

THE INTERNATIONAL
CENTRE FOR THE STUDY
OF RADICALISATION AND
POLITICAL VIOLENCE



*Developments in Radicalisation
and Political Violence*

Who Matters Online: Measuring influence, evaluating content and countering violent extremism in online social networks

J.M. Berger & Bill Strathearn

March 2013

Developments in Radicalisation and Political Violence

Developments in Radicalisation and Political Violence is a series of papers published by the International Centre for the Study of Radicalisation and Political Violence (ICSR). It features papers by leading experts, providing reviews of existing knowledge and sources and/or novel arguments and insights which are likely to advance our understanding of radicalisation and political violence. The papers are written in plain English. Authors are encouraged to spell out policy implications where appropriate.

Editors

Alexander Meleagrou-Hitchens
Dr John Bew
ICSR, King's College London

Editorial Assistant

Katie Rothman
ICSR, King's College London

Editorial Board

Prof. Sir Lawrence Freedman
King's College London
Dr. Boaz Ganor
Interdisciplinary Center Herzliya
Dr. Peter Neumann
King's College London
Dr. Hasan Al Momani
Jordan Institute of Diplomacy

Contact

All papers can be downloaded free of charge at **www.icsr.info**. To order hardcopies, please write to mail@icsr.info. For all matters related to the paper series, please write to:

ICSR
King's College London
138–142 Strand
London WC2R 1HH
United Kingdom

© ICSR 2013

Contents

Introduction and Key Findings	3
Overview	7
Methodology and terms	9
Collected Data	11
Analysis of Influence	15
Identification of Engaged Extremists	17
Content of Tweets	21
Links	21
Top Sites	23
Hashtags	25
Political Context Considerations	27
Comparison Group	29
Identification Rate	30
Content Analysis	31
CVE Implications and Recommendations	35
Measurement	37
Targeting Influence	38
Targeting Content	41
Intervening with At-Risk Individuals	43
Countering Extremist Narratives	44
Reactive Opportunities to Counter Extremist Narratives	45
Data at a glance	47
Appendices	48
A: Seed accounts	48
B: Scoring formula and glossary of metrics	48
C: Top 10 most influential	50
D: Top 10 most exposed	51
E: Top 20 most interactive	51
F: Glossary of Twitter Terms	52
G: Disclosures	53

About the Authors

J.M. Berger is a researcher and consultant working on problems related to extremism, with a particular focus on the use of the Internet by extremist movements. He has written about extremism for Foreign Policy, The New York Times, The Daily Beast, and the Sentinel, a peer-reviewed journal published by the Combating Terrorism Center at West Point. Berger is the author of “Jihad Joe: Americans Who Go to War in the Name of Islam” (Potomac Books, 2011).

Bill Strathearn is a software engineer with Google and Google.org with an interest in creating technology that improves the interface between government and citizens. He has 13 years of experience in a broad range of software systems including social graph analysis pipelines used to improve search results. Strathearn studied computer science at the University of California at Santa Barbara where he graduated with a Master’s degree.

Acknowledgement

The authors would like to acknowledge and thank *Google Ideas* for its support in commissioning this paper. Views expressed in the paper are those of the authors.

Introduction and Key Findings

It is relatively easy to identify tens of thousands of social media users who have an interest in violent ideologies, but very difficult to figure out which users are worth watching. For students of extremist movements and those working to counter violent extremism online, deciphering the signal amid the noise can prove incredibly daunting.

This paper sets out a first step in solving that problem. We have devised a scoring system to find out which social media accounts within a specific extremist circle were most influential and most prone to be influenced (a tendency we called exposure).

Our starting data centered on followers of 12 American white nationalist/white supremacist “seed” accounts on Twitter. We discovered that by limiting our analysis to interactions with this set, the metrics also identified the users who were highly engaged with extremist ideology.

Within our total dataset of 3,542 users, only 44 percent overtly identified themselves as white nationalists online. By measuring interactions alone—without analyzing user content related to the ideology—we narrowed the starting set down to 100 top-scoring accounts, of which 95 percent overtly self-identified as white nationalist. Among the top 200, 83 percent self-identified, and for the top 400, the self-identification rate was 74 percent. A comparison analysis run on followers of anarchist Twitter accounts suggests the methodology can be used without modification on any number of ideologies.

Because this approach is entirely new (at least in the public sphere), the paper spends some time discussing the methodology and findings in some detail, before concluding with a series of recommendations for countering violent extremism (CVE) based on the findings. The key terms for understanding the recommendations are:

- **Influence:** The tendency of a user to inspire a measurable reaction from other users (such as a replies or retweets).
- **Exposure:** The flip side of influence, this is the tendency of a user to respond to another user in a measurable way.
- **Interactivity:** The sum of influence and exposure scores, roughly representing how often a user interacts with the content of other users.

Our key findings include:

- Influence is highly concentrated among the top 1 percent of users in the set.
- High scores in both influence and exposure showed a strong correlation to engagement with the seed ideology (white nationalism in our primary analysis, and anarchism in a secondary analysis).
- Interactivity, the sum of influence and exposure scores, was even more accurate at identifying users highly engaged with the seed ideology.

In the course of collecting the data needed to measure influence and exposure, we incidentally collected a large amount of data on hashtags and links used by people who follow known white nationalists on Twitter. When we examined this data, we discovered that members of the dataset were highly engaged with partisan Republican and mainstream conservative politics. The paper presents a significant amount of context needed to properly evaluate this finding.

Working from these findings, the paper makes several recommendations for new CVE initiatives with a focus on NGO efforts, which was the purpose of this research, although we recognize our findings will likely have utility for government efforts in this sphere as well.

Our recommendations include:

- We believe these metrics offer ways to concretely measure which types of CVE approaches are effective and which are not, bringing some clarity to a realm where strategies are often wishful and based on assumptions, while conclusions are often anecdotal and inconclusive.
- The concentration of influence among a very few users suggests that disruptive approaches and counter-messaging should be targeted to the top of the food chain, rather than working with the larger base of users.
- Our analysis found that the seed accounts—all well-known white nationalist ideologues and activists—were not necessarily producing the most popular content and links to external Web sites. The collected data can be used to find the most important external content sources, and target them for disruption through terms-of-service violation reporting, or through counter-messaging.
- By tracking these metrics on an ongoing basis, NGO efforts to counterprogram against extremist narratives can be evaluated to measure how many users adopt or respond to counter-messaging content, and how much influence accrues to different kinds of positive messaging.
- Since the data suggests white nationalists are actively seeking dialogue with conservatives, CVE activists should enlist the help of mainstream conservatives, who may be considerably more successful than NGOs at engaging extremists with positive messaging. Further research may also suggest avenues for engagement between other kinds of extremists and other mainstream political and religious movements.

Finally, we believe that these metrics are only a starting point for the study of extremist use of social media. We believe the metrics and approaches here can be further refined, and we believe that additional research may yield substantial new techniques for monitoring and countering the promotion of violent ideologies online.



Overview

As social media grows more important to organizational activity of all types, its role in enabling radicalization and extremism has come under intense scrutiny. This paper examines one type of extremism, white nationalism, on one social network, Twitter,¹ in order to learn more about the nature of social media interactions.

We used a numerical scoring system to analyze the number and type of interactions among 3,542 Twitter users who followed the accounts of well-known white nationalists.

Not all members of the overall dataset were committed white nationalists, but we found that 95 of the 100 users with the highest scores for interactivity were also observed to be highly engaged with white nationalist (WN) ideology in their tweets and how they described themselves in their Twitter account profiles. For the top 200, the metric's accuracy in identifying users overtly engaged with WN ideology was 83 percent, and for the top 400 the accuracy rate was 74 percent.

Importantly, the scoring system did not include analysis of any terms or conditions specific to white nationalism, defined here as any ideology that prioritizes white racial identity as a fundamental component of how society should be organized. Instead, results were derived solely by measuring interactions within the dataset, regardless of content, which suggests the approach can be applied without modification to other extremist ideologies. To test that theory, we replicated the analysis of the initial dataset on a group of Twitter users who followed anarchist accounts, with encouraging but more ambiguous results.

As one component of the scoring system, we analyzed the dataset to identify the most influential users. We found that nearly 50 percent of all influence in the system (as defined by

1 A glossary of Twitter terms is included as Appendix F. For more explanation of how Twitter works see <https://support.twitter.com/articles/215585-twitter-101-how-should-i-get-started-using-twitter>

a weighted measurement of interaction described below) was concentrated among the top 1.5 percent of users.

We believe these results offer significant opportunities to improve initiatives aimed at countering violent extremism (CVE) in the online arena, including improved selection of targets for counter-messaging and disruptive tactics and the ability to quantitatively measure the success of different approaches.

In a secondary analysis, we also examined links and hashtags tweeted by users in the dataset and found a significant portion of the content related to mainstream conservative and partisan Republican politics. This finding requires a significant amount of context to be interpreted fairly, and it is important to reiterate that not all users in the dataset were committed to white nationalism. Most importantly, the timing of the data collection, during the political convention season with a major election approaching, should also be taken into consideration as it likely influenced the choice of both links and hashtags, as discussed in greater detail below.

