqqa and shows how the #IslamicState runs the daily life
2. The glorious and mighty army of the Caliphate: Young kids ready to blow themselves up
3. Health services in Islamic state
4. Wedding of an #ISIS fighter in #Raqqa
5. 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!
Syrian war
1. #IS fighters readying to fight an invasion of Yarmouk Camp by Assad's allies Jaysh Al-Islam and Liwa Sham Al-Rasool
2. Massive destruction in Douma today after one of Assad's almost daily air strikes on the city. #Syria #Damascus
3. #Syria—The evil #Assad regime lost Busra al-Harir so they tortured a 6 year old girl out of revenge…
4. Massive explosion rocked entire of #Ramadi city. No further details yet. #Iraq #ISIS
Anti-West
1. America has been at war 222 out of 239 years since 1776. Let that sink for a moment
2. If Islamic State terror is evil why would Western State war be good?
3. US-led wars on terror have killed four million Muslims since 1990
4. It's sad when I am more afraid of our government then #ISIS! At least I know #ISIS hates #America #Government = wolves
5. Why are we shocked at ISIS brutality but not shocked by US British & European brutality?

The supervised learning process works as follows. First, about a thousand human coders from two crowdsourcing platforms, Amazon Mechanical Turk and CrowdFlower,[22] labeled a random sample of posts by hand. Then, an algorithm used information on the words in each labeled post to "learn" the categorization rules and classify unlabeled posts.[23] I obtained a random sample of tweets posted by ISIS activists in English, Arabic, French, and German to create a training set for the classification model.[24] 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.[25]

[22] CrowdFlower changed its name to Figure Eight in March 2018.

[23] See Grimmer and Stewart (2013) for a review and more information on supervised machine learning methods to classify text, and James et al. (2013) for an introduction to machine learning in general. The online appendix provides more details on the supervised learning method used in this study.

[24] English, Arabic, French, and German are used in 76% of the tweets in the database. As the proportion of tweets in the database varies by language, the size of the training set accordingly varies for different languages: English (N = 9,926), Arabic (N = 10,631), French (N = 6,158), and German (N = 3,011).

[25] I took several precautionary steps to reduce the likelihood that the human coders (971 in all) might inadvertently bias the coding of radical content. First, to be sure that the coding instructions were easy to understand, I confirmed that a student research assistant was also able to correctly code tweets using these instructions (see Figure S11 in the online appendix). Second, I randomly assigned each tweet to multiple coders, which should cause idiosyncratic biases from individual coders to cancel out on average. Third, I manually checked a random sample of coded tweets to ensure that the coding reflected the correct topics.

Since Twitter textual data are very noisy and radical pro-ISIS content is rare, many tweets in the database were coded as unrelated to any of the above categories.[26] To facilitate statistical prediction, I follow King and Zeng (2001) and randomly over-sample pro-ISIS tweets and randomly under-sample 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 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. Using this method, the models were able to predict pro-ISIS content with an in-sample accuracy over 95%. More metrics on the performance of the models for each topic and language are reported in section S3 in the online appendix. 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.

To measure users' posting of radical and pro-ISIS content, I counted the number of tweets classified in these categories for each user. I also created a combined measure that counted the number of tweets falling into any of the ISIS-related categories.[27] To ensure that I capture users that post highly pro-ISIS content, I

[26] See section S3 in the online appendix for details on the classes for each outcome and language.

[27] Sympathy with ISIS, Life in ISIS territories, travel to Syria or foreign fighters, and the Syrian war.

TABLE 3. Summary Statistics

Statistic	N	Mean	St. dev.	Min	Max
A. Dependent variables					
Sympathy with ISIS (# tweets)	175,015	4.678	12.309	0	277
ISIS life, foreign fighters, or travel to Syria (# tweets)	175,015	9.103	23.479	0	479
Syrian war (# tweets)	175,015	6.725	17.288	0	343
Anti-West (# tweets)	175,015	4.429	11.961	0	286
ISIS activist	175,015	0.005	0.070	0	1
Suspended by Twitter	175,010	0.041	0.199	0	1
ISIS accounts following (# accounts)	175,010	5.445	23.827	0	3,216
B. Independent variables					
Far-right vote share (%, local level)	116,492	13.208	9.026	0	53.805
Unemployed (%, local level)	170,653	5.124	2.410	0	41
Immigrants unemployed (%, local level)	90,520	1.889	0.993	0	9
Foreigners/non-citizens (%)	171,076	10.466	7.405	0	89.026
Population	171,547	815,083.3	1,078,664	3	3,292,365

Note: This table reports summary statistics for the sample of ISIS activist and followers who are predicted to be located in France, Germany, Belgium, and the United Kingdom. Data represent content posted between 2014 and 2016.

created an indicator that is coded one for users who are at the top 1% of the distribution of radicalized content posting and zero otherwise.[28] Panel A in Table 3 provides summary statistics for these various measures of online radicalization.

While these measures only capture expressions of support for the Islamic State in the online world, they are nonetheless a plausible proxy for underlying radicalization. Social media played a central role in the radicalization process of European foreign fighters (Carter, Maher, and Neumann 2014). In the United States, the majority of individuals who attempted to travel overseas to join ISIS or planned a violent attack on the organization's behalf used social media when radicalizing. Importantly, most of these individuals have expressed their support for ISIS on social media by posting publicly viewable posts.[29] This suggests that studying radicalization using online measures of ISIS support can be a fruitful way to better understand this phenomenon.

Independent Variables

To create a local-level measure of anti-Muslim hostility, I relied on the finding presented earlier on the strong link between far-right voting and holding anti-Muslim attitudes. I created local-level measures of support for the far-right by calculating the percent of votes for parties associated with far-right positions at the electoral constituency level in France, Germany, Belgium, and the United Kingdom. Table 4 shows the elections and parties used to construct this variable. Using Twitter users' predicted geolocation data, I matched users in my database to electoral constituencies, thereby assigning users to different areas with varying degrees of far-right support. Panel B in Table 3 shows that vote share for these parties varies substantially, where some users are located in areas with zero vote share for far-right parties, and others in areas with more than 50% support for these parties.

In addition, I created variables for other indicators that might predict online support for ISIS. First, I examine whether economic grievances might be linked to radicalization by using official data on unemployment from France, Germany, the United Kingdom, and Belgium, at the lowest possible level of aggregation. In France, Germany, and Belgium, the lowest possible level was the town/municipality. In the United Kingdom, data were available at the sub-municipality level.[30] I matched users to their respective areas for which unemployment data exist. As some have hypothesized that unemployment among immigrants in particular feeds ISIS radicalization (Holland 2016), I also created a measure for the share of unemployed immigrants in each location. Panel B in Table 3 provides information on the distribution of these variables across Twitter users in the dataset.

Second, I examine whether areas that are likely to have stronger social networks have a greater number of radicalizing individuals. I use census data on the share of foreigners or noncitizens in each locality to examine the extent to which ISIS supporters on Twitter are located in areas with higher shares of noncitizen populations. To account for the recent debates over the link between refugees and support for ISIS in Europe (Marans 2015), I looked for variables that might proxy for the presence of refugees in a locality. I use information on the number of asylum seeker centers across localities in France, and the share of asylum seeker benefits receivers in localities in Germany.[31] As these two variables are measured on

[28] I chose this cutoff in order to be conservative and not erroneously classify as radicalized individuals who post less radical content. As reported in Table S25 in the online appendix, results hold in estimations with cutoffs using top 5%, 10% 15%, and 20%.

[29] See more details in section S5 in the online appendix.

[30] In the United Kingdom, statistical local data are available at the Mid-level Super Output Area (MSOA) level, in which the population ranges between 5,000 and 15,000. (Office for National Statistics 2016).

[31] These data reflect 2014 figures.

TABLE 4. Far-Right Parties in Recent European Elections

Country	Election	Far-right parties
France	2015 departmental elections	Front National (FN)
Germany	2013 Federal elections	National Democratic Party of Germany (NPD); Alternative for Germany (AfD)
United Kingdom	2015 general elections	British Democrats; British National Party; Liberty GB party; National Front party; United Kingdom Independence Party (UKIP)
Belgium	2014 Belgian federal elections	Vlaams Belang (VB)

different scales, I created a standardized measure for this combined variable. Table 3 shows the distribution of these variables across users. The online appendix provides more details on the data sources and construction of the independent variables.

DESCRIPTIVE ANALYSIS

This section presents a few examples that illustrate the kind of content that I collected and its connection to real-world events. On June 29, 2014, ISIS declared the establishment of a caliphate in an online statement distributed through Twitter and the group's media center, calling all Muslims to pledge allegiance and travel to the territories it controlled in Syria and Iraq. I calculated the daily number of tweets discussing foreign fighters or travel to Syria posted by accounts located in France, Belgium, Germany, and the United Kingdom in the month surrounding ISIS's caliphate declaration. Figure 2 shows that after the declaration, discourse on foreign fighters significantly increased among these Twitter users.

Next, I examine whether online radicalization measures correlate with Western foreign fighter figures. Figure 3 shows a map of ISIS foreign fighters from Europe (Panel A), along with a map showing the number of Twitter users flagged as ISIS activists in each country (Panel B). France, the United Kingdom, and Germany have higher numbers of foreign fighters and Twitter users flagged as ISIS activists than many other European countries. Figure 4 displays the correlation between additional online radicalization measures and the number of foreign fighters in the West. It can be seen that online measures of support for ISIS closely track official foreign fighter counts. While these scatterplots show bivariate relationships, the online appendix provides estimations controlling for population size, which show the same pattern.

CROSS-SECTIONAL STUDY

This section examines whether local-level measures of anti-Muslim hostility are linked to greater online support for ISIS. I regress the different online radicalization outcomes on the independent variables described above using a combined dataset covering all localities in France, Germany, Belgium, and the United Kingdom. The dependent variables are summarized in Panel A in Table 3 and are measured on the Twitter user level. The independent variables, summarized in Panel B in Table 3, are matched to each individual user in the dataset, but originate in local-level administrative data. To account for possible dependency across users in the same area, I cluster the standard errors at the locality level in my main regressions. I use the following least squares model in my primary estimations:

$$Y_{ijk} = \beta_1 V_{jk} + \beta_2 U_{jk} + \beta_3 F_{jk} + \beta_4 P_{jk} + \beta_5 P_{jk}^2 + \alpha_k + \varepsilon_{ijk}, \quad \textbf{(1)}$$

where i is a Twitter user in geographic area j in country k; Y_{ijk} is one of the online radicalization measures for user i in area j in country k; and V_{jk} represents the locality-level vote share for far-right parties matched to user i in area j in country k. U_{jk}, F_{jk}, and P_{jk} represent unemployment, share of foreigners, and population size matched to user i in area j in country k, respectively, and α_k is a country fixed effect.[32] The main coefficient of interest in these regressions is β_1, which estimates the relationship between the local-level vote share for far-right parties and online measures of support for ISIS. While this coefficient cannot be interpreted as evidence of a causal relationship, it provides a systematic test of the link between a context of anti-Muslim hostility and online pro-ISIS radicalization.

Far-Right Vote Share and Support for ISIS

Tables 5 and 6 report the main results. In Table 5, Column (1), the dependent variable is a text-based measure that is coded one for individuals who are at the top 1% of the distribution of posting pro-ISIS content, and zero otherwise. To ensure that this content-based measure captures individuals who frequently express sympathy with ISIS, in Column (2), the dependent variable includes only tweets sympathizing with ISIS. Regardless of the measure used, it can be seen that local-level vote share for far-right parties is positively associated with posting large numbers of pro-ISIS tweets. In substantive terms, a one percent increase in far-right vote share is associated with a 3–5% increase in the probability of being among the top 1% of posters of extremist content.

[32] Data on the share of Muslim populations in each geographic area are only available in the United Kingdom. In estimations with United Kingdom data only, shown in Table 8 below, I find that controlling for Muslim population share does not affect the results.

FIGURE 2. ISIS Declares Caliphate and Tweets Discussing Foreign Fighters or Travel to Syria

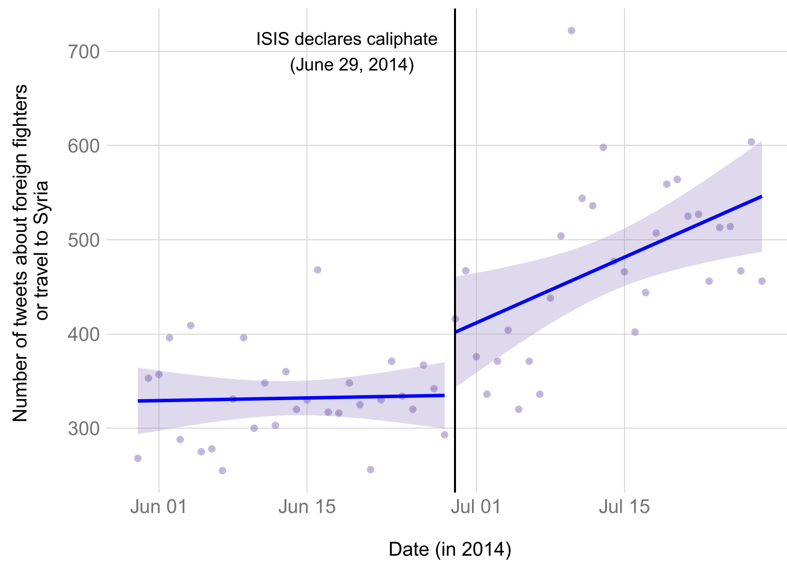

Note: The figure shows the daily number of tweets discussing foreign fighters or travel to Syria in the month surrounding ISIS's caliphate declaration on June 29, 2014.

FIGURE 3. Foreign Fighters and Online Radicalization in Europe

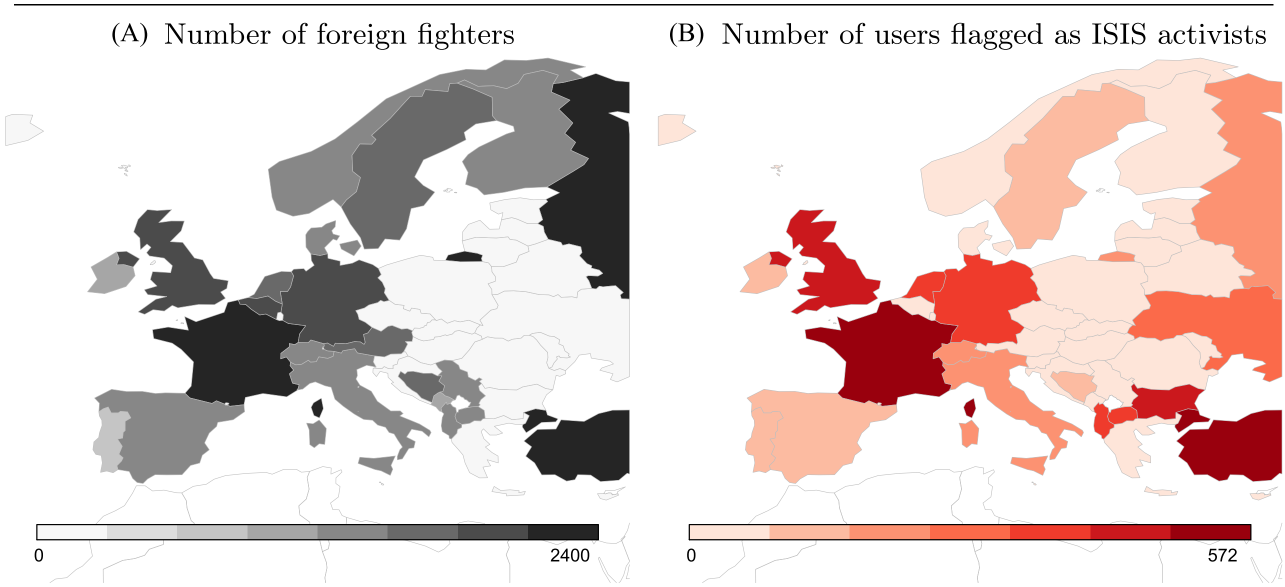

Note: Panel (A) displays official counts of ISIS foreign fighters in Europe, calculated by Barrett et al. (2015). Panel (B) shows the number of Twitter users flagged as ISIS activists, aggregated to the country level.

Columns (3)–(5) in Table 5 report the results when the dependent variable is measured as being flagged as an ISIS activist, being suspended from Twitter for association with the organization, and with a count measure of the number of ISIS accounts that a user follows. Here, as well, the results show that vote share for far-right parties is positively related to these radicalization outcomes. However, suspension and number of ISIS accounts followed are not statistically significant at conventional levels with the clustered standard error specification, although results are significant when estimating the models without clustered standard errors (see Table S27 in the online appendix). In substantive terms, vote share for far-right parties is

FIGURE 4. Foreign Fighters and Online Radicalization (Additional Measures)

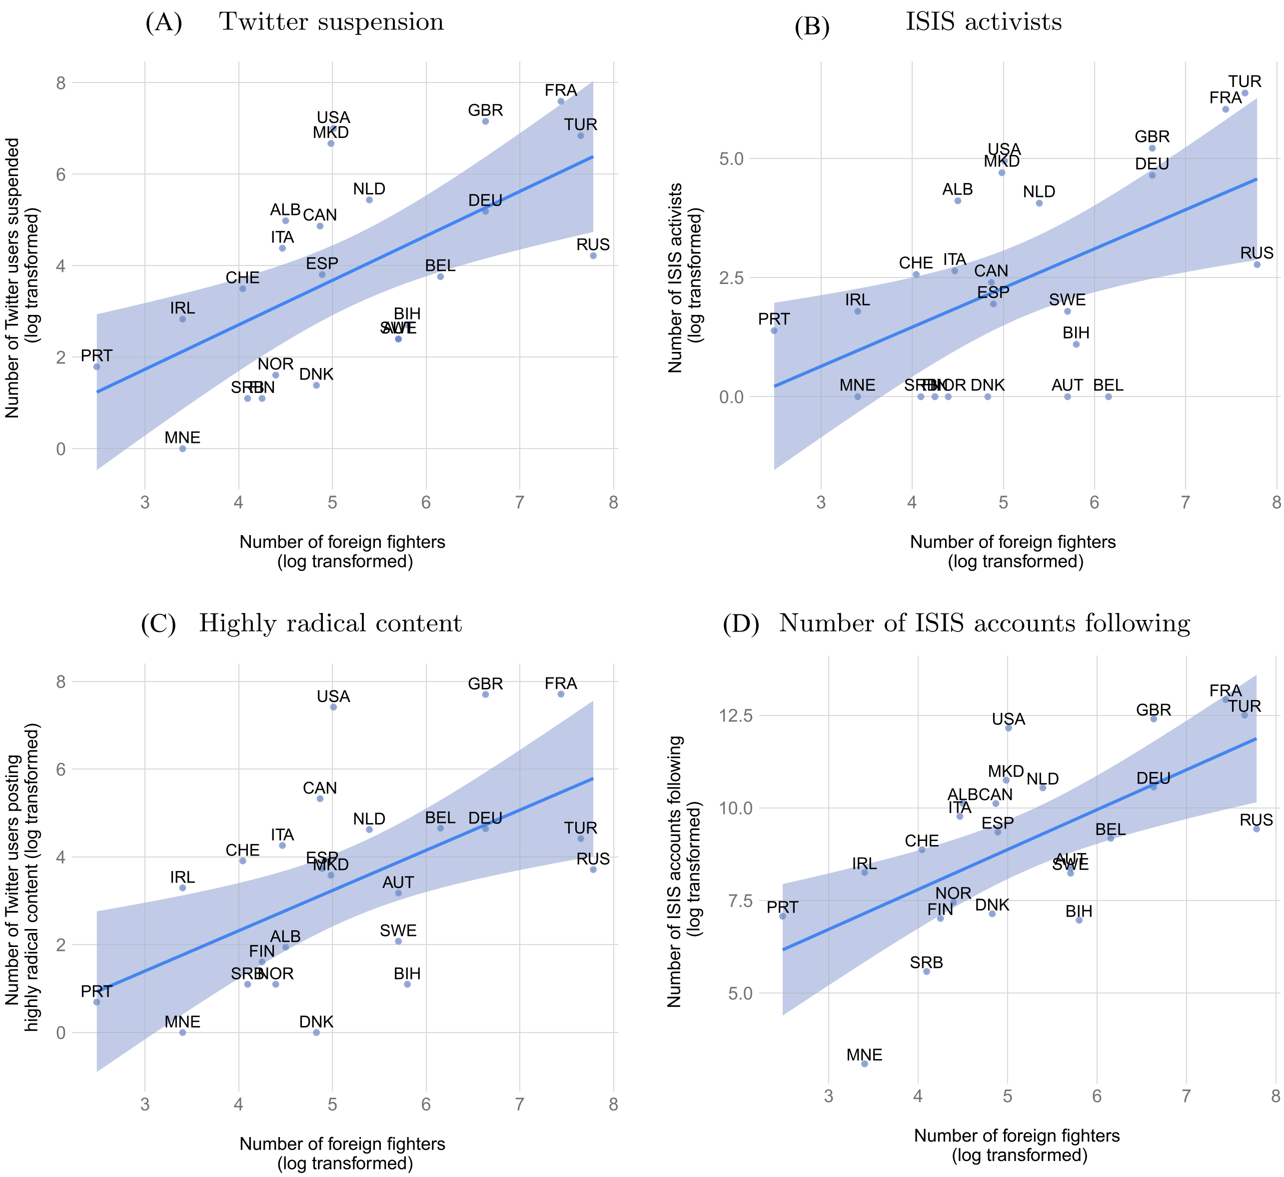

Note: The figure presents scatterplots of the relationship between the number of foreign fighters and online radicalization measures in countries that had at least one foreign fighter with ISIS. Data on foreign fighters are taken from Barrett et al. (2015). Online radicalization measures are based on data collected by the author and are aggregated to the country level. The values are log-transformed.

associated with a 6% increase in the probability of being flagged as an ISIS activist.

