


VOX

POL



DISRUPTING DAESH

**MEASURING TAKEDOWN OF
ONLINE TERRORIST MATERIAL
AND ITS IMPACTS**

Maura Conway, Moign Khawaja, Suraj Lakhani,
Jeremy Reffin, Andrew Robertson and David Weir

DISRUPTING DAESH

MEASURING TAKEDOWN OF
ONLINE TERRORIST MATERIAL
AND ITS IMPACTS

Maura Conway, Moign Khawaja, Suraj Lakhani,
Jeremy Reffin, Andrew Robertson and David Weir

About the authors

Maura Conway is Professor of International Security in the School of Law and Government at Dublin City University and Coordinator of VOX-Pol.

Moign Khawaja is a doctoral candidate in the School of Law and Government at Dublin City University.

Suraj Lakhani is Lecturer in Criminology and Sociology in the Department of Sociology at the University of Sussex.

Jeremy Reffin is a Research Fellow in the Department of Informatics and co-founder of the TAG Laboratory at the University of Sussex.

Andrew Robertson is a Research Fellow in the TAG Laboratory in the Department of Informatics at the University of Sussex.

David Weir is Professor of Computer Science in the Department of Informatics and co-founder of the TAG Laboratory at the University of Sussex.

Acknowledgements

This research was partially funded by a grant from the UK Home Office.

ISBN: 978-1-873769-70-6

© VOX-Pol Network of Excellence, 2017

This material is offered free of charge for personal and non-commercial use, provided the source is acknowledged. For commercial or any other use, prior written permission must be obtained from VOX-Pol. In no case may this material be altered, sold or rented.

Like all other VOX-Pol publications, this report can be downloaded free of charge from the VOX-Pol website: www.voxpol.eu

Designed and typeset by Soapbox, www.soapbox.co.uk

TABLE OF CONTENTS

| | |
|--|-----------|
| Executive Summary | 4 |
| INTRODUCTION | 8 |
| SOCIAL MEDIA MONITORING: METHODOLOGY | 10 |
| Caveats | 12 |
| OUR DATA | 14 |
| MEASURING EFFECTS | 18 |
| Account Longevity | 20 |
| Community Breakdown | 26 |
| BEYOND TWITTER: THE WIDER JIHADIST SOCIAL MEDIA ECOLOGY | 32 |
| Case Study: Destinations of Official IS Propaganda | 35 |
| RECOMMENDATIONS | 38 |
| FUTURE RESEARCH | 42 |
| CONCLUSION | 44 |

EXECUTIVE SUMMARY

- This report seeks to contribute to public and policy debates on the value of social media disruption activity with respect to terrorist material. We look in particular at aggressive account and content takedown, with the aim of accurately measuring this activity and its impacts.
- Our findings challenge the notion that Twitter remains a conducive space for Islamic State (IS) accounts and communities to flourish, although IS continues to distribute propaganda through this channel. However, not all jihadists on Twitter are subject to the same high levels of disruption as IS, and we show that there is differential disruption taking place.
- IS's and other jihadists' online activity was never solely restricted to Twitter. Twitter is just one node in a wider jihadist social media ecology. We describe and discuss this, and supply some preliminary analysis of disruption trends in this area.
- Our analysis rests on a dataset containing 722 pro-IS accounts (labelled Pro-IS throughout) and a convenience sample of 451 other jihadist accounts (labelled Other Jihadist throughout), including those supportive of Hay'at Tahrir al-Sham (HTS), Ahrar al-Sham, the Taliban and al-Shabaab, active on Twitter at any point between 1 February and 7 April 2017.
- The Pro-IS accounts were located and identified using three methods: the original seed set of accounts (27%) were manually identified by the research team; the second set of accounts (30%) were identified 'semi-automatically' (i.e. automatically identified by the system and manually inspected and verified); and the third group of accounts (43%) were identified using an 'advanced semi-automatic' system via IS propaganda links.