With a full consideration of the context, we believe this finding suggests that social conservatives are in a very strong position to conduct effective CVE programs aimed at neutralizing the messages and narratives promulgated by influential white nationalists.

Our comparison analysis of followers of anarchist Twitter accounts found that while those users posted a large amount of material pertaining to liberal views, they posted much less content related to mainstream and partisan politics.

Methodology and terms

Methodology

- A short glossary of Twitter-specific terms used herein is attached as Appendix F.
- **Sampled Twitter accounts:** 12 “seed” Twitter accounts for users identified as prominent individuals or organizations within the white nationalist movement
- **Followers dataset:** 3,542 Twitter accounts that follow one or more seed accounts
- **Total dataset:** The 200 most recent tweets for each seed and follower, or all available tweets if the account had less than 200. Total tweets analyzed: 342,807.
- **Sample period:** Automated collection was completed 25 August 2012. Manual inspection of accounts began 25 August 2012 and continued through 31 October 2012.
- **Metrics:** A complete explanation of all metrics may be found in Appendix B.

Key Terms

- **White nationalist:** Any ideology that prioritizes white racial identity as a fundamental component of how a given society or nation should be organized.
- **Influence:** A metric measuring a Twitter user’s ability to publicly create and distribute content that is consumed, affirmed and redistributed by other users.
- **Exposure:** A metric measuring a Twitter user’s tendency to publicly consume, affirm and redistribute content created by other users.
- **Interactivity:** A metric measuring a Twitter user’s combined influence and exposure based on their public activity.



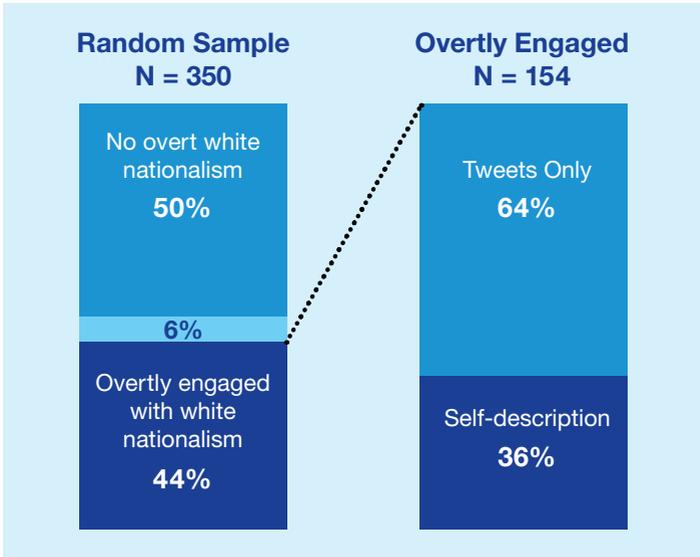
Collected Data

Collection began with 12 Twitter accounts belonging to prominent individuals and organizations that openly and unambiguously promote organized, ideological white nationalism, including the official accounts of David Duke, the Aryan Nations and the White Aryan Resistance. A full list is found in Appendix A.

These “seed” accounts were chosen based on the current and historical significance of the people and organizations to white nationalist ideology, as well as for their unambiguous status as white nationalists. We collected information that the followers of the seed accounts made publicly available between August and October 2012, using a Web application designed by co-author Bill Strathearn.

No one of the 12 seed accounts attracted a notably large number of followers. All 12 accounts together had 3,542 followers with publicly viewable profiles, many of whom followed more than one seed. For each of the 12 seeds and their 3,542 followers, we collected the 200 most recent tweets. In some cases, accounts had not yet posted 200 tweets, in which case all available tweets were collected. A total of 342,807 tweets were analyzed statistically. Tens of thousands of tweets were also manually reviewed.

Not all the follower accounts collected belonged to users who openly embraced or supported white nationalist views. A manual inspection of 350 randomly selected accounts from the follower dataset showed 44 percent were overtly engaged with white nationalism in observed tweets or by self-identification in their user profile (margin of error +/- 5.4 percent). The sample was selected using a random number function to assign a random number value to each user. The dataset was then sorted according to the random value, and the first 350 accounts were analyzed.



Sixteen percent of the dataset unambiguously self-identified with white nationalism in their profiles or usernames by using clear terms and phrases to describe themselves, such as Aryan, racist, racist or white nationalist. We took such primary self-identification to indicate a very high level of engagement with WN ideology.

Close to 50 percent of the random sample of followers of seed accounts displayed no overt sign of positive engagement with white nationalism in their tweets or profiles other than the act of following a seed account. The vast majority of these simply provided no clear evidence of their attitudes toward white nationalism, either for or against.

A tiny handful overtly expressed views incompatible with white nationalism, such as belief in racial harmony or non-discrimination. Only one account was observed to be an activist opposed to white nationalism.

Twelve percent of the random sample followed more than 500 other accounts. We judged these individuals to be indiscriminately following others, often in hopes of gaining more followers themselves (a common Twitter strategy). Twenty-two percent of the indiscriminate followers were overtly engaged with WN ideology on Twitter, about half the rate of the overall set, while 70 percent showed no overt engagement.

All of the seed accounts were based in the United States. Followers came from a number of locations. Because Twitter's profile parameters are unstructured, we could not statistically analyze the locations, but users were anecdotally observed to live primarily in the United States, the United Kingdom, Sweden and Canada. Users were observed from elsewhere in Europe, including France and Germany, Latin America, Asia and the Middle East.

Only public information was factored into this analysis. No data was collected on accounts marked as private at the time of collection. The status of some accounts changed between the start of collection and the end of analysis, including a handful of voluntary and involuntary suspensions, and some accounts which switched from public to private.



Analysis of Influence

We designed a scoring system to measure the “influence” and “exposure” of the seed accounts and their followers using their patterns of interaction with other users in the followers dataset. Interactions with the seeds and users outside the dataset of followers are not scored in the metrics presented here.

Influence and exposure are defined as Twitter interactions that mirror the real life definitions of the terms – influence defined as the ability to spread and promote content or ideas, and exposure as a tendency to consume content or ideas generated by others.

In terms of Twitter activity specifically, influence was defined as when a user did something that inspired a reaction from a second user (such as prompting the second user to reply to a tweet, or to retweet content initially tweeted by the first user). In the context of this paper, a user’s influence score may be said to determine how “influential” that user is.

Exposure was simply the other side of the coin, defined as when a user performed such an action in response to another user. A user who retweeted a second user, or replied to something the second user tweeted, received points for exposure. Throughout the paper, a user’s exposure score may be said to determine how “exposed” that user is to content generated by other users within the dataset. For instance, that a user with a high exposure score may be designated as “more exposed” than one with a low score.

Although influence and exposure represent opposite types of interactions, an account could have high or low scores in one or both categories, depending on the nature and quantity of its interactions. The complete scoring formula, with examples, is attached as Appendix B.

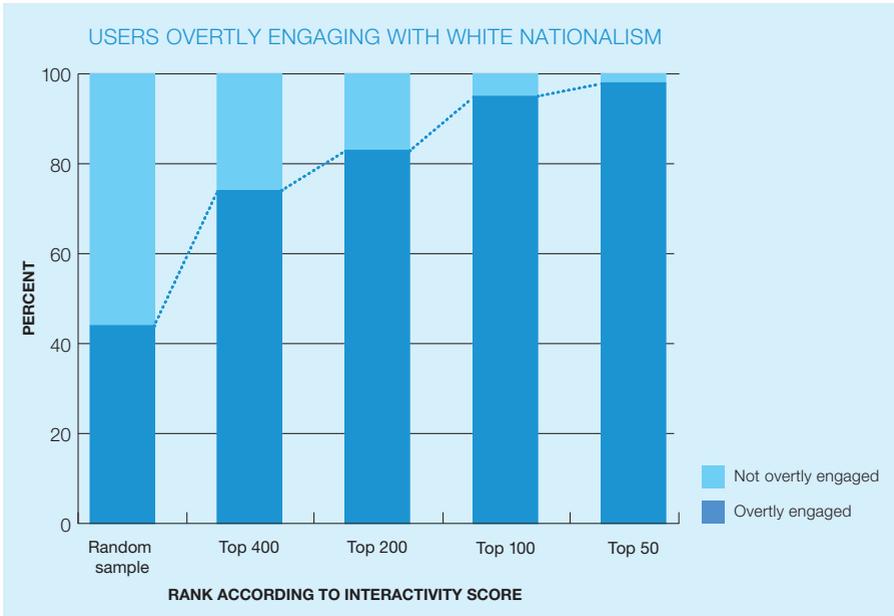
Interactions were weighted to emphasize keywords in tweets indicating they might be linked to more significant types of

communication such as “call,” “email” or “DM” (referring to Twitter’s direct messaging system for private communications between users).

After analysis of these initial results, we created a third score, which was simply the sum of the influence and exposure scores, a measurement we refer to as “total interactivity,” representing the overall quantity of meaningful interactions conducted by the account holder.

