Mutual radicalisation of opposing extremist groups via the Internet

John D. Gallacher*
Oxford Internet Institute
University of Oxford

Marc W. Heerdink
Department of Social Psychology
University of Amsterdam

Abstract

Social media has become a common arena for both far-right and Islamic extremist groups to stoke division through the spreading of propaganda and hate speech. This online hate is suggested to drive extremism online, and in some cases lead to offline hate crimes and violence. Whether this online radicalisation happens in isolation within a group, or whether there is an interdependent relationship of mutual radicalisation, is unclear. A possible process by which mutual radicalisation could occur would be if social media incite users to commit offline violence, and if this offline violence in return triggers online reactions from both the target and perpetrator groups. This however has not been tested. This study addresses these questions by investigating the nature of the online-offline relationship of extremist hate. We combine data from the social media platform Gab, variations in Internet search trends, and offline hate crimes in three countries, and test for temporal relationships between opposing extremist groups. Our findings show that online hate from far-right groups both precedes offline violence from these same groups, and spikes following offline violence from opposing Islamic extremist groups. Additionally, far-right Islamophobic violence offline is also followed by increased online interest in Islamic extremist topics. Together, these findings show that the Internet, and specifically hate speech, plays a potential key role in a cyclical process that increases mutual radicalisation.

16 February 2021

Keywords: Mutual Radicalisation | Online Hate Speech | Social Media | Social Networks | Intergroup Contact | Political Extremism | Political Violence

* **Corresponding Author:** John Gallacher – john.gallacher@oii.ox.ac.uk

Introduction

In recent decades there has been a notable rise in Islamic extremism and violence as well as an increase in far-right extremism across much of Europe and the United States of America (Global Terrorism Index, 2019). It has been argued that these two extremist trends have occurred in tandem, with an interdependent relationship of mutual escalation between opposing groups (Guhl & Ebner, 2018). These parallel trends may be an indication of mutual radicalisation, whereby for two opposing groups the actions of one group result in a negative or aggressive reaction from the other, and can lead to these groups shifting in opposing directions (Konaev & Moghaddam, 2010; Moghaddam, 2018). The narratives used by far-right and Islamic extremist groups to motivate their actions demonstrate this trend, with far-right descriptions of “*The west being at War with Islam*” and Islamist extremist narratives “*Muslims are at war with the West*” exhibiting opposing but parallel views (Ebner, 2017, p.197). This may instigate an increasing spiral of hatred and intergroup conflict, with both groups viewing themselves as the victim and the other as the perpetrator (Fielitz, Ebner, Guhl, & Quent, 2018).

The Internet is hypothesised to play an important role in promoting this spiral of intergroup hate, extremism, and violence (Torok, 2013). Social media in particular has historically given extremist groups the unparalleled ability to connect and communicate at scale, both internally and externally, and despite recent crackdowns on the prevalence of extremist material online (Conway, Khawaja, et al., 2019), the online reach of extremist groups is still considerable (e.g. Benigni, Joseph, & Carley, 2017; Berger, 2018b). The role this plays in mutual radicalisation is unclear however, and direct evidence for these effects of mutual radicalisation online is limited. In particular, whether increases in extreme derogation from one group online are associated with subsequent offline violence against these targeted groups, and whether this violence in turns radicalises the targeted groups, is undetermined. However, this role of the Internet may be key to understanding the observed increases in offline violence.

Derogation of the opposing groups is a key feature of the mutual radicalisation process. Online, extreme outgroup derogation is often framed as hate speech, i.e. messages which express hatred towards a targeted group with the intention to attack, disparage, or humiliate members of that group (Davidson, Warmley, Macy, & Weber, 2017). These messages contain deliberate attacks against (or about) a specific group of people, motivated by (or focused on) aspects of that group’s identity (Davidson et al., 2017; de Gibert, Perez, García-Pablos, & Cuadros, 2018; Gallacher, 2021). The spread of hate speech online is a growing concern. Its propagation across social media is likely to be inflicting harm on the individuals and groups who are targeted and their communities, and the wider society exposed to it (Vidgen, Harris, Cowls, & Guest, 2020). This issue has attracted increasing attention from social media platforms, governments, law enforcement agencies, and civil society. In addition, considerable improvements in automated techniques to detect and measure online hate speech have been made recently, reflecting an increased interest in the study of online hate speech (Vidgen, Tromble, et al., 2019). However, less attention has been given to the role that hate speech plays in dividing online communities, promoting intergroup conflict, and the impact that it may have on individual extremism, group radicalisation, and wider hateful extremist behaviours including offline violence (Kleinberg, Vegt, & Gill, 2020; Siegel, 2020).

There is evidence that online far-right is particularly associated with the creation and distribution of online hate speech (e.g. Berger, 2018b; Conway, Scrivens, & Macnair, 2019; Lee, 2019), and the Internet is suggested to play a central role in breeding and amplifying far-right extremist ideologies (Ischinger, 2020). Users who are most active and spend more time on far-right fringe social media

platforms are more likely to express hate speech themselves (Gallacher & Bright, 2021; Javed & Miller, 2019; Mathew, Dutt, Goyal, & Mukherjee, 2018). This effect is likely due to the increased exposure to outgroup hate from other ingroup members (Gallacher & Bright, 2021). This suggests that online hate speech can drive radicalisation within groups, however the impact that it has on relations with opposing groups, both those directly targeted and wider groups, is less clear. Furthermore, how reactions to offline attacks from opposing groups can drive this online radicalisation process requires further investigation.

In this study we investigate the relationship between online hate speech and offline hateful extremist behaviour and violence, and what role the Internet and social media platforms are playing in cycles of mutual radicalisation. A conceptual overview of this process is given in Figure 1, illustrating how the Internet may facilitate mutual radicalisation through both reaction and incitement to offline violence. Guided by this idea, we test for evidence of mutual radicalisation at three levels: firstly, whether far-right online hate speech influences far-right offline hate crime (Figure 1 step A), secondly, whether far-right offline violence affects reactionary Internet activity from opposing extremist groups (Figure 1 step B), and finally, whether offline Islamic extremist violence impacts far-right hate speech online (Figure 1 step D). By considering these three links we obtain insights into potential effects of mutual radicalisation between opposing groups via the Internet. Testing the additional stage in this cycle (Figure 1 step C) whereby interest in joining Islamic extremist groups is linked to subsequent Islamic terror attacks is beyond the scope of this study due to practical limitations around data availability and legal restrictions, the existing evidence for this stage is discussed below.

Mutual radicalisation

Online, hate speech has been linked to driving radicalisation and negative attitudes towards outgroups (Gallacher & Bright, 2021). Radicalisation is the process whereby individuals come to increasingly adopt an extreme position through a process of socialisation around shared grievances as the belief that the direct action against the outgroup is morally correct (Reicher et al., 2008; Smith, Blackwood,

Figure 1 – Proposed cycle of Mutual Radicalisation between opposing far-right and Islamic extremist groups via the Internet. In this example far-right groups are given in Blue and Islamic extremist groups are given in Orange

& Thomas, 2019). This shift towards extremism is characterised by the belief that survival of the ingroup is inseparable from direct action against the outgroup (Berger, 2018a).

For opposing groups taking increasingly extreme positions against one another, it initiates a dynamic process of mutual radicalisation (also referred to as ‘co-radicalisation’ or ‘reciprocal radicalisation’; Knott, Lee, & Copeland, 2018; Pratt, 2015; Reicher & Haslam, 2016), with the groups reacting against real or imagined threats offline, and taking positions further opposing one another (Moghaddam, 2018). This increases intergroup boundaries, perceived intergroup distance, and perceived intergroup threat (Tajfel, 1974; Tajfel & Turner, 1979; Turner, Hogg, Oakes, Reicher, & Wetherell, 1987) and can lead to a greater risk of prejudice, discrimination and conflict (Stephan & Stephan, 2000; Stephan, Ybarra, & Morrison, 2009). This process can manifest as hateful extremist behaviours, i.e. actions that can incite and amplify hate, or make the moral case for violence, while drawing on hateful, hostile or supremacist beliefs directed at an out-group (Commission for Countering Extremism, 2019). Inciting violent reactions from the outgroup can also be a deliberate strategy of terrorist organisations. McCauley & Moskaleiko (2008) identify ‘jiu-jitsu’ radicalisation whereby terror attacks are designed to elicit moral outrage amongst the targeted population and provoke a military response. This military response is then thought to increase the perceived threat posed by this nation and make the local populations more susceptible to local radicalisation.

Moghaddam (2018) identifies three stages in offline mutual radicalisation: i) group mobilisation, ii) extreme ingroup cohesion, and iii) antagonistic identity transformation. This builds on evidence from social psychology that conflict with outgroups can increase ingroup cohesion (Coser, 1956), as the real or imagined threat posed by the outgroup can cause the ingroup to overcome any internal differences and unite against the outgroup (Stein, 1976). Inside these stages of mutual radicalisation are a number of steps, a subset of which relate directly to hate speech and extreme outgroup denigration. Outgroup denigration acts to (i) increase intergroup differentiation, (ii) normalise public expression of hatred towards the outgroup, (iii) instigate a spiral of increasingly hostile responses to exaggerated outgroup threats, and (iv) shifts the groups to extreme positions against one another. This process has been demonstrated at the national level (Moghaddam, 2018), amongst opposing national and political leadership (Konaev & Moghaddam, 2010; Moghaddam, Harré, & Lee, 2008), amongst Muslims and non-Muslims in Scandinavia (Obaidi, Thompson, & Bergh, 2019), along with others.

Recent evidence of mutual radicalisation between Islamic extremist groups and far-right extremist groups across Europe is increasing (Ebner, 2017; Fielitz et al., 2018). Offline, terrorist violence against civilians predicts an increase in violent behaviour from these same civilian groups against the perceived perpetrators of the original violence (Brandsch & Python, 2020), whilst learning that Islamic extremists dehumanize Westerners increases the likelihood Muslims will be dehumanized in return (Kteily, Hodson, & Bruneau, 2016). Additionally, there is prior evidence for each of the individual links between online and offline radicalisation shown in Figure 1, which this study builds on. However, the role of the Internet in driving this overall cycle of mutual radicalisation remains a central question.

The relationship between online hate and offline violence

It is now increasingly established that online hate speech is linked with offline hate crime. In recent years, numerous high-profile offline far-right terror attacks have been preceded by online activity indicative of the future violence, and while the attackers conducted their attacks alone, they were embedded in global online communities (Ischinger, 2020). Indeed, the concept of ‘lone wolf terrorism’ is becoming progressively outdated, and evidence suggests that the Internet plays a role in connecting

attackers to wider extremist communities (Gill, Horgan, & Deckert, 2014; Schuurman, Bakker, Gill, & Bouhana, 2018; Schuurman et al., 2019).

The attackers at the 2018 Tree of life Synagogue shooting in Pittsburgh, USA, the 2019 synagogue shooting in Halle, Germany, and the 2019 El Paso shooting in the USA all displayed considerable activity on fringe far-right social media platforms prior to their attacks (Evans, 2018; Evans, 2019; Der Spiegel, 2019). Attacks have also been livestreamed over mainstream social media platforms, including the 2019 terror attack at two mosques in Christchurch, New Zealand (Evans, 2019). Similar evidence has arisen for posts to fringe far-right platforms from attackers responsible for school shootings in the USA, attacks on Black Lives Matter protests, and misogynistic attacks at University campuses (Nagle, 2017), among numerous others. In these events the online environment appears to have played a role in the radicalisation of these attackers, and these attackers then used these same platforms to distribute information of their planned attacks prior to the events (Grover & Mark, 2019).

Beyond these high-profile individual attacks, there is growing evidence that similar trends occur at the group level. Hateful comments increased on Reddit in close temporal proximity to the Unite the Right rally in Charlottesville in August 2017 (Zannettou, ElSherief, Belding, Nilizadeh, & Stringhini, 2020), which left three people dead and multiple injured. The Facebook event page used to organise this event was removed by Facebook prior to the event, due to the high level of hate speech and threats to life (Heath, 2017), and attendees had expressed support for racist neo-Nazi groups and far-right extremist movements prior to attending (Cohen-Almagor, 2018). Beyond western countries, online hate speech has played a leading role in inciting genocide of the Rohingya community in Myanmar (Stecklow, 2018), while Islamophobic hate speech has been linked with deadly mob violence in Sri Lanka (Samaratunge & Hattotuwa, 2014). Online content has also been linked to violence against health workers combatting Polio in Pakistan (Jawaid, 2013), Ebola in Central Africa (Turse, 2019), and most recently to Coronavirus conspiracy theory related violence in the UK (Kelion, 2020).

Systematic evidence that online hate speech predicts offline hate crime (Figure 1 step A)

In addition to these single event observations, systematic research on broader contexts has provided further evidence supporting the link between the online and offline hate environments. Online hate speech and offline racially and religiously aggravated hate crimes are temporally and spatially associated (Williams, Burnap, Javed, Liu, & Ozalp, 2019), even when controlling for known ‘trigger events’. In Europe, violence towards immigrants is related to the degree of hate speech expressed on social media in areas where the violence takes place (Müller & Schwarz, 2018a), while in the US, anti-Muslim messages disseminated by President Trump over social media correlate with the number of anti-Muslim hate crimes in states where social media usage is high (Müller & Schwarz, 2018b). Finally, antagonistic online discussions between opposing groups can predict subsequent offline violence between these same groups (Gallacher, Heerdink, & Hewstone, 2020).

This relationship between online activity and offline violence is also found for more indirect online metrics such as Internet search interest. A higher proportion of racially charged Google search terms is found in areas with higher racial hate crime and racial segregation, and this association increases further with greater broadband availability (Chan, Ghose, & Seamans, 2016). As a result, the Google search volumes for anti-Muslim sentiment can be used to predict anti-Muslim hate crimes (Stephens-Davidowitz, 2017).