- For the Pro-IS accounts, 57,574 tweets were collected, with 7,216 (12.5%) of these tweets containing out-links, i.e. links to wider websites, social media platforms, content hosting sites, etc. (not including links within Twitter). For the Other Jihadist accounts, 62,156 tweets were collected, of which 7,928 (13%) contained out-links.
- One of the overarching objectives of this research was to provide an up-to-date account of the effects of Twitter's disruption strategy on IS supporter accounts. We found that pro-IS accounts faced substantial and aggressive disruption, particularly those linking to official IS content hosted on a range of other platforms. The majority – around 65% – of the Pro-IS accounts in our dataset were suspended within 70 days of their establishment, with the overall suspension rate of pro-IS accounts probably being considerably higher.
- In a case study of accounts posting links to official IS content in a 24-hour period on 3 and 4 April 2017, 153 accounts were identified. A subset of 50 were 'throwaway accounts' (i.e. accounts specifically created on 3 April to disseminate IS propaganda with no expectation that they would stay online for any significant period of time). Together these accounts sent a total of 842 tweets with out-links to IS propaganda on other online platforms. Within this 24-hour period, 65% of accounts were suspended within the first 17 hours (07:00–00:00 GMT). The 50 throwaway accounts suffered even higher levels of disruption, with a 75% suspension rate during the same time period. This demonstrates that the disruption to official IS propaganda distribution was reasonably effective in the first 24 hours after linking the content.
- We also compared the suspension rates of Pro-IS accounts versus Other Jihadist accounts to check for differential disruption. We found that more than 25% of Pro-IS accounts were suspended within five days of their creation; a negligible number (less than 1%) of Other Jihadist accounts were subject to the same rapid response. Of those accounts in our dataset that were eventually suspended (i.e. 455 Pro-IS accounts and 163 Other Jihadist), more

than 30% of Pro-IS accounts were suspended within two days of their creation; less than 1% of Other Jihadist accounts met the same fate.

- As a result of this disruption, IS's ability to facilitate and maintain strong and influential communities on Twitter was found to be significantly diminished. Relationship networks were much sparser for Pro-IS accounts than Other Jihadist accounts. Other Jihadist accounts had the opportunity to send six times as many tweets, follow or 'friend' four times as many accounts and, critically, gain 13 times as many followers as Pro-IS accounts.
- Pro-IS users who persistently returned to Twitter resorted to adopting counter-measures, such as locking accounts, diluting the content of tweets, using innocuous profile pictures, and adopting meaningless Twitter handles. This situation makes it extremely difficult to maintain a strong and influential virtual community.
- Twitter is, however, just one node in a wider jihadist social media ecology. Therefore, we analysed a sample of destinations from Twitter for official IS propaganda at three time points (4–8 February, 4–8 March (excluding 7 March), and 4–8 April 2017). During these periods, Pro-IS accounts linked to 39 different third-party platforms or content hosting sites, as well as running its own server to host material. Of these, six remained prominent across the three time periods: justpaste.it, IS's own server, archive.org, sendvid.com, YouTube and Google Drive. These domains accounted for 83%, 70% and 67% of the URLs in the February, March and April sampling periods respectively. The takedown rate (as of 12 April) was 72%, 66% and 72% for the same sampling periods.
- Only 20 (or 0.04%) of all tweets from Pro-IS accounts contained a telegram.me link. The paucity of such links caused us to explore further; we found that just two of 722 Pro-IS users' biographies and two of 451 Other Jihadist users' biographies contained Telegram links. Neither group of accounts was therefore using Twitter to advertise its presence on Telegram.

- Our report makes three recommendations:
 1. Modern social media monitoring systems have the ability to dramatically increase the speed and effectiveness of data gathering, analysis and (potentially) intervention, but probably only when deployed in combination with trained human analysts.
 2. Active IS supporters who remain on Twitter, in particular content disseminators and their throwaway accounts, could probably be degraded further – though this may have both pros (e.g. detrimental impact on last remaining significant IS supporter Twitter activity) and cons (e.g. further degradation of Twitter as a source of data or open source intelligence on IS).
 3. Our focus was largely on Twitter, but we also pointed to the importance of the wider jihadist social media ecology. As our analysis was not restricted to IS users and content, we also underline the often uninterrupted online presence and activity of non-IS jihadists. We point to the usefulness of maintaining a wide-angle view of the online activity of a diversity of these, particularly HTS, across a variety of social media and other online platforms.
- For the future, we propose replicating the present research, but with a larger and more equal sample of HTS, Ahrar al-Sham, and Taliban accounts. This would allow for a more systematic and comparative analysis of the levels of disruption of a range of non-IS jihadists, the vibrancy of their contemporary Twitter communities and Twitter out-linking practices. It would also allow us to identify their other preferred online platforms. Additional research is clearly also warranted into the wider jihadist social media ecology. In particular, we suggest analysing pro-IS and other jihadist activity on Telegram, which is almost certainly where IS's online community has reconstituted, and comparing this with our present findings.

1. INTRODUCTION

USE OF THE Internet, particularly social media, by violent extremists and terrorists and their supporters is a source of concern for policy-makers and the public. This is due to apparent connections between consumption of, and networking around, violent extremist and terrorist online content. Concerns are focused on:

- adoption of extremist ideology – i.e. so-called ‘(violent) online radicalisation’;
- recruitment into violent extremist or terrorist groups or movements; and/or
- attack planning and preparation.