Importantly, the scoring system did not consider or score terms relevant to white nationalism. Our initial goal was simply to analyze the social network structure to evaluate which followers of seed accounts were extremely influential or extremely vulnerable to influence. The scoring system’s effectiveness at identifying people overtly engaged with the ideology that united the seed accounts was an unexpected finding.

Identification of Engaged Extremists

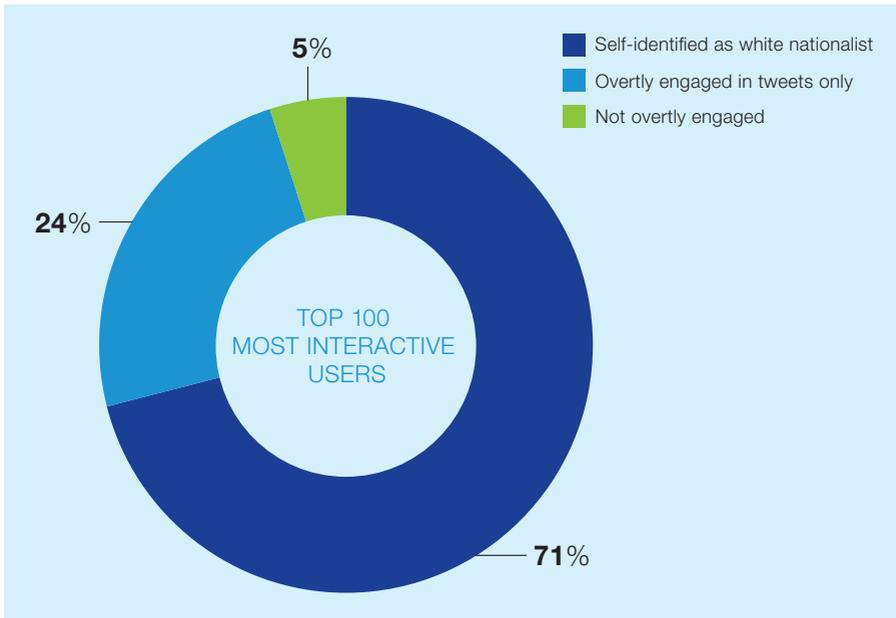


The rankings discussed throughout this paper apply to the followers of the seed accounts and exclude the seeds themselves, as our goal was to evaluate the less obvious participants in the network. Within the follower dataset, the number one most-influential account, by a wide margin, had been identified as a potentially important user in the dataset through a manual examination of the seeds and their followers prior to data collection, suggesting that the scoring system's results were not out of line with manual analysis. Manual inspection of the most influential and exposed accounts found that the ranked accounts generally conformed to the qualities we sought to identify (i.e., influence and exposure).

Among top 100 most influential, 86 percent were observed to be overtly engaged with white nationalism beyond the simple act of following a seed account (as previously defined in the dataset section, “overtly engaged” means users self-identified as white nationalists or tweeted consistently about white nationalism in a non-adversarial light). For the top 200, the percentages dropped to 73.5 percent. Ninety-three of the 100 most exposed accounts were overtly engaged with white nationalism. Of the top 200, 83 percent were overtly engaged.

After examining these results in detail and observing their unanticipated utility at identifying engaged extremists, we posed the question of whether it was possible to improve the identification process. We experimented with variations on the original scores and found that adding influence and exposure, creating a measurement of overall “interactivity,” produced improved identifications.

Manual examination showed 95 of the 100 accounts with the highest interactivity scores were overtly engaged with



white nationalism, and 71 self-identified as white nationalist in their usernames or profiles, showing extremely high levels of engagement. This compared very favorably to the random sample, where 44 percent were overtly engaged and 16 percent self-identified.

For users with the top 200 interactivity scores, 83 percent were overtly engaged and 57.5 percent self-identified. Of the top 400, 74 percent were overtly engaged.

There was substantial overlap among the highest-scoring accounts in each category, with a number of accounts scoring high in all three measurements.

Within the top 50 of all three categories, only one account appeared totally irrelevant, a person who generically interacted with accounts that began following him. Another user self-identified using neo-Nazi terminology and interacted with engaged white nationalists, but did not herself obsessively tweet on the subject.

The remaining top-50 users were observed to be highly engaged with white nationalism and spent the vast majority of their time on Twitter discussing WN ideology.

Based on these findings, we believe the scoring formula may have a number of useful applications in CVE efforts, which will be discussed below.

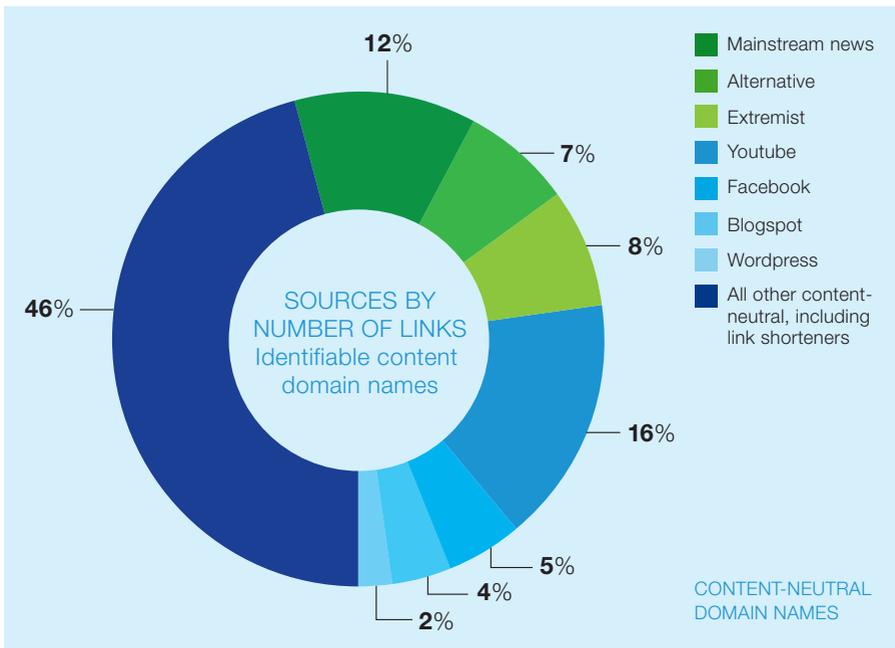


Content of Tweets

Although we set out with the primary goal of analyzing the network interactions of the followers of white nationalist seed accounts, the nature of our process also resulted in the collection of other data, including links to external content and hashtags (self-conscious content identifiers) used in tweets by members of the dataset. A casual examination of this collected data produced unexpected insights into the content being promoted and consumed by the followers dataset, and so we carried out a more detailed analysis.

Links

We detected 676 distinct domains in a total of 84,825 tweets that included links from the seed accounts and their followers.



We classified the links into four major categories: mainstream news, alternative sources, extremist sources, and content-neutral. Sites that generally adhered to journalistic standards were classified as mainstream news, even if they had a political slant, such as Fox News.

The alternative category included sites that explicitly articulated a political slant and sites that did not consistently conform to mainstream journalistic standards.

Extremist sources consisted of sites that openly advocated white nationalism or very closely related themes.

The content-neutral category included links where the specific content could not be resolved, as well as sports, entertainment, health and science news sources deemed of minimal interest to the study.

Mainstream news accounted for 12 percent of all links; alternative and extremist sources combined represented another 15 percent.

The categorization process necessarily involved a number of subjective judgments. For instance, a handful of sites included in the extremist category did not explicitly endorse white nationalism, but provided narrative support to themes white nationalists consider to be crucial. The most important of these, in terms of the number of links tweeted by users in the dataset, was Infowars.com, a conspiracy-oriented Web site. Other subjective calls included the classification of state-sponsored media and foreign nationalist sites. Overall, we estimate around 6 to 8 percent of these subjective judgments could be considered open for debate.

The single most-linked domain by an overwhelming margin was YouTube, which accounted for about 21 percent of all links detected. The content-neutral category included a number of other service providers, such as Blogger and Wordpress, where the nature of the content could not be determined. Finally, generic URLs generated by various popular Twitter clients represented 15.9 percent of all links.

We suspect many of these links would be classified as alternative or extremist content, to an extent that could substantially shift the percentages given above.