Online hate speech may therefore push some users towards offline violence, and while it is unlikely to be the sole cause of this violence, it may act as a ‘trigger’ or facilitatory mechanism (Briggs & Strugnell, 2011; Gaudette, Scrivens, & Venkatesh, 2020).

Evidence that offline hate crime drives online radicalisation (Figure 1 step B)

Online hate is also found to increase following offline violence, suggesting that offline violence can facilitate online radicalisation within extremist groups. For example, the Charlottesville ‘Unite the Right’ rally in 2017 led to a 400% increase in search terms indicating a desire to get involved with violent far-right extremist groups in the weeks following the event (Moonshot CVE, 2018), which reflected wider online interest in joining these groups (Tien, Eisenberg, Cherng, & Porter, 2020). Similar effects have been shown for Islamic extremism, and offline Jihadi-inspired terror attacks have led to increases in online posts advocating for further violence amongst Islamic extremist communities (Olteanu, Castillo, Boy, & Varshney, 2018),

This relationship also has a mutual component to it, and areas with larger far-right communities and greater anti-Muslim hostility offline are also linked with greater levels of pro-Islamic State content online (Mitts, 2019), suggesting that offline hostility against a group leads to greater sympathy from members of this group with extremist movements. Similarly, the geographical localisation of anti-Muslim Internet searches is strongly associated with pro-ISIS searches (Bail, Merhout, & Ding, 2018).

As such, the geographical connection between the online and offline environments may, in this context, result in greater offline impacts caused by online radicalisation in areas where use of far-right Internet platforms is high. This warrants further research.

Evidence that offline terror attacks provoke online hate speech (Figure 1 step D)

Offline violence can also cause a reactionary response from opposing groups online. Online Islamophobia has been shown to increase amongst the far-right following Islamic extremist terror attacks (Burnap et al., 2014; Kaakinen, Oksanen, & Räsänen, 2017; Vidgen, Yasseri, & Margetts, 2019; Williams & Burnap, 2016). Similar effects have also been shown whereby online hate speech increases after certain offline ‘trigger events’ (Awan & Zempi, 2016) including local and national political events (Faris, Ashar, & Gasser, 2016; Saleem, Dillon, Benesch, & Ruths, 2017), and both domestic and overseas terror attacks (Siegel, Tucker, Nagler, & Bonneau, 2017).

This phenomena is further demonstrated by evidence that offline intergroup conflict between far-right extremist groups and opposing groups can lead to increases in the levels of ingroup allegiance for these same far-right groups online, including stronger outgroup sentiments (Bliuc, Betts, Vergani, Iqbal, & Dunn, 2019). However, it is unclear whether these reflexive actions following offline conflict are targeted solely towards the perceived preparators of the violence, or whether they also transfer across to wider outgroups, a process similar to the secondary transfer effect (Pettigrew, 2009). Online, it has been shown that exposure to hate speech against one target group can encourage hate expressions against wider groups (Gallacher & Bright, 2021). Whether this is also true for reactions to offline violence deserves investigation, as if so, it could lead to a worsening of intergroup relations across multiple dimensions.

Evidence that online search is related to joining terror groups (Figure 1 step C)

There is growing evidence that Internet searches for how to join terror and extremist groups is related to genuine interest in doing so, and predictive of future recruitment success for these organisations. Bail et al (2018) report a high geographical association between the Google search phrase ‘*How to join ISIS*’ and offline intention to join the terror group. Similarly, online interest for ISIS is shown to be

predictive of the onset of offline violence from the group (Johnson et al., 2016). Conversely, campaigns which have aimed to prevent or hinder potential extremists from discovering ISIS-related content on Google have been found to help prevent recruitment into the terror group, providing additional evidence that this is an active route for those looking to join the group (Helmus & Klein, 2018). Although testing this mechanism is beyond the scope of this paper, we include it as a step in our model (Figure 1) for completeness and to help elucidate the entire mutual radicalisation cycle.

The current research

These previous studies demonstrate how the dynamics of online and offline intergroup hate are interconnected, and how the relationship between hate speech and offline violence may be a composite and self-reinforcing process. What remains unclear however is the role of the Internet in facilitating the full-cycle process of mutual radicalisation, whether by allowing for online hate which precedes offline violence or/and by amplifying the impact of offline events from opposing groups by giving users space to express shared grievances against the perceived perpetrator group. Additionally, it remains unclear whether these effects are specific to a particular target group, or whether they are more universal and transfer across the range of minority groups targeted by online hate speech. Determining the role of social media in this process is necessary to understand the impact of increasing online hate speech, and to mitigate for its negative impacts on intergroup conflict.

This study addresses these questions by analysing a combination of hate speech messages from Gab, a fringe social media platform popular with the far-right, data on offline extremist attacks, and Internet search data for extremist terms, in the UK, US and Germany.

Overall analytical approach

We construct our analysis in three steps. First, we identify hate speech in online conversations using a supervised machine learning approach (Gallacher, 2021), and determine the group targeted with this hate (Gallacher & Bright, 2021). Secondly, we use Granger causality models to test whether the proportion of online hate speech over the study period relates to offline events including hate crimes and terror attacks from across multiple prejudices, hate targets, and countries, and investigate whether one type of activity consistently precedes the other. Finally, we use Google Trends information to investigate the relationship between offline islamophobia and online Islamic extremist interest. We also use Google Trends data to make local level estimates of interest in Gab and test whether the overall effects between the online and offline environments are stronger at in regions with higher estimated Gab usage.

Gab: a useful platform to study online hate

Fringe social media platforms likely play a key role in the distribution and spread of hate across wider online communities (Nagle, 2017; Zannettou et al., 2017). One of these platforms is Gab, a micro-blogging platform designed as a replica of Twitter, but with lax hate speech moderation policies and an emphasis on tolerance of ‘freedom of speech’. It is known for its high proportion of messages containing explicit hate terms (Lima et al., 2018; Zannettou et al., 2018).

Gab was taken offline in 2018 after the shooter at the Tree of Life Synagogue attack was found to have posted anti-Semitic messages on the platform immediately prior to the attack (Mathew, Dutt, Goyal, & Mukherjee, 2019). It has since returned in a distributed format, making it even less resistant to moderation, and is more popular than ever (Gilbert, 2019; Tech Against Terrorism, 2019). The

platform is heavily focused on political content and topics closely follow current affairs, particularly around ideology, race, and terrorism (Zhou, Dredze, Broniatowski, & Adler, 2019).

This makes Gab a useful platform to study the relationship between online and offline violence. Firstly, the known prevalence of far-right activity and connection with offline violence (Evans, 2018) means that any signals of future violence are likely to be present on the platform. Secondly, Gab data is available historically and the platform is not ephemeral unlike other fringe platforms popular with the far-right (e.g. 4Chan, 8Chan), so it is possible to investigate historical trends. Finally, due to the relatively small userbase, compared to mainstream platforms, we are able to measure hate on the platform as a whole without the need to down-sample or focus on any particular subset from the platform which could lead to a biased estimate of overall hate.

This smaller userbase however, has the limitation that much other far-right extremist content is hosted on other fringe platforms, such as Parler (Aliapoulos, Bevensee, Blackburn, & Cristofaro, 2021), Voat (Papasavva, Blackburn, Stringhini, Zannettou, & de Cristofaro, 2020), Telegram (Baumgartner, Zannettou, Squire, & Blackburn, 2020), and others. While it is beyond the scope of this study to include measurements of these platforms, an awareness that similar parallel conversations likely also take place elsewhere is important.

Countries studied

For this analysis we focus on three countries: the United States of America (USA), the United Kingdom (UK) and Germany. These countries were selected as they have the highest levels of Gab usage (Zannettou et al., 2018). Importantly however they differ in primary language (English vs German). The hate speech detection models we use are trained on English language data and so will not detect German language hate speech. We therefore expect the relationship between the online and offline spaces in the USA and UK to be stronger than in Germany, because in the latter country only a small proportion of the online conversation will take place in English. As such, including Germany will allow us—if indeed we find a weaker association in Germany than in English-speaking countries—to demonstrate more clearly that we are measuring a direct link between online conversation and offline space, rather than an effect driven by an external global variable, which would result in similar associations in all three countries.

Research Aims

This study aims to improve our understanding of how online hate speech relates to offline hateful extremist violence. By measuring the temporal relationship between online hate speech and offline hate crimes, and measuring the reactionary responses, we aim to test how the Internet affects the cycle of mutual radicalisation (Figure 1) between far-right and Islamic extremist groups. By focusing on the entire cycle of mutual radicalisation our study aims to provide a more complete picture of the process, complimenting previous studies which have focused on single stages in this cycle.

First, we test whether far-right online hate speech predicts far-right offline hate crimes (Figure 1 step A). Then, we investigate mutual radicalisation between far-right and Islamic extremism, first by testing whether Islamophobic offline hate crimes are followed by greater online interest in Islamic extremist groups (Figure 1 step B), then by testing the relationship between offline Islamic extremist terror attacks and subsequent far-right online hate (Figure 1 step D).

Based on the theory of mutual radicalisation, we hypothesise that all three relationships described above take place and that we will find that (i) increases in far-right online hate speech precede

increases in far-right offline violence, (ii) increases in Islamophobic offline hate crime precede increases in Google search trends for terms linked with Islamic extremism, and (iii) increases in attacks linked with or inspired by Islamic extremist groups lead to increases in far-right online hate speech..

In addition, we explore whether this first potential relationship between far-right online hate speech and offline violence varies with usage of fringe social media platforms in the local geographical area. We also explore whether this relationship is specific to certain hate types, i.e. whether the groups targeted online are also targeted offline, or whether the effect is more general and transfers across hate types. If this transfer occurs, we would expect to observe increases in online hate against one group to precede increases in hate against wider groups offline.

Methods

Data collection

We utilise data collected from fringe micro-blogging social media platform Gab. We analyse all the messages from the entire platform posted between its formation on 10th August 2016 and the 29th October 2018 when the platform was taken offline following a terror attack committed by one of its users. This dataset contains 33,089,208 messages posted from 259,598 accounts. This period consists of an amalgamation of data shared by (Zannettou et al., 2018) which covers August 2016 – January 2018 and data from the online repository Pushshift which covers the remainder of this period.

Natural language classification

For each message in the whole corpus we collect two natural language classifications: whether the message contains hate speech, and (if so) what type of hate speech (see Gallacher & Bright, 2021 Supplementary Information (SI) Figure 1 for the hate speech processing pipeline).

Hate speech detection

To detect hate speech in the Gab dataset we used a previously developed supervised machine learning approach (see Gallacher, 2021, for details) trained on datasets from Facebook, Twitter, Stormfront and Gab in order to detect various presentations of hate speech against multiple target groups. This approach uses pre-trained contextualised word embedding models (BERT) to identify the semantic features of online hate, along with syntactic features (such as message length and linguistic complexity) and non-linguistic features (such as hate symbols) to identify hate when it is presented in a more nuanced or subtle way. The model classifies all messages into one of three ordered categories; ‘clean’, ‘offensive’, or ‘hate speech’. The inclusion of this intermediate category is found to improve performance over binary classifiers (Davidson et al., 2017).

This approach has been shown to have a good level of performance for Gab data, with an overall accuracy of 90.9%. The precision and recall values for hate speech in particular are 0.75 and 0.83, indicating that this approach will detect more than 8 out of 10 of the hate speech messages on the platform. Conversely, one out of four of the messages flagged as hate speech will be highlighted incorrectly. Whilst not perfect, these levels of accuracy allow us to make a good approximation of the overall hate speech on Gab over time. The overall proportion of hate speech within the Gab conversations over the period studied is shown in SI Figure 2.

Target of hate detection

In order to identify the type of hate speech expressed in the messages, i.e. the group they target, we use guided topic modelling (Anoop, Asharaf, & Deepak, 2016; Li, Chen, Xing, Sun, & Ma, 2019). Specifically, we used Guided Latent Dirichlet Allocation (LDA) (Blei, Ng, & Jordan, 2003), and to ensure consistency we used training data from the same four social media platforms used to train the earlier hate speech detection model outlined above. The details of this approach are described in Gallacher & Bright, 2021, and we use the same LDA model here.

Topic modelling identified six distinct types of hate speech present in the dataset: anti-Semitism, anti-immigration, anti-Black racism, Islamophobia, misogyny, and homophobia, along with an ‘other’ category for messages which could not be classified reliably or contained derogation against other less common groups. These six hate types have been previously highlighted as the most prominent online by a UN report on online hate (Ischinger, 2020), and as the primary groups targeted by the online far-right (Conway, 2018). For each hate speech message, we gathered the type of hate it contained, along with a confidence judgement. Low confidence judgements were re-classified into the ‘other’ category. Accuracy across the six retained topics was high (average 79.7%) but varied between topics (anti-Semitism 90.0%, anti-Black racism 89.6%, anti-immigration 81.0%, Islamophobia 73.7%, homophobia 72.2%, misogyny 71.8%). Further details including inter-coder reliability checks, equivalence zones, and the most salient words in each of these six hate topics are detailed in Gallacher & Bright, 2021.

Offline extremist violence data collection

In order to explore the relationship between online hate speech and offline hateful extremist behaviour we collected data on far-right terror attacks and hate crime from official Government statistics and police reports (e.g. Federal Bureau of Investigation, FBI), civil society crowdsourced datasets (e.g. Anti-defamation league) and academic organisations (e.g. Centre for Research on Extremism Threats / Global Terrorism Database), to build the fullest publicly available picture available of offline hate over the period of interest (August 2016 – October 2018). Details on the sources of offline events along with numbers of events are shown in SI Table 1. Overall, we collected 21,151 offline events from across all countries. The spatial and temporal distributions of these events are shown in Figure 3.