Particular concerns have been raised regarding easy access to large volumes of potentially influential violent extremist and terrorist content on prominent and heavily trafficked social media platforms. The micro-blogging platform Twitter has been subject to particular scrutiny, especially regarding their response to use of their platform by the so-called ‘Islamic State’ (hereafter IS), also known as ‘Daesh’.

One of the major aims of this analysis is to supply an up-to-date account of the effects of Twitter’s disruption strategy on IS-supporter accounts. Twitter continues to be ‘called out’ in the media and by policy-makers for the use of their platform by a variety of violent extremists. However, Twitter is not alone among social media companies and other online platforms in hosting extremist accounts and content.¹ The company has taken significant steps over the last three years to disrupt IS activity on their platform. Detailed description and analysis of the precise nature of this disruption activity

1 See, for example, UK House of Commons Home Affairs Committee, *Hate Crime: Abuse, Hate and Extremism Online*, London: House of Commons, 2017.

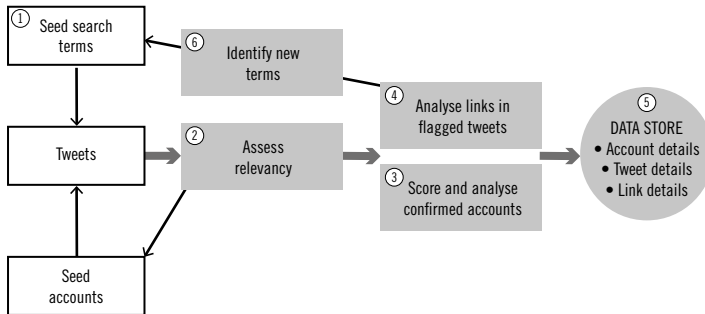
and, importantly, its effects is sparse. Therefore, this report aims to contribute to public and policy debates on the value of disruption activity, particularly aggressive account and content takedown, by seeking to accurately measure this activity and its impacts. Our findings challenge the notion that Twitter remains a conducive space for IS accounts and communities to flourish, although IS continue to distribute propaganda through the platform. Not all jihadists on Twitter are subject to the same high levels of disruption as IS, however, and we show that there is differential disruption taking place. An important related point is that the social media presence of IS and other jihadists has never been solely restricted to Twitter. Twitter is just one node in a wider jihadist social media ecology. We describe and discuss this, and supply some preliminary analysis of disruption trends in this area.



2. SOCIAL MEDIA MONITORING: METHODOLOGY

FOR THIS PROJECT we developed a semi-automated methodology for identifying pro-jihadist accounts on Twitter. Figure 1 illustrates this methodology, which was implemented using the social media analysis platform known as Method52.²

Figure 1. Detailed flow diagram for semi-automated social media analysis



The first step was to identify candidate accounts of interest. Our approach was based on finding tweets that had specific terms of interest in them (i.e. ‘seed search terms’) and/or finding accounts that were in some way related to other accounts known to be of interest (i.e. ‘seed accounts’). See step 1 in Figure 1.

When a tweet matched these search criteria, it was automatically analysed to see if it was actually relevant, using a machine-learning classifier trained to mimic the classification decisions of a human analyst.³ A key task of the relevancy classifier was to separate target Twitter accounts from other Twitter accounts using similar language (e.g. journalist or researcher accounts). If the tweet was deemed relevant, then further historic tweets were automatically extracted for the candidate account and assessed for relevancy (see step 2 in Figure 1).

- 2 Method52 was developed by the TAG Laboratory at the University of Sussex. For more information, see www.taglaboratory.org.
- 3 Classifiers were trained using semi-supervised machine learning approaches. Method52 provides components that enable this to be done swiftly and in a manner that is bespoke to a project.

This provided the system with an aggregate view of the tweet history of the account. This overview of the tweet history was combined with other account metadata that could be extracted; these pieces of information were scored automatically and candidate Twitter accounts that exceeded the set thresholds were presented to a human analyst for decision (step 3).

If the analyst confirmed that the account was pro-jihadist, then the out-links found in all the account's tweets were automatically analysed (step 4) and details of the account, its tweets and its links stored (step 5).

Information from new confirmed accounts was used by the system in a feedback loop to continually improve the efficiency of the system, thereby identifying new seed search terms (step 6) and providing additional seed accounts (step 2).