Table 6 reports the results when the dependent variables reflect the number of tweets posted by a user across the content outcomes. Here, a one percent increase in the vote share for far-right parties is positively and statistically significantly associated with increases in the number of tweets sympathizing with ISIS, discussing life in ISIS-controlled territories and foreign fighters, discussing the Syrian civil war, and expressing anti-West sentiment. Substantively, these reflect an average increase of 4,000–10,000 pro-ISIS tweets across the entire sample. Note that these measures are calculated from content generated in English, Arabic, French, and German and are measured across thousands of individuals in four countries. The consistency of the results across these text-based measures suggests that this association did not occur by random chance.

Hate Crimes and Support for ISIS

One might wonder whether the findings are driven by greater levels of animosity against Muslims in areas

TABLE 5. Far-Right Vote Share and Support for ISIS on Twitter

	(1) Top 1% radical content	(2) Top 1% sympathy with ISIS only	(3) Flagged as an ISIS activist	(4) Suspended from Twitter	(5) Number of ISIS accounts following
Far-right vote share (%)	0.25**	0.20**	0.30**	0.10	0.09
	(0.10)	(0.09)	(0.14)	(0.40)	(0.08)
Unemployment (%)	0.25	0.23	−0.21	−1.25*	−0.11
	(0.24)	(0.22)	(0.52)	(0.69)	(0.14)
Foreigners (%)	0.11	0.14*	0.27*	−0.06	0.08
	(0.09)	(0.08)	(0.15)	(0.32)	(0.07)
Constant	7.87*	4.34	−9.82	35.03**	1.12
	(4.43)	(4.17)	(6.36)	(15.28)	(3.73)
Population controls	✓	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓	✓
R^2	0.0003	0.0004	0.006	0.002	0.006
Number of clusters	2,655	2,655	2,655	2,654	2,654
Number of observations	112,271	112,271	112,271	112,267	112,267

Robust standard errors in parentheses, clustered at the locality level. Base category is Belgium.
Coefficients in columns 1–4 are × 1,000 to account for the skewed distribution of the dependent variables.
$^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

TABLE 6. Far-Right Vote Share and Posting Pro-ISIS and Anti-West Content on Twitter

	(1) Sympathy with ISIS	(2) ISIS life/Foreign fighters	(3) Syrian war	(4) Anti-West
Far-right vote share (%)	0.05**	0.09**	0.07**	0.04*
	(0.02)	(0.04)	(0.03)	(0.02)
Unemployment (%)	0.12**	0.24**	0.15**	0.13**
	(0.05)	(0.10)	(0.08)	(0.05)
Foreigners (%)	0.02	0.04	0.03	0.02
	(0.02)	(0.04)	(0.03)	(0.02)
Constant	3.53***	7.31***	5.77***	3.21***
	(0.99)	(1.97)	(1.44)	(0.91)
Population controls	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓
R^2	0.001	0.002	0.002	0.003
Number of clusters	2,655	2,655	2,655	2,655
Number of observations	112,271	112,271	112,271	112,271

Robust standard errors in parentheses, clustered at the locality level. Base country is Belgium.
$^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

where far-right parties are popular. Earlier in this article, I showed that at the individual level voting for far-right parties strongly correlates with anti-Muslim sentiment. In this section, I test if the relationship found in the previous section holds when using hate crimes as a proxy for anti-Muslim hostility. I also examine whether hate crimes moderate the link between far-right vote share and pro-ISIS radicalization.

As systematic local-level data on hate crimes is not publicly available in most countries, this section only uses data from the United Kingdom. Using official data from the U.K. police, I matched accounts of Twitter users in the United Kingdom with information on hate crimes motivated by religion in each police force area,[33] as well as granular geo-spatial data on public order crimes.[34] Public order crimes are incidents that "cause fear, alarm, or distress" and subsume most hate crimes

[33] Hate crime data in each police force area cover the years 2015–2017. See https://www.gov.uk/government/statistics/hate-crime-england-and-wales-2015-to-2016 and https://www.gov.uk/government/statistics/hate-crime-england-and-wales-2016-to-2017.
[34] The data can be downloaded at https://data.police.uk/data/.

FIGURE 5. Far-Right Vote Share, Hate Crimes, and Support for ISIS in the U.K.

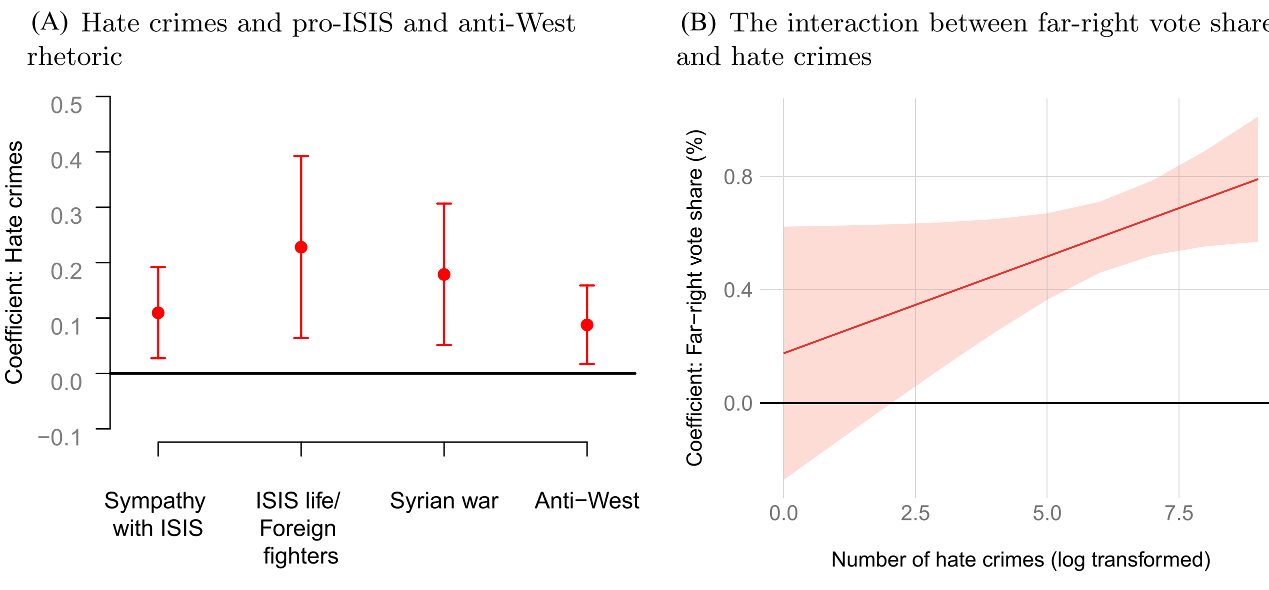

Note: Panel (A) shows the correlation between the number of hate crimes and pro-ISIS tweeting in the U.K. Panel (B) presents the interaction between far-right vote share and hate crimes in a regression where the dependent variable combines all pro-ISIS tweets posted by users in the U.K.

in the United Kingdom.[35] To estimate the link between hate crimes and pro-ISIS rhetoric on social media, I regressed online radicalization outcomes on local-level hate crime data, controlling for other variables that might explain support for ISIS and the likelihood of hate crimes in each area.[36]

Figure 5A shows the correlation between the number of hate crimes and pro-ISIS tweeting in the United Kingdom. It can be seen that individuals located in areas with a higher number of hate crimes are more likely to post greater number of tweets expressing sympathy with ISIS, discussing life in ISIS territories and foreign fighters, and expressing interest in the Syrian civil war. In addition, individuals located in these areas tend to voice greater anti-West sentiment. Substantively, a unit increase in the number of hate crimes is linked to an increase of about 10–11% in the number of pro-ISIS tweets posted by each user, or an average increase of 7,000–18,000 tweets across the entire sample.

To examine whether hate crimes moderate the relationship between far-right vote share and online support for ISIS, I interact far-right vote share with hate crimes in a regression where the dependent variable combines all pro-ISIS tweets.[37] Figure 5B presents this interaction, showing that the relationship between far-right vote share and pro-ISIS tweets is stronger in areas with more hate crimes. These findings suggest that anti-Muslim hostility likely drives the relationship between far-right vote share and support for ISIS on Twitter, at least among users located in the United Kingdom.

OTHER CORRELATES OF ONLINE RADICALIZATION

Unemployment

Next, I investigate other correlates of online radicalization. As can be seen in Table 6, the unemployment rate at the local level is also strongly associated with online support for ISIS when considering the content-based outcomes. A one percent increase in the level of unemployment is associated with a 1–3% increase in posting tweets sympathizing with ISIS, discussing life in ISIS territories, or expressing an interest in traveling to Syria to become foreign fighters. Unemployment, however, is not positively related to other radicalization outcomes, such as being flagged as an ISIS activist, being suspended from Twitter, or the number of ISIS accounts followed (Table 5). In Table 7, I find that the share of unemployed immigrants is not significantly related to online measures of pro-ISIS radicalization.

Since levels of unemployment can drive both far-right vote share and support for ISIS, I run additional estimations to further rule out the confounding effect of unemployment. First, as presented in the next section, I conduct high frequency studies around events that may mobilize support for ISIS and examine whether pro-ISIS rhetoric increases after these events more strongly in areas with higher levels of far-right vote share. In particular,

[35] See: https://www.police.uk/about-this-site/faqs/#what-do-the-crime-categories-mean. Since official police-force data on hate crimes is reported at a very aggregate level, I use incident-level, geo-tagged data on public order crimes. A test of the correlation between public order crimes and religiously motivated hate crimes, at the Twitter user level, shows a very strong relationship: the correlation coefficient is 0.9 with a p-value <0.01. See online appendix for more details.
[36] I control for far-right vote share, unemployment, share of foreigners, Muslims, and Arabs, and population size.
[37] See online appendix section S6 for details.

TABLE 7. Unemployed Immigrants, Asylum Seekers and Support for ISIS on Twitter

	(1) Top 1% radical content	(2) Flagged as an ISIS activist	(3) Suspended from Twitter	(4) Number of ISIS accounts following
Far-right vote share (%)	0.24*	0.52**	0.62	0.23*
	(0.13)	(0.26)	(0.54)	(0.14)
Unemployed immigrants (%)	0.70*	0.38	0.08	0.36
	(0.36)	(1.05)	(2.12)	(0.63)
Asylum seekers (%, sd units)	−0.40	−11.86***	−14.40***	−2.62*
	(1.07)	(4.01)	(4.38)	(1.45)
Constant	−4.25	−64.04**	−42.15	−14.51
	(7.50)	(28.79)	(39.70)	(14.66)
Population controls	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓
R^2	0.001	0.012	0.004	0.005
Number of clusters	1,135	1,135	1,135	1,135
Number of observations	30,383	30,383	30,382	30,382

Robust standard errors in parentheses, clustered at the locality level. Data available only for France and Germany. Base category is Germany.
Coefficients in columns 1–3 are ×1,000 to account for the skewed distribution of the dependent variables.
$^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

TABLE 8. Pro-ISIS and Anti-West Content in the United Kingdom, Additional Controls

	(1) Sympathy with ISIS	(2) ISIS life/foreign fighters	(3) Syrian war	(4) Anti-West
Far-right vote share (%)	0.06***	0.12***	0.09***	0.05***
	(0.02)	(0.04)	(0.03)	(0.02)
Muslims (%)	−0.07**	−0.13**	−0.09**	−0.06***
	(0.03)	(0.06)	(0.04)	(0.02)
Males (%)	−0.04	−0.11	−0.10	−0.05
	(0.04)	(0.09)	(0.07)	(0.04)
Pakistanis (%)	0.04	0.08*	0.05	0.03
	(0.02)	(0.05)	(0.04)	(0.02)
Bangladeshis (%)	0.01	0.04	0.01	0.01
	(0.02)	(0.05)	(0.04)	(0.02)
Arabs (%)	0.09	0.17	0.13	0.06
	(0.07)	(0.15)	(0.11)	(0.06)
Foreigners (%)	0.02	0.03	0.03*	0.02*
	(0.01)	(0.02)	(0.02)	(0.01)
Unemployed (%)	−0.04	−0.09	−0.08	−0.03
	(0.04)	(0.08)	(0.06)	(0.03)
Population	0.00***	0.00***	0.00***	0.00***
	(0.00)	(0.00)	(0.00)	(0.00)
Population^2	−0.00***	−0.00***	−0.00***	−0.00***
	(0.00)	(0.00)	(0.00)	(0.00)
Constant	2.14	5.71	5.26	2.44
	(2.34)	(4.67)	(3.71)	(2.02)
R^2	0.001	0.001	0.001	0.001
Number of observations	62,081	62,081	62,081	62,081

Heteroskedasticity robust standard errors in parentheses.
$^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

when I examine heterogeneous changes following these events for both far-right vote share and unemployment, it is clear that these high-frequency changes are linked to the former and not the latter. The results show systematic evidence that ISIS followers express greater support for the organization after these events in localities where far-right parties are more popular.

Second, in the online appendix, I carry out a more comprehensive examination using a matching design. In the matching approach, I compare users located in areas with high and low far right vote share that are matched on levels of unemployment, the proportion of foreigners, population size, and the country in which they are located. I find almost identical results.[38] This suggests that unemployment does not confound the relationship between far-right vote share and pro-ISIS radicalization.

FOREIGNERS, REFUGEES, AND OTHER MINORITIES

In addition, I examine whether support for ISIS on Twitter relates to the share of foreigners or noncitizens in a locality. The third row in Tables 5 and 6 shows that a greater number of foreigners in a locality is positively associated with online radicalization, but the relationships are not statistically significant for almost all outcomes. I also investigate in Table 7 whether the share of refugees in a locality relates to greater support for ISIS on Twitter. I find that the share of asylum seeker and/or asylum seeker centers in a locality is negatively related to being flagged as an ISIS activist, being suspended from Twitter, and to the number of ISIS accounts followed.

To examine whether these results might be driven by a common third variable linked to both radicalization and far-right support, I use data on possible omitted variables that are available only in the United Kingdom, such as the share of Muslims, Arabs, Pakistanis, Bangladeshis, and foreign-born in each local area. Table 8 shows that when controlling for these variables, vote share for far-right parties remains strongly correlated with posting pro-ISIS content on Twitter. The findings also show that the local proportion of Muslims is negatively correlated with posting pro-ISIS content. This is an important finding in light of recent debates on Muslim populations in the West, as it casts doubt on the argument that areas with larger Muslim populations are more likely to be prone to jihadi radicalization.

Overall, these results are consistent with the hypothesis that exposure to anti-Muslim animosity, measured as the local-level vote share for far-right parties, might lead individuals to radicalize and support the Islamic State on social media. The results hold across various dependent variables, in a large number of locations in four European countries. However, since the findings are based on cross-sectional comparisons, it is possible that these relationships are driven by reverse causality or omitted variables. In the remaining parts of the article, I investigate these relationships using high frequency Twitter data surrounding events that likely stir support for ISIS, such as terrorist attacks, releases of propaganda materials on the Internet, and anti-Muslim protests.

HIGH FREQUENCY EVENT STUDIES

The relationship between anti-Muslim hostility and pro-ISIS radicalization is complex and may also run in the other direction or be driven by omitted variables. To further investigate the link between far-right vote share and support for ISIS on social media, I take advantage of the high-frequency nature of Twitter data and examine whether events that likely stir sympathy with ISIS among potential recruits are immediately followed by increased posting of pro-ISIS content in areas with high far-right support. While this design cannot completely rule out reverse causality, it strengthens the inference that anti-Muslim hostility is indeed linked to pro-ISIS content and not reflecting a spurious cross-sectional correlation.

It is important to note that observing high frequency changes in pro-ISIS content does not imply that people are radicalizing in such short amount of time. While the process of radicalization often unfolds slowly and gradually, it is possible to identify events that ‘trigger’ this dynamic. Prior research from other conflicts showed that individuals’ support for extremism can be strongly shaped by exposure to violent events (Crone 2016). I seek to examine whether changes in the voicing of radical sentiments are stronger among individuals in areas with greater levels of anti-Muslim hostility.

In this section, I carry out several high-frequency analyses around three types of events. First, I examine the impact of the terrorist attacks in Paris (11/13/2015) and Brussels (3/22/2016) on support for ISIS on Twitter. Second, I study the effect of a widespread propaganda release by ISIS, which was distributed on the organization’s Twitter networks on June 29, 2014. Third, I evaluate how individuals responded to a high profile anti-Muslim event, the Patriotic Europeans Against the Islamization of the West (PEGIDA) marches across Europe on February 6, 2016. If a local hostile context has any influence on support for ISIS on social media, then we would expect to find a significantly different pattern in the responses to these events in areas where far-right parties are popular.

Events that Can Increase Support for ISIS

Terrorist Attacks

Terrorist attacks perpetrated by individuals associated with the Islamic State might inspire individuals sympathetic to the group to voice their support for the organization. I examine whether the Paris attacks of November 2015[39] and the Brussels attacks of March 2016[40] were immediately followed by increased radical, pro-ISIS content among Islamic State followers on Twitter, especially in locations with high support for far-right parties.

ISIS Propaganda Releases

Propaganda materials distributed on the Internet can also increase pro-ISIS rhetoric among potential

[38] See section S7 in the online appendix for more details.

[39] On November 13, 2015, several perpetrators identified with the Islamic State launched several attacks in Paris, including suicide bombings and mass shootings. The attacks killed 130 people and injured hundreds of others, becoming the deadliest atrocities in France since the Second World War.

[40] On March 22, 2016, ISIS-affiliated suicide bombers detonated explosive devices in Brussels Airport and at a train station nearby, killing 32 civilians and injuring over 300.

FIGURE 6. "The End of Sykes Picot" Propaganda Video Disseminated by ISIS on Twitter

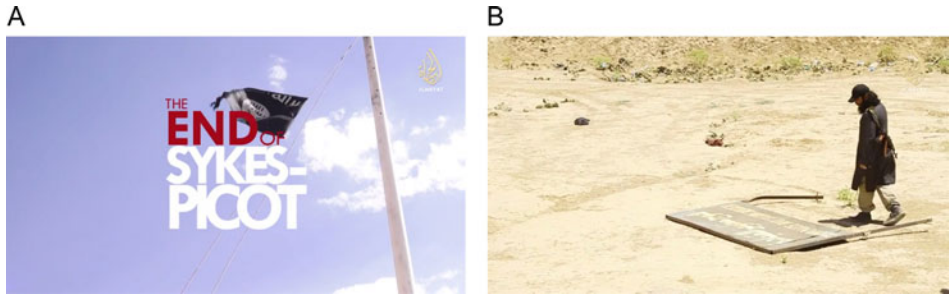

Note: Screenshots from the propaganda video "The End of Sykes-Picot," disseminated by ISIS on Twitter on June 29, 2014. Source: http://jihadology.net

supporters. On June 29, 2014, ISIS disseminated on Twitter a video called "The End of Sykes Picot," which showcased the territory captured by the organization and stated that it will eliminate all "so-called borders" created by Western powers in the Middle East. Figure 6 shows screenshots from the video, in which an Islamic State soldier steps over a sign that used to mark the border between Syria and Iraq. I examine how ISIS followers on Twitter responded to this video, and especially whether individuals located in areas with high-far right support were more responsive to the propaganda release.

Anti-Muslim Marches

Events that exhibit animosity to Muslims may also lead individuals to voice their support for ISIS. On February 6, 2016, PEGIDA organized large marches in multiple cities in Germany, Britain, France, Netherlands, Austria, Ireland, Poland, Czech Republic, and Slovakia, to protest

TABLE 9. Terrorist Attacks, ISIS Propaganda, and Changes in Pro-ISIS Rhetoric

	Paris attacks		Brussels attacks		ISIS propaganda release	
	(1) Sympathy with ISIS	(2) ISIS topics	(3) Sympathy with ISIS	(4) ISIS topics	(5) Sympathy with ISIS	(6) ISIS topics
A. Changes in pro-ISIS content (standard deviation units)						
After event = 1	0.118***	0.126***	0.024***	0.029***	0.067**	0.018
	(0.033)	(0.026)	(0.006)	(0.010)	(0.028)	(0.022)
Constant	0.099	0.245***	0.053*	0.110**	−0.285***	0.031
	(0.064)	(0.040)	(0.028)	(0.051)	(0.059)	(0.125)
R^2	0.008	0.009	0.003	0.002	0.007	0.0002
Number of clusters	409	409	609	609	150	150
Number of observations	35,176	35,176	67,438	67,438	5,502	5,502
B. Changes in pro-ISIS content (standard deviation units), by far-right support						
After event = 1	0.031	0.036	0.040***	−0.009	−0.091*	−0.069
	(0.027)	(0.032)	(0.015)	(0.016)	(0.054)	(0.079)
Far-right vote share (%)	−0.003	−0.005**	0.003*	−0.000	−0.009	−0.012**
	(0.002)	(0.002)	(0.001)	(0.002)	(0.006)	(0.005)
After event = 1 × far right vote share (%)	**0.003***	**0.005****	**−0.001**	**0.002****	**0.009*****	**0.005**
	(0.002)	(0.002)	(0.001)	(0.001)	(0.004)	(0.004)
Constant	0.255*	0.355***	-0.037	0.086	-0.340	0.038
	(0.134)	(0.132)	(0.071)	(0.077)	(0.221)	(0.203)
Controls	✓	✓	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓	✓	✓
R^2	0.005	0.006	0.002	0.002	0.010	0.008
Number of clusters	362	362	529	529	140	140
Number of observations	21,459	21,459	46,460	46,460	3,216	3,216

Standard errors in parentheses, clustered by location.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

against the "Islamization of Europe" (Reuters 2016). The marches drew thousands who came to express their opposition to the arrival of millions of migrants from Middle Eastern and North African countries, and to warn about Europe "being overrun by Muslims" (Reuters 2016). Since the anti-Muslim hostility expressed in the PEGIDA marches was likely to most resonate in areas where far-right parties are popular, and therefore stir hostility in these areas, I examine whether individuals located in these areas responded differently to the PEGIDA marches, compared to individuals located in areas with little far-right support.