Here we use a definition of hate crime similar to that of the FBI: *“crimes in which the perpetrators acted based on a bias against the victim’s race, colour, religion, or national origin”* (FBI Crime Data Explorer, 2020), but we expand it to include any targeted group based on aspects of that group’s characteristics to align with our earlier definition of online hate speech (see additional discussion in Gallacher, 2021). In all cases we collected the date and location of the events, as well as the motivation of each attack and the details on the group targeted to infer the type of hate crime committed (Islamophobia, anti-Semitism etc.). We also collected information on the number of perpetrators of each attack and the number of victims.

In addition to far-right violence, we also collected data on Islamic extremist-inspired terror attacks perpetrated against the West during the study period (August 2016 – October 2018). Here, we take a definition of ‘the West’ as Europe, North America and Australasia. We focused on these regions as terror attacks there are likely to have received high media and press coverage in English speaking media (Chalabi, 2018; Wendling, 2016), whereas attack in non-western countries do not receive the same level of attention in the West. This data is collected from the Global Terrorism Database (The National Consortium for the Study of Terrorism and Responses to Terrorism (START), 2020) and Wikipedia. As with far-right violence we collected the location and date of the attack, the number of victims and the number of perpetrators. Unlike with far-right attacks whose specific target we identify, we do not differentiate Islamic extremist attacks by victim/target as these attacks are

typically carried out against the secular/Christian West as a whole, and inspired by a complex mix of religious and historical motivations, as well as perceived military operations and territorial incursions (Institute for Global Change, 2016).

To measure offline violence, it is important to include a measure of the severity of the violence rather than simply the presence/absence of it (Gallacher, Heerdink & Hewstone, 2020). As such, for each event we computed an ‘offline violence’ metric calculated as the number of perpetrators of the attack multiplied by the number of victims. In cases of attacks against property rather than individuals, this was treated as if the number of victims was equal to one. This effectively treats each interaction between an attacker and a victim as a separate sub-event (1 person attacking 10 victims is given the same violence score as 2 people attacking 5 victims, as in each case there are 10 potential connections in the violence). This approach is similar to violence indices used in other studies, e.g. using the number of victims (Ritchie, Hasell, Appel, & Roser, 2013), while also placing importance on the number of perpetrators as this also reflects the severity of the intention to cause harm (Dugan, LaFree, Cragin, & Kasupski, 2008).

Google Trends data collection

We collected Google Trends data to infer local estimates of Gab interest and to measure local and national interest trends in Islamic extremist topics. Google Trends is an underutilised data resource for the study of online radicalisation and intergroup conflict. The predictive power of search data has been demonstrated in diverse fields (Jun, Yoo, & Choi, 2018). For example, increases in Google search terms associated with influenza symptoms can be a leading indicator of subsequent outbreaks (Ginsberg et al., 2009) and while these results are sometimes overstated (Lazer, Kennedy, King, & Alessandro, 2014), they indicate the power of using search trends to infer wider behaviours.

Measuring relationships between online and offline activity with Granger causality

To investigate the relationship between online hate speech and offline hate crime we constructed time series of both online hate speech and offline hate crime activity and compared the dynamics of these timeseries using Granger causality tests (Granger, 1969). These test whether changes in time series A are predictive of changes in time series B, and construct a test of significance by comparing the predictive power of this relationship between A and B in comparison to using timeseries A to predict itself. All timeseries were normalised to remove any overall trends (upwards / downwards) and any effects of seasonality.

To better understand the directionality of the relationship between online and offline activity, we compared the relationship in both directions. If online information predicts offline events, but not the reverse, then this indicates that the former time series ‘Granger causes’ the latter. Insight into the directionality of this relationship is the primary benefit of this approach over more simple correlation measures, although we should note that this is not a true causal statement, as other confounding variables may also drive this relationship. This Granger causality is therefore better considered as an approach for forecasting one timeseries using another. This approach has previously been successful in identifying the relationship between online activity and civil unrest (Bastos, Mercea, & Charpentier, 2015).

To control for effects of changing platform popularity over time, we used the proportion of hate speech messages over the whole daily activity on Gab rather than simply the number of hate

messages. As such, more active periods on Gab will not necessarily appear as more hateful, but increases in the density of hate on the platform will.

We compared this relationship between online and offline activity at multiple levels and in different countries separately to ensure robustness of results. Initially we compared how the daily proportion of hate speech on Gab relates to the overall daily offline far-right violence metric for the same period, then to the offline Islamic extremist violence metric, and calculated the two-way Granger causality for each. For each Granger causality test, we ascertained the appropriate lag between the two timeseries by computing the relationship for lags ranging from 1-7 days and retaining the optimum selected as that which minimised Akaike Information Criterion (AIC) (Giles, 2011; Thornton & Batten, 1985). We selected seven days as the upper limit of this lag as effects beyond this range are unlikely to be reliable given the fast turnover of online content.

We then performed more targeted analysis, comparing the online activity around specific hate types to the offline hate crimes of the same type. We look at these effects separately for the USA, UK, and Germany. Statistical analysis was performed using the `lmtest` package (Hothorn et al., 2020) and significance values were adjusted with Holm’s correction to control for the effects of multiple comparisons.

Estimates of local Gab usage

We used local Gab Google search interest as a geographical correction factor for our offline violence data. Gab data is not geolocated, and so to take an approximation of Gab usage by country and region over this period, we collected the Google search interest for ‘Gab’ (taking the specific social media search term rather than keyword) over the period using the Google Trends tool (Google, 2020). This gives us the relative interest in the search term over a given period, and the distribution of this interest over states (USA), and sub-national regions (UK and Germany) (see Figure 2 for an example).

We make this estimate in order to provide additional validation for the potential relationship between online hate and offline violence. If there is an instrumental relationship between the two, we would expect a more significant association (Granger causality F-statistic) in geographical areas where Gab usage is higher. This is to say that following spikes in online hate, any subsequent offline violence would be more likely to occur in areas with a lot of Gab users present, compared to areas with very few Gab users. As such, we would expect the relationship of online hate on offline violence to be more significant, i.e. a greater predictive power of online hate on offline violence, after weighting by local Gab interest.

Google Trends data is only available at the daily level for periods shorter than 9 months and so in order to collect a correctly normalised dataset over the entire period we collect overlapping 9-month periods (1-month overlap) and re-scale the subsequent dataset to match the trailing/overlapping final month of the previous collection period. This approach accurately estimates daily Google search interest over longer than 9-month periods (Bleher & Dimpfl, 2019; Tseng, 2019). We perform this calculation separately for each geographical area, giving us daily trends in the relative google search interest at the country level. Subsequently, to estimate Google search interest for each country’s sub-regions we rescaled the local daily search interest by the country-wide search interest on that day from the national timeseries data. Data collection was performed using the `gtrendsR` package (Massicotte & Eddelbuettel, 2020) to connect to the Google trends API. We used this Google trends data to weight the violence of each offline event by the relative level of Google search interest in the

Figure 2 – Relative Google search interest for ‘Gab’ over the period of interest for the USA, UK and Germany. The map shows the geographical distribution of search interest in the USA for a given month (August 2017) for demonstration purposes.

state/sub-region on the day of the event, to obtain a timeseries of offline violence weighed by local Internet interest in the Gab platform.

We re-ran the Granger Causality analysis outlined above (on the general hate measure only, not separated by type) with this new corrected timeseries, for each country, and compare the significance of that relationship to that of the original model (without correction for local Gab interest).

Estimates of interest in joining Islamic extremist groups

We used the approach from Bail et al (2018) to estimate the local interest in Islamic extremism and in joining Islamic extremist groups, and make use of the search phrase ‘How to join ISIS’ which is highly correlated with offline intention to join the terror group. Robustness checks performed by Bail et al show that the level of false positive associations with this term are low, and few people enter this term while searching for unrelated topics. This indicates it is a reasonable proxy for broader Islamic extremist sentiment (Ceron, Curini, & Iacus, 2019). It should be noted however that the relative power and influence of ISIS as a terror organisation was declining globally over the study period (Burke, 2017), and was of varying relevance to citizens in Europe vs the USA (Byman, 2016) due to

geographical distance from the group and the prevalence of terror attacks. We accounted for these effects by ensuring the timeseries for ISIS search interest was de-trended prior to conducting the Granger causality to account for reducing global sway, and by conducting separate analyses for the USA, UK, and Germany.

As with the Gab interest outlined above, we collected country-wide interest over the entire period for the USA, UK and Germany. We then collected search term volumes at the local daily level and combined these with long-term estimates to generate local long-term trends. Due to low search volumes in the UK and Germany at the local daily level, we only collected the local daily trends for the USA. These trends are shown in SI figure 1.

Data analysis

All analysis was done in R (version 3.6.1) using the tidyverse (Wickham et al., 2019) and tidymodels (Kuhn, Wickham, & RStudio, 2020) collections of packages. Where particular additional packages have been used, they have been referenced in the text.

Ethics

All research was conducted in accordance with the University of Oxford Ethics Committee (Reference: SSH_OII_CIA_19_062). All data collection was conducted using open source methods and publicly available data, and hence, informed consent was not explicitly obtained. In order to preserve anonymity, we took a cryptographic hash of all usernames prior to analysis, and real account usernames were not used in the analysis at any point.

Figure 3 – Offline events used in the Granger causality analysis. Upper panel gives the spatial distribution of events coloured by motivation. Lower panel gives the temporal distribution of events in each country. Larger circles indicate greater levels of violence

Results

Overall, 8.9% of messages on Gab during the study period were classified as hate speech.

Islamophobia was the most common type of hate speech with 24.1% of all hate messages classified into this category, followed by anti-immigration sentiment (22.1%) and then anti-Semitism (9.5%). anti-Black racism (2.2%), homophobia (1.9%) and misogyny (1.6%) were less common in this dataset, but still with substantial quantities given the dataset size. The remaining 38.5% of hate speech messages were contained within the ‘other’ category.

Online hate speech precedes far-right offline violence

When testing for two-way Granger causality between the timeseries in online hate and offline violence, we observed significant Granger causality in the USA and the UK between overall levels of hate speech on Gab preceding far-right hate crime offline (USA; $df = 747$, $F = 6.80$, $p = 0.008$, UK; $df = 747$, $F = 4.93$, $p = 0.026$), while there was no observed effect in the reverse direction. (USA; $df = 733$, $F = 1.31$, $p = 0.270$, UK; $df = 738$, $F = 1.250$, $p = 0.273$). This means that the proportion of hate speech on Gab was predictive of future offline hate crimes, in both countries, with the strength of this effect larger in the USA. In Germany we observed no relationship in either direction, see Table 1.

Looking at these effects for specific hate types, we found significant Granger causality in the USA for anti-Semitism, Islamophobia and anti-immigration, whereby in all three cases spikes in online hate of this type preceded increases in offline hate crimes against the same specific target groups. In the UK, this effect was most significant for Islamophobia, less but still significant for homophobia and anti-immigration, but not significant for anti-Semitism or other hate types. Looking at the reverse direction, we found a relationship where offline Islamophobic hate crimes preceded online hate speech on Gab in the UK. Again, no specific relationships were observed in Germany. Full results across all hate types and countries are given in Table 1.

Offline Islamic extremist violence precedes online far-right hate speech

With regard to extremist terror attacks, spikes in the level of offline Islamic extremist violence preceded spikes in online far-right hate speech ($df = 741$, $F = 82.78$, $p < 0.001$) – the reverse to the effect observed for far-right violence, where online hate speech preceded offline violence. This effect occurred predominately in terms of Islamophobia and anti-immigration sentiment online, but also spread to include misogyny and homophobia (Table 1). Figure 4 demonstrates this effect for Islamophobia in the USA, whereby spikes in Islamic extremist terror attacks preceded increases in Islamophobia on Gab, which in turn preceded increases in offline Islamophobia hate crime. To check whether this primary Granger causality effect (online behaviour preceding offline events) was driven purely by a relationship between offline events we also ran Granger causality tests directly between offline far-right and Islamophobic events, and found a Granger causality effect of Islamic extremist terror attacks preceding increases in Islamophobic hate crimes in the UK ($df = 741$, $F = 9.41$, $p < 0.001$) and Germany ($df = 741$, $F = 22.0$, $p < 0.001$), but not in the USA ($df = 743$, $F = 0.290$, $p = 0.590$).

Figure 4 – Example of Granger causality between Islamic extremist offline violence (yellow), Islamophobia on Gab (blue), and offline Islamophobia hate crime (red) in the USA. Vertical dotted lines highlight temporal differences in patterns, with spikes in offline Islamic extremist violence preceding spikes in online islamophobia, which themselves precede spikes in Islamophobic hate crimes.

Table 1 – Two-way Granger causality effects for online hate speech and offline violence across combined and specific online far-right hate types. Sections 1:3 relate to far-right attacks in the US, UK, and Germany; in these sections hate types relate to both the motivation of the attack and the types of online hate speech. The final section relates to Islamic extremist attacks against the ‘west’, the hate types relate only to online hate types. Hate types are ranked by *F* statistic for each country and test, darker yellow indicates a larger estimate of *F* (stronger signal) while darker blue indicates a lower *p*-value.