2.1 CAVEATS

There are a number of caveats attached to the data-collection process:

- The bulk of data gathering was undertaken over two months in early 2017. The system to implement our semi-automated methodology was created, tested and evolved throughout this period. Online accounts it returned were integrated with those found via traditional, manual search for accounts of interest. The overall approach was, therefore, a combination of automated and manual, and snowball and purposive sampling methods.
- Not all the available data was captured. There were various periods of downtime for the semi-automated system throughout this period as we developed and modified the methodology. Further, we were unable to include some accounts found via automated means because they were taken down before the

human analyst could assess and confirm their affiliation.⁴ By the project's end, when the system was working optimally, 100% of these accounts were identified by the software as pro-IS, reflecting the high level of disruption of IS-related accounts. We discuss this further below.

- The semi-automated system primarily focused on pro-IS accounts operating in English and Arabic or some combination of these languages. There is a possibility that accounts using, for example, Bahasa,⁵ Russian or Turkish were overlooked. We believe this possibility is worth mentioning but is negligible, as the system's effectiveness improved as we learned more about pro-IS users' contemporary Twitter activity and refined the methodology accordingly. By early April, for example, the software was able to identify accounts directly distributing IS propaganda with very high precision, no matter what language was used. At the same time, we believe it also identified the majority of accounts linking to that propaganda.
- Our data from the latter stages of this project suggests that around 50 or more throwaway IS accounts were produced daily. These accounts appear to be set up solely to distribute propaganda, typically have no followers and send only IS propaganda tweets until they are suspended. If we had been gathering all of these throwaway IS accounts over the whole research period, we would have had many hundreds more – perhaps as many as 2,000 to 3,000 accounts – in our sample data set. The analysis that follows is however based on pro-IS accounts with at least one follower and thus excludes these throwaway accounts unless otherwise stated.

4 This is a perennial issue in this type of research. It is also mentioned, for example, in J.M. Berger and Jonathon Morgan, *The ISIS Twitter Census: Defining and Describing the Population of ISIS Supporters on Twitter*, Washington DC: Brookings, 2015, p.41 and p.44.

5 J.M. Berger and Heather Perez, *The Islamic State's Diminishing Returns on Twitter: How Suspensions are Limiting the Social Networks of English-speaking ISIS Supporters*, Washington DC: George Washington University Program on Extremism, 2016, p.6.



3. OUR DATA

THE RESEARCH DATASET comprised 722 pro-IS accounts (labelled Pro-IS throughout) and 451 other jihadist accounts (labelled Other Jihadist throughout) with at least one follower active on Twitter at any time between 1 February and 7 April 2017 (see Table 1). Accounts were defined as pro-IS if their avatar or carousel image contained explicitly pro-IS imagery and/or text, and/or they had at least one recent tweet by the user (i.e. not a retweet) that contained explicitly pro-IS images and/or text, such as referring to IS as ‘Dawlah’ or their fighters as ‘lions’. Accounts maintained by journalists and others who tweeted, for example, Amaq News Agency content for informational purposes, were manually excluded. The Other Jihadist accounts included, among others, those supportive of HTS, Ahrar al-Sham, the Taliban and al-Shabaab. The same parameters were used to categorise these accounts.

Table 1. Description of final dataset

| | PRO-IS | OTHER JIHADIST |
|---|---------------|-----------------------|
| Number of accounts | 722 | 451 |
| Number of tweets | 57,574 | 62,156 |
| Number of out-links | 7,216 | 7,928 |
| Percentage of tweets containing out-links | 12.5% | 13% |

The accounts in our dataset were located and identified in three different ways (see Table 2). One set of accounts was manually identified by the research team, principally by looking at known jihadi accounts (or those known to be of interest to jihadi supporters) and inspecting accounts that were following or being followed by them. A second group of accounts was identified semi-automatically – that is, automatically by the above-described social-media monitoring system and then manually inspected by a human analyst who confirmed: (i) whether they were jihadist accounts or not; and (ii) if they were, of what type. Several approaches were used to identify seed

Accounts were defined as Pro-IS if their avatar or carousel image contained explicitly pro-IS imagery and/or text, and/or they had at least one recent tweet by the user (i.e. not a retweet) that contained explicitly pro-IS images and/or text, such as referring to IS as ‘Dawlah’ or their fighters as ‘lions’.

accounts or generate seed accounts. This included analysing the vocabulary being used in known jihadi accounts that were currently or had recently been active, determining which terms were being used much more often than would be expected statistically, and searching for tweets that contained these terms. These candidates were then winnowed based on the relevancy of their tweets in general (see above) and other metadata. A third group of

accounts was identified automatically by the social-media monitoring system based on the presence of known IS propaganda links. These links were first identified through other tracking procedures, including (but not limited to) being spotted in confirmed IS tweets.