Top Sites

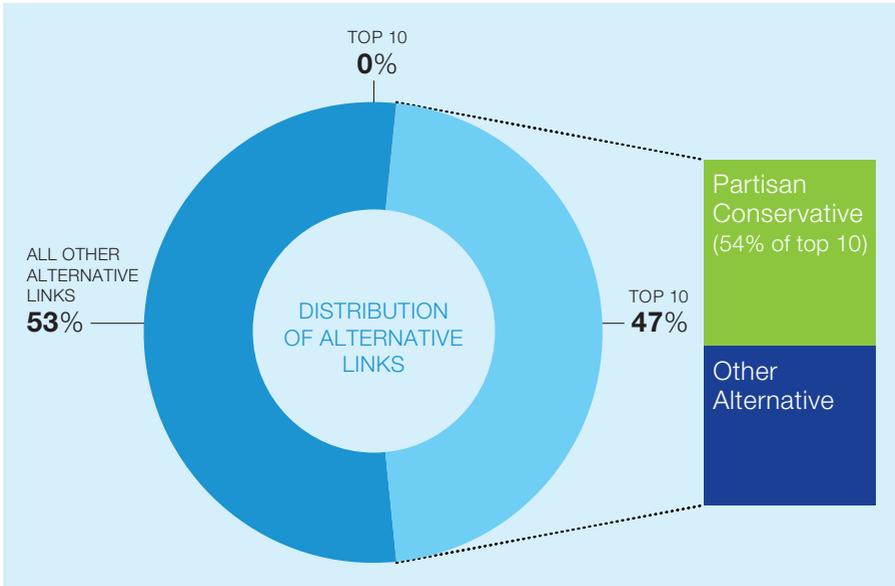
The top 10 sites linked in each category were:

Extremist	Alternative	Mainstream	Content-neutral
whiteresister.com	rt.com	dailymail.co.uk	youtube.com
cofcc.org	examiner.com	foxnews.com	bit.ly
realisten.se	theblaze.com	bbc.co.uk	facebook.com
infowars.com	wnd.com	nytimes.com	youtu.be
vnnforum.com	breitbart.com	guardian.co.uk	blogspot.com
amren.com	presstv.ir	huffingtonpost.com	twitter.com
northwestfront.org	dailycaller.com	telegraph.co.uk	tumblr.com
malevolentfreedom.org	hotair.com	thehill.com	wordpress.com
thenewamerican.com	globalresearch.ca	cnn.com	fb.me
vdare.com	thegatewaypundit.com	reuters.com	instagram.com
Top 10 as a percentage of all links			
46.3	46.8	36.1	58.5

Mainstream conservative information sources were found throughout the links, mostly in the alternative category, which by definition included political content.

The top 10 alternative URLs accounted for nearly half of all links in the alternative category. Among the top 10, 54 percent of links went to sources espousing clear partisan Republican viewpoints (as opposed to simply expressing conservative views).

The number of links per site dropped sharply in the lower rankings. Many sites in the lower 112 URLs were oriented so ambiguously among mainstream conservative, fringe conservative and libertarian views as to defeat easy classification.



Among the top 10 mainstream news sites, 31 percent of links went to the top two sources, the Daily Mail (UK) and Fox News (US), both of which are strongly oriented toward mainstream political conservatism. The remainder went to a mix of top mainstream news sources, including sites that were reasonably considered apolitical, such as the New York Times, or had a liberal orientation, such as the Huffington Post.

The most-linked extremist site was WhiteResister.com, but more than half the links to that site originated with two Twitter accounts, both affiliated with the site. Discounting links from those two accounts, WhiteResister.com would have dropped to sixth.

Significantly, the accounts promoting WhiteResister.com were the two most followed by other users in the dataset, pointing to influence above and beyond the bias introduced by their self-promotional tweets (or at least highlighting the effectiveness of self-promotion at garnering links).

The top 10 list also illustrates an important point. While the followers of white nationalist Twitter accounts are clearly engaged with the ideology, that doesn't mean they are involved in promoting the seed accounts of leading WN figures and organizations.

Of the top 10 most-linked extremist sites, only two – amren.com and northwestfront.org – were associated with seed accounts. Our influence rankings excluded the seed accounts. If they had been included in the dataset, three of the seed accounts would have ranked among the top 10 most influential. But the Web sites associated with those three accounts did not appear among the 100 sites most-linked by the followers dataset.

This demonstrates that followers were far more likely to promote their own related agendas than to promote well-known leaders and groups and suggests that these well-known leaders of white nationalism in the United States may be losing touch with their constituents. While they still carry significant weight in the movement, they are not generating daily buzz, on Twitter at least.

Connecting the most influential Twitter users with online content hosted elsewhere also expands the universe of options for CVE practitioners, which will be discussed in more detail in the CVE Implications section of this paper.

Hashtags

Hashtags included in tweets by the seed accounts and their followers pointed more clearly to engagement with mainstream politics. The top 10 hashtags used by the seeds and followers combined were:

1. tcot (top conservatives on Twitter)
2. svpol (Swedish police)
3. teaparty
4. p2 (used to address comments to progressives)
5. gop
6. tlot (top libertarians on Twitter)

7. ff (used to recommend Twitter accounts to other followers)
8. Obama
9. news
10. ows (Occupy Wall Street)

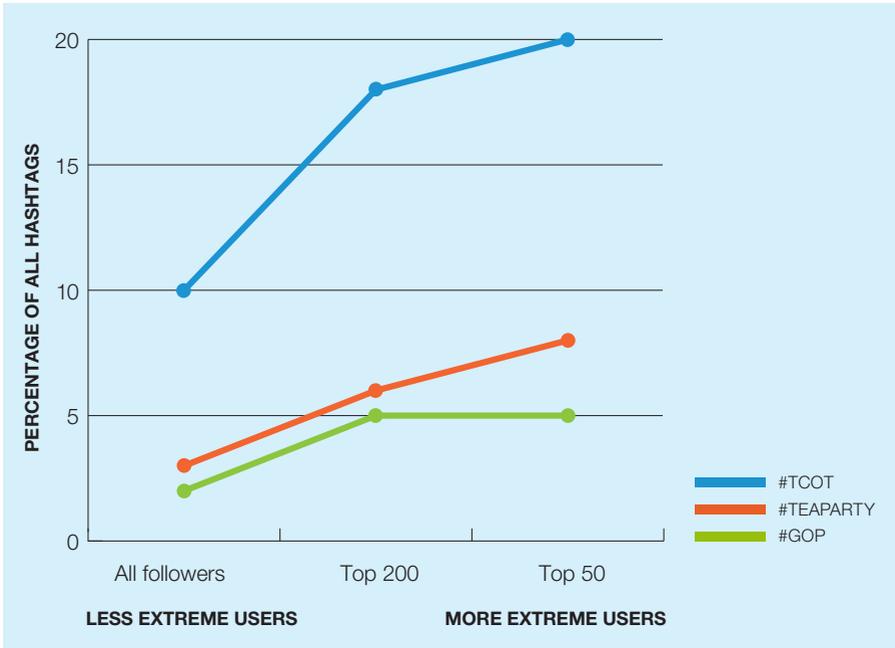
Given that not all accounts overtly engaged with white nationalism, we analyzed the top 10 hashtags tweeted by the 200 and 50 most influential followers of seed accounts, who were more highly engaged with white nationalism.

For hashtags, narrowing the dataset down to the most influential and extreme users surfaced specific tags relating to extremism, including #wpww (White Power World Wide), #nwo (New World Order) and #edl (English Defense League).

However, the mainstream political tags remained significant and in many cases represented a *higher* percentage as the user base became more extreme. The tags #tcot, #gop and #teaparty remained in the top 10, in the same order as they appeared for the whole dataset. #tcot was the most-used hashtag in all subsets examined.

All followers		200 most influential		50 most influential	
TCOT	10%	TCOT	18%	TCOT	20%
SVPOL	3%	SVPOL	11%	SVPOL	18%
TEAPARTY	3%	TEAPARTY	6%	WPWW	14%
P2	3%	GOP	5%	EDL	11%
GOP	2%	P2	5%	TEAPARTY	8%
TLOT	2%	WPWW	4%	FF	6%
FF	2%	TLOT	4%	GOP	5%
OBAMA	1%	OCRA	3%	OBAMA	5%
NEWS	1%	NATPOL	3%	NWO	4%
OWS	1%	NOW	3%	P2	3%

TOP 10 HASHTAGS, PERCENTAGE OF ALL TAGS



Political Context Considerations

In our random sample, 16 percent of users self-identified as white nationalists, but only 4 percent self-identified as mainstream conservatives. That suggests the link and hashtag results are driven more by white nationalists feeling an affinity for conservatism than by conservatives feeling an affinity for white nationalism.

The timing of the data collection, during the political convention season with a major election approaching, likely influenced the choice of both links and hashtags. President Obama's race almost certainly motivated white nationalists to take an enhanced interest in mainstream Republican politics.

A number of the 10 hashtags most often tweeted by the follower dataset (including those most oriented toward Republicans) were recommended for use in an organized

program of tweets by a social media group outside the dataset with a stated intention to “discredit the Obama campaign.” This effort began no later than August and may have tilted the results.²

Many of the same tags were recommended on a non-racial militia movement Web site in September 2012 as a way to broadcast that movement’s message to those in the political mainstream.³

This tactic could also have been adopted by white nationalists, although explicit evidence of such an effort was not found. These online campaigns were found by searching for several of the hashtags together in a database of extremist Web sites routinely maintained by co-author J.M. Berger.

Finally, it is important to note that engagement does not necessarily mean positive or two-way engagement. For instance, the ranking of #p2 and #ows among the top 10 suggests the mere presence of a hashtag does not necessarily represent an endorsement.