Online Hate → Offline Hate				Offline Hate → Online Hate			
HATE TYPE	DF	ESTIMATE	P	HATE TYPE	DF	ESTIMATE	P
USA - Far-right violence							
Anti-semitism	741	7.107	0.008	Homophobia	740	3.656	0.056
All hate types	747	6.800	0.009	Misogyny	733	1.838	0.120
Anti-immigration	741	5.210	0.023	Anti-black racism	733	1.399	0.233
Islamophobia	733	3.095	0.003	All hate types	745	1.313	0.270
Misogyny	739	2.277	0.132	Islamophobia	735	1.306	0.252
Homophobia	736	1.307	0.271	Anti-semitism	739	0.877	0.416
Anti-black racism	737	0.732	0.481	Anti-immigration	739	0.636	0.530
UK - Far-right violence							
All hate types	747	4.943	0.026	Islamophobia	740	3.021	0.017
Islamophobia	746	4.320	0.038	Misogyny	731	1.971	0.081
Homophobia	732	3.701	0.003	Anti-semitism	728	1.894	0.068
Anti-immigration	738	2.995	0.030	Anti-black racism	738	1.560	0.211
Anti-semitism	734	1.621	0.167	All hate types	735	1.250	0.273
Anti-black racism	740	0.759	0.384	Anti-immigration	732	1.025	0.408
Misogyny	733	0.497	0.738	Homophobia	740	0.160	0.689
Germany - Far-right violence							
Islamophobia	745	3.664	0.392	Anti-black racism	737	4.661	0.068
Anti-immigration	737	1.984	0.690	Misogyny	727	1.768	0.544
All hate types	746	1.195	1.000	Anti-immigration	729	1.581	0.689
Anti-black racism	727	1.126	1.000	Anti-semitism	732	0.940	1.000
Homophobia	740	1.104	1.000	Islamophobia	741	0.706	1.000
Anti-semitism	740	0.719	1.000	Homophobia	740	0.141	1.000
Misogyny	737	0.279	1.000	All hate types	746	0.035	1.000
Global - Islamic extremist attacks against the 'West'							
Islamophobia	737	2.925	0.089	Islamophobia	741	82.784	0.000
Homophobia	728	2.276	0.161	All hate types	742	41.982	0.000
All hate types	738	2.038	0.357	Anti-immigration	741	9.932	0.008
Mysogyny	739	1.481	0.896	Mysogyny	735	4.819	0.010
Anti-immigration	729	1.174	0.947	Homophobia	734	3.068	0.048
Anti-semitism	732	1.007	0.947	Anti-semitism	740	1.461	0.454
Anti-black racism	739	0.415	0.947	Anti-black racism	735	0.810	0.489

Offline-online relationship adjusted for local Gab interest

To test whether these effects were more significant in areas with greater Gab usage, we used data on Google search interest for ‘Gab’ to make local-level adjustments to the ‘violence’ timeseries. Locally-adjusted effects of Granger Causality between online hate speech and far-right offline hate crime were more significant in regions with higher Gab interest in the USA and UK, however again we observed no effects in Germany (Table 2).

Both Islamophobic violence and Islamic extremist violence precede online Islamic extremist interest

When testing for two-way Granger causality between the offline far-right hate crime and Google interest (search trends) for Islamic extremist content we observed significant Granger causality in the USA and the UK. In both countries overall hate crime levels preceded Google interest, while there was no relationship for the reverse (Table 3). When looking at hate crimes against specific targeted groups we found that this effect was specific to Islamophobic hate crimes in the USA, while in the UK the effect was most significant for Islamophobic hate crimes but also significant for anti-immigration hate crimes and anti-Semitic hate crimes. The results for all hate crime types are given in SI table 2. Investigations of local state-wide effects in the USA pointed towards potential state-specific relationships, but trends were not sufficiently clear to make a robust conclusion (SI Section 5).

We also found significant Granger causality between Islamic extremist terror attacks against the West and Google Islamic extremism interest. In all three countries (USA, UK, and Germany) increases in Islamic extremism interest occurred following increases in offline Islamic extremist violence (Table 3). These effects were most significant for the UK but significant for all three countries.

In Germany (but not elsewhere), we also found a significant Granger causality whereby the Islamic extremist Google search volumes preceded offline far-right hate crimes overall, and specifically Islamophobic hate crimes, anti-immigration attacks, anti-black racism and misogynistic hate crimes (Table 3, SI Table 2).

Table 2 – Google adjusted granger causality effects for online hate and offline violence for the USA, UK, and Germany. The difference in model performance for these internet-Gab-search-adjusted models from the unadjusted model is given in the ‘Delta F’ column. Darker yellow indicates a larger estimate of F (stronger signal) while darker blue indicates a lower p-value.

TEST DIRECTION	DF	ESTIMATE	P	DELTA F
USA - Far-right violence				
Offline Hate → Online Hate	747	11.215	0.001	4.415
Online Hate → Offline Hate	745	1.941	0.144	0.628
UK - Far-right violence				
Offline Hate → Online Hate	747	6.018	0.014	1.075
Online Hate → Offline Hate	735	1.005	0.426	-0.245
Germany - Far-right violence				
Offline Hate → Online Hate	747	2.088	0.149	0.894
Online Hate → Offline Hate	735	0.648	0.716	0.613

Table 3 – Granger causality effects between offline hate and offline Islamic extremist google search volumes for the USA, UK, and Germany. Darker yellow indicates a larger estimate of F (stronger signal) while darker blue indicates a lower p-value. Results are given for far-right hate overall and Islamophobic hate specifically, and for Islamic extremist Terror attack violence

Offline Hate → ISIS Internet Searches				ISIS Internet Searches → Offline Hate			
HATE TYPE	DF	ESTIMATE	P	HATE TYPE	DF	ESTIMATE	P
USA - Far-right violence							
All hate types	744	8.175	0.001	Islamophobia	747	2.534	0.224
Islamophobia	745	4.702	0.019	All hate types	738	2.397	0.072
UK - Far-right violence							
Islamophobia	743	6.555	0.001	Islamophobia	739	0.999	0.417
All hate types	743	5.685	0.002	All hate types	747	0.681	0.410
Germany - Far-right violence							
Islamophobia	742	1.975	0.116	All hate types	734	4.873	0.000
All hate types	742	1.267	0.285	Islamophobia	734	3.932	0.001
Global - Islamic extremist attacks against the 'West'							
UK	738	159.296	0.000	Germany	738	1.246	0.286
Germany	744	34.545	0.000	UK	734	0.926	0.485
USA	738	12.685	0.000	USA	746	0.126	0.723

Discussion

In this study we analysed the dissemination of online hate speech on a fringe social media platform popular with the far-right, and investigated how online hate speech relates to offline hate crime, hateful extremist violence, and mutual radicalisation.

We found a positive association between online hate speech on Gab and far-right offline hate crimes, with the former preceding the latter. The effect was most significant within specific hate types—in other words online hate speech against a specific group more strongly predicted offline attacks against that group, and was more significant in areas with higher levels of interest in fringe social media platforms. Conversely, we found that hate speech against specific groups spiked following offline violence associated with connected groups. Finally, we found that offline Islamophobic violence preceded increases in online interest in Islamic extremism. Together, these results highlight the key role of the Internet in processes of mutual radicalisation and support the cyclical process presented in Figure 1, and is congruent with the idea that the online environment mediates both the instigation of offline violence, and the response to it.

Online hate predicts offline violence

We find a robust relationship between online hate speech and offline hate crimes in the USA and the UK. In both countries, increases in the proportion of online hate speech on Gab temporally preceded increases in offline far-right violence targeting minority groups. The reverse effect did not occur however, and offline far-right hate crimes did not reliably precede increases in online hate speech,

demonstrating that this is a unidirectional relationship. This supports prior studies which have shown a relationship between online hateful content and offline violence (J. D. Gallacher et al., 2020; Müller & Schwarz, 2018a, 2018b; Williams et al., 2019). Our results help confirm this relationship, while also providing new insight into how this effect is strongest for certain hate types (Islamophobia & anti-Semitism) but also present for a wider range of targeted groups (homophobia, anti-immigration sentiment).

Importantly, while our use of Granger causality models to demonstrate this directionality provides stronger evidence for than simple correlative approaches would have done, this directionality itself does not necessarily imply direct causality between the online and offline hate, nor does it imply that individual users active in the online conversations were the same individuals involved in the offline violence. Instead, this result indicates that there is something in the hate speech of online conversations which is indicative and predictive of the future violence.

There are three likely explanations for this online-offline relationship. Firstly, it has been shown that there are some individuals radicalised in these types of online conversations, for whom the conversations may act as a trigger for existing views. In these cases, the attackers express indicators of their attacks before the event (possibly instigating wider hate and encouragement from other users), then subsequently commit the attack. In this way, the online hate speech is a direct signal of the future offline violence (Meloy, Hoffman, Guldemann, & James, 2011). This pattern has been shown in a small number of high-profile and high-violence cases (e.g. Minnesota (USA), 2015; Toronto (Canada), 2017; Christchurch (New-Zealand), 2019; Poway (USA), 2019; El Paso (USA), 2019, among others), but these events are rare and this pattern is unlikely to apply to the majority of offline violence. Indeed, this is highly improbable given the relatively small userbase on Gab (~250,00 users for the period studied) when compared to the populations size of the USA and UK.

A second more plausible explanation is that Gab is a representative sample of much wider far-right communities both online and offline. Therefore, measuring the hate on Gab at any one point in time gives an estimate of the overall hateful sentiment from these wider communities. Any individual from these wider communities may engage in violence completely separately from the online conversations, but the motivating factors are mirrored in the online space. This would make hate crime occurrence similar to that of stochastic terrorism (Hamm & Spaaij, 2018; May, 2020; Woo, 2002), where random acts of ideologically-motivated violence are statistically predictable globally (from a fixed point of collection) but individually unpredictable.

A third, and related, potential explanation for this relationship is that the online environment and offline world may be separate leading and lagging indicators of a third instigation to violence – such as a political decision, a terror attack etc (e.g. Awan & Zempi, 2016). In this way, following a trigger event, hate increases quickly online, and then latterly occurs offline against the same targets. This gives the impression of the online spike causing the offline, but in reality, they are both responses to the third variable. While this third explanation is likely in some situations, a relationship between online and offline hate has been shown previously even when controlling for known trigger events (Williams et al., 2019). Similarly, in the current study we found no direct relationship between offline Islamic extremist terror attacks and offline far-right Islamophobic hate crimes in the USA, but we do find that these events are linked by online activity (discussed below), which supports a more substantial involvement of the Internet in processes of online and offline hate.

The second of these explanations, that Gab is a representative sample of much wider far-right communities, is bolstered by the results for specific hate types individually. In addition to the overall

effect of online hate speech preceding offline violence, we also found that increases in specific types on hate of Gab predicted increases in hate crimes against these same groups offline. In the USA, this relationship was significant for anti-Semitism, Islamophobia and anti-immigration, while in the UK the effect was strongest for Islamophobia but also significant for homophobia and anti-immigration. Additionally, the use of local Google search interest in Gab showed that online hate was more predictive of offline overall hate in geographical areas with higher Gab interest. Similar local effects have been shown previously in the relationship between anti-refugee sentiment on mainstream social media platforms and offline violence in Germany (Muller & Schwartz, 2018a), and the relationship between online hate and offline violence on Twitter in the USA (Muller & Schwartz, 2018b). Taken together, these results support the argument that measuring the density of hate speech on Gab can be used to gauge the violent sentiment of wider far-right groups that have similar geographical distributions.

While this relationship between online hate speech and offline hate crime was shown in both the USA and UK, we did not find it in Germany. This lack of significant relationship in Germany compared to its presence in English-speaking countries provides additional support that there is a meaningful relationship between online hate and offline violence in the USA and UK. If the effects in the UK and USA were due to hidden third variables causing both online and offline activity, we would expect to observe a similar relationship in Germany. Instead, given that our hate speech classifier was trained on English language content and will therefore not have performed for German language messages, we have not detected German users exposure to hate speech messages in German but only in English, which likely only represents a small proportion. Additionally, while Germany is the third most common country of origin for Gab users (Zannettou et al., 2018), it trails the USA and UK substantially, with only 5% of accounts belonging to German users (Zhou et al., 2019) compared to 88% from the USA or UK. This means that the sample may simply be too small to detect a relationship, and that much meaningful German fringe far-right conversations are occurring elsewhere (Guhl, Ebner, & Rau, 2020). Finally, there may be high levels of inauthentic accounts or automated activity for German accounts on Gab (Zhou et al., 2019). In future, developing cross-language hate speech detection approaches would bring additional insight into international variation in these relationships (Vidgen & Derczynski, 2020). The fact that we observe an offline-to-offline effect of global Islamic extremist violence preceding offline far-right hate crime in Germany, as we do in the USA and UK, supports the idea that mutual radicalisation is a more global phenomenon. However, the lack of predictive power of online hate on these German attacks shows that this effect is not mediated by English language Gab content in Germany.

Offline violence predicts online hate

A key finding of our study is that in addition to far-right online hate speech predicting offline hate crimes against minority groups, we also found a significant relationship whereby offline Islamic extremist terror attacks preceded increases in online hate speech on Gab. This effect was most strongly observed for Islamophobia, indicating a reactionary response to the violence, but one which attacks the perceived outgroup as a whole (in this case Muslims) rather than focusing on the attackers themselves. In addition to Islamophobia, we also observed increases in anti-immigration sentiment, misogyny and homophobia. This supports prior evidence that offline violence can be a ‘trigger event’ for spikes in online hate speech (Burnap et al., 2014; Kaakinen et al., 2017; Williams & Burnap, 2016), and builds on this by demonstrating that while this reaction may be primarily focused on the perceived instigators of the violence, it also spreads across to target wider minority groups. This is suggestive of a transitive effect for intergroup conflict (Pettigrew, 2009; Gallacher & Bright, 2021), with offline terror attacks leading to wider conflict with other groups which may be perceived to be similar or associated in some way to the perpetrator group, such as non-Muslim immigrants. In this

way, Islamic terror attacks may lead to a reduction in intergroup relations between far-right groups and a range of perceived outgroups.