It is important to underline that our pro-IS Twitter account dataset is as close as possible – taking into account the caveats in section 2.1 above – to a full dataset of explicitly IS-supportive accounts with at least one follower for the period studied. On the other hand, the Other Jihadist category is a convenience sample of non-IS jihadist Twitter accounts collected for comparison purposes and in no way reflects the true number of these accounts on Twitter.

Table 2. Location and identification of Twitter accounts

| | PRO-IS | | OTHER JIHADIST | |
|-------------------------|------------|----|----------------|----|
| | NO. | % | NO. | % |
| Manually identified | 193 | 27 | 332 | 74 |
| Semi-automated | 218 | 30 | 119 | 26 |
| Advanced semi-automated | 311 | 43 | – | – |
| TOTAL | 722 | | 451 | |



4. MEASURING EFFECTS

TWITTER HAS BEEN one of the most preferred online spaces for IS and their supporters, even prior to the establishment of their so-called 'caliphate' in June 2014. It was estimated that there were between 46,000 and 90,000 pro-IS Twitter accounts active in the period September to December 2014.⁶ However, their activity was subject to disruption by Twitter from mid-2014 and, although initially low level and sporadic, significantly increasing levels of disruption were instituted throughout 2015 and 2016. From mid-2015 through January 2016, for example, Twitter claimed to have suspended in the region of 15,000 to 18,000 IS-supportive accounts per month.⁷ From mid-February to mid-July 2016, this increased to an average of 40,000

It was estimated that there were between 46,000 and 90,000 pro-IS Twitter accounts active in the period September to December 2014.

IS-related account suspensions per month, according to the company.⁸ Despite the growing costs attached to remaining on Twitter (such as greater effort to maintain a public presence while relaying diffused messages and deflated

morale), during this period IS supporters routinely penned online missives exhorting 'Come Back to Twitter'.⁹ In 2017, is it worthwhile for pro-IS users to do so?

Until now, the small amount of publicly available research on the online disruption of IS has focused on the impact of Twitter's suspension activities on follower numbers for re-established accounts.¹⁰ We also looked at the longevity or survival time of accounts, and compared Pro-IS to Other Jihadist accounts on both measures (i.e. follower numbers and longevity). Our overall finding

6 Berger and Morgan, *The ISIS Twitter Census*, 2015, p.9.

7 Twitter, 'Combating Violent Extremism.' *Twitter Blog*, 5 February 2016, <https://blog.twitter.com/2016/combating-violent-extremism>.

8 Twitter, 'An Update on our Efforts to Combat Violent Extremism.' *Twitter Blog*, 18 August 2016, <https://blog.twitter.com/2016/an-update-on-our-efforts-to-combat-violent-extremism>.

9 Cole Bunzel, "'Come Back to Twitter': A Jihadi Warning Against Telegram." *Jihadica*, 18 July 2016, www.jihadica.com/come-back-to-twitter/.

10 Berger and Perez, *The Islamic State's Diminishing Returns on Twitter*, 2016.

was that pro-IS accounts are being significantly disrupted and this has effectively eliminated IS's once vibrant Twitter community. Differential disruption is taking place, however, which means that other Jihadist accounts are subject to much less pressure.

4.1 ACCOUNT LONGEVITY

This section addresses the survival time of accounts. All the Twitter accounts in our database were active at the time they were identified and classified as Pro-IS or Other Jihadist. Once an account was entered in the database, we monitored its status and recorded when

Once an account was entered in the database, we monitored its status and recorded when it was suspended, if this subsequently occurred. This allowed us to measure the age of each account (i.e. the time elapsed since the account's creation) at the date of suspension.

it was suspended, if this subsequently occurred. This allowed us to measure the age of each account (i.e. the time elapsed since the account's creation) at the date of suspension.

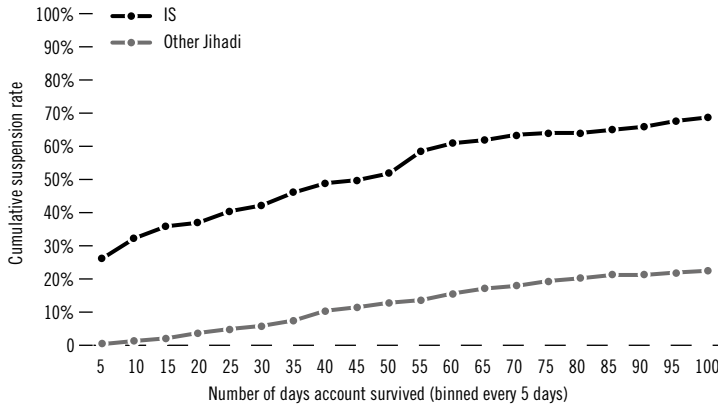
Worth underlining here is that the below-described survival rates of Pro-IS accounts would likely have been considerably shorter if the analysis included those

accounts suspended – often within minutes of creation – before they could be captured by the research team for inclusion in our dataset.