In light of these factors, we recommend revisiting this research after the election and caution against drawing overbroad conclusions about the relationship between mainstream conservative politics and white nationalism.

Despite all this, we believe these results are neither invalid nor insignificant. The context and timing may have inflated numbers regarding mainstream political engagement, but the social network’s involvement in electoral politics is itself interesting and significant.

Additional research over time may help illuminate the strength and persistence of these patterns, but we suspect it would not erase them. We discuss the implications of these findings further in the CVE recommendations below.

2 <https://www.facebook.com/groups/TwitterAttack/>, retrieved October 2012

3 “HASHTAGS # for Waging Political Warfare: Use these tags if you’re using twitter” posted on a militia Web site monitored by Intelwire.com on September 15, 2012

Comparison Group

After reviewing the first set of results, we performed supplemental data collection in October using a seed group of eight well-known and overtly anarchist Twitter accounts in order to further test the scoring system and try to capture a better picture of how mainstream political discourse relates to extremist movements aside from white nationalism.

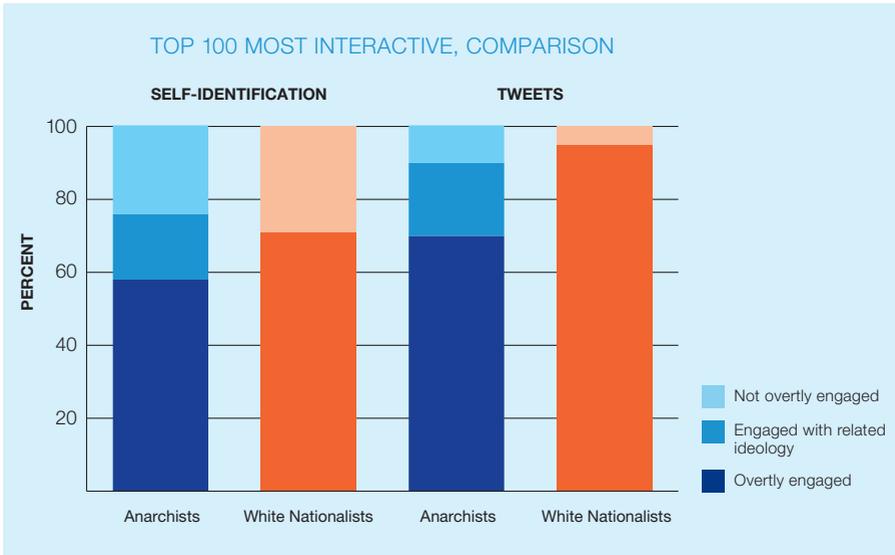
Anarchism is by definition less leadership-oriented, which may skew comparisons in various ways. It is also, by its nature, less clearly defined, making the categorization of accounts more difficult. For instance, the Occupy movement has strong anarchistic components, but Occupy is not generally considered anarchist by definition.

We collected from fewer accounts than we did WNs, but the eight anarchist seeds had more followers than the 12 white national accounts (5,409 versus 3,542). In part, this likely reflects the lesser social stigma associated with anarchism compared to white nationalism, as well as an anecdotally observed tendency of anarchists to be younger and more technologically proficient than white nationalists.

Based on anecdotal observation, we would also argue it is easier to be casually involved with anarchism than it is to be casually involved with ideological white nationalism. Users who self-identified with a variation of the word anarchist often included several other interests in their self-identification. Users who self-identified as white nationalist were much more prone to identify solely with that ideology.

For all of these reasons, we believe that anarchism likely represents the most difficult ideological extremism challenge for metrics-based identification.

Identification Rate



Of the 100 followers of anarchist seeds with the highest interactivity scores, 70 percent overtly engaged on Twitter with anarchism, 70 percent overtly engaged with anarchism; engaged with very closely related movements, including the Black Bloc segment of the Occupy movement and Anonymous; or self-described as radical or revolutionary.

Another 20 percent overtly engaged in tweets with Occupy but not Black Bloc, or overtly engaged with strongly anti-government leftist movements, making their presence clearly relevant, while not precisely matching the definition of the seeds. Depending on how you classify that 20 percent, the identification rate falls between 70 and 90 percent for the top 100, compared to 95 percent for white nationalists.

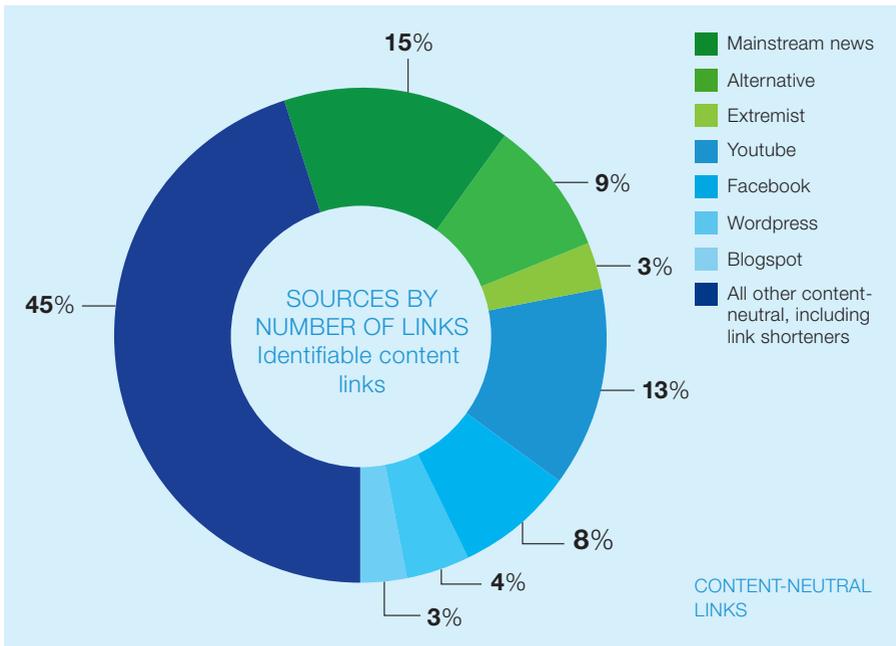
Fifty-eight of the top 100 accounts self-identified as anarchist, Black Bloc, Anonymous or something clearly within the same ideological space. Another 18 self-identified with Occupy, communism or socialism, but did not self-identify as anarchist.

This self-identification rate of 58 percent to 76 percent (depending how tightly one defines the target category) compared to a self-identification rate of 71 percent for the top 100 most interactive members of the white nationalist follower dataset.

Given that anarchism is extremely diverse and resistant to consistent labeling and movement cohesion, and given the lower threshold for casual involvement, we believe these results are extremely encouraging, especially given that only about 10 percent of the entire dataset self-identified with a variation of the word “anarchist.”

In light of these findings, we believe the scoring system can be effective at identifying followers who are highly engaged in ideological extremism in social networks revolving around seed accounts with a clear ideological orientation, regardless of what specific ideology is being examined.

Content of Tweets



Links and hashtags used by the anarchist seed accounts and their followers somewhat predictably skewed to the left side of the political spectrum. However, they tweeted considerably less content related to partisan politics than WN seeds and followers.

In addition to being a larger dataset, the anarchist followers were also prolific tweeters. We analyzed 188,573 links (the top 85 percent of the 220,214 collected) and found that followers of anarchist seeds relied somewhat more on mainstream news than white nationalists (15 percent compared to 13 percent), while relying much less on clearly extremist sites (3 percent versus 8 percent). The number of content neutral links was the same, 73 percent, with the remaining links going to alternative sources.

Categorizations were even more subjective for this dataset than for the original. For instance, it was difficult to determine whether Occupy, Wikileaks, Anonymous and hacker sites should be considered as extremist sites relevant to the seed ideology.

For purposes of the analysis above, all of these were counted as “extremist.” If they were discounted, extremist sites would have represented only 1.4 percent of the total mix, far less than the 7 percent tweeted by the white nationalist dataset.

Alternative sites were overwhelmingly liberal in their orientation, although debatably less partisan. For instance, Democracy Now, Mother Jones and Salon were among the top 10 alternative sites, but the latter two especially have frequently and sharply criticized the Obama administration on topics such as law enforcement, surveillance, foreign policy and the use of drones.