Importantly, our results also show that offline far-right hate crime is predictive of subsequent increases in Google search volumes for Islamic extremist material. This supports prior suggestions that offline discrimination can be an initial driver of the process of radicalisation (Dechesne, 2009; Moghaddam, 2005; Sageman, 2008). Given that the sharing of common grievances, such as experiences of prejudice and discrimination, is suggested to be a factor driving group extremism (Reicher & Haslam 2008; Smith et al, 2019) this result is particularly concerning. In the USA, this association was found only for Islamophobic hate crimes and no other type of hate, supporting the specificity of this relationship, and suggesting that it is the occurrence or perception of Islamophobic hate crime itself which drives the relationship rather than broader trends of intolerance against multiple groups. In the UK, this association was strongest following Islamophobic hate crimes, but also occurred after anti-immigration and anti-Semitic hate crimes, suggesting a degree of overlap between targeted groups, or a general sentiment of discrimination against minority groups. Overall, this result supports prior evidence that offline Islamophobia is linked with interest in online Islamic Extremist content in a given location (Mitts, 2019) and that the geographical localisation of Islamophobic Internet searches and Islamic extremist searches is strongly associated (Bail, Merhout, & Ding, 2018). We build on this and show that this association is likely directional, with offline discrimination preceding online interest in these extremist topics.

In addition, we find a strong relationship between Islamic extremist terror attacks and Islamic extremist Google search volumes in all three study countries. This suggests that in addition to prompting reactionary hate from far-right groups, Islamic extremist terror attacks may also inspire sympathisers to investigate how to join the terror group themselves (Berger, 2014; Berger & Morgan, 2015). This result should be considered with caution however, due to the risk of false positive searches caused by interest in the terror group from across the spectrum being detected in the Google trends result (see limitations and future directions).

Role of the Internet in the Mutual Radicalisation cycle

Our results shine new light on the role of the Internet in processes of mutual radicalisation. While the process of mutual radicalisation has been demonstrated for offline intergroup conflict (e.g. Konaev & Moghaddam, 2010; Moghaddam, Harré, & Lee, 2008), we show that the Internet may have a catalyst function which facilitates this process further. Previously, users on fringe far-right platforms have been shown to respond to real world political events with greater online activity (Scrivens, 2020; Scrivens et al., 2020; Zannettou et al., 2018). Our findings expand on this and show that offline Islamic extremist violence predicts online Islamophobic hate and demonstrate a degree of mutual radicalisation between far-right and Islamic extremist groups following Islamic extremist terror attacks (Figure 1 step D). External prompts which are perceived to legitimise radical beliefs have been shown to stoke outgroup derogation (Thomas, McGarty, & Louis, 2014). In this way offline terror attacks may be perceived to legitimise Islamophobic attitudes leading to an increase in online hate speech towards the perceived perpetrator group.

This increase in online hate following offline events can itself instigate processes of group radicalisation and extremism as other ingroup members are exposed to this negative sentiment from their ingroup contacts and adopt this behaviour themselves through a process of ‘social contagion’ (Gallacher & Bright, 2021; Ferrara, 2017; Valente, Dyal, Chu, Wipfli, & Fujimoto, 2015). This is reflected in our finding that online hate speech is associated with subsequent increases in offline hate crime across multiple targeted groups. This suggests that any reactionary online responses from far-

right groups to perceived outgroup actions may also lead to further offline violence, thereby potentially creating a self-perpetuating cycle of violence (Moghadam, 2018). Prior evidence suggests that over time individual events such as verbal abuse or physical confrontation can lead to ‘micro-radicalisations’ (Bailey & Edwards, 2017), where negative intergroup contact between far-right extremist group members of Islamic extremist groups leads to mutual escalation. Additionally, direct antagonistic contact between opposing groups in unstructured and relatively unmoderated online spaces may also contribute to mutual radicalisation (Gallacher, Heerdink and Hewstone, 2020). Our results build on these findings and suggest that even in the absence of direct online contact between opposing groups, the Internet can play a role in facilitating the effects of mutual radicalisation, both in reaction to, and an indication of, offline violence (Figure 1 step A).

Finally, we find that Islamophobic offline violence is associated with subsequent online interest in joining Islamic extremist groups (Figure 1 step B). Whether this interest translates directly into the desire to join terror groups is somewhat unclear (Sageman, 2008). Regardless, this evidence helps us to better understand the role of the Internet in driving intergroup conflict by providing easy access to global communities to discuss shared grievances (Smith et al, 2019), and demonstrate the interplay between online and offline activity. By responding with hate speech to terror attacks, online far-right groups increase the chance that the two groups will mutually escalate, increasing the likelihood of further conflict (Figure 1 & Figure 4).

Limitations and future directions

Although we employed statistical methods to provide more robust temporally directional associations than simple correlation (Granger, 1969), this work remains observational and associative in nature, so caution should be taken not to infer strong causal mechanisms. In addition, the statistical and machine learning approaches we employ have limitations which need to be kept in mind when interpreting results.

Topic models for hate speech identification

The topic modelling approach that we use in this work to classify hate speech into distinct types assumes that each message can be classified into exactly one topic, although it may in fact contain multiple types of hate (Burnap & Williams, 2016). The conflation and combination of hate targets may play a role our observed effects of the transfer of hate from one target group to another, but whether the simultaneous expression of hate against multiple targets assists in these transfer effects of hate is unknown. As such, future research should explore this phenomenon in greater depth and include more detailed and nuanced hate type distinctions, for example through a combination of qualitative and quantitative work (e.g., measuring topic overlap in topic modelling approaches in combination with manual inspection).

Data limitations

In this study we analyse data from fringe social media platform Gab, mainly because of the availability of Gab data and how the platform historical data is stored, allowing us to study messages from the entire platform over a >2-year period. However, considering a wider range of social media platforms, and the interplay between them, would allow to for more generalisable results across topics and geographical locations, and should be an interesting avenue for future research. A challenge to overcome to work with fringe platforms popular with the online far-right such as 4chan and 8chan, or closed messaging apps such as Discord and Telegram, relates to the more ephemeral nature of these platforms. Messages and conversation threads within these platforms are deleted after a short period of time, and therefore studying these platforms and the role they play in spreading hate and division

will require forward thinking data collection plans, and data contracts which preserve user anonymity and platform user agreements. Progress has been made recently in curating public datasets covering these platforms (Baumgartner et al., 2020; Papasavva, Zannettou, De Cristofaro, Stringhini, & Blackburn, 2020), but as online hate groups migrate between fringe platforms this will present an ongoing challenge. These platforms have been linked with numerous offline terror attacks (Evans, 2019), and so understanding their role in the wider online ecosystem is vital.

Another caveat in our study relates to our use Internet search trends to measure phenomena which would otherwise be very difficult to gauge, such as interest in joining specific extremist groups or local geographic distribution of Gab usage. A primary drawback of this method is that Internet search trends can have a high degree of noise or false positive association (Lazer et al., 2014). The Google Trends platform provides estimates for phrases based on all possible combinations of keywords in each search phrase (Bail, Merhout, & Ding, 2018), so a query for the phrase ‘are they evil’ will return the same results as a search for ‘they are evil’ despite substantial differences in meaning. It is therefore possible that results will include reactionary searches from those interested in the topics from an external or observer position, as well as those genuinely interested in the content. This limitation does not affect our results linked with Google search data for “Gab” as these were for the specific social media platform, but it may affect our results linked with searches for Islamic extremist content. This could potentially explain in part the Granger causality results for an association between ISIS search volume following terror attacks, as these attacks may cause greater interest in these terror groups as well as a greater radicalising effect. Similarly, it is likely that our observed results whereby Islamic extremism search volumes predict subsequent far-right hate crimes in Germany may be caused by false positives being detected from Internet searches on Islamic extremist topics by the far-right. Future work should consider collecting multiple search topics across orthogonal dimensions to create an unbiased estimate of the search topic of interest.

Finally, there are also limitations regarding our dataset of offline hate crimes, including systematic under reporting of certain hate types. This may hinder research efforts to understand how the online environment affects these types of attacks and ultimately the implementation of adequate protective measures. For example, misogynistic attacks on women are underreported (Scott, 2020), and while domestic abuse statistics could be used as a proxy, this is an imperfect solution (Craanen, Berntsson, Ging, & DiBranco, 2020). We identified misogynistic hate speech as one of the top six most common on Gab, but it is not currently considered a protected characteristic under UK legislation (Grierson, 2020). Given a recent increase in popularity of the so-called ‘manosphere’ online, a conglomerate of online misogynist movements focused on “men’s issues” including ‘involuntary celibate’ groups and ‘pick-up artists’, this is especially concerning. These groups have been directly linked with at least six large scale hateful extremist violent events since 2014 (Hoffman, Ware, & Shapiro, 2020) as well as numerous large-scale online harassment campaigns (Nagel, 2017). Efforts are therefore needed to address these gaps so that the negative impacts of online hate across the entire spectrum of targeted groups can be understood and then mitigated.

Conclusion

This study provides evidence for the damaging effects of online hate speech in fringe platforms popular with the far-right on intergroup conflict, extremist violence and mutual radicalisation. We show that the online environment does not occur in isolation from the offline world, instead we find a high degree of association between online and offline hate, with online hate speech preceding offline hate crimes of the same motivation, while offline events can both predict subsequent online hate speech against perceived perpetrators and interest in extremist topics from the target groups. This

negative feedback loop has serious implications for intergroup relations, intergroup conflict and extremist violence. Together, our results emphasise the potential central role that Internet activity plays in the process of mutual radicalisation, which both drives far-right extremism and goes hand-in-hand with increased Islamic extremism worldwide.

References

- Aasland, J. R. (2016). Right-wing terrorism and violence in Western Europe: Introducing the RTV dataset. *Perspectives on Terrorism*, 10(3), 2–15. Retrieved from <http://www.terrorismanalysts.com/pt/index.php/pot/article/view/508>
- Aliapoulos, M., Bevensee, E., Blackburn, J., & Cristofaro, E. De. (2021). An early look at the Parler online social network. *ArXiv*, 1–8. Retrieved from <https://idrama.science/papers/parler-2021-01-07.pdf>
- Anoop, V. S., Asharaf, S., & Deepak, P. (2016). Unsupervised concept hierarchy learning: A topic modeling guided approach. *Procedia Computer Science*, 89, 386–394. <https://doi.org/10.1016/j.procs.2016.06.086>
- Awan, I., & Zempi, I. (2016). The affinity between online and offline anti-Muslim hate crime: Dynamics and impacts. *Aggression and Violent Behavior*, 27, 1–8. <https://doi.org/10.1016/j.avb.2016.02.001>
- Bail, C. A., Merhout, F., & Ding, P. (2018a). Using Internet search data to examine the relationship between anti-Muslim and pro-ISIS sentiment in U.S. counties. *Science Advances*, 4(6), 1–10. <https://doi.org/10.1126/sciadv.aao5948>
- Bail, C. A., Merhout, F., & Ding, P. (2018b). Using Internet search data to examine the relationship between anti-Muslim and pro-ISIS sentiment in U.S. counties. *Science Advances*, 4(6), 1–10. <https://doi.org/10.1126/sciadv.aao5948>
- Bailey, G., & Edwards, P. (2017). Rethinking ‘radicalisation’: Microradicalisations and reciprocal radicalisation as an intertwined process. *Journal for Deradicalization*, (10), 255–281.
- Bastos, M. T., Mercea, D., & Charpentier, A. (2015). Tents, tweets, and events: The interplay between ongoing protests and social media. *Journal of Communication*, 65(2), 320–350. <https://doi.org/10.1111/jcom.12145>
- Baumgartner, J., Zannettou, S., Squire, M., & Blackburn, J. (2020). The Pushshift Telegram dataset. *ArXiv*, 1–7. Retrieved from <https://arxiv.org/pdf/2001.08438.pdf>
- Benigni, M. C., Joseph, K., & Carley, K. M. (2017). Online extremism and the communities that sustain it: Detecting the ISIS supporting community on Twitter. *PLoS ONE*, 12(12), 1–23. <https://doi.org/10.1371/journal.pone.0181405>
- Berger, J. M. (2014). How ISIS games Twitter. *The Atlantic*. Retrieved from <https://www.theatlantic.com/international/archive/2014/06/isis-iraq-twitter-social-media-strategy/372856/>
- Berger, J. M. (2018a). *Extremism*. Cambridge, Massachusetts: The MIT Press.
- Berger, J. M. (2018b). *The Alt-Right Twitter census: Defining and describing the audience for alt-right content on Twitter*. VOX-Pol Network of Excellence. Retrieved from https://www.voxpol.eu/download/vox-pol_publication/AltRightTwitterCensus.pdf
- Berger, J. M., & Morgan, J. (2015). *The ISIS Twitter census: Defining and describing the population of ISIS supporters on Twitter*. *The Brookings Project on U.S. Relations with the Islamic World*. Retrieved from <https://www.brookings.edu/research/the-isis-twitter-census-defining-and-describing-the-population-of-isis-supporters-on-twitter/>
- Bleher, J., & Dimpfl, T. (2019). Knitting multi-annual high-frequency Google trends to predict inflation and consumption. *SSRN Electronic Journal*. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3357424
- Blei, D., Ng, A., & Jordan, M. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 12(4), 421. <https://doi.org/10.2307/3050792>
- Bliuc, A. M., Betts, J., Vergani, M., Iqbal, M., & Dunn, K. (2019). Collective identity changes in far-right online communities: The role of offline intergroup conflict. *New Media and Society*, 21(8), 1770–1786. <https://doi.org/10.1177/1461444819831779>
- Brandsch, J., & Python, A. (2020). Provoking ordinary people: The effects of terrorism on civilian violence. *Journal of Conflict Resolution*, 1–31. <https://doi.org/10.1177/0022002720937748>
- Briggs, R., & Strugnell, A. (2011). *Radicalisation: The role of the Internet. A working paper of the PNN*. Retrieved from <https://www.isdglobal.org>
- Burke, J. (2017). Rise and fall of Isis: its dream of a caliphate is over, so what now? *The Guardian*. Retrieved from <https://www.theguardian.com/world/2017/oct/21/isis-caliphate-islamic-state-raqqa-iraq-islamist>
- Burnap, P., & Williams, M. L. (2016). Us and them: Identifying cyber hate on Twitter across multiple protected characteristics. *EPJ Data Science*, 5(1). <https://doi.org/10.1140/epjds/s13688-016-0072-6>
- Burnap, P., Williams, M. L., Sloan, L., Rana, O., Housley, W., Edwards, A., ... Voss, A. (2014). Tweeting the terror: modelling the social media reaction to the Woolwich terrorist attack. *Social Network Analysis and Mining*, 4(1), 1–14. <https://doi.org/10.1007/s13278-014-0206-4>