Estimation

I estimate several heterogeneous event study models, where I examine whether the difference in the number of pro-ISIS tweets three days after the event is larger in areas that have higher vote-share for far-right parties. A 'pro-ISIS tweet' is coded one if its predicted value of belonging to any of the content categories is above the mean of the predicted values for that category, and zero if not. To measure content that explicitly sympathizes with ISIS, I also create a variable for the Sympathy with ISIS topic only. For each event, I estimate the following least squares model:

$$Y_{ijk} = \beta_1 T_i + \beta_2 V_{jk} + \beta_3 (T_i \times V_{jk}) + \boldsymbol{\delta}\mathbf{X}_{jk} + \alpha_k + \varepsilon_{jk}, \quad \textbf{(2)}$$

where Y_{ijk} represents the level of radical content in tweet i posted in area j and country k, T_i is an indicator coded one for tweets appearing after the event (Paris attacks, Brussels attacks, ISIS propaganda release, and the PEGIDA marches) and zero if before, V_{jk} is the locality-level vote share for far-right parties in area j in country k, $\mathbf{X}_{jk}$ represents other independent variables described in equation (1), α_k represents country fixed effects, and ε_{jk} are standard errors clustered at the locality level.

RESULTS

Terrorist Attacks and ISIS Propaganda

Table 9 presents the findings for the terrorist attacks and the ISIS propaganda release. Panel A reports the pooled results for each of these events. In the first few days after the terrorist attacks in Paris and Brussels, individuals posted significantly more pro-ISIS content on Twitter. We find a similar result after the release of ISIS's propaganda video, where ISIS followers posted more tweets sympathizing with ISIS after the video's dissemination.

Panel B in Table 9 reports the heterogeneity results, where I interact far-right vote share with the timing of the events. It can be seen that the largest changes in the number of pro-ISIS posts were concentrated among individuals in high far-right support areas. For the Paris attacks, this finding holds for all ISIS-related topics, as well

TABLE 10. PEGIDA Marches and Changes in Pro-ISIS and Anti-West Rhetoric

	PEGIDA marches		
	(1) Sympathy with ISIS	(2) ISIS topics	(3) ISIS topics + anti-West
A. Changes in pro-ISIS content (sd units)			
After event = 1	0.003	−0.022*	−0.015
	(0.009)	(0.012)	(0.009)
Constant	0.141***	0.243***	0.273***
	(0.038)	(0.030)	(0.032)
R^2	0.004	0.001	0.002
Number of clusters	577	577	577
Number of observations	56,402	56,402	56,402
B. Changes in pro-ISIS and far-right content (sd units), by far-right support			
After event = 1	−0.003	−0.035	−0.042**
	(0.021)	(0.023)	(0.020)
Far-right vote share (%)	0.001	0.000	0.000
	(0.002)	(0.001)	(0.001)
After event = 1 × far right vote share (%)	**0.001**	**0.002**	**0.002****
	(0.002)	(0.001)	(0.001)
Constant	0.093	0.135*	0.156*
	(0.100)	(0.078)	(0.087)
Controls	✓	✓	✓
Country fixed effects	✓	✓	✓
R^2	0.003	0.001	0.002
Number of clusters	508	508	508
Number of observations	38,527	38,527	38,527

Standard errors in parentheses, clustered by location.
$^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

as when restricting the analysis to sympathy with ISIS only. For the Brussels attack, the difference is only significant for the content outcome combining all ISIS-related topics. For the ISIS propaganda video, the changes are significant with the variable capturing sympathy with ISIS.

Anti-Muslim Marches

Table 10 presents the findings for the PEGIDA marches. It can be seen that the marches did not lead to an increase in pro-ISIS tweets when considering the sample as a whole. However, when examining how the responses varied by areas with different levels of far-right vote share, we find that pro-ISIS tweets increased in areas with high levels of far-right support. However, the difference is not statistically significant when restricting the analysis only to topics that explicitly discuss ISIS.

Since qualitative data revealed that ISIS supporters generate high levels of anti-West content, I also examine how tweeting patterns changed when considering a measure including the anti-West topic. I find that anti-West discourse significantly increased after the PEGIDA marches among individuals located in areas with high far-right support. While anti-West sentiment does not necessarily imply support for violent extremism, it is certainly part of the discourse of ISIS supporters, especially around incidents carried out by Western actors (Cunningham, Everton, and Schroeder 2017). The difference between the ISIS-initiated events (terrorist attacks and propaganda release) and the PEGIDA marches could be driven by the different kinds of actors involved (Western vs. non-Western). Nonetheless, the pattern across these events is strikingly similar.

Figure 7 presents these patterns visually, showing that across all four events, the increase in pro-ISIS and anti-West tweeting is significantly larger in areas with higher far right vote share. The figures also illustrate why the PEGIDA marches did not lead to a positive change in tweeting in the pooled analysis: in areas with low far-right

FIGURE 7. Responses to Terrorism, Propaganda, and Anti-Muslim Marches, by Far-Right Vote Share

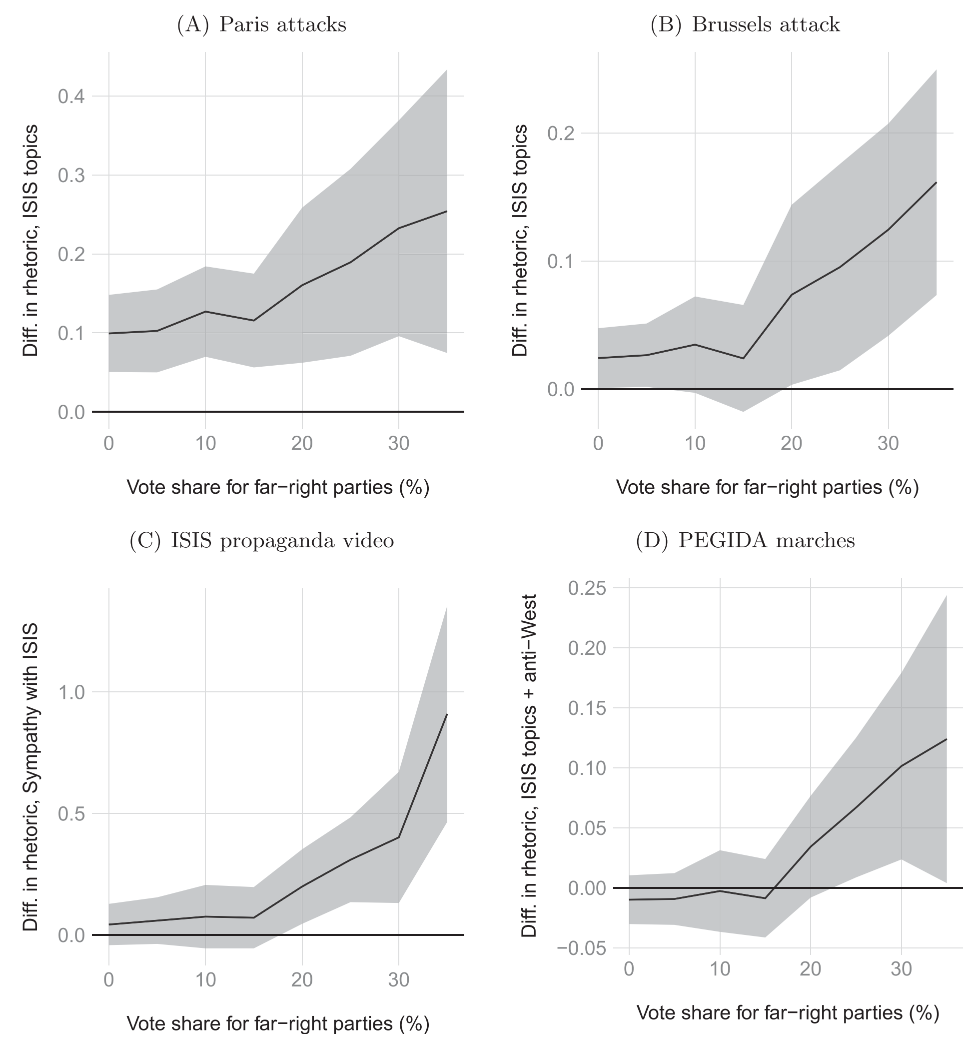

Note: The figure plots the difference in the frequency of pro-ISIS and anti-West tweets after various events for areas with different levels of far-right vote share. The differences are reported in standard deviation units.

vote share, the marches were followed by a *decrease* in the number of ISIS-related and anti-West tweets. Indeed, on the day of the PEGIDA marches, many counter-protests took place in opposition to the movement's positions (Worley 2016). This illustrates the powerful role of local context in facilitating support for ISIS.

As a robustness test, in the online appendix, I run the same estimation when interacting the events' timing with local-level unemployment. I find no difference in the tweeting patterns after the events in areas with high levels of unemployment.[41] This suggests that there is something unique about areas with high support for far-right parties that might shape support for extremism. As the results reflect responses to various types of events taking place in different (even arbitrary) points in time, these consistent findings provide further support for the hypothesis that a local context of anti-Muslim hostility can facilitate support for ISIS.

CONCLUSION

This study seeks to shed light on what drove so many to support the Islamic State in the West in the past several years. By collecting data on thousands of Twitter users affiliated with or following ISIS accounts, classifying millions of tweets along various dimensions of ISIS sympathy, and mapping Twitter users to geographic locations in France, Germany, Belgium, and the United Kingdom, I showed that those located in areas that voted for far-right, anti-Muslim parties were more likely to show signs of radicalization than others in less hostile areas. While some have noted that there might be a link between the rise of far-right parties and support for ISIS in Europe (Van Zeller 2016), this article has provided the first systematic, rigorous study of this proposition.

The findings stress the importance of understanding the process of radicalization and support for extremist movements in the age of social media. The ability to directly reach potential recruits on the Internet, interact with them through online platforms, and persuade them to embrace extremist ideology is changing how we think about recruitment in subnational conflicts. As the Internet and mobile technology continue to spread across the world, radicalization through the Internet is likely to continue, given the ongoing conflicts in the Middle East, North Africa, and other parts of the world. Studying how the online and offline worlds interact in this setting suggests that hostility in one's offline world might lead to the consumption of online radical content.

Looking forward, research on radicalization in the West would benefit from more localized studies aiming to causally identify the mechanisms by which hostility can facilitate support for extremism. Does an environment of anti-Muslim hostility increase support for jihadi ideologies through a process of identity-seeking? Or is it driven by lack of opportunity to integrate into the surrounding society, e.g., by finding employment or increasing social status? In addition, studies could unpack the role of inter-group contact in this setting, especially in light of the finding in prior research that support for far-right parties tends to be stronger in areas where minority communities are smaller (Biggs and Knauss 2012).[42]

Future work can also study the determinants of ISIS radicalization in non-Western countries. While some of the same mechanisms might be at play, descriptive evidence suggests that recruits' motivations, as well as ISIS's recruitment strategy have been different in Middle Eastern and North African countries (Raghavan 2016; Wilson 2015). Nonetheless, it is certainly possible that institutional exclusion in authoritarian settings might create a similar dynamic, where individuals who feel alienated from the regime become attracted to propaganda disseminated by ISIS and other extremist groups.

Finally, future studies might examine ways to de-radicalize potential recruits. With the rise of Islamic State recruitment on social media, there has been a dramatic increase in initiatives to counter extremism around the world: in just a few years, over forty countries have announced official national strategies to fight violent extremism, and many more initiatives have been launched by nongovernmental organizations.[43] While policy efforts such as the U.S. Department of State's "Think Again Turn Away" campaign have had limited impact (Fernandez 2015), other, more local de-radicalization efforts have reportedly been more successful in this and prior conflicts (Frenett and Dow 2016; Horgan 2015; Rabasa et al. 2010). The findings of this study, especially those showing differential results across localities around the PEGIDA marches, suggest that it might be fruitful to investigate the role of an inclusive local context in quelling support for extremism. A better understanding of what drives individuals to sympathize with foreign extremist groups could guide policymakers in responding to this troubling phenomenon.

SUPPLEMENTARY MATERIAL

To view supplementary material for this article, please visit https://doi.org/10.1017/S0003055418000618.

Replication material can be found on Dataverse at: https://doi.org/10.7910/DVN/5QUCW7.

REFERENCES

Adida, Claire L., David D. Laitin, and Marie-Anne Valfort. 2016. *Why Muslim Integration Fails in Christian-Heritage Societies*. Cambridge, MA: Harvard University Press.

Allport, Gordon W. 1979. *The Nature of Prejudice*. New York, NY: Basic Books.

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*. Raleigh, NC: ACM, 61–70.

41 See section S7 in the online appendix.

42 One might argue that the findings in this article support the 'contact hypothesis' that frequent interactions between groups can reduce inter-group hostility (Allport 1979).

43 Information collected by the author. See Figure S21 in the online appendix.

Barrett, Richard, Jack Berger, Lila Ghosh, Daniel Schoenfeld, Mohamed el-Shawesh, Patrick M. Skinner, Susan Sim, and Ali Soufan. 2015. "Foreign Fighters: An Updated Assessment of the Flow of Foreign Fighters into Syria and Iraq." *The Soufan Group*. http://soufangroup.com/wp-content/uploads/2015/12/TSG_ForeignFightersUpdate3.pdf.

Bartlett, Jamie, and Jonathan Birdwell. 2013. "Cumulative Radicalisation between the Far-Right and Islamist Groups in the UK: A Review of Evidence." *Demos* (3–15).

Bass, Loretta. 2014. *What Motivates European Youth to Join ISIS? Syria Comment: University of Oklahoma*. https://bit.ly/2AsjhFH.

Bayman, Daniel, and Jeremy Shapiro. 2014. "Be Afraid. Be a Little Afraid: The Threat of Terrorism from Western Foreign Fighters in Syria and Iraq." *Foreign Policy at Brookings*.

Benmelech, Efraim, and Esteban F. Klor. 2018. "What Explains the Flow of Foreign Fighters to ISIS?" *Terrorism and Political Violence*, 1–24. Published online 31 October 2018.

Berger, J. 2015. "The Metronome of Apocalyptic Time: Social Media as Carrier Wave for Millenarian Contagion." *Perspectives on Terrorism* 9 (4): 61–71.

Berger, J., and Jonathan Morgan. 2015. "The ISIS Twitter Census: Defining and Describing the Population of ISIS Supporters on Twitter." *The Brookings Project on US Relations with the Islamic World* 3: 20.

Biggs, Michael, and Steven Knauss. 2012. "Explaining Membership in the British National Party: A Multilevel Analysis of Contact and Threat." *European Sociological Review* 28 (5): 633–46.

Borum, Randy. 2011. "Radicalization into Violent Extremism I: A Review of Social Science Theories." *Journal of Strategic Security* 4 (4): 7.

Bryan, Jennifer. 2005. "Constructing 'The True Islam' in Hostile Times: The Impact of 9/11 on Arab Muslims in Jersey City." In *Wounded City: The Social Impact of 9/11*, ed. Nancy Foner. New York, NY: Russel Sage Foundation, 133–162.

Burrows, Thomas. 2016. "Islamophobic Hate Crimes in London Have Nearly Doubled in the Last Two Years as Muslims Say They Are 'Very Conscious of Keeping Your Head Down'." The Daily Mail. http://www.dailymail.co.uk/news/article-3358535/Islamophobic-hatecrimes-London-nearly-doubled-two-years-Muslims-say-conscious-keeping-headdown.html.

Carter, Joseph A., Shiraz Maher, and Peter R. Neumann. 2014. "# Greenbirds: Measuring Importance and Influence in Syrian Foreign Fighter Networks." *The International Centre For the Study of Radicalisation and Political Violence*.

Crone, Manni. 2016. "Radicalization Revisited: Violence, Politics and the Skills of the Body." *International Affairs* 92 (3): 587–604.

Cunningham, Daniel, Sean F. Everton, and Robert Schroeder. 2017. "Social Media and the ISIS Narrative." White Paper. Monterey, CA: Defense Analysis Department, Naval Postgraduate School.

Dabiq. 2015. "The Extinction of the Grayzone." *Dabiq* 7: 54–66.

Dalgaard-Nielsen, Anja. 2010. "Violent Radicalization in Europe: What We Know and what We Do Not Know." *Studies in Conflict & Terrorism* 33 (9): 797–814.

Dancygier, Rafaela M., and David D. Laitin. 2014. "Immigration into Europe: Economic Discrimination, Violence, and Public Policy." *Annual Review of Political Science* 17: 43–64.

della Porta, Donatella. 2018. "Radicalization: A Relational Perspective." *Annual Review of Political Science* 21: 461–74.

Eatwell, Roger. 2006. "Community Cohesion and Cumulative Extremism in Contemporary Britain." *The Political Quarterly* 77 (2): 204–16.

European Social Survey Round 7. 2014. "European Social Survey Round 7 Data (2014)." *Data File Edition 2.1*. NSD - Norwegian Centre for Research Data. http://www.europeansocialsurvey.org/data/download.html?r=7.

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*.

Flaig, Joseph. 2016. "Khola Hasan from Theydon Garnon Says She Has Experienced Increased Islamophobia in Epping Forest." *The Guardian*. http://bit.ly/2AFAAya.

Frenett, Ross, and Moli Dow. 2016. "One to One Online Interventions: A Pilot CVE Methodology." *Institute for Strategic Dialogue*. https://www.isdglobal.org/wp-content/uploads/2016/04/One2One%5C_Web%5C_v9.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. http://www.jstatsoft.org/v33/i01/.

Front National. 2017. "The National Front Platform (In French)." http://www.frontnational.com/le-projet-de-marine-le-pen/.

Goldenberg, Jacob, and Moshe Levy. 2009. "Distance Is Not Dead: Social Interaction and Geographical Distance in the Internet Era." arXiv preprint arXiv:0906.3202.

Golder, Matthew. 2016. "Far Right Parties in Europe." *Annual Review of Political Science* 19 (1): 477–97.

Gould, Eric D., and Esteban F. Klor. 2016. "The Long-run Effect of 9/11: Terrorism, Backlash, and the Assimilation of Muslim Immigrants in the West." *The Economic Journal* 126 (597): 2064–114.

Greenberg, Karen J. 2016. "Case by Case: ISIS Prosecutions in the United States." *Center on National Security*, Fordham University School of Law.

Grimmer, Justin, and Brandon M. Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21 (3): 267–97.

Gurr, Ted R. 1970. *Why Men Rebel*. Princeton, NJ: Princeton University Press.

Hainmueller, Jens, and Daniel J. Hopkins. 2014. "Public Attitudes toward Immigration." *Annual Review of Political Science* 17: 225–49.

Hall, Ellie. 2015. "Gone Girl: An Interview with an American in ISIS." *Buzzfeed News*. https://www.buzzfeed.com/ellievhall/gone-girl-an-interview-with-an-american-inisis?utm%5C_term=.grBaBZdyN%5C#.iv22gLK9N.

Hegghammer, Thomas. 2010. "The Rise of Muslim Foreign Fighters: Islam and the Globalization of Jihad." *International Security* 35 (3): 53–94.

Holland, Joshua. 2016. "Islamic Extremism Is Not the Root Cause of Europe's Terror Problem." *The Nation* (March 30). https://www.thenation.com/article/islamic-extremism-isnot-the-cause-of-europes-terror-problem/.

Horgan, John G. 2015. "De-Radicalization Programs Offer Hope in Countering Terrorism." *Los Angeles Times*. http://www.latimes.com/opinion/op-ed/la-oe-0215-horgan-terrorist-deradicalization-20150215-story.html.

Horowitz, Donald L. 1985. *Ethnic Groups in Conflict*. Oakland, CA: University of California Press.

Humphreys, Macartan, and Jeremy M. Weinstein. 2008. "Who Fights? the Determinants of Participation in Civil War." *American Journal of Political Science* 52 (2): 436–55.

Inge, Sophie. 2013. "Nazi Graffiti and Pig's Head Left at Mosque Site." *The Local*. http://www.thelocal.fr/20130415/nazi-graffiti-and-pigs-head-found-on-mosque-site.