The hashtags offer a better window for comparing the political proclivities of followers of white nationalist seeds versus followers of anarchist seeds. The top 10 hashtags used in the anarchist dataset were:

1. ows (Occupy Wall Street)
2. occupy
3. oo (Occupy Oakland)

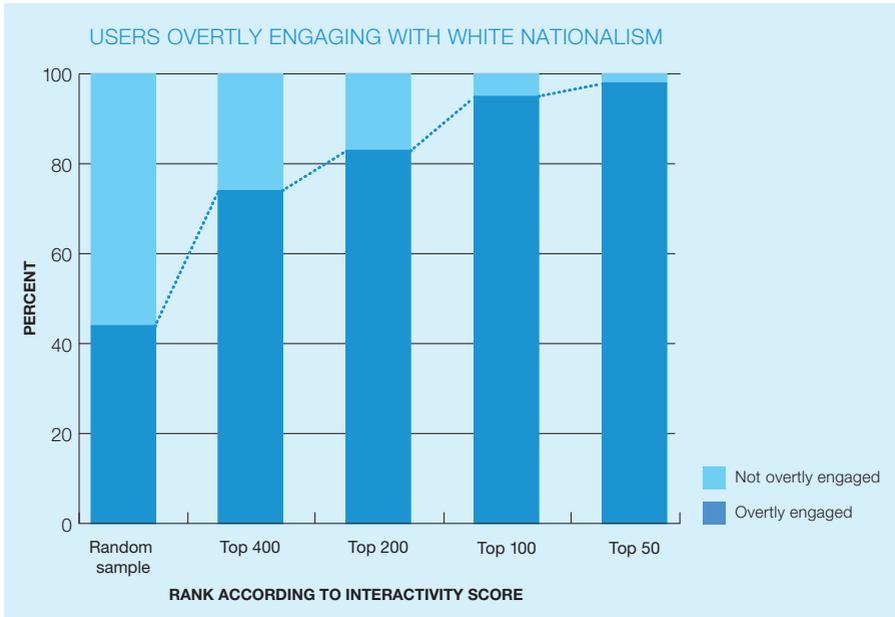
4. occupydenver
5. anonymous
6. p2 (used to engaged with progressives)
7. ff (Follow Friday)
8. s17 (September 17, the anniversary of the founding of Occupy Wall Street)
9. occupywallstreet
10. occupyoakland

Compared to the top 10 hashtags used by white nationalists, these tags suggest a much higher disenfranchisement from mainstream politics. Again, this may relate to the basic nature of anarchism, which is fundamentally opposed to political institutions, compared to white nationalism, which is not opposed to institutions per se.

Overall, we interpreted these results to indicate anarchists, while clearly embracing many liberal values, were far less engaged with partisan liberal politics than white nationalists with partisan conservative politics.



CVE Implications and Recommendations



The first challenge in conducting any type of Countering Violent Extremism (CVE) program is finding the target audience. Given the ease of access to extremism online and the allure of large numbers, many CVE efforts have focused on social media as a hunting ground for extremists. But while it is relatively simple to collect large datasets of potential extremists online, it is much more challenging to identify people who are sufficiently engaged to be suitable targets for CVE. Online activity often involves a lot of bad behavior and indiscriminate sampling of ideas. Shielded by physical

distance and anonymity, people often express views online that they would rarely, if ever, express in their offline lives.

In short, the vast majority of people taking part in extremist talk online are unimportant. They are casually involved, dabbling in extremism, and their rhetoric has a relatively minimal relationship to the spread of pernicious ideologies and their eventual metastasization into real-world violence. Any CVE program must begin by sifting the wheat from this chaff.

While this study is not an end-state solution to the sorting problem, we believe it represents a significant step forward. As shown in the graph above, a higher interactivity score clearly corresponded to a higher probability that a user was overtly engaged with the ideology espoused by the seed accounts. Ranked by the interactivity metric, we found that 95 of the 100 highest-scoring members of the followers dataset were highly engaged with white nationalist ideology.

These results are not a substitute for manual examination of user activity, but they provide a way to take a snapshot of a large user base and almost instantly prioritize the dataset, isolating fruitful avenues of investigation. For instance, a manual examination of the followers of seed white nationalist accounts took several days and identified only two or three follower accounts that subjectively appeared to be very influential. The metrics identified the same accounts and many more while requiring much less direct investment of time and effort by a human analyst.

CVE initiatives are not limited to one type of extremism. They can cover diverse ideologies such as jihadism, white nationalism, black nationalism, anti-Muslim, sovereign citizen, anarchism, eco-terrorism and more.

Based on the comparison results for anarchists, we believe this is an off-the-shelf approach that can be applied to any ideology without customization. However we do believe the metrics and analysis can be improved, perhaps substantially, with further research and testing.

We further believe that the paradigm of counting and weighting interactions can be adapted to any social network where an adequate amount of data can be scraped using automated tools.

In addition to identification, we believe the scoring system applied in this study has several other quantitative applications to CVE programs.

Measurement

CVE takes various forms. The most-discussed avenues involve persuasion (convincing potential extremists not to pursue violent ideologies) and disruption (preventing networks that promote radicalization from functioning effectively). These efforts can be conducted by government or by non-governmental organizations.

Critics of current CVE initiatives, including co-author J.M. Berger, have questioned the effectiveness of many CVE programs, pointing to the fact that there are no meaningful methods to measure success, especially as it pertains to persuasion programs. CVE strategies to date have been guided largely by intuition and anecdotal observation rather than by clearly relevant metrics.⁴

The influence, exposure and interactivity scores showed remarkable success at identifying highly engaged white nationalists within the starting set. Manual inspection of the top-ranked accounts also suggests that the scores themselves can be interpreted as they were intended – as quantitative measures of engagement-related activity.

Furthermore, the scores are not proportions or percentages; they measure the temperature of a given community, which means that a drop in score represents an absolute drop in the property (such as influence) being measured. So if the sum of all influence scores in the dataset drops, it suggests

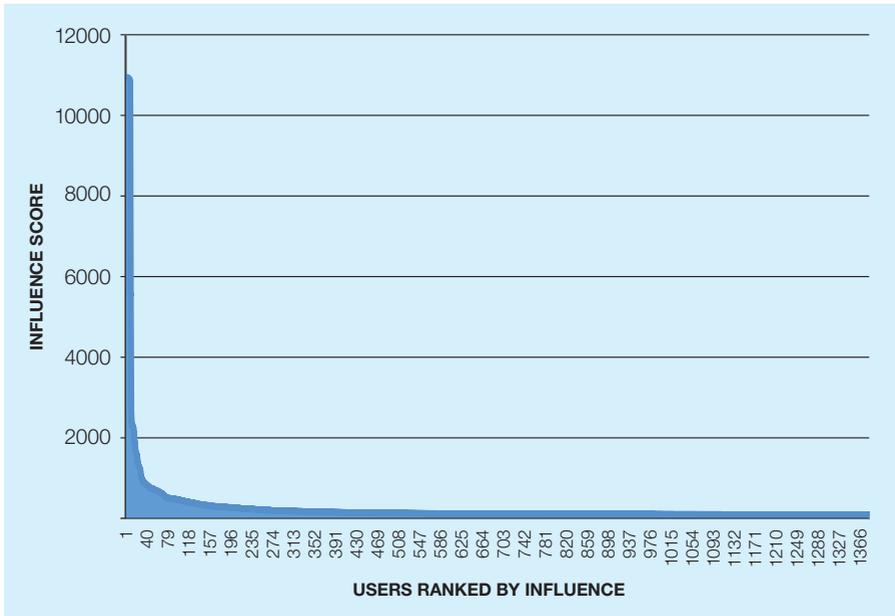
4 J.M. Berger, Intelwire.com, "Finding a Way Forward for CVE," 19 August 2011, http://www.foreignpolicy.com/articles/2011/12/09/monsters_and_children; Will McCants and Clint Watts, Foreign Policy Research Institute, "U.S. Strategy for Countering Violent Extremism: An Assessment," December 2012

that the leaders of the dataset have become less influential. If interactivity scores drop, it suggests that engagement with the target ideology by members of the dataset has been reduced.

Put simply, we believe the scoring system can be used to measure the absolute amount of engagement within a targeted online community in real time and over extended periods. This opens the door to quantitative testing of different CVE approaches. For instance, the scoring system can measure efforts to diminish the influence of key leaders, to inoculate a target audience against exposure, or simply to depress all interaction in the dataset.

Several CVE approaches are suggested by these conclusions.

Targeting Influence



Influence was disproportionately distributed throughout users with a nonzero score, as shown above. More than 50 percent of all the influence in the system was attributed to just 54 accounts, or 1.5 percent of the total number of accounts. The chart above shows the curve for all users with nonzero influence scores from highest to lowest.

Our understanding of influence in the real world is fairly simple. A relative few individuals wield a tremendous amount of influence over much larger numbers of people. Our findings put some specificity on that intuitive concept.

The vast majority of influence in the WN dataset was disproportionately concentrated at the top, as seen in the chart above. The top 50 most-influential users, looking this time at both followers and seeds – just 1.4 percent of the total – accounted for nearly 46 percent of all the influence in the system. (Because of their prominence, the seeds become more relevant when targeting accounts for CVE efforts.)

These findings generally conform to other studies of online activity, including the well-known 90-9-1 rule, which states 90 percent of users in most online social networks are passive, 9 percent are somewhat engaged and 1 percent drive most of the engagement and discussion.⁵

We attempted to enhance measures of simple engagement by weighting interactions to emphasize users with greater reach and who provide content that is widely redistributed and deemphasize users whose engagement carried less weight in group dynamics. We expect the scoring system can be refined to improve these measurements further, but it performed very well out of the gate.