- Byman, D. (2016). *Europe vs. America: Comparing the terrorism threat*. Brookings Institute. Retrieved from <https://www.brookings.edu/blog/order-from-chaos/2016/04/05/europe-vs-america-comparing-the-terrorism-threat/>
- C-REX - Center for Research on Extremism. (2020). *Codebook: The right-wing terrorism and violence (RTV) dataset, 1990 - 2019*. Retrieved from <https://www.sv.uio.no/c-rex/english/groups/rtv-dataset/rtv-codebook-revised-5.5.2020.pdf>
- Ceron, A., Curini, L., & Iacus, S. M. (2019). ISIS at Its apogee: The Arabic discourse on Twitter and what we can learn from that about ISIS support and foreign fighters. *SAGE Open*, 9(1), 215824401878922. <https://doi.org/10.1177/2158244018789229>
- Chalabi, M. (2018). Terror attacks by Muslims receive 357% more press attention, study finds | US news | The Guardian. *The Guardian*. Retrieved from <https://www.theguardian.com/us-news/2018/jul/20/muslim-terror-attacks-press-coverage-study>
- Chan, J., Ghose, A., & Seamans, R. (2016). The internet and racial hate crime: Offline spillovers from online access. *MIS Quarterly: Management Information Systems*, 40(2), 381–403. <https://doi.org/10.25300/MISQ/2016/40.2.05>
- Cohen-Almagor, R. (2018). Taking North American white supremacist groups seriously: The scope and challenge of hate speech on the Internet. *International Journal for Crime, Justice and Social Democracy*, 7(2), 38–57. <https://doi.org/10.5204/ijcjsd.v7i2.517>
- Commision for Countering Extremism. (2019). Challenging Hateful Extremism, (October).
- Conway, M. (2018). *Violent Extremism and Terrorism Online in 2018*. Retrieved from <https://www.voxpol.eu/year-in-review-2018/>
- Conway, M., Khawaja, M., Lakhani, S., Reffin, J., Robertson, A., & Weir, D. (2019). Disrupting Daesh: Measuring takedown of online terrorist material and its impacts. *Studies in Conflict & Terrorism*, 42(1–2), 141–160. Retrieved from <https://www.tandfonline.com/doi/full/10.1080/1057610X.2018.1513984>
- Conway, M., Scrivens, R., & Macnair, L. (2019). Right-wing extremists' persistent online presence: History and contemporary trends. *ICCT Policy Brief*. <https://doi.org/10.19165/2019.3.12>
- Coser, L. A. (1956). *The functions of social conflict*. Free Press. <https://doi.org/10.4324/9780203714577>
- Craanen, A., Berntsson, J., Ging, D., & DiBranco, A. (2020). *Tech Against Terrorism: Incels, online misogyny and gender-based terrorism*. Retrieved from <https://www.techagainstterrorism.fm/incels-online-misogyny-and-gender-based-terrorism/>
- Davidson, T., Warmesley, D., Macy, M., & Weber, I. (2017). Automated hate speech detection and the problem of offensive language. *Proceedings of the Eleventh International AAAI Conference on Web and Social Media (ICWSM 2017)*, 512–515. Retrieved from <http://arxiv.org/abs/1703.04009>
- de Gibert, O., Perez, N., García-Pablos, A., & Cuadros, M. (2018). Hate speech dataset from a white supremacy forum. *Proceedings Of the Second Workshop on Abusive Language Online (ALW2)*, 11–20. <https://doi.org/10.18653/v1/w18-5102>
- Dechesne, M. (2009). Explorations in the experimental social psychology of terrorism: The struggle-violence link and its predictors. *Revue Internationale de Psychologie Sociale*, 22(3–4), 87–102.
- Dugan, L., LaFree, G., Cragin, K., & Kasupski, A. (2008). *Building and analyzing a comprehensive open source database on global terrorist events*. Retrieved from <https://www.ncjrs.gov/App/AbstractDB/AbstractDBDetails.aspx?id=245203>
- Ebner, J. (2017). *The rage: The vicious circle of Islamist and far-right extremism*. London: I.B.Tauris.
- Evans, R. (2018). *How the MAGAbomber and the synagogue shooter were likely radicalized*. *Bellingcat*. Retrieved from <https://www.bellingcat.com/news/americas/2018/10/31/magabomber-synagogue-shooter-likely-radicalized/comment-page-2/>
- Evans, R. (2019). *Shitposting, inspirational terrorism, and the Christchurch mosque massacre*. *Bellingcat*. Retrieved from <https://www.bellingcat.com/news/rest-of-world/2019/03/15/shitposting-inspirational-terrorism-and-the-christchurch-mosque-massacre/>
- Faris, R., Ashar, A., & Gasser, U. (2016). Understanding harmful speech online. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.2882824>
- FBI Crime Data Explorer. (2020). Downloads & Documentation. Retrieved from <https://crime-data-explorer.fr.cloud.gov/downloads-and-docs>
- Ferrara, E. (2017). Contagion dynamics of extremist propaganda in social networks. *Information Sciences*, 418–419, 1–12. <https://doi.org/10.1016/j.ins.2017.07.030>

- Fielitz, M., Ebner, J., Guhl, J., & Quent, M. (2018). *Loving hate. Anti-Muslim extremism, radical Islamism and the spiral of polarization*. Retrieved from <https://www.isdglobal.org/isd-publications/hassliebe-muslimfeindlichkeit-islamismus-und-die-spirale-gesellschaftlicher-polarisierung-deutsch/>
- Fuchs, B. C. (2017). Advocates warn of possible underreporting in FBI hate crime data. *NBC News*. Retrieved from <https://www.nbcnews.com/news/asian-america/advocates-warn-possible-underreporting-fbi-hate-crime-data-n830711>
- Gallacher, J. D., Heerdink, M. W., & Hewstone, M. (2020). Online contact between opposing political protest groups via social media is linked to physical violence of offline encounters. *Social Media + Society*, 1–44.
- Gallacher, J. D. (2021). Leveraging cross-platform data to improve automated hate speech detection. *ArXiv*, 1–34. Retrieved from <https://arxiv.org/abs/2102.04895v1>
- Gallacher, John D., & Bright, J. (2021). Hate Contagion: Measuring the spread and trajectory of hate on social media. *PsyArXiv*, 1–47. Retrieved from <https://psyarxiv.com/b9qhd/>
- Gaudette, T., Scrivens, R., & Venkatesh, V. (2020). The role of the Internet in facilitating violent extremism: Insights from former right-wing extremists. *Terrorism and Political Violence*, 1–18. <https://doi.org/10.1080/09546553.2020.1784147>
- Gilbert, D. (2019). Here's how big far right social network Gab has actually gotten. Retrieved from https://www.vice.com/en_us/article/pa7dwg/heres-how-big-far-right-social-network-gab-has-actually-gotten
- Giles, D. (2011). *Testing for Granger Causality. Econometrics Beat*. Retrieved from <https://davegiles.blogspot.com/2011/04/testing-for-granger-causality.html>
- Gill, P., Horgan, J., & Deckert, P. (2014). Bombing alone: Tracing the motivations and antecedent behaviors of lone-actor terrorists. *Journal of Forensic Sciences*, 59(2), 425–435. <https://doi.org/10.1111/1556-4029.12312>
- Ginsberg, J., Mohebbi, M. H., Patel, R. S., Brammer, L., Smolinski, M. S., & Brilliant, L. (2009). Detecting influenza epidemics using search engine query data. *Nature*, 457(7232), 1012–1014. <https://doi.org/10.1038/nature07634>
- Global Terrorism Index. (2019). *Global terrorism index 2019: Measuring the impact of terrorism*. Retrieved from <https://www.visionofhumanity.org/wp-content/uploads/2020/11/GTI-2019-web.pdf>
- Google. (2020). Google Trends. Retrieved November 10, 2020, from <https://trends.google.com/trends/?geo=US>
- Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3), 424–438. <https://doi.org/10.1017/ccol052179207x.002>
- Grierson, J. (2020). Misogyny “should become a hate crime in England and Wales.” *The Guardian*. Retrieved from <https://www.theguardian.com/law/2020/sep/23/misogyny-hate-crime-england-wales-law-commission>
- Grover, T., & Mark, G. (2019). Detecting potential warning behaviors of ideological radicalization in an alt-right subreddit. *Proceedings of the 13th International Conference on Web and Social Media, ICWSM 2019*, 193–204.
- Guhl, J., & Ebner, J. (2018). *Islamist and far-right extremists: Rhetorical and strategic allies in the digital age*. Retrieved from <https://www.radicalisationresearch.org/debate/ebner-islamist-far-right-extremists-rhetorical-digital-age/>
- Guhl, J., Ebner, J., & Rau, J. (2020). *The online ecosystem of the German far-right*. Retrieved from <https://www.isdglobal.org/wp-content/uploads/2020/02/ISD-The-Online-Ecosystem-of-the-German-Far-Right-English-Draft-11.pdf>
- Hamm, M. S., & Spaaij, R. (2018). The age of lone wolf terrorism. *The British Journal of Criminology*, 58(6), 1521–1523. <https://doi.org/10.1093/bjc/azy022>
- Heath, A. (2017). Facebook removed the event page for white nationalist “Unite the Right” rally in Charlottesville one day before it took place. *Business Insider*. Retrieved from <https://www.businessinsider.com/facebook-removed-unite-the-right-charlottesville-rally-event-page-one-day-before-2017-8?op=1&r=US&IR=T>
- Helmus, T. C., & Klein, K. (2018). *Assessing outcomes of online campaigns countering violent extremism: A case study of the redirect method*. Retrieved from https://www.rand.org/pubs/research_reports/RR2813.html
- Hoffman, B., Ware, J., & Shapiro, E. (2020). Assessing the threat of incel violence. *Studies in Conflict and Terrorism*, 43(7), 565–587. <https://doi.org/10.1080/1057610X.2020.1751459>
- Hothorn, T., Zeileis, A., Farebrother, R. W., Cummins, C., Millo, G., & Mitchell, D. (2020). Package “lmtree”: Testing Linear Regression Models. *CRAN*. Retrieved from <https://cran.r->

- project.org/web/packages/lmtest/index.html
- Institute for Global Change. (2016). *In their own words: Why ISIS hates the West*. Retrieved from <https://institute.global/policy/their-own-words-why-isis-hates-west>
- Ischinger, W. (2020). *Munich Security Report 2020*. Retrieved from <https://securityconference.org/en/>
- Javed, J., & Miller, B. (2019). When content promotes hate: Moral-emotional content, outgroup cues, and attitudes toward violence and anti-muslim policies. *ArXiv*, 1–29. Retrieved from http://www.jeffreyjaved.com/uploads/8/6/1/2/86128854/jm_when_content_promotes_hate_041919.pdf
- Jawaid, A. (2013). Pakistan’s polio workers targeted for killing. *Aljazeera*. Retrieved from <https://www.aljazeera.com/features/2013/12/17/pakistans-polio-workers-targeted-for-killing/>
- Johnson, N. F., Zheng, M., Vorobyeva, Y., Gabriel, A., Qi, H., Velasquez, N., ... Wuchty, S. (2016). New online ecology of adversarial aggregates: ISIS and beyond. *Science*, 352(6292), 1459–1463. <https://doi.org/10.1126/science.aaf0675>
- Jun, S. P., Yoo, H. S., & Choi, S. (2018). Ten years of research change using Google Trends: From the perspective of big data utilizations and applications. *Technological Forecasting and Social Change*, 130, 69–87. <https://doi.org/10.1016/j.techfore.2017.11.009>
- Kaakinen, M., Oksanen, A., & Räsänen, P. (2017). Did the risk of exposure to online hate increase after the November 2015 Paris attacks? A group relations approach. *Computers in Human Behavior*, 78, 90–97.
- Kelion, L. (2020). Mast fire probe amid 5G coronavirus. *BBC News*. Retrieved from <https://www.bbc.co.uk/news/uk-england-52164358>
- Kleinberg, B., Vegt, I. Van Der, & Gill, P. (2020). The temporal evolution of a far-right forum. *Journal of Computational Social Science*. <https://doi.org/10.1007/s42001-020-00064-x>
- Knott, K., Lee, B., & Copeland, S. (2018). *Briefings: Reciprocal radicalisation*. Retrieved from <https://crestresearch.ac.uk/resources/reciprocal-radicalisation/>
- Konaev, M., & Moghaddam, F. M. (2010). Mutual Radicalization: Bush, Ahmadinejad, and the “Universal” cycle of out-group threat/in-group cohesion. In *Words of Conflict, Words of War*. Praeger.
- Kteily, N., Hodson, G., & Bruneau, E. (2016). They see us as less than human: Metadehumanization predicts intergroup conflict via reciprocal dehumanization. *Journal of Personality and Social Psychology*, 110(3), 343–370. <https://doi.org/10.1037/pspa0000044>
- Kuhn, M., Wickham, H., & RStudio. (2020). Package ‘tidymodels’: Easily Install and Load the “Tidymodels” Packages. *CRAN*, 1–5. Retrieved from <https://cran.r-project.org/web/packages/tidymodels/index.html>
- Lazer, D., Kennedy, R., King, G., & Alessandro, V. (2014). The parable of Google flu: Traps in big data analysis. *Science*, 343, 1203–1205. Retrieved from www.sciencemag.org/SCIENCEVOL34314MARCH2014
- Lee, B. (2019). *Overview of the far-right*. Centre for Research and Evidence on Security Threats (CREST). Retrieved from <https://www.gov.uk/government/publications/overview-of-the-far-right>
- Li, C., Chen, S., Xing, J., Sun, A., & Ma, Z. (2019). Seed-guided topic model for document filtering and classification. *ACM Transactions on Information Systems*, 37(1). <https://doi.org/10.1145/3238250>
- Lima, L., Reis, J. C. S., Melo, P., Murai, F., Araujo, L., Vikatos, P., & Benevenuto, F. (2018). Inside the right-leaning echo chambers: Characterizing gab, an unmoderated social system. *Proceedings of the 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2018*, 515–522. <https://doi.org/10.1109/ASONAM.2018.8508809>
- Massicotte, P., & Eddelbuettel, D. (2020). Package “gtrendsR”: Perform and Display Google Trends Queries. *CRAN*. Retrieved from <https://cran.r-project.org/web/packages/gtrendsR/gtrendsR.pdf>
- Mathew, B., Dutt, R., Goyal, P., & Mukherjee, A. (2018). Spread of hate speech in online social media. Retrieved from <http://arxiv.org/abs/1812.01693>
- Mathew, B., Dutt, R., Goyal, P., & Mukherjee, A. (2019). Spread of hate speech in online social media. *Proceedings of the 10th ACM Conference on Web Science - WebSci ’19*, 173–182. <https://doi.org/10.1145/3292522.3326034>
- May, R. P. (2020). Early warning? Opportunities and limits of automated internet monitoring. *Global Network on Extremism and Technology*. Retrieved from <https://gnet-research.org/2020/05/14/early-warning-opportunities-and-limits-of-automated-internet-monitoring/>
- McCauley, C., & Moskaleiko, S. (2008). Mechanisms of political radicalization: Pathways toward terrorism. *Terrorism and Political Violence*, 20(3), 415–433. <https://doi.org/10.1080/09546550802073367>
- Meloy, R. J., Hoffman, J., Guldemann, A., & James, D. (2011). The role of warning behaviors in threat assessment: An exploration and suggested typology. *Behavioral Sciences & the Law*, 28(2), 211–223.