Isaac, Mike. 2016. "Twitter Steps up Efforts to Thwart Terrorists' Tweets." *New York Times*. http://www.nytimes.com/2016/02/06/technology/twitter-account-suspensions-terrorism.html.

Jamal, Amaney. 2008. "Civil Liberties and the Otherization of Arab and Muslim Americans." In *Race and Arab Americans before and after 9/11: From Invisible Citizens to Visible Subjects*, eds. Amaney A. Jamal and Nadine Naber. Syracuse, NY: Syracuse University Press, 114–30.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2013. *An Introduction to Statistical Learning*. Vol. 112. New York: Springer.

Jurgens, David. 2013. "That's what Friends Are for: Inferring Location in Online Social Media Platforms Based on Social Relationships." *ICWSM* 13: 273–82.

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)*, eds. Daniele Quercia and Bernie Hogan. Palo Alto: AAAI Press, 188–97.

Karam, Nicoletta. 2012. *The 9/11 Backlash: A Decade of US Hate Crimes Targeting the Innocent*. Pennsauken, NJ: BookBaby.

Kent Online. 2015. "MP Gareth Johnson Contacts Police over Britain First "Anti-halal Operation" in Dartford." *Kent Online*. http://www.kentonline.co.uk/dartford/news/mp-contactspolice-over-far-right-38953/.

King, Gary, and Langche Zeng. 2001. "Logistic Regression in Rare Events Data." *Political analysis* 9 (2): 137–63.

Lubbers, Marcel, and Peer Scheepers. 2002. "French Front National Voting: A Micro and Macro Perspective." *Ethnic and Racial Studies* 25 (1): 120–49.

Lyons-Padilla, Sarah, Michele J. Gelfand, Hedieh Mirahmadi, Mehreen Farooq, and Marieke van Egmond. 2015. "Belonging Nowhere: Marginalization and Radicalization Risk Among Muslim Immigrants." *Behavioral Science & Policy* 1 (2): 1–12.

Malet, David. 2013. *Foreign Fighters: Transnational Identity in Civil Conflicts*. Oxford, UK: Oxford University Press.

Marans, Daniel. 2015. "What You Need to Know about Syrian Refugees Who Have 'Positive' Views of ISIS." *The World Post*. http://www.huffingtonpost.com/entry/syrian-refugeespositive-views-isis%5C_us%5C_564e2c72e4b08c74b734fc83.

McCaul, Michael, Bennie Thompson, John Katko, Loretta Sanchez, Barry Loudermilk, Filemon Vela, John Ratcliffe, Donald Payne, Will Hurd, and Martha McSally. 2015. "Final Report of the Task Force on Combating Terrorist and Foreign Fighter Travel." *Homeland Security Committee*.

McGee, Jeffrey, James A. Caverlee, and Zhiyuan Cheng. 2011. "A Geographic Study of Tie Strength in Social Media." In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*, eds. Bettina Berendt, Arjen de Vries, Wenfei Fan, Craig Macdonald, Iadh Ounis, and Ian Ruthven. Glasgow: ACM, 2333–6.

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*. eds. Qi He and Arun Iyangar. San Francisco: ACM, 459–468.

Mousseau, Michael. 2011. "Urban Poverty and Support for Islamist Terror: Survey Results of Muslims in Fourteen Countries." *Journal of Peace Research* 48 (1): 35–47.

Naber, Nadine. 2008. "Introduction: Arab Americans and U.S. Racial Formations." In *Race and Arab Americans before and after 9/11: From Invisible Citizens to Visible Subjects*, eds. Amaney A. Jamal and Nadine Naber. Syracuse, NY: Syracuse University Press, 1–45.

Neumann, Peter R. 2013. "The Trouble with Radicalization." *International Affairs* 89 (4): 873–93.

Norris, Pippa. 2005. *Radical Right: Voters and Parties in the Electoral Market*. New York, NY: Cambridge University Press.

Nussio, Enzo. 2017. "The Role of Sensation Seeking in Violent Armed Group Participation." *Terrorism and Political Violence*, 1–19.

Office for National Statistics. 2016. *Census Geography*. https://www.ons.gov.uk/methodology/geography/ukgeographies/censusgeography.

Petersen, Roger D. 2001. *Resistance and Rebellion: Lessons from Eastern Europe*. New York, NY: Cambridge University Press.

Rabasa, Angel, Stacie L. Pettyjohn, Jeremy J. Ghez, and Christopher Boucek. 2010. Deradicalizing Islamist Extremists. Technical Report, RAND Corp Arlington VA National Security Research Div.

Raghavan, Sudarsan. 2016. "Two of Her Daughters Joined ISIS. Now She's Trying to Save Her Two Younger Girls." *The Washington Post*. https://www.washingtonpost.com/world/middle%5C_east/two-of-her-daughters-joined-isis-now-shes-trying-to-save-her-twoyounger-girls/2016/05/25/700ea478-1d1e-11e6-82c2-a7dcb313287d%5C_story.html.

Reuters. 2016. "Anti-islam Movement PEGIDA Stages Protests across Europe." *Reuters*. http://www.reuters.com/article/us-europe-migrants-protests-germany-idUSKCN0VF0P4.

Robinson, Belinda. 2015. "ISIS Supporter, 17, Is Sentenced to 11 Years in Prison for Tweeting 7,000 Messages Backing the Terror Group and Helping a Friend to Join Fighters in Syria." *Daily Mail*. http://www.dailymail.co.uk/news/article-3214436/Virginia-teenage-rran-pro-ISIS-Twitter-account-sentenced-11-years-prison.html.

Rydgren, Jens. 2008. "Immigration Sceptics, Xenophobes or Racists? Radical Right-Wing Voting in Six West European Countries." *European Journal of Political Research* 47 (6): 737–65.

Scacco, Alexandra. 2018. *Anatomy of a Riot: Why Ordinary People Participate in Ethnic Violence*. WZB Berlin Social Science Center.

Sedgwick, Mark. 2010. "The Concept of Radicalization as a Source of Confusion." *Terrorism and Political Violence* 22 (4): 479–94.

Shane, Scott, Matt Apuzzo, and Eric Schmitt. 2015. "Americans Attracted to ISIS Find an 'Echo Chamber' on Social Media." *New York Times*. http://www.nytimes.com/2015/12/09/us/americans-attracted-to-isis-find-an-echo-chamber-on-social-media.html?%5C_r=0.

Stack, Liam. 2015. "American Muslims under Attack." *New York Times*. http://www.nytimes.com/interactive/2015/12/22/us/Crimes-Against-Muslim-Americans.html?%5C_r=0.

Takhteyev, Yuri, Anatoliy Gruzd, and Barry Wellman. 2012. "Geography of Twitter Networks." *Social Networks* 34 (1): 73–81.

The Atlantic. 2017. "The 'First Lady of ISIS'." *The Atlantic*. https://www.theatlantic.com/video/index/544688/after-jihad-documentary-film/.

Van Zeller, Mariana. 2016. "Radicals Rising." http://interactive.fusion.net/radicalsrising/.

Victoroff, Jeff, Janice R. Adelman, and Miriam Matthews. 2012. "Psychological Factors Associated with Support for Suicide Bombing in the Muslim Diaspora." *Political Psychology* 33 (6): 791–809.

Vidino, Lorenzo, and Seamus Hughes. 2015. "ISIS in America: From Retweets to Raqqa." *George Washington University: Program on Extremism*.

Walter, Barbara F. 2017. "The Extremist's Advantage in Civil Wars." *International Security* 42 (2): 7–39.

Wiktorowicz, Quintan. 2005. *Radical Islam Rising: Muslim Extremism in the West*. Lanham, MD: Rowman & Littlefield Publishers.

Wildman, Sarah. 2017. "Meet the Far-Right Party That's Bringing Racism and Xenophobia Back to Germany." *Vox*. https://www.vox.com/world/2017/9/26/16360916/afd-german-far-right-racism-xenophobia.

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.

Wilson, Lydia. 2015. "What I Discovered from Interviewing Imprisoned ISIS Fighters." *The Nation*. http://www.thenation.com/article/what-i-discovered-from-interviewing-isisprisoners/.

Wood, Elisabeth Jean. 2003. *Insurgent Collective Action and Civil War in El Salvador*. New York, NY: Cambridge University Press.

Worley, Will. 2016. "Thousands Take Part in Anti-islam Pegida Protests across Europe." *The Independent*. https://www.independent.co.uk/news/world/europe/thousands-take-part-in-anti-islam-pegida-protests-across-europe-a6857911.html.

From Isolation to Radicalization: Anti-Muslim Hostility and Support for ISIS in the West

Online Appendix

Contents

List of Tables

List of Figures

S1 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,[1] I designed an algorithm that continually monitored and recorded ISIS accounts identified by the hacktivist group @CtrlSec.[2] Immediately upon observing a new account in the @CtrlSec list, I downloaded 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 S1. The database contains "snapshots" of each user's profile at various points in time. In particular, between December 2015 and 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. Between May 2016 and January 2017, 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 S2.

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 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 all of these accounts. Thus, I began by downloading the full historical tweet timelines of all @CtrlSec-identified "ISIS activist" accounts ($N = 15,088$), as well as of all the friends of a sub-sample 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-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. Moreover, as noted previously, user profile information was 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, were 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 a current profile snapshot for a user, which contains their most

[1] https://dev.twitter.com/overview/documentation
[2] Lists are available in these handles: @ctrlsec, @ctrlsec0, @ctrlsec1, @ctrlsec2, @ctrlsec9.

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 over 1.6 million accounts on a daily basis, requesting latest profile and tweet, which led 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 100 million as of January 2017.

Table S1: 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 S2: 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

Figure S1: Scraping ISIS accounts

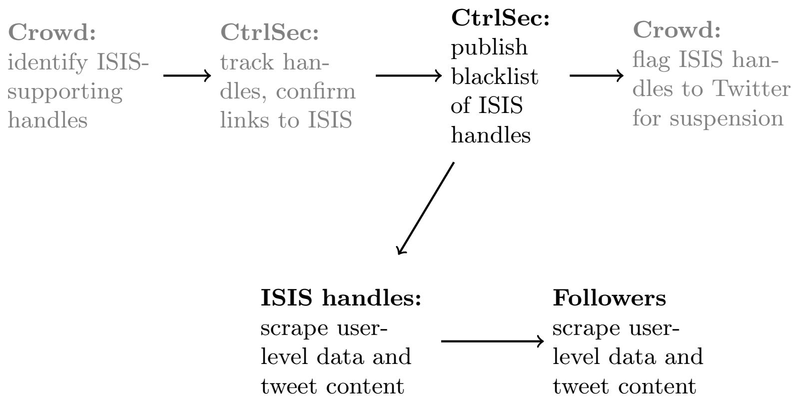

Table S3: Number of tweets posted by all users in database, by year

Year	# tweets
2007	849
2008	4,740
2009	42,667
2010	113,625
2011	376,627
2012	1,299,006
2013	3,285,090
2014	6,552,219
2015	17,887,290
2016	69,900,477
2017	4,903,609

Note: The number of tweets is accurate to 1/30/2017.