The top 10 most influential accounts (less than 0.3 percent of all sampled users) accounted for 24.8 percent of the total influence in an ecosystem of 3,554 Twitter accounts (followers plus seeds) – a massively disproportionate distribution. This suggests the system is extremely vulnerable to disruption by focusing CVE efforts at the top of the food chain.

Possible methods for disruption include submitting complaints over violations of terms of service, where applicable; challenging and discrediting influencers through investigative

5 Jakob Nielsen, Nielsen Norman Group, "Participation Inequality: Encouraging More Users to Contribute," 9 October 2006, http://www.useit.com/alertbox/participation_inequality.html; Ben McConnell, The Church of the Customer Blog, "The 1% Rule: Charting citizen participation," 3 May 2006, http://www.churchofthecustomer.com/blog/2006/05/charting_wiki_p.html

reporting and research; or challenging and discrediting influencers through open debate in the online sphere.

Influence was disproportionately distributed even after removing the top 50 users. The chart above shows the curve for all remaining users with nonzero influence scores.

After the top 50, a sharp curve remained even among the users who contributed the lower 50 percent of influence to the system. While this suggests there would be a diminishing benefit to removing influencers over time, the distribution remains top-heavy, with the second 50 (1.4 percent of the remaining sampled users) comprising slightly more than 27 percent of the remaining influence. This suggests targeting the top of the curve can be effective even after the ecosystem has suffered significant attrition.⁶

While it is impossible to know for certain without direct experimentation, we believe these curves suggest disruptive CVE aimed at the top can be very effective over the course of a long engagement.

Another CVE option is to aim disruptive techniques at those with the highest exposure or interactivity scores. We originally suspected that high exposure scores might correlate to people at risk of radicalization who could be targeted for personalized interventions by nonprofit organizations. What we discovered instead was that users with high exposure scores were extremely committed to white nationalism and very active participants in the network.

While interventions seem unlikely to be effective with such users, disruption might be effective. The distribution of both exposure and interactivity was less concentrated at the top than influence (with interactivity slightly outperforming exposure), but the distribution remained disproportionately top-heavy for both metrics.

6 These calculations do not account for the redistribution of influence when users are removed from the system, with some of the remaining accounts becoming more prominent and thus more influential. Nor do they account for the reemergence of removed users under new Twitter handles.

Our evaluation of users with high exposure but medium to low influence is that they were mainly reactive, functioning as a supportive cheering section but not necessarily directly driving the content of discussions (which is not the same as saying their role in the system is insignificant). However, the interactivity metric was significantly more effective than any other method at identifying highly engaged users. Since users cannot be influential without other users who are exposed to them, we suspect that targeting the highest interactivity scores might have benefits that are not obvious but could become visible when scores are tracked over time.

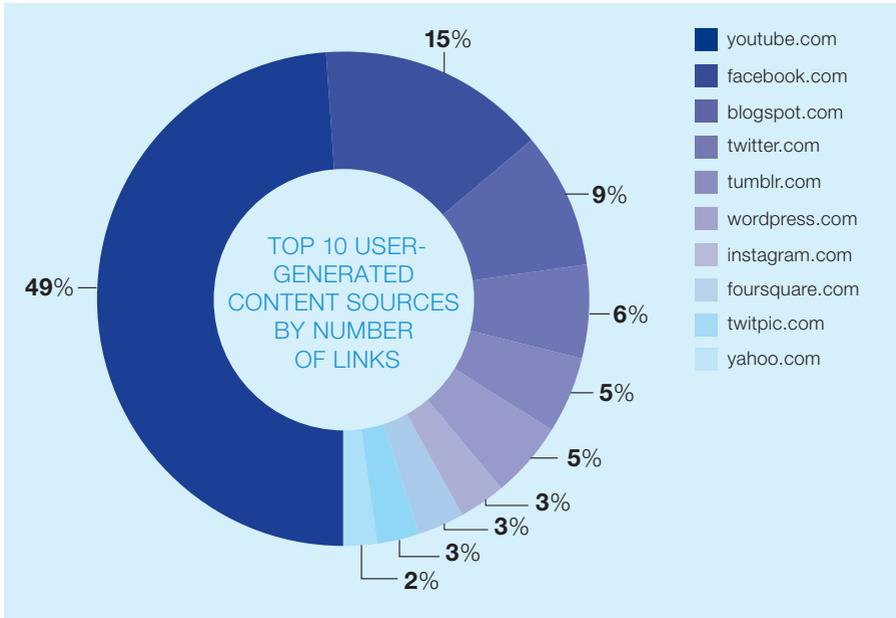
Finally, we should note that the influence distribution was different for anarchists than for white nationalists. The distribution of influence for the anarchists was much flatter, but still quite top heavy, with half the influence in the system residing among the top 106 accounts. This suggests the approach may need to be fine-tuned for different ideologies, but that the basic principles will remain effective.

Targeting Content

Disruptive CVE approaches on Twitter are limited by broadly permissive Terms of Service. Social media platforms in general are biased towards freedom of speech, although it should be noted that freedom of speech is not synonymous with the freedom to use a specific company's services. Providers of Internet publishing platforms are not obligated to host content in violation of their published rules, and they have wide latitude in deciding what those rules should be.

Twitter's TOS as of October 2012 banned "direct, specific threats of violence against others" and the use of Twitter to conduct or promote "illegal activities" and loosely defined "abusive behavior."

These terms are sufficiently ambiguous to allow designated global terrorist groups like Al Shabab to maintain active accounts on Twitter (although Shabab's account was recently suspended for a direct threat of violence, it has since opened



a new account).⁷ The Taliban and many other extremists also use Twitter extensively. However several prominent extremist accounts have been suspended for unknown reasons as of December 2012. This may reflect a shift in policies, or it may be due to other unknown reasons.

Regardless, disruption need not be limited to Twitter, thanks in part to the content analysis, which shows the most significant external sources of content. A large part of Twitter's value as a social media tool involves the ability to provide links and steer users to more detailed material on other Web sites.

For instance, YouTube emerged as the most-linked source of content in the dataset, representing more than 21 percent of all links and 49 percent of the top 10 links. Facebook, Blogger and Wordpress also ranked among the top 10 sources, even

7 BBC News, "Somalia's al-Shabab opens new Twitter account," 4 February 2013, <http://www.bbc.co.uk/news/world-africa-21321687>

without accounting for sites that use those services under a custom domain. As noted previously, we could not determine the nature of the content being linked to under these domains, but it's likely a significant percentage of these links point to user-generated content related to white nationalism.

Terms of service vary for each of these online publishing platforms. YouTube has relatively robust tools for reporting and removing hate speech, including the ability to report users for “violent or repulsive,” and “hateful or abusive” content, with specific subcategories for promoting terrorism and violence and hatred.

CVE practitioners can also make note of which subsets of a sampled dataset are linking to content, for instance targeting links provided by the most influential, exposed or interactive users, and report TOS violations, if any exist, with the content publisher.

Intervening with At-Risk Individuals

Manual inspection of the highest influence, exposure and interactivity scores found that almost all of the top 50 in each category were highly committed white nationalists unlikely to be swayed by intervention.

The percentages dropped steadily and consistently as the view expanded to the top 100 and 200 accounts. We believe manually examining users whose interactivity rank falls between 200 and 500 might be fruitful in identifying people vulnerable to radicalization or in its early stages, whose interactions indicate an interest in an extremist ideology but not a single-minded obsession with it.

Additional research is required to evaluate these findings in order to capture the characteristics of at-risk users and evaluate possible approaches to online intervention. Approaches can be field-tested and the scoring system can be used to help evaluate the results, although manual inspection and contextualization will likely be required. For instance, if an at-risk user is questioning the wisdom of pursuing white nationalism as a personal ideology, he or she

might temporarily become more interactive while seeking to test assumptions and challenge influencers.

Since radicalization is a process, the rankings are only relevant to this dataset during the time period of collection, so scores should be recalculated prior to experimentation.

Additional research might eventually point to ranges of influence, exposure and interactivity where the highest percentage of at-risk individuals can be found. Some adjustments to the scoring system would be necessary in order to support a universal range-based approach.

CVE partners who undertake this research should keep detailed notes on the scores, ranks and profiles of individuals they approach and the results of their intervention in order to further refine the most effective target ranges.

Countering Extremist Narratives

A popular approach to countering violent extremism involves creating and promoting narratives that oppose extremist ideologies, in the hopes of undermining existing extremist messages and competing for the hearts and minds of potential recruits.

Although such efforts can be evaluated by how many participants they attract, their utility in actually countering extremism remains unclear. For instance, accumulating 1,000 social media users who promote messages opposed to racism is an achievable goal, but it's not clear how that effort lands in at-risk communities and whether people already attracted to white nationalism are positively influenced by such efforts or if they may even be further alienated by them.

The metrics developed in this paper offer the possibility of performing quantitative evaluations for online messaging efforts. For instance, if the effort is centered around a Web site, CVE practitioners can track whether the links to the site are being tweeted by members of a targeted dataset, how the site is being characterized and whether influential users are discussing it.