- <https://doi.org/10.1002/bsl>
- Miller, E., LaFree, G., & Dugan, L. (2018). Global Terrorism Database (GTD). *National Consortium for the Study of Terrorism and Responses to Terrorism*. <https://doi.org/10.5260/chara.19.3.14>
- Mitts, T. (2019). From isolation to radicalization: Anti-muslim hostility and support for ISIS in the west. *American Political Science Review*, 113(1), 173–194. <https://doi.org/10.1017/S0003055418000618>
- Moghaddam, F. M. (2005). The staircase to terrorism a psychological exploration. *American Psychologist*, 60(2), 161–169. <https://doi.org/10.1037/0003-066X.60.2.161>
- Moghaddam, F. M. (2018). *Mutual radicalization: How groups and nations drive each other to extremes*. Washington: American Psychological Association. <https://doi.org/10.1037/0000089-000>
- Moghaddam, F. M., Harré, R., & Lee, N. (2008). *Global conflict resolution through positioning analysis*. Springer New York. https://doi.org/10.1007/978-0-387-72112-5_1
- Moonshot CVE. (2018). *Searching for hate in America*. Retrieved from <http://moonshotcve.com/searching-for-hate-in-america/>
- Müller, K., & Schwarz, C. (2017). Fanning the Flames of Hate: Social Media and Hate Crime. *Ssrn*. <https://doi.org/10.2139/ssrn.3082972>
- Müller, K., & Schwarz, C. (2018a). Fanning the flames of hate: Social media and hate crime. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3082972>
- Müller, K., & Schwarz, C. (2018b). Making America hate again? Twitter and hate crime under Trump. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3149103>
- Nagle, A. (2017). *Kill All Normies: Online Culture Wars from 4chan and Tumblr to Trump and The Alt-Right*. Zero Books.
- Obaidi, M., Thompson, L., & Bergh, R. (2019). “They think we are a threat to their culture”: Meta-cultural threat fuels willingness and endorsement of extremist violence against the cultural outgroup. *International Journal of Conflict and Violence (IJCIV)*, 12. <https://doi.org/10.4119/UNIBI/ijcv.647>
- Olteanu, A., Castillo, C., Boy, J., & Varshney, K. R. (2018). The effect of extremist violence on hateful speech online. *ArXiv*. Retrieved from <http://arxiv.org/abs/1804.05704>
- Papasavva, A., Blackburn, J., Stringhini, G., Zannettou, S., & de Cristofaro, E. (2020). Is it a coincidence?: A first step towards understanding and characterizing the qanon movement on Voat.co. *ArXiv*. Retrieved from <https://arxiv.org/pdf/2009.04885.pdf>
- Papasavva, A., Zannettou, S., De Cristofaro, E., Stringhini, G., & Blackburn, J. (2020). Raiders of the Lost Kek: 3.5 Years of Augmented 4chan Posts from the Politically Incorrect Board. *ICWSM*. Retrieved from <http://arxiv.org/abs/2001.07487>
- Pettigrew, T. F. (2009). Secondary transfer effect of contact: Do intergroup contact effects spread to noncontacted outgroups? *Social Psychology*, 40(2), 55–65. <https://doi.org/10.1027/1864-9335.40.2.55>
- Pratt, D. (2015). Islamophobia as reactive co-radicalization. *Islam and Christian–Muslim Relations*, 26(2), 205–218. <https://doi.org/10.1080/09596410.2014.1000025>
- Reicher, S. D., & Haslam, S. A. (2016). Fueling Extremes. *Scientific American Mind*, 27(3), 34–39. <https://doi.org/10.1038/scientificamericanmind0516-34>
- Reicher, S., Haslam, S. A., & Rath, R. (2008). Making a virtue of evil: A five-step social identity model of the development of collective hate. *Social and Personality Psychology Compass*, 2(3), 1313–1344. <https://doi.org/10.1111/j.1751-9004.2008.00113.x>
- Ritchie, H., Hasell, J., Appel, C., & Roser, M. (2013). Terrorism. Retrieved from <https://institute.global/policy/their-own-words-why-isis-hates-west>
- Sageman, M. (2008). *Leaderless Jihad: Terror networks in the twenty-first century*. University of Pennsylvania Press.
- Saleem, H. M., Dillon, K. P., Benesch, S., & Ruths, D. (2017). A web of hate: Tackling hateful speech in online social spaces. *ArXiv*. Retrieved from <http://arxiv.org/abs/1709.10159>
- Samaratunge, S., & Hattotuwa, S. (2014). *Liking Violence: A study of hate speech on Facebook in Sri Lanka*. Retrieved from <https://www.cpalanka.org/wp-content/uploads/2014/09/Hate-Speech-Final.pdf>
- Schuurman, B., Bakker, E., Gill, P., & Bouhana, N. (2018). Lone Actor Terrorist Attack Planning and Preparation: A Data-Driven Analysis. *Journal of Forensic Sciences*, 63(4), 1191–1200. <https://doi.org/10.1111/1556-4029.13676>
- Schuurman, B., Lindekilde, L., Malthaner, S., O'Connor, F., Gill, P., & Bouhana, N. (2019). End of the lone wolf: The typology that should not have been. *Studies in Conflict and Terrorism*, 42(8), 771–778.

- <https://doi.org/10.1080/1057610X.2017.1419554>
- Schwencke, K. (2017). Why America fails at gathering hate crime statistics. *ProPublica*. Retrieved from <https://www.propublica.org/article/why-america-fails-at-gathering-hate-crime-statistics>
- Scott, J. (2020). Misogyny: Women “should be protected” under hate crime laws. *BBC News*. Retrieved from https://www.bbc.co.uk/news/uk-politics-54254541?fbclid=IwAR0vLQhcheQkKVcEHGlj2NzEo2h9Pd7IYE_IZ65AIHqTPh-athCv6n7rWRM
- Scrivens, R. (2020). Exploring radical right-wing posting behaviors online. *Deviant Behavior*, 1–15. <https://doi.org/10.1080/01639625.2020.1756391>
- Scrivens, R., Burruss, G. W., Holt, T. J., Chermak, S. M., Freilich, J. D., & Frank, R. (2020). Triggered by defeat or victory? Assessing the impact of presidential election results on extreme right-wing mobilization online. *Deviant Behavior*, (August). <https://doi.org/10.1080/01639625.2020.1807298>
- Siegel, A. A. (2020). Online hate speech. In *Social media and democracy: The state of the field and prospects for reform* (pp. 56–88). Cambridge University Press.
- Siegel, A., Tucker, J., Nagler, J., & Bonneau, R. (2017). Socially mediated sectarianism. *Unpublished Manuscript*. Retrieved from https://alexandra-siegel.com/wp-content/uploads/2019/05/Siegel_Sectarianism_January2017.pdf
- Smith, L. G. E., Blackwood, L., & Thomas, E. F. (2019). The need to refocus on the group as the site of radicalization. *Perspectives on Psychological Science*, 15(2), 327–352. <https://doi.org/10.1177/1745691619885870>
- Stecklow, S. (2018). Why Facebook is losing the war on hate speech in Myanmar. *Reuters*. Retrieved from <https://www.reuters.com/investigates/special-report/myanmar-facebook-hate/>
- Stein, A. A. (1976). Conflict and cohesion: A review of the literature. *Journal of Conflict Resolution*, 20(1), 143–165. <https://doi.org/10.18574/nyu/9780814786390.003.0007>
- Stephan, W. G., & Stephan, C. W. (2000). An integrated threat theory of prejudice. In *Reducing prejudice and discrimination* (pp. 23–45). Lawrence Erlbaum Associates Publishers.
- Stephan, W. G., Ybarra, O., & Morrison, K. R. (2009). *Handbook of prejudice, stereotyping, and discrimination*. Psychology Press. Taylor and Francis Group. Retrieved from <https://archive.org/details/handbookprejudic00nels>
- Stephens-Davidowitz, S. (2017). *Everybody lies: Big data, new data, and what the internet can tell us about who we really are*. Dey Street Books.
- Tajfel, H. (1974). Social identity and intergroup behaviour. *Social Science Information*, 13(2), 65–93. <https://doi.org/10.1177/053901847401300204>
- Tajfel, H., & Turner, J. (1979). An integrative theory of intergroup conflict. In M. A. Hogg & D. Abrams (Eds.), *Key readings in social psychology. Intergroup relations: Essential readings* (pp. 94–109). Psychology Press.
- Tech Against Terrorism. (2019). *Analysis: The use of open-source software by terrorists and violent extremists*. Retrieved from <https://www.techagainstterrorism.org/2019/09/02/analysis-the-use-of-open-source-software-by-terrorists-and-violent-extremists/>
- The National Consortium for the Study of Terrorism and Responses to Terrorism (START). (2020). Global Terrorism Database. Retrieved November 10, 2020, from <https://project-iris.app-staging.cloud/>
- Thomas, E. F., McGarty, C., & Louis, W. (2014). Social interaction and psychological pathways to political engagement and extremism. *European Journal of Social Psychology*, 44(1), 15–22. <https://doi.org/10.1002/ejsp.1988>
- Thornton, D. L., & Batten, D. S. (1985). Lag-length selection and tests of Granger Causality between money and income. *Journal of Money, Credit and Banking*, 17(2), 164. <https://doi.org/10.2307/1992331>
- Tien, J. H., Eisenberg, M. C., Cherng, S. T., & Porter, M. A. (2020). Online reactions to the 2017 ‘Unite the right’ rally in Charlottesville: Measuring polarization in Twitter networks using media followership. *Applied Network Science*, 5(1). <https://doi.org/10.1007/s41109-019-0223-3>
- Torok, R. (2013). Developing an explanatory model for the process of online radicalisation and terrorism. *Security Informatics*, 2(1), 1–10. <https://doi.org/10.1186/2190-8532-2-6>
- Tseng, Q. (2019). Reconstruct Google trends daily data for extended period. *Towards Data Science*. Retrieved from <https://towardsdatascience.com/reconstruct-google-trends-daily-data-for-extended-period-75b6ca1d3420>
- Turner, J. C., Hogg, M. A., Oakes, P. J., Reicher, S. D., & Wetherell, M. S. (1987). *Rediscovering the social group: A self-categorization theory*. Basil Blackwell.