S2 Predicting geographic location of ISIS activists and followers

S2.1 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).

```
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 ∈ (U − domain(L)) (i.e., users who do not currently have location information) do
        Let M be a list of locations;
        for n ∈ N(u) (i.e., friends of user u) do
            if E(n) ≠ ∅ (i.e., if the friend n has location information) then
                add E(n) to M;
            end
        end
        if M ≠ ∅ (i.e., user u's friends have location information) then
            E'(u) = arg min_{x∈L} Σ_{y∈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)

Figure S2 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 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.[3]

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 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).[4] 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

[3] I employed three iterations, which predicted locations for 1,676,419 users in the database.

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

Figure S2: Spatial Label Propagation Algorithm

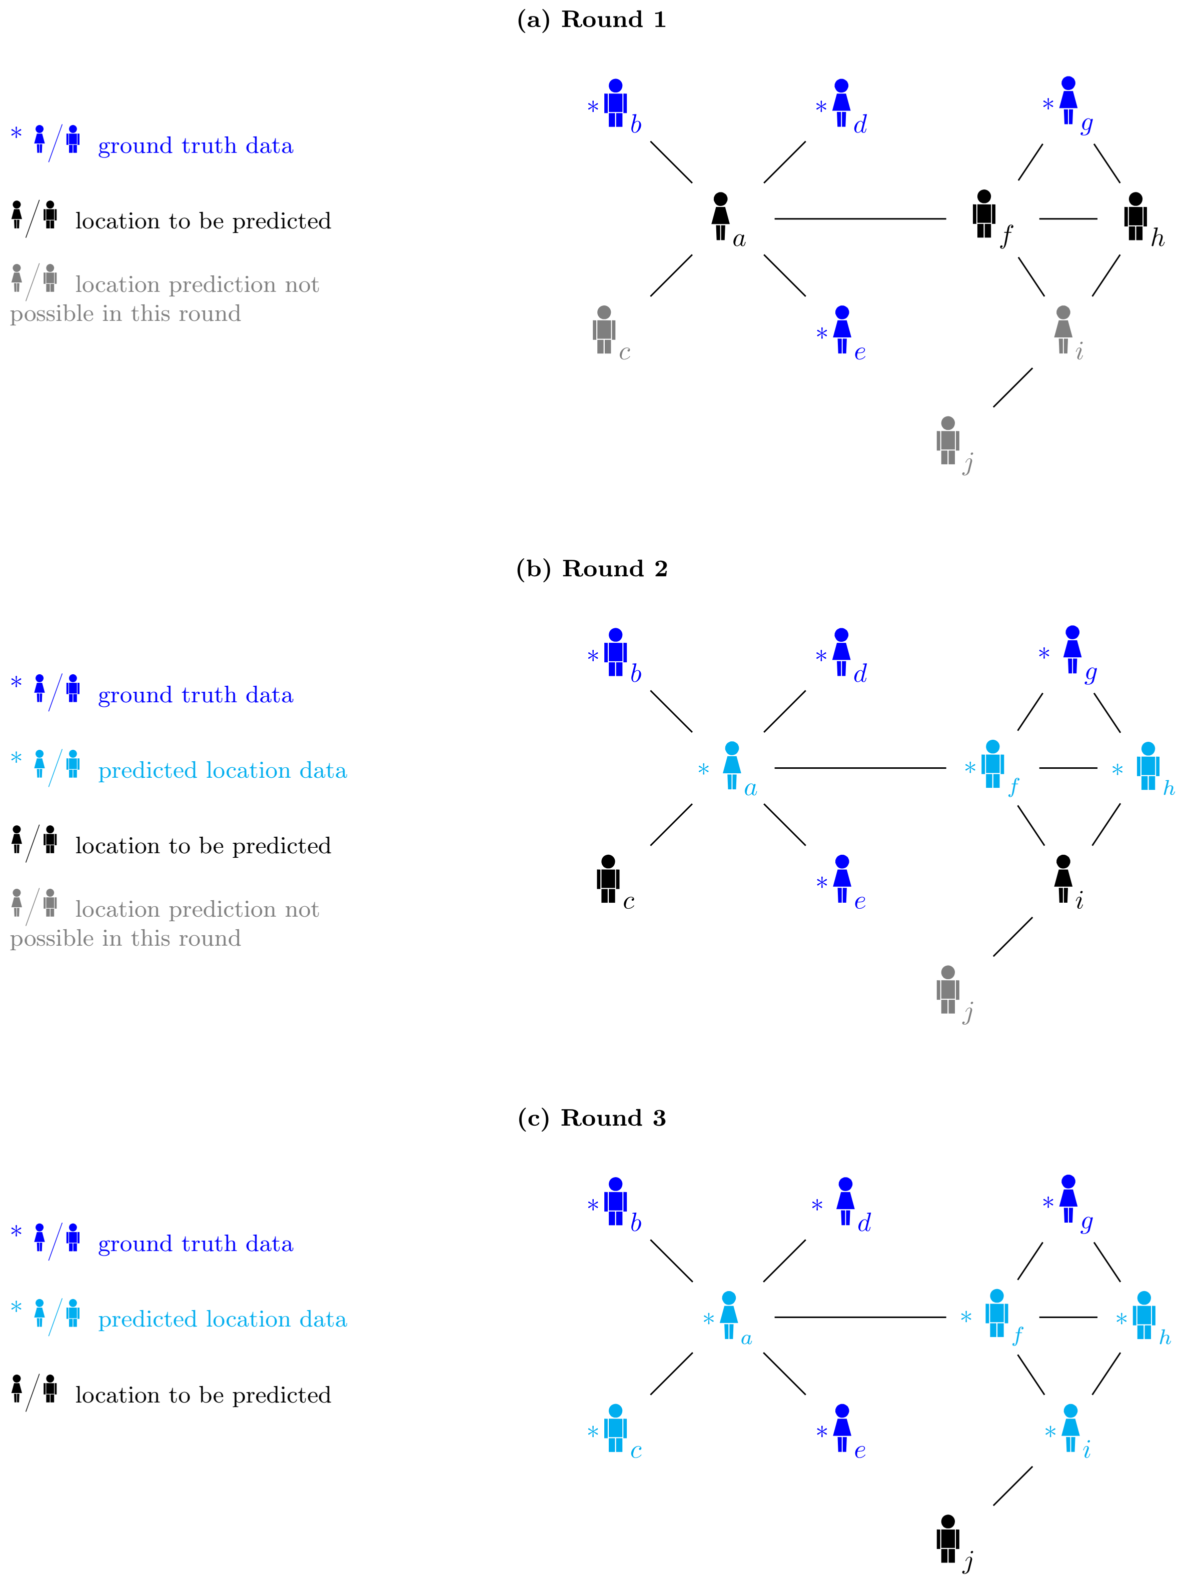

truth data ($N = 287,482$), these errors should be inconsequential.

S2.2 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 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 S3 shows the performance for the ISIS activists' accounts. Figure S4 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.

Figure S3: 10-Fold out-of-sample stability test (ISIS activists' accounts)

(a) Mean stability (b) Median stability

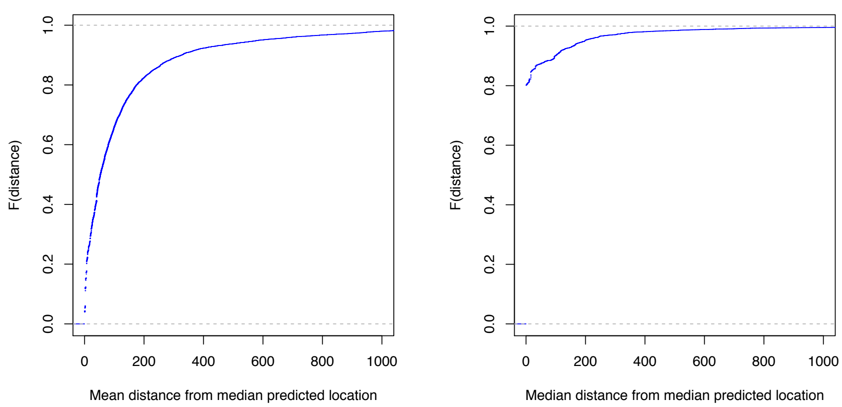

Note: The figure plots the cumulative distribution function of the stability of location predictions of ISIS activists across ten prediction estimations when leaving out one-tenth of the training data each time. In Panel (a), the stability is calculated as the mean of the predicted locations' deviations 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.

Figure S4: 10-Fold out-of-sample stability test (ISIS followers' accounts)

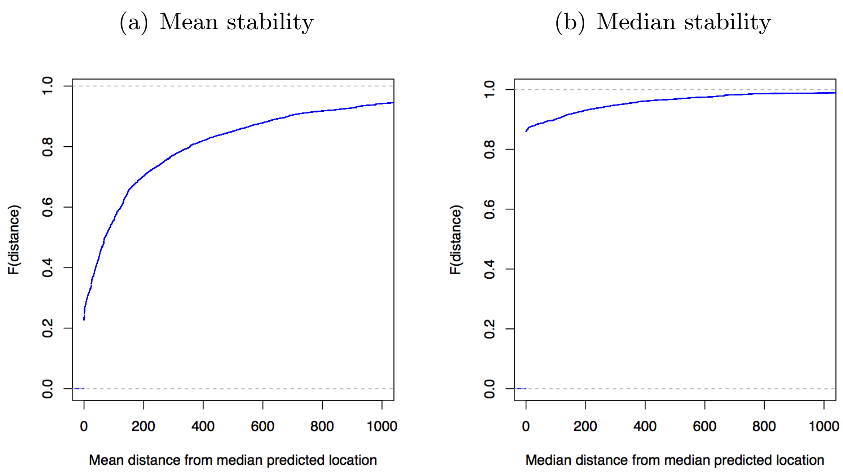

Note: The figure plots the cumulative distribution function of the stability of location predictions of ISIS followers across ten prediction estimations when leaving out one-tenth of the training data each time. In Panel (a), the stability is calculated as the mean of the predicted locations' deviations 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.

S2.3 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 the accounts are geo-enabled, the number of friends, followers, and twitter posts, as well as the language used by the users.

Table S4 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.

Table S4: 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 S5: 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

S2.4 Location prediction for individuals whose friends traveled to Syria

Another concern that may arise with the location prediction approach described above is that predictions will be biased for individuals whose friends have traveled to Syria. As the algorithm relies on the network of friends and their locations to predict geo-location, a person who has many friends that traveled to Syria is likely to be predicted to be in Syria. In the analysis in this paper, such an individual would be excluded from the sample, as this study only analyzes users whose locations are predicted to be in France, Germany, Belgium, and the UK.

It is important to note that the algorithm predicts locations by calculating the *geometric median* of the coordinates of a user's friends. Using the geometric median is crucial, since it predicts locations using the distribution of friends' actual locations. If I were to use the geometric mean — which I am not doing in this project — a user's location would be predicted to lie in places where they have no friends, or even in meaningless locations like the middle of the ocean. In addition, in the context of this study, using the geometric mean could bias the results by pulling out individuals located in cities to more rural areas where far-right parties might be more popular. This problem does not occur with the geometric median, where predicted locations are never pulled out of cities into rural areas if there are no friends in rural areas.

To visually show how this works, consider Figures S5 and S6, which display results from simulations using the geometric mean and the geometric median to predict a hypothetical user's location. In the simulation, user i is located in Paris, France, and has 100 friends. In each iteration, the distribution of user i's friends' locations changes, such that in the beginning most friends are located in Paris, and as the simulation progresses more and more move to Syria, Turkey, or Iraq. The simulation parameters are set such that out of the friends that travel abroad (whose number increases in each iteration), 60% are located in Raqqa, Syria, 30% in Mosul, Iraq, and 10% in Gaziantep, Turkey. The simulation shows what happens to the predicted location of user i as more of his or her friends travel to the Syrian civil war.

Figure S5 shows the results from the simulation using the geometric mean. Each point represents the predicted location of user i in each of the 100 simulation iterations. The color of the points changes from blue to red with each iteration, as more friends move out of Paris to Syria, Turkey or Iraq. The figure shows that the geometric mean introduces a lot of bias. In the early phases of the simulation, user i is still predicted to be in France, but having friends who traveled abroad pulls the user's predicted location out of Paris into more rural areas in France. Furthermore, as the proportion of user i's friends who travel abroad increases, user i's location is predicted to be outside of France, sometimes in arbitrary places like in the waters of the Mediterranean Sea. This example illustrates how serious the bias can be when using the geometric mean to predict a users' geo-locations.

However, we find a very different result when using the geometric median. Figure S6 shows that as more friends move out of Paris, user i's location shifts from Paris to Raqqa in Syria, but is never predicted to be in arbitrary locations outside of these two points. Specifically, with the parameters set in this simulation, user i's predicted location moves from Paris to Raqqa in the 52^{nd}

iteration, when 48 of the user's friends are in Paris, 31 in Raqqa (Syria), 15 in Mosul (Iraq), and 5 in Gaziantep (Turkey). Figure S6 shows that the blue points (which mark the earlier phases of the simulation) are located in Paris, and the red points (which mark the later phases of the simulation) are located in Raqqa in Syria. For easier visualization, in both figures I jittered the coordinates of user i's predicted location.

Figure S5: Location prediction with geometric mean (NOT used in this paper)

Figure S6: Location prediction with geometric median (the method used in this paper)

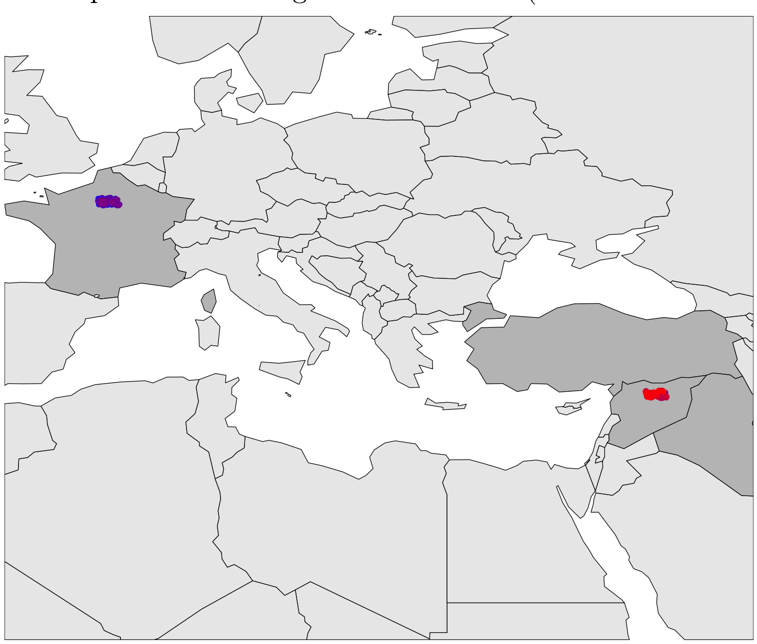

S2.5 Location prediction error and far-right vote share

Finally, one might worry that location prediction errors might have the effect of spreading out users in cities to rural areas outside of cities, where voters may be more inclined to vote for far-right parties. If measurement errors diffuse users in this way, then the results might be biased by erroneously having more users in areas with greater far-right support.

I examine this possibility in two ways. First, in Table S6 I test if there is in-sample correlation between location prediction errors and far-right vote share. Using the mean and median location prediction stability measures described in section S2.2 as the dependent variable, I estimate regressions on far-right vote share. The results show that there is no systematic relationship between far-right support and in-sample location prediction errors.

Second, I test if there is evidence of out-of-sample prediction error by examining the spatial distribution of predicted locations. As the Spatial Label Propagation algorithm predicts locations on the basis of the number of friends in each area, we would intuitively expect prediction errors to be biased in favor of cities (rather than rural areas), simply due to population size. Indeed, this can be seen in Columns (1) and (3) in Table S7, which regresses the number of users in each locality on the population and far-right support. These columns show that, as expected, more users are predicted to be located in areas with larger populations.

Nonetheless, if the effect of the prediction error is to spread out users outside of cities to rural areas where far-right parties are more popular, we would expect to observe a link between far-right vote share and the number of accounts in each location, even after adjusting for population size. There are two possible patterns that might emerge in the data. One possibility is that location prediction errors place *more users* in certain localities with high levels of far-right support. If this were the case, then we would observe a positive relationship between far-right vote share and the number of users in each locality. Another possibility is that measurement errors place users in a *greater number of areas* with high levels of far-right support. In the second scenario, we will observe a negative relationship between far-right vote share and the number of accounts, as diffusion will result in fewer users in each location.

As can be seen in Columns (2) and (3) in Table S7, the relationship between the number of users in each locality and far-right vote share is statistically insignificant. This finding holds whether population size is accounted for or not, as well as when adjusting for country fixed effects. Overall, these findings suggest that prediction errors are not spreading out users into areas with greater voting for far-right parties. In addition, it is worth recalling that the dependent variable in this study does not examine the number of users in each locality, but the correlation between the content that they produce and far-right vote share. A greater number of users in high far-right areas does not necessarily imply any such correlation.

Table S6: Location prediction errors and far-right vote share

	Mean error (km)	Median error (km)
Far-right vote share (%)	3.25	1.68
	(3.60)	(1.28)
Constant	248.22***	53.69***
	(32.68)	(12.24)
Country fixed effects	✓	✓
R^2	0.021	0.004
Number of clusters	3,136	3,136
Number of observations	116,465	116,465

Robust standard errors in parentheses, clustered at the locality level.
Base category is Belgium.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table S7: Predicted locations and far-right vote share

	Dependent variable: Number of accounts		
	(1)	(2)	(3)
Population	0.01***		0.00***
	(0.00)		(0.00)
Far-right vote share (%)		-1.89	-1.43
		(1.66)	(1.76)
Constant	3.44	21.95	19.95
	(71.96)	(54.79)	(58.59)
Country fixed effects	✓	✓	✓
R^2	0.191	0.001	0.007
Number of observations	2987	3140	2803

Standard errors in parentheses. Base country is Belgium.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

S3 Classifying Twitter content

To generate the textual content outcomes in this study, I used supervised machine learning to classify tweets into several categories. The categories classified by the model included (1) Anti-West, (2) Sympathy with ISIS, (3) Life in ISIS territories, (4) Travel to Syria or foreign fighters, and (5) Syrian war. When developing my training set, I coded content into additional categories, including references to Islam (expressions of faith, Islamic quotes, and prayers and/or requests for prayers), as well as Islamophobia (content describing discrimination against Muslims). These two categories are not utilized in my analysis, which focuses on radicalization. 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 Crowdflower 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 S11 for an example of instructions for the classification task in the Crowdflower 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 the Arabic language, I similarly removed punctuation and numbers. To pre-process Arabic tweets, I used the R package arabicStemR to stem Arabic text (Nielsen, 2017). See `https://CRAN.R-project.org/package=arabicStemR` for more details.

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 database were coded as unrelated to any of the above categories. Class proportions for each language in the training set are shown in Tables S8 – S11. 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 λ by cross-validation to maximize the area under the ROC curve. Figures S7 – S10 show the cross-validation curves for each language and topic. 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.

S3.1 Model performace

Model performance statistics from 10-fold cross validation for each topic and language are shown in Tables S12 – S11. It can be seen that the models were able to predict the content categories with high levels of in-sample accuracy. For example, for the Sympathy with ISIS topic in English, the accuracy rate is over 99.3%. This means that the misclassification rate is less than 1% for this topic and language. For the same topic in Arabic, the in-sample accuracy is 99.4%, for French it is 99.4% and for German it is 96.2%. As can be seen below, we find similar metrics for other topics and languages.

These high accuracy rates are driven by the fact that tweets labeled as these pro-ISIS topics are extremely different from tweets on other topics. The difference in content is related to the rare frequency of these categories: in the entire population of tweets, there may very well be content that has similar words and phrases to these pro-ISIS topics, but occurs so infrequently that it was not included in my training set. Those population tweets may be incorrectly classified as belonging to one of these pro-ISIS categories as a result. It is thus reasonable to suppose that my sample may contain more false positives than false negatives.

However, it is unlikely that my sample contains many such false positives because the proportion of tweets containing these topics in the sample is extremely small (see Tables S8 - S11 for an illustration of the distribution of these topics in the training set). Further, if there are a small number of false positives, there is little reason to think they would be concentrated in far-right areas. The consistency of my text-based results with non-text measures like being flagged as an ISIS activist, suspension, and the number of activist accounts followed suggests that false positives in the textual variables are not biasing my estimates.

Table S8: Class proportions by topic (English)

	0	1
Anti-West	0.984577	0.015423
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 S9: Class proportions by topic (Arabic)

	0	1
Anti-West	0.998104	0.001896