Figure 2 shows the estimated cumulative suspension rate for all Twitter accounts in our dataset, outlining the probability of an account being suspended against its age (represented in days) for the 722 Pro-IS accounts and 451 Other Jihadist accounts. Figure 2 shows that the majority – around 65% – of Pro-IS accounts were suspended before they reached 70 days since inception. At the same time point, less than 20% of Other Jihadist accounts had been suspended. In fact, as regards differential disruption, more than 25% of Pro-IS accounts were suspended within five days of inception; a negligible number (less than 1%) of Other Jihadist accounts were subject to the same rapid response.

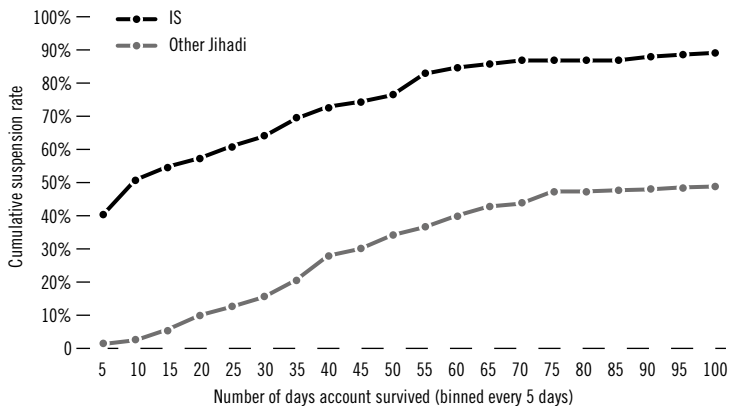
Our categorisation of these accounts as being jihadist in orientation was necessarily subjective. It is possible that others may disagree.

Figure 2. Cumulative suspension rate for all accounts in database



To address this possibility, Figure 3 focuses on those accounts in our dataset that were eventually suspended: 455 Pro-IS accounts and 163 Other Jihadist accounts. The rationale is that these accounts were judged independently to have breached Twitter’s terms of use. Again, as regards differential disruption, our data illustrates that 85% of Pro-IS accounts were suspended within the first 60 days of their life, compared to 40% of accounts falling into the Other Jihadist category. More than 30% of Pro-IS accounts were suspended within two days of their creation; less than 1% of Other Jihadist accounts met the same fate.

Figure 3. Cumulative suspension rate for accounts eventually suspended



In addition to the differences in longevity of Pro-IS and Other Jihadist accounts, the three subsets of Pro-IS accounts (i.e. those identified manually, semi-automatically based on general tweet content, and semi-automatically as a result of linking to official IS propaganda) also displayed different survival and activity patterns. From the 722 Pro-IS accounts in our dataset, the manually identified accounts (27%) survived disruption for longer periods and were predominantly tweeting about general IS and non-IS related news. The ‘general content’ semi-automated accounts (30%) had a somewhat shorter lifespan and were tweeting content generically related to the conflict (e.g. daily battle updates from several IS frontlines such as Mosul, Al-Bab, Deir Ez-Zor, eastern Aleppo, etc.). The advanced semi-automated group (43%) experienced the shortest lifespans. They were initially identified as a result of sending at least one tweet specifically disseminating ‘official’ IS propaganda (e.g. from the Amaq News Agency). Many were then found to be exclusively tweeting links to official IS propaganda.

4.1.1. Case Study: Intervention Effectiveness

Throughout the period of data collection, IS operated a 24-hour ‘news cycle,’ disseminating a new batch of propaganda on a daily basis via Twitter and other online platforms, using links to content hosted elsewhere on the Internet. These may be so-called ‘ghazwa’

or social media ‘raids’ orchestrated using some other online platform, potentially Telegram¹¹ and/or the dark web. The rapid takedown of Twitter accounts sending tweets containing links to official IS propaganda is seen in greater detail in this case study, which shows the effectiveness of intervention over a single 24-hour period. Figure 4 depicts survival

The rapid takedown of Twitter accounts sending tweets containing links to official IS propaganda is seen in greater detail in this case study, which shows the effectiveness of intervention over a single 24-hour period.