- Turse, N. (2019). How this pastor of a megachurch is fueling Ebola conspiracy theories. *Time Magazine*. Retrieved from <https://time.com/5703662/ebola-conspiracy-theories-congo/>
- Valente, T. W., Dyal, S. R., Chu, K. H., Wipfli, H., & Fujimoto, K. (2015). Diffusion of innovations theory applied to global tobacco control treaty ratification. *Social Science and Medicine*, 145, 89–97. <https://doi.org/10.1016/j.socscimed.2015.10.001>
- Vidgen, B., & Derczynski, L. (2020). Directions in abusive language training data: Garbage in, garbage out. *ArXiv*, 1–26. Retrieved from <http://arxiv.org/abs/2004.01670>
- Vidgen, B., Harris, A., Cowls, J., & Guest, E. (2020). *An agenda for research into online hate*. Retrieved from https://www.turing.ac.uk/sites/default/files/2020-10/an_agenda_for_research_into_online_hate.pdf
- Vidgen, B., Tromble, R., Harris, A., Hale, S., Nguyen, D., & Margetts, H. (2019). Challenges and frontiers in abusive content detection. *3rd Workshop on Abusive Language Online*, 1–14. Retrieved from <https://www.aclweb.org/anthology/W19-3509/>
- Vidgen, B., Yasseri, T., & Margetts, H. (2019). Trajectories of Islamophobic hate amongst far right actors on Twitter. *ArXiv*, 1–20. Retrieved from <https://arxiv.org/pdf/1910.05794>
- Wendling, M. (2016). Do terror attacks in the Western world get more attention than others? *BBC Trending*. Retrieved from <https://www.bbc.co.uk/news/blogs-trending-35886051>
- Wickham, H., Averick, M., Bryan, J., Chang, W., D’L., McGowan, A., ... Yutani, H. (2019). Welcome to the Tidyverse. *Journal of Open Source Software*, 4(43), 1686. <https://doi.org/10.21105/joss.01686>
- Williams, M. L., & Burnap, P. (2016). Cyberhate on social media in the aftermath of Woolwich: A case study in computational criminology and big data. *British Journal of Criminology*, 56(2), 211–238. <https://doi.org/10.1093/bjc/azv059>
- Williams, M. L., Burnap, P., Javed, A., Liu, H., & Ozalp, S. (2019). Hate in the machine: Anti-black and anti-muslim social media posts as predictors of offline racially and religiously aggravated crime. *The British Journal of Criminology*, 60(1), 1–25. <https://doi.org/10.1093/bjc/azz049>
- Woo, G. (2002). Quantitative terrorism risk assessment. *Journal of Risk Finance*, 4(1), 7–14. <https://doi.org/10.1108/eb022949>
- Zannettou, S., Bradlyn, B., De Cristofaro, E., Kwak, H., Sirivianos, M., Stringhini, G., & Blackburn, J. (2018). What is Gab? A bastion of free speech or an alt-right echo chamber? *ArXiv*. <https://doi.org/10.1145/3184558.3191531>
- Zannettou, S., Caulfield, T., De Cristofaro, E., Kourtellis, N., Leontiadis, I., Sirivianos, M., ... Blackburn, J. (2017). The web centipede: Understanding how web communities influence each other through the lens of mainstream and alternative news sources. *Proceedings of the 2017 Internet Measurement Conference*. Retrieved from <http://arxiv.org/abs/1705.06947>
- Zannettou, S., ElSherief, M., Belding, E., Nilizadeh, S., & Stringhini, G. (2020). Measuring and characterizing hate speech on news websites. *ArXiv*. Retrieved from <http://arxiv.org/abs/2005.07926>
- Zhou, Y., Dredze, M., Broniatowski, D. A., & Adler, W. D. (2019). Elites and foreign actors among the alt-right: The Gab social media platform. *First Monday*. <https://doi.org/10.5210/fm.v24i9.10062>

Supplementary Information (SI) for:

Mutual radicalisation of opposing extremist groups via the Internet

- 1 – Offline violence data collection
- 2 - Google search volumes for ‘How to join ISIS’
- 3 - Overall density of hate on Gab
- 4 - Wider effects of violence on ISIS search volumes
- 5 - USA estimates of Islamic extremist search interest and offline Islamophobia

1 - Offline violence data collection

In order to estimate the level of offline hate crime over the study period, we collected data from a number of sources and combined these datasets together. SI Table 1 shows the data sources used, along with the number of events contributed from each dataset. In all cases the events were de-duplicated across sources to ensure that no event was double counted.

Far-right hate crime and violence

For the USA, we collected data on far-right terror attacks from the Global Terrorism Database (GTD, Miller, LaFree, & Dugan, 2018) and from Wikipedia. The Global Terrorism Database (GTD) is an open-source database including information on Global terrorist events since 1970 and is updated through to the end of 2018, therefore covering the period of interest for the current study (August 2016 – October 2018). For Wikipedia, we used an automated collection technique to collect all events listed on the global ‘List of terrorist incidents’ page. In both cases we filtered the datasets by location and our time period of interest.

USA hate crime data was collected primarily from the Federal Bureau of Investigation (FBI) dataset. This contains all hate crime reported to the FBI in the USA and gives detailed information on the type of attack, motivation or the attacker, and details of the victim(s). As this dataset is known to under-report certain hate crimes (Fuchs, 2017; Schwencke, 2017), we additionally collected community-reported data on hate crimes from the Anti-defamation league (ADL) and the Organization for Security and Co-operation in Europe (OSCE), which collates numerous reports from civil society organisations including the United Nations and the Council of Europe amongst others. All data is available at the daily level.

For the United Kingdom and Germany, we collected data from a similar range of sources, with the GTD and Wikipedia providing information on larger scale terror attack and more notable hate crimes. We also included data from the Right-wing terrorism and violence in Western Europe dataset (RTV) collated by the Centre for Research on Extremism (CREX) (Ravndal, 2016). In the case of the RTV data, this also included plots which were foiled prior to the attack being carried out, which adds to the richness of the data. As with the USA, we supplement this data with community-reported data on hate crimes from the OSCE.

Far-right motivation classification

Ensuring that all the events included in this analysis had a far-right motivation is challenging. In the case of the larger and more notable events, we included an attack in the analysis if the prosecuting authority, police, or media indicated a far-right, nationalist or white supremacist motivation. For the community-reported data, we included an attack if the victim type matched a typical far-right

targeted outgroup. It is therefore possible that this latter data includes attacks without a far-right motivation. To mitigate for this in the UK, we made sure to remove any entries which referred to Irish separatist / Irish sectarian violence, while in Germany we also removed any left-wing attacks and attacks attributed to the Kurdistan Workers' Party.

In order to classify the events into specific hate types (e.g. Islamophobia, anti-Semitism etc) we used a combination of official designation and inference. In the case of the FBI dataset for example, the hate motivation was provided through both the offender bias and the victim designation and so these classifications were used. The OCSE and ADL data was also broken down by hate motivation. In the case of the GTDB, C-Rex and Wikipedia data, the hate type was extracted from the description of the event, and details on offender background, motivation and history, along with details on the victims. This process of classification of events by hate type was performed manually.

Hate crimes could be classified as belonging to multiple categories and these are not mutually exclusive (e.g. an attack on women at a Mosque would be classified as Islamophobia and Misogyny). Matching with the Centre for Research on Extremism codebook, victims were also coded reflecting perpetrator intention (e.g. if a Sikh individual was targeted because the perpetrator believed the victim was a Muslim, target group will be coded as Muslim) (C-REX - Center for Research on Extremism, 2020).

We manually checked all datasets to ensure that the descriptions of hate type were consistent across all.

Islamic extremist violence and terror attacks

We collected Islamic extremism data from a combination of the GTD and Wikipedia. Islamic extremist terrorism consists of either direct attacks claimed by Al-Qaeda and Islamic State (ISIS), or events where the attacker either claimed to have been inspired by these groups or where authorities have formally made this connection. We restricted these events to those perpetrated in the ‘West’ (Europe, North America or Australasia) for reasons discussed in the main text.

SI Table 1 – Sources for offline violence collection and the number of events each source contributed to the collection after de-duplication

<i>Country</i>	<i>Source</i>	<i>Number of Events</i>	<i>Link to source</i>
<i>United States of America</i>	Global Terrorism Database	188	https://www.start.umd.edu/gtd/
	Wikipedia	2	https://en.wikipedia.org/wiki/List_of_terrorist_incidents
	FBI Hate Crime	18,534	https://crime-data-explorer.fr.cloud.gov/downloads-and-docs
	ADL	1,190	https://www.adl.org/education-and-resources/resource-knowledge-base/adl-heat-map
	OCSE	89	https://hatecrime.osce.org/united-states-america
<i>United Kingdom</i>	Global Terrorism Database	54	https://www.start.umd.edu/gtd/
	Wikipedia	1	https://en.wikipedia.org/wiki/List_of_terrorist_incidents
	OCSE	401	https://hatecrime.osce.org/united-kingdom
	C-Rex RTV	14	https://www.sv.uio.no/c-rex/english/topics/online-resources/rtv-dataset/index.html
<i>Germany</i>	Global Terrorism Database	43	https://www.start.umd.edu/gtd/
	Wikipedia	7	https://en.wikipedia.org/wiki/List_of_terrorist_incidents
	OCSE	470	https://hatecrime.osce.org/germany
	C-Rex RTV	41	https://www.sv.uio.no/c-rex/english/topics/online-resources/rtv-dataset/index.html
<i>Islamic extremist Terrorism</i>	Global Terrorism Database	92	https://www.start.umd.edu/gtd/
	Wikipedia	25	https://en.wikipedia.org/wiki/List_of_terrorist_incidents

2 - Google search volumes for ‘How to join ISIS’

SI Figure 1 shows the Google search trends for the phrase ‘How to join ISIS’ for Germany, the United Kingdom and the United States of America (USA). In each case the data was collected for a series of 9-month overlapping periods and then normalised across the entire period (see methods section), which gives rise to the differential upper bounds for each country.

SI Figure 1 – Google search trends for the phrase ‘How to join ISIS’ for Germany, the United Kingdom and the United States of America (USA)

3 - Overall density of hate speech on Gab

SI Figure 2 gives the overall density of hate speech within the Gab conversations over the period August 2016 – October 2018. The density rose quickly over the first 4 months before remaining fairly level but with significant variation from day-to-day.

SI Figure 2 – Daily hate density for Gab over the period August 2016 – October 2018

4 - Wider effects of violence on ISIS search volumes

SI Table 2 shows the Granger causality effects for offline hate and offline Islamic extremist Internet interest for the USA, UK, and Germany across specific hate types not outlined in the main results. Darker yellow indicates a larger F statistic (stronger signal) while darker blue indicates a lower p-value.

We observed no effects for any hate types other than Islamophobia in the United States of America. In the UK we observed effects for anti-immigration sentiment and anti-Semitism, but these were less significant than for Islamophobia.

In Germany we observed an association between offline anti-Black hate crimes and subsequent Islamic extremism Google interest, but not for any other offline hate types. We do not read too far into this result and consider that it might be a spurious false positive. We do however observe interesting results whereby the Google search volumes for Islamic extremist content in Germany consistently preceded hate crimes across anti-immigration, misogynistic and anti-Black racism motivations.

SI Table 2 – Granger causality effects for offline hate and offline Islamic extremist google search volumes for the USA, UK, and Germany across specific hate types. Darker yellow indicates a larger estimate of F (stronger signal) while darker blue indicates a lower p-value.

Offline Hate → ISIS Google Searches				ISIS Google Searches → Offline Hate			
HATE TYPE	DF	ESTIMATE	P	HATE TYPE	DF	ESTIMATE	P
USA - Far-right violence							
Anti-immigration	746	3.236	0.145	Misogyny	736	1.681	0.245
Anti-black racism	746	2.809	0.188	Anti-semitism	738	1.404	0.441
Anti-semitism	746	1.262	0.523	Homophobia	736	1.377	0.443
Homophobia	740	0.780	1.000	Anti-black racism	742	1.168	0.642
Misogyny	740	0.475	1.000	Anti-immigration	742	0.896	0.886
UK - Far-right violence							
Anti-immigration	742	5.913	0.002	Anti-immigration	734	0.915	0.886
Anti-semitism	742	5.327	0.004	Anti-semitism	734	0.909	0.499
Homophobia	746	0.069	1.000	Misogyny	746	0.397	0.529
Misogyny	746	0.067	1.000	Homophobia	746	0.373	0.541
Anti-black racism	746	0.038	0.847	Anti-black racism	746	0.021	0.884
Germany - Far-right violence							
Anti-black racism	740	4.149	0.007	Anti-immigration	734	4.597	0.000
Anti-immigration	742	1.824	0.145	Misogyny	734	3.280	0.006
Misogyny	742	1.656	0.525	Anti-black racism	734	2.607	0.035
Homophobia	742	1.381	0.742	Homophobia	734	2.372	0.063
Anti-semitism	743	0.893	0.523	Anti-semitism	737	2.213	0.120

5 - Local USA estimates of Islamic extremist search interest and offline Islamophobia

In order to investigate the local effects of Islamophobic hate crime on Islamic extremist search interest we separately compared the timeseries data from the five states with the highest number of Islamophobic hate crimes in the USA over the study period. For each state we ran two-way Granger Causality tests looking at the predictive power of offline Islamophobic hate crimes on Islamic extremist Internet search interest in that state and vice-versa. We then compared the hate volumes in each of these top five states to the Islamic extremist search volumes in each of the other five states. We limited the analysis to the top five states for two reasons; firstly, because with the highest offline crime volumes these states will lead to the most reliable statistical measures, and secondly because including more states in the analysis increases the likelihood of spurious effects due to the high number of combinations of states that would be need to be calculated. In all cases we corrected for multiple comparisons using Holms corrections.

Here, we follow the method of Bail et al (2018) and excluded Washington DC due to the high likelihood that investigative journalists or intelligence personnel may type skew the Google search volumes for ISIS related content as part of research into ISIS recruitment tactics. This left the following states included in the analysis; California, Massachusetts, Michigan, New jersey and New York.

Results showed a degree of local specificity in the associations between offline hate crime and online search activity (SI Figure 3). Panel A shows how offline Islamophobic hate crime in Michigan predicted Google search volumes in that state, and only in that state, while offline crime in California predicted the Google search volume in California. The activities in these two states were mutually independent. However, the offline activity in California also predicted Google search volumes in New Jersey and New York suggesting that the results are not completely clear cut. We observed no effects in Massachusetts, and in none of the states did Google search volumes predict offline violence.

SI Figure 3 – Local state-wide effects of Islamophobic hate crime in the USA on Islamic extremist Google search volumes for the top 5 states by Islamophobic hate. Panel A shows the Granger Causality effects of offline violence on online searches across, while panel B shows the reverse. Darker Orange indicates a larger estimate of F (stronger signal) while darker blue indicates a smaller value. Significance values are given by * for $p < 0.05$ and ** for $p < 0.01$.