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 S10: Class proportions by topic (French)

	0	1
Anti-West	0.971370	0.028630
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 S11: Class proportions by topic (German)

	0	1
Anti-West	0.959585	0.040415
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 S12: 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

Table S13: 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

Table S14: 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

Table S15: 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

Figure S7: Cross validation for model choice (English tweets)

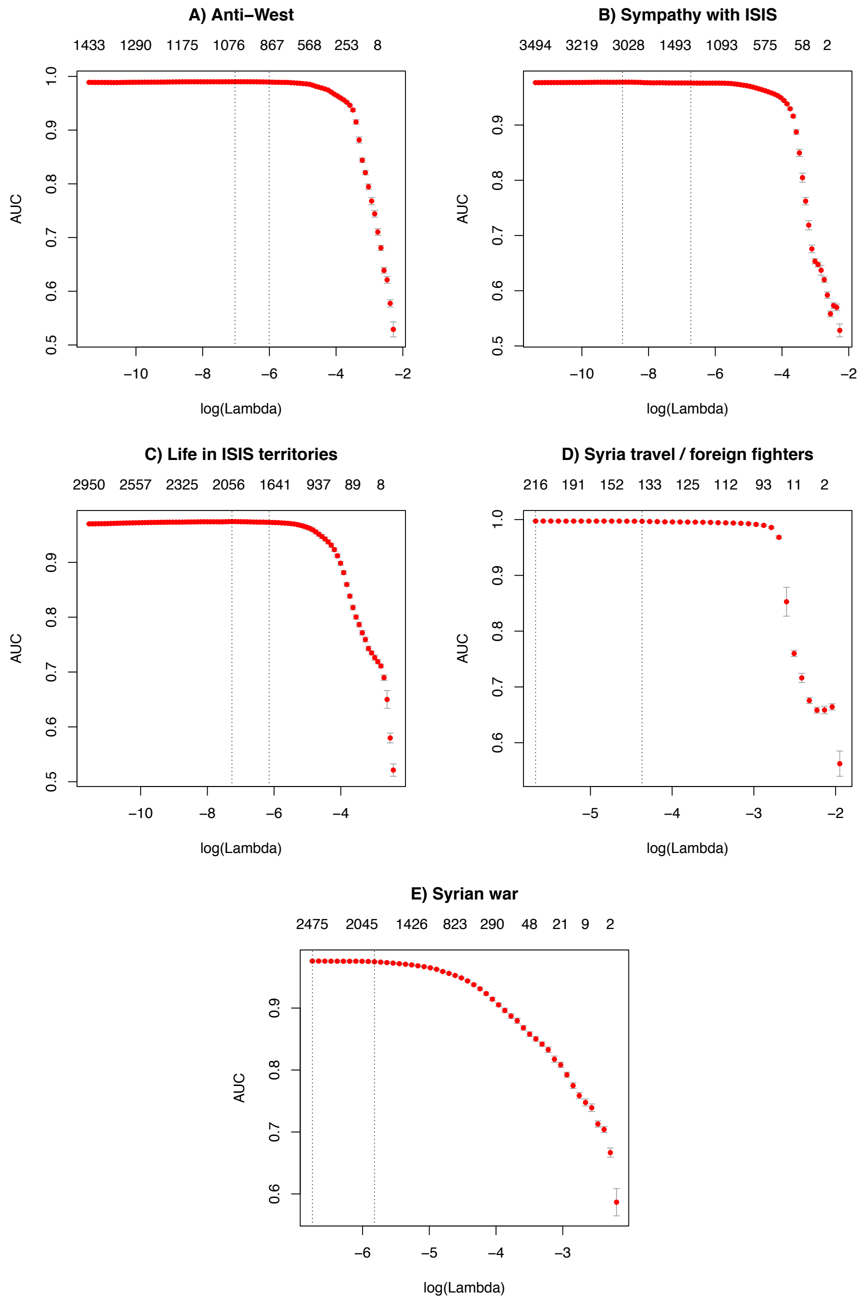

Note: The figure shows cross-validation curves for model choice in text classification of English language tweets for six topics. The cross-validation estimates for each model are shown in red dots, surrounded by error bars, plotted against the λ sequence. The y axis marks the Area Under the ROC Curve (AUC). Two selected λs 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 S8: Cross validation for model choice (Arabic tweets)

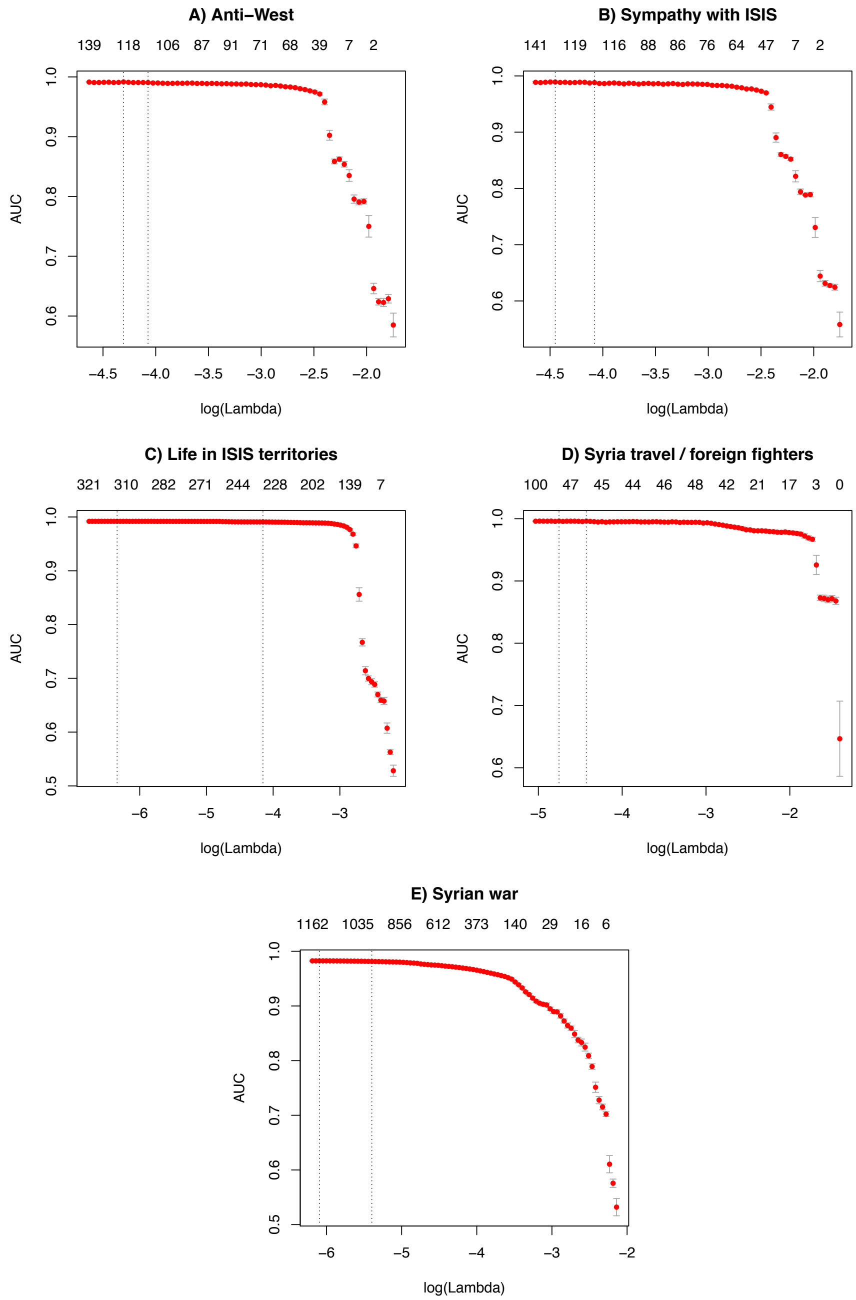

Note: The figure shows cross-validation curves for model choice in text classification of Arabic language tweets for six topics. The cross-validation estimates are shown in red dots, surrounded by error bars, plotted against the λ sequence. The y axis marks the Area Under the ROC Curve (AUC). Two selected λs 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 S9: Cross validation for model choice (French tweets)

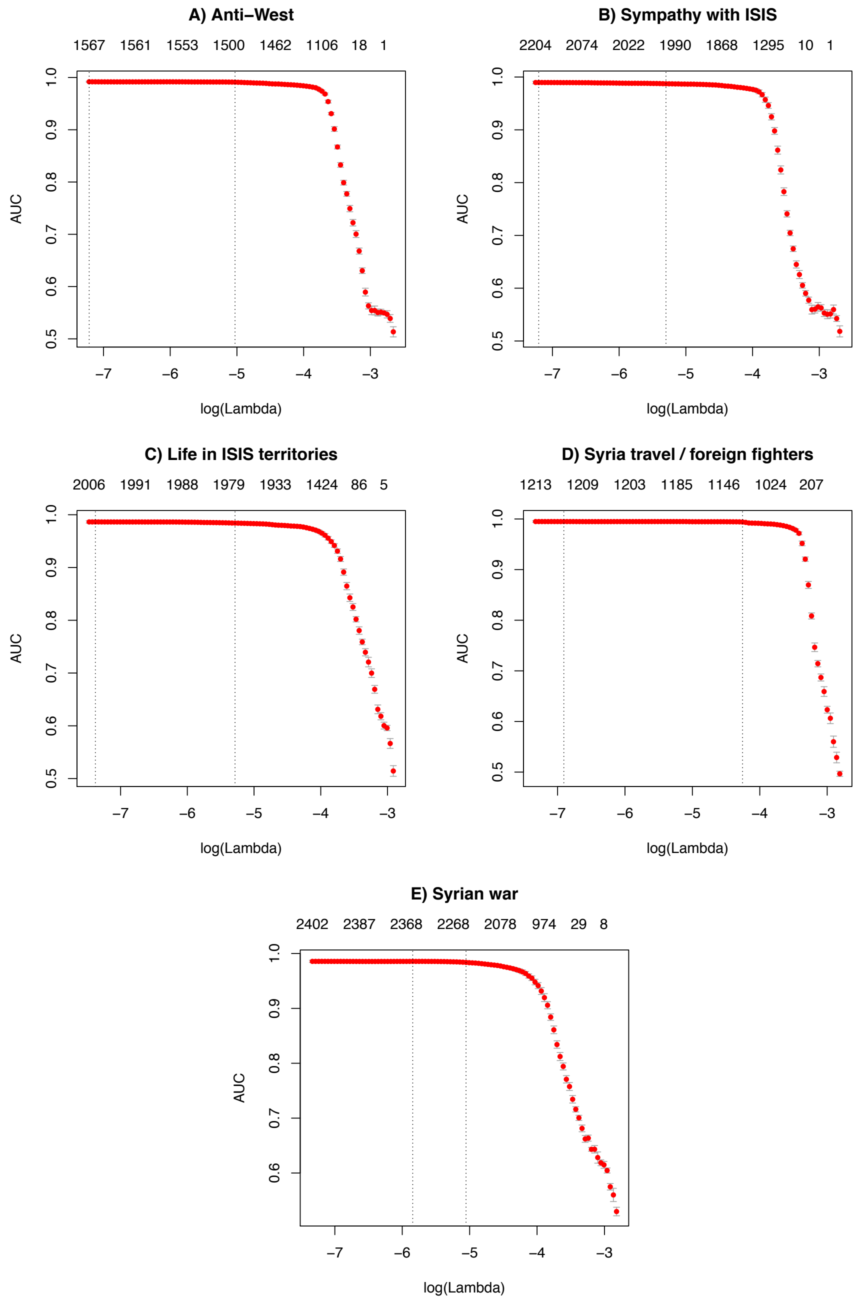

Note: The figure shows cross-validation curves for model choice in text classification of French language tweets for six topics. The cross-validation estimates are shown in red dots, surrounded by error bars, plotted against the λ sequence. The y axis marks the Area Under the ROC Curve (AUC). Two selected λs 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 S10: Cross validation for model choice (German tweets)

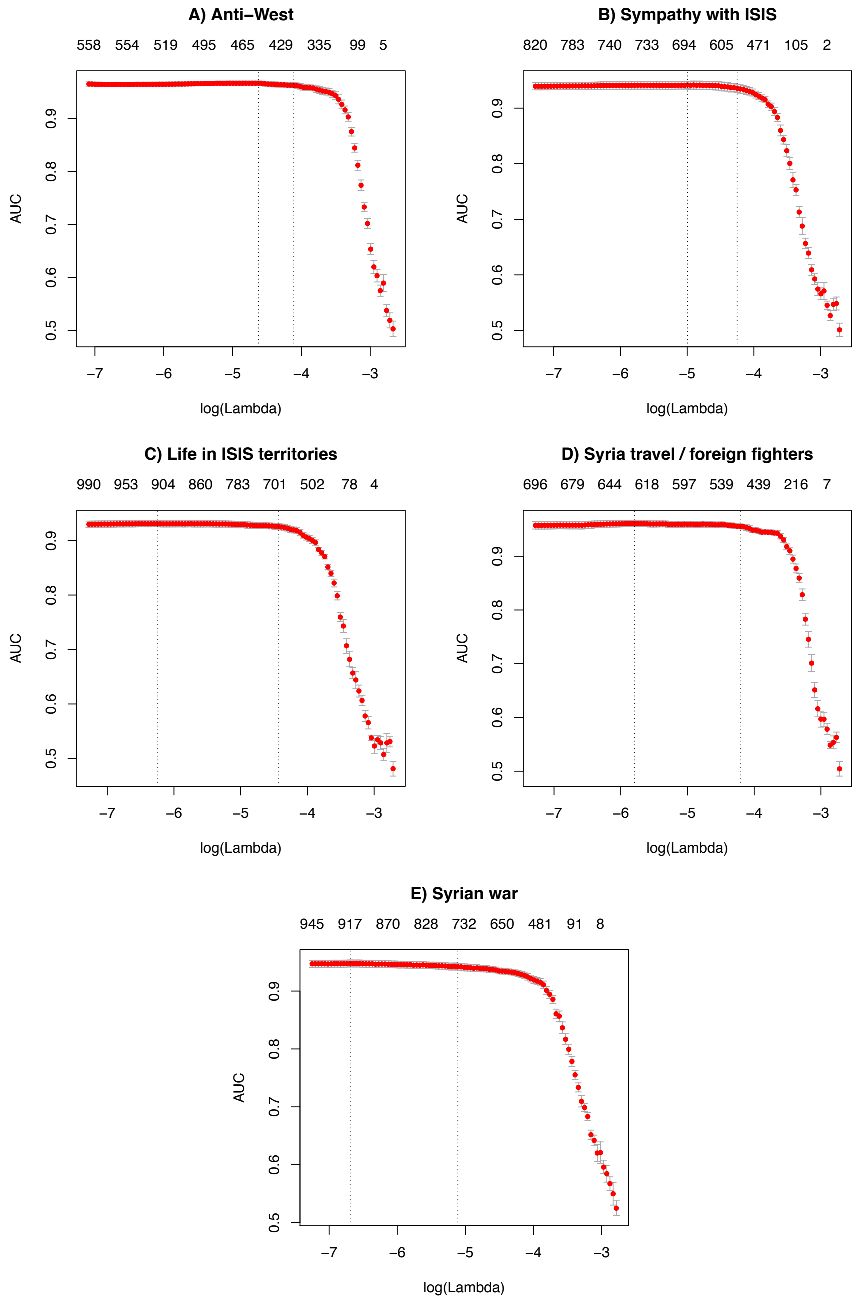

Note: The figure shows cross-validation curves for model choice in text classification of German language tweets for six topics. The cross-validation estimates are shown in red dots, surrounded by error bars, plotted against the λ sequence. The y axis marks the Area Under the ROC Curve (AUC). Two selected λs 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 S11: 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.

Figure S12: Supervised machine learning

$$\begin{tikzcd} \textbf{Random sample of tweets}^* \arrow[r] & \textbf{Label tweets} \arrow[r] \arrow[d] & \textbf{Train models} \arrow[r] \arrow[d] & \textbf{Predict unlabeled tweets' content} \\ & \text{971 coders (Crowdflower, Mturk)}^{**} & \text{Using labeled tweets}^{***} & \end{tikzcd}$$

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.

S4 Collecting administrative data from European countries

To assign independent variables to each user in my database, I collected administrative data from France, Germany, Belgium and the United Kingdom on far-right vote share, percent unemployment, share of foreigners, population size, and additional variables described below. I matched each variable to its corresponding spatial polygon using shape files from official government databases. Then, I used Twitter users' predicted geo-location data and the shape files of local administrative areas to assign users to areas with local-level socioeconomic data. This process was done in `R`, and the code to replicate the point-to-polygon matching is available upon request.

S4.1 Far-right vote share

France I obtained data on voting results in the 2015 French Departmental Elections at the polling station level from France's open platform of public data.[5]. The data contain information on the votes for each party in each polling station, the total eligible votes, as well as the electoral canton in which each polling station is located, among other variables. I aggregated the votes for the Front National party to the electoral canton level, and then divided the raw vote total for the party by the total eligible votes in each electoral canton. I used the electoral canton level vote share because of the availability of shape files at that level.

Germany I obtained data on voting results in the 2013 Federal Elections in Germany at the constituency level from Germany's Federal Returning Officer's Office.[6] For each constituency, I calculated the percent vote share in the Second Vote for the National Democratic Party of Germany (NPD) and the Alternative for Germany (AfD) party.

United Kingdom I obtained information on the vote share of the United Kingdom Independence Party (UKIP), British Democrats, British National Party, Liberty GB party, and the National Front party in the United Kingdom's 2015 General Elections from the country's Electoral Commission website.[7] For each constituency, I calculated the percent vote share for these parties.

Belgium I downloaded voting results from the 2014 Belgian Federal Elections at the municipality level from the country's Election Board website.[8] I calculated the vote share for Vlaams Belang for each constituency.

[5] `https://www.data.gouv.fr/fr/datasets/elections-departementales-2015-resultats-par-bureaux-de-vote/`
[6] `https://www.bundeswahlleiter.de/en/bundestagswahlen/BTW_BUND_13/ergebnisse/wahlkreisergebnisse/index.html`
[7] `http://www.electoralcommission.org.uk/our-work/our-research/electoral-data`
[8] `http://www.elections.fgov.be/index.php?id=3265&L=1`

S4.2 Socioeconomic data

France I obtained data on unemployment, share of foreigners, number of asylum seeker centers, and population size from the National Institute of Statistic and Economic Studies (INSEE).

1. *Unemployment (2011).* Unemployment at the municipality level the 2011 census.[9]
2. *Share of foreigners (2011).* The share of non-nationals in each municipality from the 2011 census.[10]
3. *Asylum seekers (2014).* The number of asylum seeker centers in each municipality as of 2014.[11]
4. *Population (2011).* Population size in each municipality from the 2011 census.[12]

Germany I downloaded data on unemployment, immigration, asylum seeker benefit receivers, and population size at the municipality level from The Regional Database Germany.[13] In order to access the data, it is necessary to create an account. Thus, I provide the names of the tables that I downloaded from the database.

1. *Unemployment (2015).* Unemployed individuals by selected groups of persons (Arbeitslose nach ausgewählten Personengruppen)
2. *Share of foreigners (2014).* Immigration and emigration by gender and age groups, over municipal boundaries, yearly total (Zu- und Fortzüge nach Geschlecht und Altersgruppen, über Gemeindegrenzen, jahressumme)
3. *Asylum seeker benefits receivers (2014).* Recipients of asylum seekers standard benefits, by gender, type of service, and age groups (Empfänger von Asylbewerberregelleistungen, Geschlecht, Art der Leistung, Altersgruppen)
4. *Population size (2011).* Population size at the municipality level from the 2011 census.

United Kingdom I obtained data from the 2011 census on unemployment, immigration, population size, religion, and ethnicity at the level of the Mid-layer super output area (MSOA), which is roughly equal to the size of a neighborhood, from the United Kingdom's Office of National Statistics.[14] I provide the names and numbers of tables that I downloaded from the database.

1. *Unemployment (2011).* KS601UK – Economic activity
2. *Share of foreigners (2011).* QS803EW – length of residence in the UK
3. *Population (2011).* KS101EW – Usual resident population

[9] http://www.insee.fr/fr/themes/detail.asp?reg_id=99&ref_id=td-population-13
[10] http://www.insee.fr/fr/themes/detail.asp?reg_id=99&ref_id=td-nationalite-13
[11] http://www.insee.fr/fr/themes/detail.asp?reg_id=99&ref_id=equip-serv-action-sociale
[12] http://www.insee.fr/fr/themes/detail.asp?reg_id=99&ref_id=td-population-13
[13] https://www.regionalstatistik.de/genesis/online/online;jsessionid=EE45147898822814978BE734145275C4?operation=sprachwechsel&option=en
[14] https://www.ons.gov.uk

4. *Religion (2011).* LC1202EW – Household composition by religion of Household Reference Person (HRP)

5. *Ethnic group (2011).* KS201EW – Ethnic group

Belgium I downloaded data on unemployment, immigration, and population at the statistical sector (sub-municipality) level from the 2011 Belgian census.[15] I provide the names of the tables that I downloaded from the database.

1. *Unemployment (2011).* Employed population by gender and age group - Total population - Statistical Sector (Werkende bevolking naar geslacht en leeftijdsklasse - Totale bevolking - Statistische sector)

2. *Share of foreigners, population (2011).* Population of Belgian and foreign nationality by gender – Statistical sector (Bevolking van Belgische en vreemde nationaliteit naar geslacht - Statistische sector)

S4.3 Stability of socioeconomic data over time

Since the Twitter data in this study covers content posted between 2014–2016, and the local administrative data captures socioeconomic conditions in earlier years (2011–2015), one might wonder how this gap might affect the results. As long as local-level socioeconomic data stay stable over time, the results should hold. To test the stability of these data, I collected information on every variable on which I could find over time information. Since yearly data in the relevant years is only available for France and Germany, I present results for these countries.[16] Table S16 presents the over-time correlations in unemployment and share of foreigners for each locality in France and Germany. In the analysis, I regressed each year's local data at time t on the data at time $t-1$. It can be seen that local-level socioeconomic data are highly stable over time.

Table S16: Stability of local-level socioeconomic data over time

	France		*Germany*	
	(1)	(2)	(3)	(4)
	Foreigners	Unemployment	Foreigners	Unemployment
	(2011-2014)	(2009-2014)	(2013-2015)	(2013-2015)
$t-1$	0.964***	0.607***	1.066***	0.143***
	(0.001)	(0.005)	(0.011)	(0.003)
Constant	0.002***	0.042***	0.002**	0.020***
	(0.00004)	(0.0004)	(0.001)	(0.0001)
R^2	0.958	0.291	0.326	0.102
Number of observations	106,170	35,900	20,503	20,415
Note:			*p<0.1; **p<0.05; ***p<0.01	

[15] http://census2011.fgov.be/download/statsect_nl.html

[16] The latest local-level socioeconomic data from the U.K. and Belgium comes from the 2011 census.

S4.4 Shape files

France I obtained shape files for the electoral cantons in France's 2015 Departmental Elections from the country's open platform of public data.[17]. For other administrative data, I obtained shape files of the contours of France's municipalities from France's open platform for public data.[18]

Germany I downloaded shape files of electoral constituencies in the 2013 German Federal Elections from Germany's Federal Returning Officer's Office.[19] For other socioeconomic variables, I used shape files from the contours of Germany's administrative boundaries.[20]

United Kingdom I obtained shape files for UK parliamentary constituencies from *MapIt*, a charity that provides data on contours of administrative areas in the United Kingdom.[21] I then matched the constituency-level vote share of far-right parties to the relevant polygon. For census data at the MSOA level, I used shape files from the Office of National Statistics.[22]

Belgium I downloaded the shape files of the contours of Belgium's statistical sectors (sub-municipality level) from Statistics Belgium, the official website of national statistics.[23]

[17] https://www.data.gouv.fr/fr/datasets/contours-des-cantons-electoraux-departementaux-2015/
[18] https://www.data.gouv.fr/fr/datasets/geofla-communes/
[19] https://www.bundeswahlleiter.de/en/bundestagswahlen/BTW_BUND_13/wahlkreiseinteilung/kartographische_darstellung.html
[20] https://www.zensus2011.de/DE/Infothek/Begleitmaterial_Ergebnisse/Begleitmaterial_node.html
[21] https://mapit.mysociety.org/areas/WMC.html
[22] http://webarchive.nationalarchives.gov.uk/20160105160709/http://www.ons.gov.uk/ons/guide-method/geography/products/census/spatial/2011/index.html
[23] http://statbel.fgov.be/nl/statistieken/opendata/datasets/tools/geografisch/

S5 Social media usage by ISIS supporters in the United States

Table S17 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 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 62%). Among those who used social media, the vast majority (about 86%) posted publicly their support for ISIS.

Table S17: 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	Nicholas Michael Teausant	Acampo, CA	1	1
8	Adam Dandach	Orange County, CA	0	0
9	Enrique Marquez Jr.	Riverside, CA	0	0
10	Aws Mohammed Younis al-Jayab	Sacramento, CA	1	0
11	Mahamad Saeed Koadimati	San Diego, CA	1	0
12	Shannon Maureen Conley	Denver, CO	1	0
13	James Gonzalo Medina	Hollywood, FL	0	0
14	Harlem Suarez	Key West, FL	1	1
15	Gregory Hubbard	West Palm Beach, FL	1	0
16	Dayne Antani Christian	Lake Park, FL	0	0
17	Darren Arness Jackson	West Palm Beach, FL	0	0
18	Miguel Moran Diaz	Miami-Dade, FL	1	1
19	Robert B. Jackson	Pensacola, FL	1	1
20	Leon Nathan Davis	Augusta, GA	0	0
21	Hasan R. Edmonds	Aurora, IL	1	1
22	Jonas M. Edmonds	Aurora, IL	0	0
23	Mhammed Hamzah Khan	Bolingbrook, IL	0	0
24	Ramiz Zijad Hodzic	Saint Louis, MO	1	1
25	Sedina Unkic Hodzic	Saint Louis, MO	1	1
26	Nihad Rosic	Utica, NY	1	1
27	Mehida Medy Salkicevic	Schiller Park, IL	1	1
28	Armin Harcevic	Saint Louis, MO	1	1
29	Jasminka Ramic	Rockford, IL	1	1
30	Abdullah Ramo Pazara	Saint Louis, MO	1	0
31	Akrami I. Musleh	Brownsburg, IN	1	1
32	Alexander E. Blair	Topeka, KS	0	0
33	John T. Booker	Topeka, KS	1	1
34	Alexander Ciccolo	Adams, MA	1	1
35	David Wright	Everett, MA	0	0
36	Mohamed Elshinaway	Edgewood, MD	1	1
37	Khalil Abu Rayyan	Dearborn Heights, MI	1	1
38	Sebastian Gregerson	Detroit, MI	0	0
39	Al-Hamzah Mohammad Jawad	East Lansing, MI	0	0
40	Abdirizak Mohamed Warsame	Eagan, MN	0	0
41	Abdul Raheem Habil Ali-Skelton	Glencoe, MN	0	0
42	Mohamed Abdihamid Farah	Minneapolis, MN	0	0
43	Adnan Abdihamid Farah	Minneapolis, MN	1	1
44	Abdurahman Yasin Daud	Minneapolis, MN	0	0
45	Zacharia Yusuf Abdurahman	Minneapolis, MN	0	0
46	Hanad Mustafe Musse	Minneapolis, MN	0	0
47	Guled Ali Omar	Minneapolis, MN	0	0
48	Hamza Ahmed	Minneapolis, MN	1	1
49	"H.A.M"	Burnsville, MN	1	1
50	Abdullahi Yusuf	Inver Grove Heights, MN	1	1
51	Abdi Nur	Minneapolis, MN	1	1
52	Yusra Ismail	St. Paul, MN	0	0
53	Safya Roe Yassin	Bolivar, MO	1	1
54	Jaelyn Delshaun Young	Starkville, MS	1	1
55	Muhammad Oda Dakhlalla	Starkville, MS	0	0

Note: The table provides details on the social media usage d 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.

Social media usage by ISIS supporters in the United States

	Name	Location	Used social media	Posted public posts
56	Justin Nojan Sullivan	Burke County, NC	1	0
57	Erick Jamal Hendricks	Charlotte, NC	1	1
58	Avin Marsalis Brown	Raleigh, NC	0	0
59	Akba Johad Jordan	Raleigh, NC	0	0