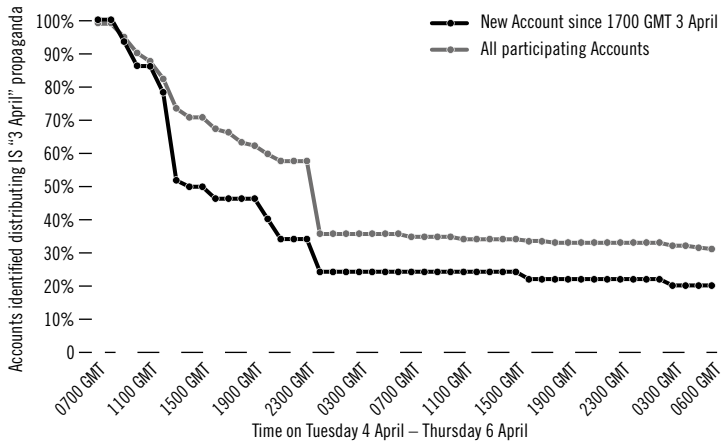
11 Nico Prucha, ‘IS and the Jihadist Information Highway: Projecting Influence and Religious Identity via Telegram.’ *Perspectives on Terrorism*, 10(6), 2016, pp.51–52.

curves for Twitter accounts that disseminated links to one or more pieces of official IS propaganda produced on Monday 3 April¹² (based on data collected on Monday 3 April and Tuesday 4 April 2017).¹³

On Monday 3 April 2017, IS uploaded its daily propaganda content to a variety of social media and online content-hosting platforms. This content generally included videos (in daily news format and other propaganda videos), 'picture stories' (a photo montage that tells a story), brief pronouncements similar to short press releases, radio podcasts and other documents (e.g. magazines). Over the course of Monday afternoon and evening, 153 unique Twitter accounts were identified that sent a total of 842 tweets with links to external (non-Twitter) web pages, each loaded with an item or items of IS propaganda. We identified only 10 of those Twitter accounts (7%) as being independent, third-party 'mainstream' accounts. The balance of accounts were identified as pro-IS. Fifty of these accounts appeared to be throwaway accounts created on Monday evening.

- 12 By early April 2017, the research reached the stage where there was complete access to IS's main Twitter propaganda apparatus. This enabled the semi-automated system to determine what IS and supporter tweets would be linking to before those tweets were sent. It is thought that this occurred several hours before Twitter themselves became aware of these accounts and their tweets. Much of this may have been due to the research team being able to access data and intelligence across multiple sites, allowing early prediction of tweet material, where Twitter's disruption team were likely restricted to monitoring their own platform only. The system was thus able to immediately identify when an account disseminated one of these propaganda links on Twitter. It was then possible to capture the rate and speed of suspension.
- 13 It should be noted that this date was chosen at random and thus propaganda represented in this graph had no relation to the chemical attack on the town of Khan Shaykhun, also on 4 April, as the propaganda was produced by IS on Monday 3 April 2017.

Figure 4. Case study of intervention effectiveness: survival of IS disseminator accounts 4–6 April 2017



The semi-automated system tracked all the accounts disseminating this propaganda – those sending one or more tweets with a 3 April propaganda link at some point prior to 06.00 GMT on the morning of Tuesday 4 April 2017. Figure 4 shows the survival curves for all 153 Twitter accounts tweeting IS propaganda from Monday 3 April and for the subset of 50 accounts specifically created on the Monday evening. The data shows that, at 07.00 GMT on Tuesday 4 April 2017, 100% of these accounts were active. However, by 13.00 GMT, only 73% of the 153 accounts were still active, falling to 58% by 23.00 GMT. This then dropped sharply to 35% surviving un-suspended by midnight on Tuesday. Very few of these surviving accounts were suspended over the subsequent 48 hours that we tracked. The 50 throwaway accounts created on Monday evening specifically to disseminate propaganda were suspended or deleted even faster: by 13.00 GMT only 52% were still active, falling to 34% by 23.00 GMT and 24% by midnight on Tuesday.

Figure 5 illustrates which accounts were responsible for posting original links (i.e. links that had not been sent before by another Twitter account). The diagram shows the account that sent the tweet (labelled Account 1, Account 2, etc. in order to retain user anonymity), the domain the link pointed at (e.g. sendvid.com), the time the