If the effort is being promoted by specific Twitter accounts, those accounts can be introduced into the dataset and their influence can be tracked. Similarly, efforts to sway sentiment among specific subsets of users (or even individuals) can be tracked and scored to measure their success or failure in the online context. For instance, if the goal is to diminish the reach of the seed accounts, it is possible to track the influence of those accounts, whether users have exited the dataset by unfollowing them, and so forth.

It should be noted that it is possible to game the system, so CVE practitioners must be coached on how to achieve honest results, although it is also possible that gaming the system could in some cases actually achieve the desired result.

Further study and rigorous analysis of future outcomes will be required to evaluate how the scoring system can best be applied to such initiatives.

Reactive Opportunities to Counter Extremist Narratives

The political content observed during data collection strongly suggests that white nationalist seed accounts, their Twitter followers and the most engaged subsets of their followers believe that mainstream political conservatives represent potential recruits, potential allies, or at minimum, persons with whom they share some interests.

Regardless of exactly why, the data indisputably shows that during the collection period, people engaged with WN ideology consumed and distributed content related to mainstream political conservatism and tried to engage with mainstream conservatives.

Many mainstream conservatives likely find these approaches offensive and choose to ignore them, but this outreach represents opportunities to intervene with people at various stages in the radicalization process. If someone reaches out to influence you, you gain an opportunity to influence them back.

Conservatives who are politically active online may therefore be better positioned than anyone to attempt CVE interventions with white nationalists. We strongly recommend that politically active conservatives consider ways they can spearhead or aggressively contribute to such CVE efforts by directly challenging white nationalists who attempt to get their attention through different styles of interaction online.

Data at a glance

Data collected	
Seed accounts	12
Follower dataset	3,542
Tweets collected	342,807
Total number of links	84,825
Distinct domains linked to	676
Total number of hashtags	69,476
Unique hashtags	1,240

Identification rates by score based on random sample of 350 accounts (margin of error +/- 5.4 percent)	% overtly engaged with WN as observed in manual inspection of accounts
All followers	44
Influence (Top 200)	73.5
Influence (Top 100)	86
Influence (Top 50)	98
Exposure (Top 200)	83
Exposure (Top 100)	93
Exposure (Top 50)	100
Interactivity (Top 200)	83
Interactivity (Top 100)	95
Interactivity (Top 50)	98

Link content by percentage		
Mainstream news	12	
Alternative	7	
Extremist	8	
Content neutral	73	15.9 (could not be determined)

Appendices

A: Seed accounts

The following seed accounts were manually identified for collection based on co-author J.M. Berger's previous research on white nationalist movements. Each account was associated with a well-known ideological white nationalist organization or individual. The organizational affiliation for each account is given in parentheses.

1. @HeimdallsGhost (White Aryan Resistance)
2. @American3rdP (American 3rd Position)
3. @ANP14 (American Nazi Party)
4. @AryanNations (Aryan Nations, Morris Gullett branch)
5. @david_duke (David Duke)
6. @jartaylor (Jared Taylor, American Renaissance)
7. @TheNewNaziParty (American Nazi Party)
8. @nwfront (Northwest Front)
9. @UKANW (Northwest chapter of The United Klans of America)
10. @UKASOUTHFLA (Southern chapter of The United Klans of America)
11. @TheHoodedone88 (Affiliated with the United Klans)
12. @NSFM_Commander (Edward McBride, National Socialist Freedom Movement)

The number of seed accounts was designed to be large enough to produce substantial amounts of data, but small enough to keep collection from becoming prohibitively time- and resource-consuming.

B: Scoring formula and glossary of metrics

For each interaction, points were awarded for influence and exposure. For instance, if Account A retweeted a post by Account B, Account B received points for influencing Account A, while Account A received points for exposure to Account B.

If account C tweeted many links that were frequently retweeted, but never directly replied to or retweeted other users, Account C would accrue a high influence score but a low exposure score. If Account D sent replies to many users, but only a few responded, Account D would accrue a high exposure score, but a low influence score.

If Account E often engaged in two-way conversations with other users, Account E would register high scores in both influence and exposure.

Interactions were scored as follows:

- **at-replies:** 10
- **retweets (RT):** 5
- **hat-tip (HT):** 20
- **modified tweet (MT):** 15
- **non-reply mention:** 2
- **any of the above actions in a tweet with a high-value keyword:** 25. The list of high-value keywords included references to private communications, and offline and real-life interactions such as “call,” “email,” or “DM” (direct message). These terms were chosen on the theory that the most influential members of the network would be those who extend their influence beyond Twitter.
- **following:** For Account A following Account B, scores were awarded between one and 10 based on the average number of tweets per day. For instance, if Account B tweeted an average of one or fewer times per day, Account B received one point for influence and Account A received one point for exposure. If Account B tweeted an average of 10 times per day or more, the score was capped at 10 in order to avoid unduly weighting scores toward users who were simply prolific.

We experimented with several different parameters for the social network we would measure – influence and exposure to the seeds, influence and exposure among followers of the seeds, and influence and exposure to followers of the followers. Based on manual examination, the second category was easily the most interesting and effective. Except where explicitly noted, all references to influence and exposure are limited by the parameter, interactions only among followers of the seeds.

A quick overview of the measurements produced:

- Influence: A score that measures the ability of one twitter user to prompt other Twitter users to redistribute or engage with content.
- Influential: A Twitter user who has a high influence score and prompts many other users to engage with content.
- Exposure: A score that measures the tendency of one Twitter user to engage with or redistribute content created by other users.
- Exposed: The measured tendency of a given Twitter user to engage with or redistribute content.
- Interactivity: The sum of a user's influence and exposure score, reflecting quantity and quality of the user's overall interactions with other users.
- Interactive: The tendency of a user to interact with other users.

C: Top 10 most influential

1. Eaglesnest1488
2. MAFreedom
3. GILLY1488
4. MarmiteMan4
5. WARRIOR33_6
6. Aryanliving

7. Par_Sjogren
8. emilyscreams18
9. johnnywhiter
10. whiteunity14

D: Top 10 most exposed

1. WARRIOR33_6
2. danielb1714
3. Eaglesnest1488
4. karlhanke32
5. whiteunity14
6. HellWar88
7. DogofWarN
8. xdtac
9. craig25081989
10. viking14ranger

E: Top 20 most interactive

1. Eaglesnest1488
2. MAFreedom
3. WARRIOR33_6
4. GILLY1488
5. whiteunity14
6. Aryanliving
7. MarmiteMan4
8. johnnywhiter
9. danielb1714
10. WhiteResister
11. KevinGoudreau
12. Par_Sjogren
13. emilyscreams18
14. IvanaWAU
15. HVRabbit
16. DogofWarN
17. B14USA
18. xdtac
19. PaulinaForslund
20. craig25081989

F: Glossary of Twitter Terms

Account: An account registered with Twitter. These are usually, but not always, managed by a single user. Accounts may also be configured automatically tweet content from a Web site.

Handle: The username chosen by a Twitter account holder.

Tweet: A short message posted on Twitter, limited to 140 characters in length. Tweets may be simple comments or replies to other users, and may include links to content from other Web sites.

Follow: When a Twitter user subscribes to another user's tweets.

Followers: A list of all users that follow a specific Twitter account.

Followed: A list of all users followed by a specific Twitter account, sometimes referred to as "friends."

Timeline: A user's view of all tweets posted by the people they follow.

Reply or @reply: When a user directs a comment publicly to another user. This is done by placing an "@" sign in front of the target user's handle.

Direct Message: When two users follow each other, Twitter allows them to communicate privately through this service.

Retweet or RT: When a user republishes another user's tweet, with attribution.

Modified Tweet or MT: A retweet in which the user modifies the text of the original tweet.

Hat-Tip or HT: A tweet that acknowledges another user as the source of a link.

Non-Reply Mention: Including another user's handle in a tweet that is not a response to a previous tweet by that user. This often functions as the equivalent of CCing someone on an e-mail message.

Hashtag: When a Twitter users highlights a term in a tweet by placing the “#” sign in front of a word or phrase. Hashtags can be used for emphasis but are more often intended to make the tweet show up in thematic searches by other users.

G: Disclosures

This paper was commissioned by Google Ideas, but the opinions expressed are those of the authors.

A minimum of 11 accounts in the set were noted in the data to be following co-author J.M. Berger's Twitter account (@intelwire) at the time of collection. Berger's Twitter feed focuses on violent extremism and national security.

The most-linked URL in the white nationalist dataset was YouTube.com, which is owned by Google, of which Google Ideas is a part. The sixth most-linked URL was Blogspot.com, another Google property. We considered a number of online platforms for this study, including YouTube, but Twitter was ultimately chosen due to its more accessible API for developers, and the relative transparency of its social transactions.

About ICSR

ICSR is a unique partnership of King's College London, the University of Pennsylvania, the Interdisciplinary Center Herzliya (Israel) and the Regional Center for Conflict Prevention at the Jordan Institute of Diplomacy. Its aim and mission is to bring together knowledge and leadership to counter the growth of radicalisation and political violence. For more information, see www.icsr.info

KING'S
College
LONDON

University of London



www.icsr.info