60	Donald Ray Morgan	Rowan County, NC	1	1
61	Nader Saadeh	Rutherford, NJ	1	1
62	Alaa Saadeh	West New York, NJ	0	0
63	Samuel Rahamin Topaz	Fort Lee, NJ	1	1
64	Tairod Nathan Webster Pugh	Neptune, NJ	0	0
65	Sajmir Alimehmeti	Bronx, NY	1	0
66	Abdursasul Hasanovich Juraboev	Brooklyn, NY	1	1
67	Akhror Saidakhmetov	Brooklyn, NY	1	1
68	Arbor Habibov	Brooklyn, NY	0	0
69	Dilkhayot Kasimov	Brooklyn, NY	0	0
70	Almal Zakirov	Brooklyn, NY	0	0
71	Mohimanul Bhuiya	Brooklyn, NY	0	0
72	Noelle Velentzas	Queens, NY	0	0
73	Asia Siddiqui	Queens, NY	1	1
74	Arafat M. Nagi	Lackawanna, NY	1	1
75	Ali Saleh	Fort Wayne, IN	1	1
76	Munther Omar Saleh	Queens, NY	1	1
77	Emanuel L. Luchtman	Rochester, NY	1	0
78	Mufid A. Elfgeeh	Rochester, NY	1	1
79	Farred Mumuni	Staten Island, NY	0	0
80	Terrence Joseph Mcneil	Akron, OH	1	1
81	Christopher Lee Cornell	Cincinnati, OH	1	1
82	Amir Aid Abdul Rahman Al-Ghazi / Robert C. McCollum	Sheffield Lake, OH	1	1
83	Munir Abdulkader	West Chester, OH	1	1
84	Jalil Ibn Amer Aziz	Harrisburg, PA	1	1
85	Keonna Thomas	Philadelphia, PA	1	1
86	David Wright	Everett, MA	0	0
87	Nicholas Rovinski	Warwick, RI	1	1
88	Usama Rahim	Roslindale, MA	0	0
89	Michael Todd Wolfe	Houston, TX	0	0
90	Omar Faraj Saeed Al Hardan	Houston, TX	0	0
91	Asher Abid Khan	Spring, TX	1	0
92	Sixto Ramiro Garcia	Houston, TX	1	1
93	Bilal Abood	Mesquite, TZ	1	1
94	Mohamad Jamal Khweis	Alexandria, VA	1	1
95	Haris Qamar	Burke, VA	1	1
96	Nicholas Young	Fairfax, VA	0	0
97	Amine El Khalifi	Fairfax, VA	1	1
98	Yusuf Abdirizak Wehelie	Failfax, VA	0	0
99	Heather Elizabeth Coffman	Richmond, VA	1	1
100	Mohamed Bailor Jalloh	Sterling, VA	1	1
101	Ali Shukri Amin	Woodbridge, VA	1	1
102	Joseph Hassan Farrokh	Woodbridge, VA	0	0
103	Mhamoud Amin Mohamed Elhassan	Woodbridge, VA	0	0
104	Daniel Seth Franey	Montesano, WA	1	1
105	Joshua Van Haften	Madison, WI	1	1
Proportion using social media			0.62	
Proportion posting public posts (among those using social media)				0.86

Note: The table provides details on the social media usage d 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.

S6 Hate crimes and far-right vote share

This study proxies anti-Muslim hostility with local-level vote share for far-right parties in the United Kingdom, France, Germany, and Belgium. However, since support for the far-right is driven by various factors, such as unemployment and rising levels of immigration (Golder, 2016), it is important to examine whether locations with high far-right support have greater levels of hate towards Muslims. Empirically, this is a challenging task, since systematic local-level data on hate crimes is not publicly available in most countries. Nonetheless, local data on hate crimes are available in the United Kingdom. This section examines the relationship between far-right support, hate crimes motivated by religion, and support for ISIS in the U.K.

Using official data from the U.K. police, I matched Twitter accounts of ISIS activists and followers in the U.K. with information on hate crimes motivated by religion in each police force area,[24] as well as granular geo-spatial data on public order crimes.[25] Public order crimes include incidents that "cause fear, alarm or distress" and subsume most hate crimes in the U.K.[26] Since official police-force area data on hate crimes is reported at a very aggregate level that includes both areas with high and low support for far-right parties,[27] I use incident-level, geo-tagged data on public order crimes that are reported at more granular levels. A test of the correlation between public order crimes and religiously motivated hate crimes, at the Twitter user level, shows a very strong relationship: the correlation coefficient is 0.9 with a p-value < 0.01. This means that Twitter users in areas with higher levels of public order crimes are also located in police force areas with higher levels of hate crimes.

Tables S18 and S19 show the relationship between hate crimes, public order offneses, far-right support and and pro-ISIS discourse in the U.K. Both tables report the same specifications, but vary in the outcome variable. In Table S18, the dependent variable is a composite measure of all pro-ISIS topics: sympathy with ISIS, life in ISIS territories, foreign fighters or travel to Syria, and the Syrian war; in Table S19 the dependent variable includes only sympathy with ISIS.

Hate crimes motivated by religion. Columns (1), (3), and (4) in both tables show that users located in police force areas with greater levels of hate crimes motivated by religion significantly tweet more pro-ISIS content. This result holds even when controlling for a battery of other variables, including far-right support, unemployment, the share of foreigners, Muslims, and Arabs in each local area.

Public order offenses. Columns (2) and (3) show a very similar relationship when using public order incidents to proxy for hate crimes. Users located in local-areas (Mid-layer super output area

[24] Hate crime data in each police force area cover the years 2015-2017. See `https://www.gov.uk/government/statistics/hate-crime-england-and-wales-2015-to-2016` and `https://www.gov.uk/government/statistics/hate-crime-england-and-wales-2016-to-2017`

[25] See `https://data.police.uk/data/`

[26] See `https://www.police.uk/about-this-site/faqs/#what-do-the-crime-categories-mean`

[27] These data are reported at the police force area level; there are 45 police force areas in the UK. See `https://www.police.uk/forces/`

(MSOA), which is roughly equal to the size of a neighborhood) that have greater levels of public order offenses are also more likely to tweet more pro-ISIS content.

Far-right vote share. As found in the main paper, all models show that far-right vote share at the local level is strongly associated with posting greater pro-ISIS content. Column (4) interacts far-right vote share with the number of offenses in each local area to examine whether users located in areas with higher far-right support post greater pro-ISIS content if they are exposed to more public order crimes (which, as mentioned above, are a plausible proxy for hate crimes). Both tables show that this is the case. The interaction term *Number of offenses in local area* × *Far-right vote share* is positive and significant at the 10% level. This evidence suggests that exposure to hate crimes is a mechanism that might be driving ISIS support in areas with greater support for far-right parties. The data also show that exposure to hate crimes in and of itself has a strong relationship with pro-ISIS support, which provides further support for the hypothesis tested in this paper, that anti-Muslim hostility might be driving pro-ISIS radicalization in Europe.

Table S18: Hate crimes and pro-ISIS discourse in the UK

Dependent variable: *Number of tweets on pro-ISIS topics**	(1)	(2)	(3)	(4)
Number of hate crimes motivated by religion†	0.59***		0.60***	0.65***
	(0.18)		(0.18)	(0.19)
Far-right vote share (%)	0.49***	0.57***	0.60***	0.22
	(0.06)	(0.06)	(0.06)	(0.23)
Number of offenses in local area‡		1.30***	1.28***	0.49
		(0.25)	(0.25)	(0.52)
Number of offenses in local area x Far-right vote share				0.06*
				(0.03)
Muslims (%)	−0.05	−0.08	−0.07	−0.12**
	(0.05)	(0.05)	(0.06)	(0.06)
Arabs (%)	−0.01	0.23	0.02	0.13
	(0.26)	(0.26)	(0.27)	(0.27)
Unemployment (%)	0.29**	0.13	0.08	0.02
	(0.14)	(0.14)	(0.14)	(0.15)
Foreigners (%)	0.15*	−0.07	−0.15	−0.09
	(0.08)	(0.09)	(0.10)	(0.10)
Constant	11.05***	6.05***	3.34*	8.66**
	(1.34)	(1.83)	(2.01)	(3.66)
R^2	0.001	0.001	0.001	0.001
Observations	80,058	79,134	79,132	79,132

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
* Pro ISIS topics include tweets sympathizing with ISIS, discussing life in ISIS territories or foreign fighters, and describing the Syrian civil war.
† Hate crimes motivated by religion reflect the logged number of hate crimes reported in each police force area in the UK during 2014-15.
‡ Number of offenses in local area reflects the logged number of public order crimes, which subsume most hate crimes, in each local area (middle layer super output areas).

Table S19: Hate crimes and sympathy with ISIS in the UK

Dependent variable: *Number of tweets sympathizing with ISIS*	(1)	(2)	(3)	(4)
Number of hate crimes motivated by religion†	0.13^{***}		0.13^{***}	0.14^{***}
	(0.04)		(0.04)	(0.04)
Far-right vote share (%)	0.11^{***}	0.13^{***}	0.13^{***}	0.04
	(0.01)	(0.01)	(0.01)	(0.05)
Number of offenses in local area‡		0.28^{***}	0.28^{***}	0.10
		(0.05)	(0.05)	(0.12)
Number of offenses in local area x Far-right vote share				0.01^{*}
				(0.01)
Muslims (%)	−0.02	-0.03^{**}	-0.02^{*}	-0.03^{**}
	(0.01)	(0.01)	(0.01)	(0.01)
Arabs (%)	0.01	0.06	0.02	0.04
	(0.06)	(0.06)	(0.06)	(0.06)
Unemployment (%)	0.07^{**}	0.03	0.02	0.01
	(0.03)	(0.03)	(0.03)	(0.03)
Foreigners (%)	0.03^{**}	−0.01	−0.03	−0.02
	(0.02)	(0.02)	(0.02)	(0.02)
Constant (%)	2.34^{***}	1.25^{***}	0.68	1.89^{**}
	(0.30)	(0.40)	(0.44)	(0.81)
R^2	0.001	0.001	0.001	0.001
Observations	80,058	79,134	79,132	79,132

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
† Hate crimes motivated by religion reflect the logged number of hate crimes reported in each police force area in the UK during 2014-15.
‡ Offenses in local area reflects the logged number of public order crimes, which subsume most hate crimes, in each local area (middle layer super output areas).

S7 Unemployment, far-right vote share, and support for ISIS on Twitter

One concern that may arise with the analysis presented in the paper is that far-right vote share and pro-ISIS rhetoric may both be driven by unemployment. While all specifications control for unemployment at the local level, this might not be enough to rule out the confounding effect of unemployment. To address this issue, I carry out several additional tests. First, as presented in the main paper, I conduct high frequency studies around events that may mobilize support for ISIS, and examine whether pro-ISIS rhetoric increases after these events more strongly in areas with higher levels of far-right vote share. The idea is that if far-right areas make people more likely to support ISIS, then we should also observe this pattern in the high frequency time dimension. The results show systematic evidence that across various events, including terrorist attacks, anti-Muslim marches, and ISIS propaganda releases, users express greater support for ISIS after these events in localities where far-right parties are more popular. In particular, when I examine heterogeneous changes following these events for both far-right vote share and unemployment (see Table S23), it is clear that these high-frequency changes are linked to the former and not the latter.

Second, I carry out a more comprehensive examination using a matching design. In the matching approach, I compare users located in areas with high and low far right vote share that are matched on levels of unemployment, the proportion of foreigners, population size, and the country in which they are located.[28] I created a binary variable for areas with high far-right support that is coded 1 when a location is at or above the median far-right vote share, and 0 otherwise. I then estimated a logistic regression of the high-far right variable on these covariates, choosing the single nearest neighbor as a control. I use propensity scores from this matching procedure as a weight in a regression comparing the difference in ISIS support between users located in areas with low and high levels of far-right vote share, as well as around events that mobilize support for ISIS.

Before discussing the results, I examine whether the matching method was able to achieve balance. Table S20 shows results from regressions of local-level far-right vote share (measured with the binary variable described above) on the covariates used in the matching. Columns (1) and (2) show that in the unbalanced model ("UB"), greater levels of unemployment are significantly correlated with high far-right vote share. This is expected, as the popularity of far-right parties in Europe is driven to a great extent by unemployment. However, this correlation disappears in the balanced model ("B") presented in Column (2). I find the same results when adding covariates to the model in Columns (3) – (6). Interestingly, in Column (7), which presents the unbalanced regression when adding country fixed effects, the relationship between unemployment and far-right vote share also goes to zero. This suggests that this model successfully accounts for the confounding effect of unemployment. As these are the covariates used in the paper's main specifications, it reduces the concern that unemployment drives the results.

[28] I use these covariates since they are available for all countries in the study.

Next, I examine the results from the matching design. Table S21 shows the relationship between far right vote share and pro-ISIS support when comparing users in high and low far-right areas. In Column (1) the variable is coded 1 for individuals who are at the top 1% of the distribution of posting pro-ISIS content, and 0 otherwise. Column (2) is measured similarly, but uses only sympathy with ISIS to measure radical content. In Columns (3) and (4) the dependent variable is a binary measure of being flagged as an ISIS activist and being suspended from Twitter, respectively. Column (5) uses the number of ISIS accounts that a user follows. Overall, the matching results show very similar findings to those found in the main paper. Moving from areas that are below the median far-right vote share to matched areas that are above the median significantly increases the probability that a user is at the top 1% posters of tweets sympathizing with ISIS, is flagged as an ISIS activist, suspended from Twitter, and follows a greater number of ISIS accounts.

Table S22 presents results from the event studies using matching. Panel A shows the impact of the events on pro-ISIS content in all areas, using data from three days before and after the events. Panel B examines whether this effect differs between areas with low and high support for far-right parties. I find that in most models, users in areas with greater far-right vote share post significantly more pro-ISIS content after terrorist attacks, the ISIS propaganda release, and the anti-Muslim marches. These results, together with the cross-sectional matched design described above, suggest that the link between far-right vote share and support for ISIS on Twitter are not driven by unemployment.

Table S20: Balance test

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
High far-right vote share	UB	B	UB	B	UB	B	UB	B
Unemployment (%)	0.03***	0.00	0.03***	-0.00	0.03***	-0.00	0.00	-0.00
	(0.00)	(0.02)	(0.00)	(0.02)	(0.00)	(0.02)	(0.00)	(0.02)
Foreigners (%)			-0.01***	0.00	-0.01***	0.00	-0.01***	0.01
			(0.00)	(0.01)	(0.00)	(0.01)	(0.00)	(0.01)
Population					-0.00***	-0.00	0.00	-0.00
					(0.00)	(0.00)	(0.00)	(0.00)
Constant	0.58***	0.49***	0.60***	0.49***	0.60***	0.49***	0.68***	0.46***
	(0.02)	(0.11)	(0.02)	(0.10)	(0.02)	(0.11)	(0.05)	(0.10)
Country fixed effects	✗	✗	✗	✗	✗	✗	✓	✓
R^2	0.025	0.000	0.064	0.001	0.071	0.003	0.312	0.003
Number of observations	2790	2,367	2,790	2,367	2,786	2,367	2,786	2,367

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table S21: Far-right vote share and support for ISIS on Twitter (Matched design)

	(1) Top 1% radical content	(2) Top 1% sympathy with ISIS only	(3) Flagged as an ISIS activist	(4) Suspended from Twitter	(5) Number of ISIS accounts following
High far-right = 1	1.97	3.75***	2.16*	10.16***	1.86***
	(1.29)	(1.29)	(1.21)	(3.78)	(0.55)
Constant	8.69***	7.48***	2.90***	33.77***	4.02***
	(1.15)	(1.14)	(0.78)	(3.04)	(0.27)
R^2	0.000	0.000	0.000	0.001	0.002
Number of clusters	2,367	2,367	2,367	2,367	2,367
Number of observations	157,873	157,873	157,873	157,872	157,872

Standard errors in parentheses
Coefficients in columns 1-4 are × 1,000 to account for the skewed distribution of the dependent variables.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table S22: Events and changes in pro-ISIS rhetoric (matched design)

	Paris attacks		Brussels attacks		ISIS propaganda release		PEGIDA marches		
	(1) Sympathy with ISIS	(2) ISIS topics	(3) Sympathy with ISIS	(4) ISIS topics	(5) Sympathy with ISIS	(6) ISIS topics	(7) Sympathy with ISIS	(8) ISIS topics	(9) ISIS topics + Anti-West
A. Changes in pro-ISIS content (standard deviation units)									
After event = 1	0.146***	0.118***	0.022*	0.024***	0.017	0.063**	0.008	-0.020	-0.011
	(0.030)	(0.034)	(0.013)	(0.009)	(0.025)	(0.031)	(0.010)	(0.014)	(0.011)
Constant	0.081	-0.082	0.166	0.002	-0.641***	-0.478***	-0.059	0.316***	0.283***
	(0.156)	(0.069)	(0.171)	(0.113)	(0.222)	(0.029)	(0.050)	(0.064)	(0.083)
R^2	0.011	0.008	0.002	0.003	0.001	0.005	0.004	0.001	0.002
Number of clusters	293	293	445	445	113	113	417	417	417
Number of observations	32,451	32,451	60,224	60,224	5,016	5,016	50,874	50,874	50,874
B. Changes in pro-ISIS content (standard deviation units), by far-right support									
After event = 1	0.016	0.070	0.041*	-0.041	-0.128*	-0.074	0.015	-0.030	-0.030
	(0.031)	(0.062)	(0.021)	(0.037)	(0.075)	(0.100)	(0.021)	(0.028)	(0.020)
Far-right vote share (%)	-0.003	-0.006*	0.005**	-0.002	-0.007	-0.009	0.003*	0.001	0.001
	(0.002)	(0.003)	(0.002)	(0.002)	(0.008)	(0.006)	(0.002)	(0.002)	(0.002)
After event = 1 × Far-right vote share (%)	0.004**	0.004	-0.001	0.004**	0.011**	0.004	-0.000	0.002	0.002*
	(0.002)	(0.003)	(0.001)	(0.002)	(0.004)	(0.005)	(0.002)	(0.001)	(0.001)
Constant	0.028	-0.172	0.141	0.452***	-0.400***	-0.312***	-0.371***	0.061	-0.029
	(0.202)	(0.379)	(0.216)	(0.101)	(0.108)	(0.110)	(0.052)	(0.194)	(0.193)
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.007	0.011	0.002	0.002	0.008	0.008	0.004	0.001	0.003
Number of clusters	281	281	428	428	110	110	405	405	405
Number of observations	19,235	19,235	40,419	40,419	2,777	2,777	34,391	34,391	34,391

Standard errors in parentheses, clustered by location. Base country is Belgium.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table S23: Events and changes in pro-ISIS rhetoric, by far-right vote share and unemployment

	Paris attacks		**Brussels attacks**		**ISIS propaganda release**		**PEGIDA marches**		
	(1) Sympathy with ISIS	(2) ISIS topics	(3) Sympathy with ISIS	(4) ISIS topics	(5) Sympathy with ISIS	(6) ISIS topics	(7) Sympathy with ISIS	(8) ISIS topics	(9) ISIS topics + Anti-West
Changes in pro-ISIS content (standard deviation units), by far-right vote share and unemployment									
After event = 1	-0.006	-0.008	0.071***	-0.009	-0.069	-0.085	-0.019	-0.039	-0.053*
	(0.041)	(0.043)	(0.020)	(0.024)	(0.078)	(0.088)	(0.028)	(0.028)	(0.027)
Unemployment (%)	-0.018**	-0.010	0.006*	-0.003	0.012	-0.009	-0.003	-0.006	-0.004
	(0.008)	(0.009)	(0.004)	(0.004)	(0.017)	(0.015)	(0.005)	(0.004)	(0.005)
After event = 1 × Unemployment (%)	0.012	0.015*	-0.010***	-0.000	-0.008	0.005	0.005	0.001	0.003
	(0.009)	(0.008)	(0.004)	(0.005)	(0.019)	(0.011)	(0.006)	(0.004)	(0.004)
Far-right vote share (%)	-0.002	-0.004*	0.002	-0.000	-0.009	-0.012**	0.001	0.000	0.000
	(0.002)	(0.002)	(0.001)	(0.002)	(0.006)	(0.005)	(0.002)	(0.001)	(0.001)
After event = 1 × Far-right vote share (%)	0.002	0.003*	-0.000	0.003**	0.010**	0.004	0.000	0.002	0.002**
	(0.002)	(0.002)	(0.001)	(0.001)	(0.005)	(0.004)	(0.002)	(0.001)	(0.001)
Constant	0.279**	0.384***	-0.055	0.086	-0.354	0.048	0.106	0.138*	0.163*
	(0.132)	(0.133)	(0.072)	(0.074)	(0.223)	(0.200)	(0.099)	(0.078)	(0.085)
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.006	0.006	0.002	0.002	0.010	0.008	0.003	0.001	0.002
Number of clusters	362	362	529	529	140	140	508	508	508
Number of observations	21,459	21,459	46,460	46,460	3,216	3,216	38,527	38,527	38,527

Standard errors in parentheses, clustered by location. Base country is Belgium.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

S8 Additional figures

Figure S13: @CtrlSec request to expose ISIS members on Twitter

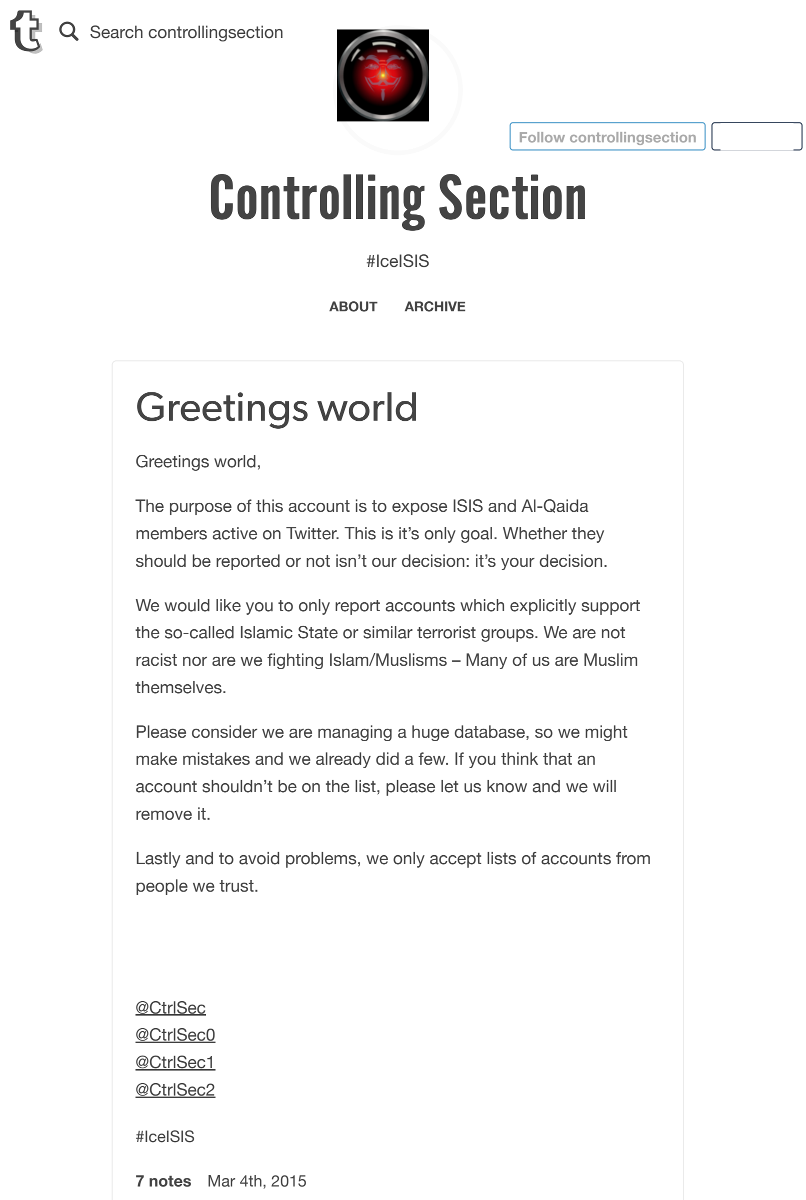

Search controllingsection

Follow controllingsection

Controlling Section

#IceISIS

ABOUT ARCHIVE

Greetings world

Greetings world,

The purpose of this account is to expose ISIS and Al-Qaida members active on Twitter. This is it's only goal. Whether they should be reported or not isn't our decision: it's your decision.

We would like you to only report accounts which explicitly support the so-called Islamic State or similar terrorist groups. We are not racist nor are we fighting Islam/Muslisms – Many of us are Muslim themselves.

Please consider we are managing a huge database, so we might make mistakes and we already did a few. If you think that an account shouldn't be on the list, please let us know and we will remove it.

Lastly and to avoid problems, we only accept lists of accounts from people we trust.

@CtrlSec
@CtrlSec0
@CtrlSec1
@CtrlSec2

#IceISIS

7 notes Mar 4th, 2015

Source: `http://controllingsection.tumblr.com/post/112703617620/greetings-world`

Figure S14: Example of @CtrlSec real-time flagging of ISIS acounts

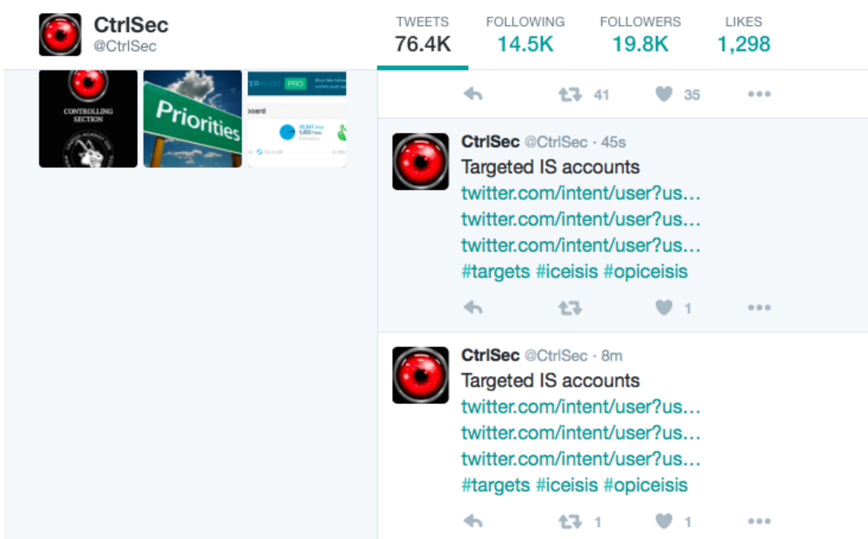

Figure S15: Example of ISIS accounts

Figure S16: Example of a Western fighter tweeting from Syria

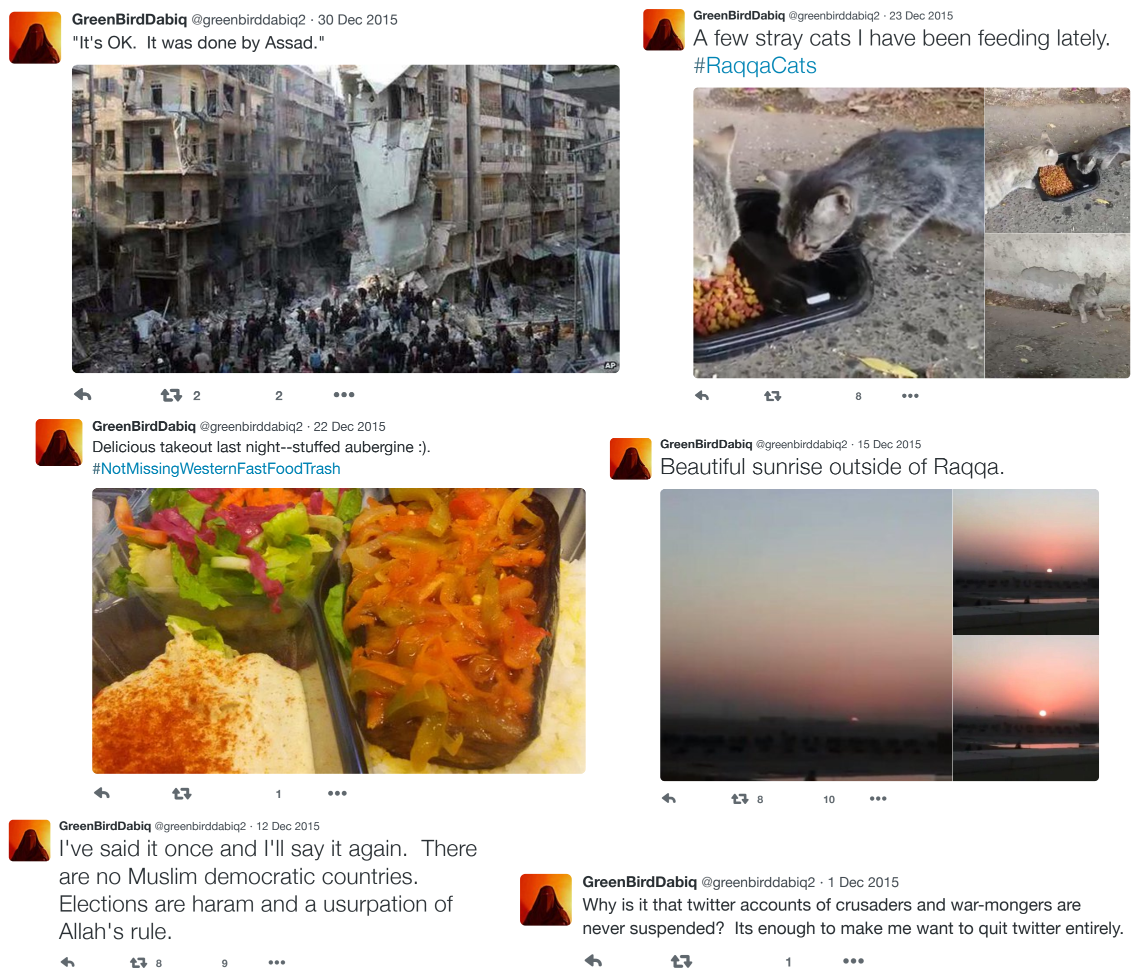

Note: This account has already been suspended as of February 2016.

Figure S17: Example of a suspended account

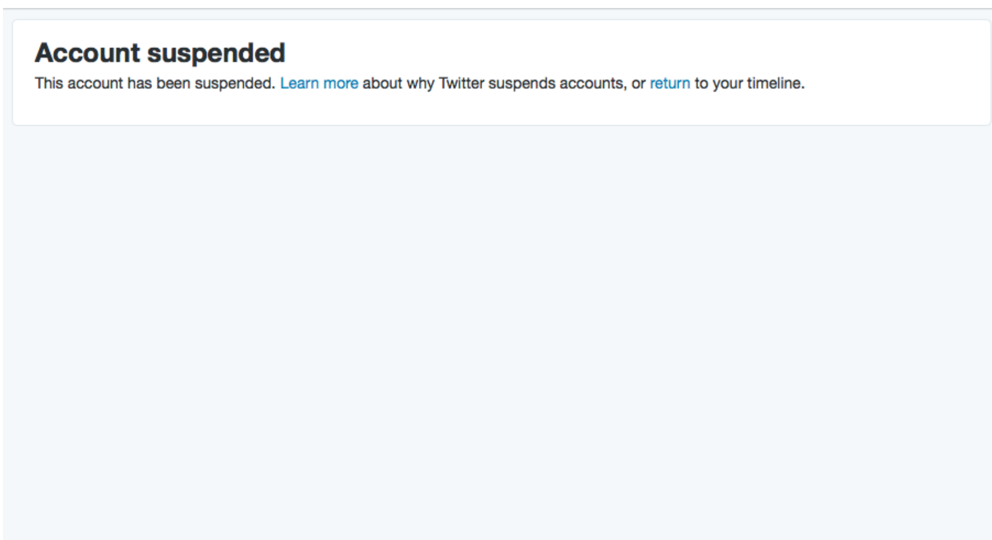

Figure S18: Vote share for far-right parties

(a) France

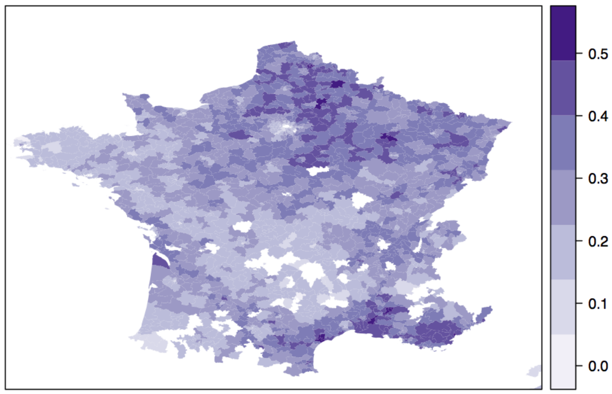

(b) Germany

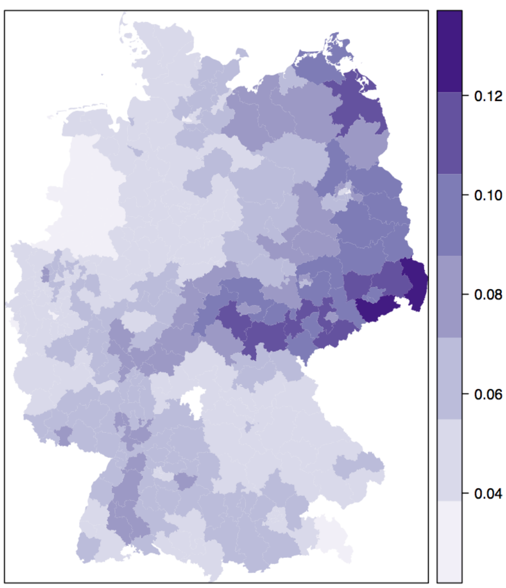

(c) United Kingdom

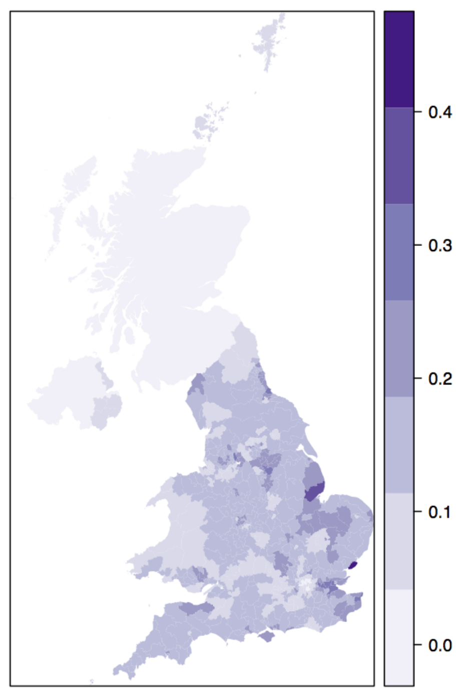

Note: For France, the map displays the vote share for the Front National party in the 2015 departmental elections at the electoral canton level. For Germany, the map displays the vote share for the Alternative for Germany (AfD) party and the National Democratic Party (NPD) in the 2013 federal elections. For the UK, the map represents the vote share for the British Democrats, British National Party, Liberty GB party, National Front party, and United Kingdom Independence Party in the 2015 UK parliamentary general elections.

Figure S19: The cumulative distribution functions for the distance to a user's geographically closest friend (Figure taken from Jurgens (2013))

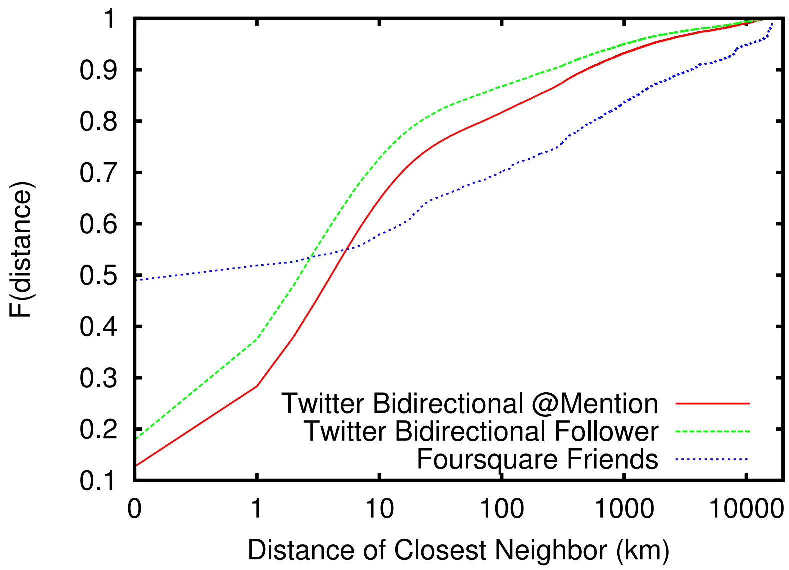

Note: The figure, taken from the study of Jurgens (2013), shows cumulative distribution functions (CDFs) of users' geographical distance to their closest neighbor in three social media networks. In the figure, the x axis shows distance in kilometers, and the y axis shows the probability that the closest neighbor for each user is located x distance or less from that user. It can be seen that more than half of the users in these three networks had neighbors that were located within 4 kilometers from them, thereby allowing location prediction within 4-kilometer bounds.

Figure S20: Anti-Muslim marches organized by PEGIDA across Europe

Note: Photos credit: *Radio Free Europe Radio Liberty* (2016) and Malm (2015)

Figure S21: National action plans to counter violent extremism

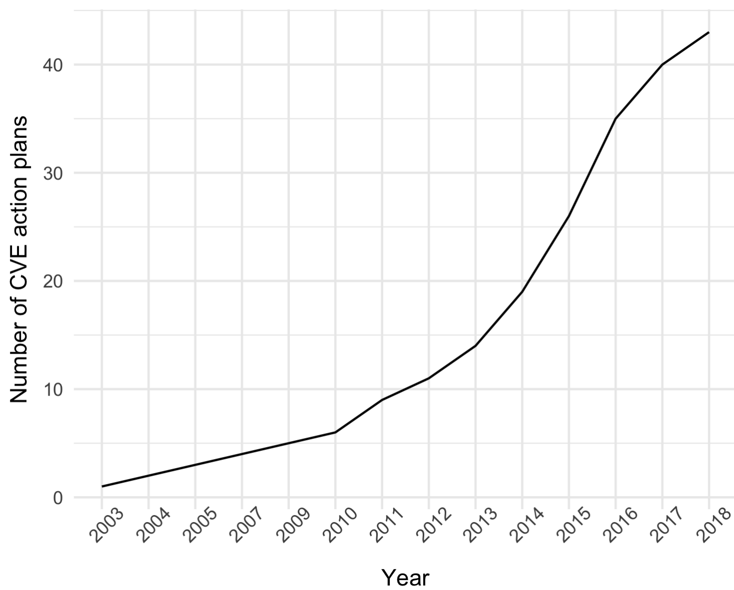

Note: The figure presents the number of official national action plans to counter violent extremism by year. National action plans to counter violent extremism are official policies adopted by countries, and are reflected in formal documents collected by the author. It can be seen that official strategies to counter extremism have dramatically increased in recent years.

S9 Additional results

Table S24: Western foreign fighters and online radicalization by country

	(1) Number of Twitter users flagged as ISIS activists	(2) Number of Twitter users posting highly radical content	(3) Number of ISIS accounts followed	(4) Number of Twitter users suspended from Twitter
Number of foreign fighters (official count)	0.135***	0.159	79.253***	0.305***
	(0.030)	(0.133)	(20.192)	(0.107)
Constant	11.916	3.446	4,734.094	42.960
	(15.882)	(71.194)	(10,773.890)	(56.946)
Population controls	✓	✓	✓	✓
R^2	0.396	0.336	0.434	0.360
Number of observations	46	46	46	46

Note: The table reports the correlation between online radicalization measures and foreign fighter counts in European countries, controlling for population size. It can be seen that all online radicalization variables are positively correlated with the number of foreign fighters in each country, with the number of users flagged as ISIS activists, number of ISIS accounts followed, and the number of users suspended from Twitter significant at the 5% level.
*p<0.1; **p<0.05; ***p<0.01

Table S25: Different cutoffs for classifying top posters of radical content

	(1) Top 5%	(2) Top 10%	(3) Top 15%	(4) Top 20%	(5) Top 25%
Far-right vote share (%)	0.81**	0.88*	1.00*	1.63*	1.99
	(0.35)	(0.51)	(0.60)	(0.88)	(1.25)
Unemployment (%)	1.19	3.15**	3.66**	5.38**	7.66**
	(0.75)	(1.27)	(1.60)	(2.54)	(3.49)
Foreigners (%)	0.40	0.48	-0.05	-0.68	-1.03
	(0.29)	(0.46)	(0.54)	(0.77)	(1.06)
Constant	45.45***	75.29***	149.63***	215.52***	281.94***
	(16.97)	(23.13)	(28.12)	(40.00)	(55.40)
Population controls	✓	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓	✓
R^2	0.001	0.002	0.001	0.002	0.003
Number of clusters	2,654	2,654	2,654	2,654	2,654
Number of observations	112,253	112,253	112,253	112,253	112,253

Robust standard errors in parentheses, clustered at the locality level.
Base country is Belgium.
All coefficients are × 1,000 to account for the skewed distribution of the DV.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table S26: Correlates of activists

	(1) Top 1% radical content	(2) Top 1% sympathy with ISIS only	(3) Suspended from Twitter	(4) Number of ISIS accounts following
Flagged as an IS activist	0.15***	0.13***	0.45***	128.22***
	(0.01)	(0.02)	(0.04)	(15.25)
Constant	0.01**	0.01	0.04***	2.37
	(0.00)	(0.00)	(0.01)	(3.11)
Controls	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓
R^2	0.010	0.009	0.028	0.131
Number of clusters	2,654	2,654	2,653	2,653
Number of observations	112,253	112,253	112,249	112,249

Standard errors in parentheses, clustered at the locality level. Base country is Belgium.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table presents the relationship between various radicalization outcomes and being an ISIS activist on Twitter. The regressions control for local-level vote share for far-right parties, unemployment, the share of foreigners, and population size, and include country fixed effects. It can be seen that ISIS activists on Twitter are significantly more likely to show signs of radicalization, when compared to ISIS followers.

Table S27: Far-right vote share and support for ISIS on Twitter

	(1) Top 1% radical content	(2) Top 1% sympathy with ISIS only	(3) Flagged as an ISIS activist	(4) Suspended from Twitter	(5) Number of ISIS accounts following
Far-right vote share (%)	0.25***	0.20***	0.30***	0.09	0.09***
	(0.06)	(0.07)	(0.04)	(0.13)	(0.02)
Unemployment (%)	0.25	0.23	-0.20*	-1.24***	-0.11***
	(0.17)	(0.17)	(0.12)	(0.32)	(0.03)
Foreigners (%)	0.11*	0.14**	0.26***	-0.06	0.08***
	(0.06)	(0.06)	(0.04)	(0.12)	(0.02)
Constant	7.89**	4.35	-9.78***	35.10***	1.12
	(3.74)	(3.48)	(1.91)	(6.70)	(0.74)
Population controls	✓	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓	✓
R^2	0.000	0.000	0.006	0.002	0.006
Number of observations	112,253	112,253	112,253	112,249	112,249

Robust standard errors in parentheses. Base category is Belgium.
Coefficients in columns 1– 4 are × 1,000 to account for the skewed distribution of the dependent variables.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table S28: Far-right vote share and posting pro-ISIS content on Twitter

	(1) Sympathy with ISIS	(2) ISIS life/ Foreign fighters	(3) Syrian war	(4) Anti-West
Far-right vote share (%)	0.05***	0.09***	0.07***	0.04***
	(0.01)	(0.02)	(0.01)	(0.01)
Unemployment (%)	0.12***	0.24***	0.15***	0.13***
	(0.02)	(0.04)	(0.03)	(0.02)
Foreigners (%)	0.02***	0.04**	0.03***	0.02**
	(0.01)	(0.01)	(0.01)	(0.01)
Constant	3.54***	7.32***	5.79***	3.21***
	(0.43)	(0.82)	(0.61)	(0.42)
Population controls	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓
R^2	0.001	0.002	0.002	0.003
Number of observations	112,253	112,253	112,253	112,253

Robust standard errors in parentheses. Base country is Belgium.

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

Table S29: Unemployed immigrants, asylum seekers and support for ISIS on Twitter

	(1) Top 1% radical content	(2) Flagged as an ISIS activist	(3) Suspended from Twitter	(4) Number of ISIS accounts following
Far-right vote share (%)	0.24**	0.52***	0.62***	0.23***
	(0.09)	(0.07)	(0.18)	(0.02)
Unemployed immigrants (%)	0.70*	0.39	0.09	0.36***
	(0.42)	(0.26)	(0.77)	(0.09)
Asylum seekers (%, sd units)	-0.40	-11.77***	-14.21***	-2.62***
	(0.93)	(1.15)	(1.87)	(0.23)
Constant	-4.27	-63.72***	-41.47***	-14.52***
	(5.81)	(6.86)	(12.03)	(2.93)
Population controls	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓
R^2	0.001	0.012	0.003	0.005
Number of observations	30,373	30,373	30,372	30,372

Robust standard errors in parentheses. Data available only for France and Germany. Base category is Germany.

Coefficients in columns 1-3 are × 1,000 to account for the skewed distribution of the dependent variables.

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

References

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/*

Golder, Matthew. 2016. "Far Right Parties in Europe." *Annual Review of Political Science* 19(1).

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.

Malm, Sara. 2015. "Muslim is stabbed to death while shopping in Dresden, where thousands have gathered for right wing rallies." Daily Mail.
URL: *http://www.dailymail.co.uk/news/article-2909950/PEGIDAs-anti-Muslim-calls-shake-German-politics.html*

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

Radio Free Europe Radio Liberty. 2016. "Prague Refugee Center Fire-Bombed After Anti-Muslim Protests.".
URL: *http://www.rferl.org/content/pegida-refugee-crisis-protests-prague-attack/27536604.html*

Ramos, Juan. 2003. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*.