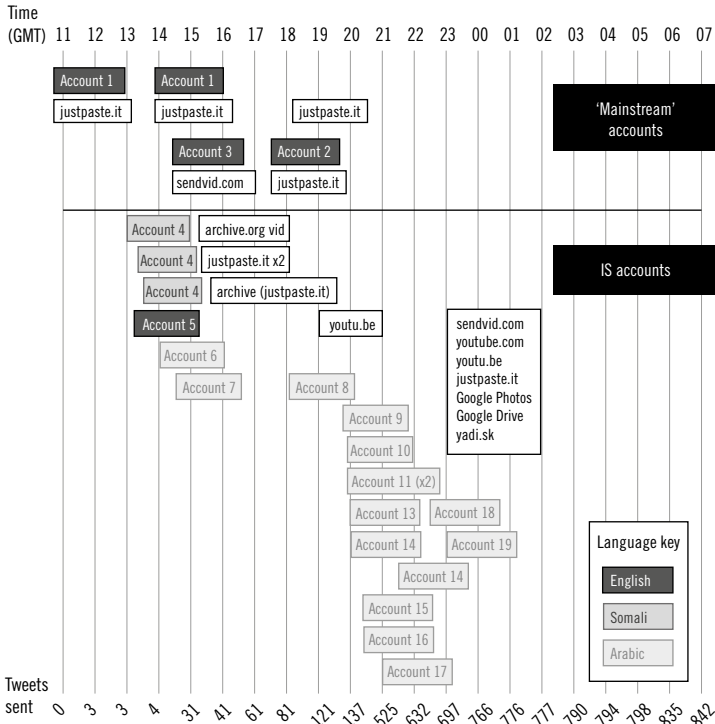
tweet was sent, the language used, and whether the account was a 'mainstream' third-party account or a pro-IS account. Overall, we identified 19 accounts sending original links to a total of 24 different URLs (destinations). Two of the accounts (identified as Account 1 and Account 2 in Figure 5) were mainstream, independent third-party accounts. A third account (Account 3) is a 'wolf in sheep's clothing', a fake account, a duplicate of the account of a widely followed, US-based new media journalist. This was an old account, unused since mid-2013, which was hacked (presumably by IS) and used to transmit IS propaganda since 26 March 2017.¹⁴ Account 1 is not a pro-IS user, but more of a citizen journalism-type account that albeit outside of the region was nonetheless in a position to supply three 'exclusives' (first releases) of pieces of official IS propaganda during the day. The mainstream accounts (including the fake one) and one known IS supporter account (i.e. Account 4) all tweeted in English. It is likely, however, that there were other English language accounts taken down before 06.00 GMT on Tuesday 4 April when we first identified the propaganda accounts. Most of the tweets and accounts in this case study were in Arabic, with an additional small number in Somali.

What this shows is that the response to official IS propaganda being distributed via Twitter was reasonably effective in terms of identifying and taking down such disseminator accounts in the first 24 hours after linking to official IS content. Comparing these rates to the rates across our entire Pro-IS dataset, it was also clear that those accounts disseminating official IS propaganda were taken down at a higher rate, compared to other pro-IS accounts that were not disseminating this propaganda. However, it must be borne in mind that pro-IS users were operating on a 24-hour 'news cycle' and creating a large number of accounts every day to disseminate daily propaganda. As these accounts were being taken down during Tuesday, a similar number of fresh accounts were being created and used to distribute the next day's official IS content. Therefore, it could be argued that, while efforts to remove

14 The account was still live as of 21 April 2017. For more on such accounts, see Berger and Perez, *The Islamic State's Diminishing Returns on Twitter*, 2016, p.16.

permanent traces of IS propaganda links from Twitter were relatively successful, IS was still able to broadcast links to its daily propaganda using Twitter in 24-hour bursts during the research period.

Figure 5. Case study of intervention effectiveness: IS disseminator account types, tweet timings, languages and URL destinations, 4 April 2017



4.2 COMMUNITY BREAKDOWN

What are the effects of this disruption of IS accounts? The truncated survival rates for Pro-IS accounts meant that their relationship networks were much sparser than for the Other Jihadist accounts in our dataset and compared to previously mapped IS-supporter networks on Twitter. From a more qualitative perspective, this means that the IS Twitter community was virtually non-existent in the research period.

Table 3. Median number of tweets, followers and friends for accounts not yet suspended

| | TWEETS | FOLLOWERS | FRIENDS |
|----------------|---------------|------------------|----------------|
| Pro-IS | 51 | 14 | 33 |
| Other Jihadist | 320 | 189 | 122 |

Table 3 compares the median number of tweets, followers and friends of Pro-IS accounts versus those of Other Jihadists. The short lifespan of the Pro-IS accounts meant that many had only a small window in which to tweet, gain followers and follow other accounts. This resulted in the Other Jihadist accounts enjoying the opportunity to: send six times as many tweets; follow or ‘friend’ four times as many accounts; and importantly, gain 13 times as many followers as the Pro-IS accounts. An even more stark comparison is between median figures for contemporary Pro-IS accounts versus those recorded for similar accounts in 2014. The median number of followers for contemporary Pro-IS accounts was 14 versus 177 in 2014,¹⁵ a decrease of 92%. The median number of accounts followed by IS supporters in 2014 was 257, while we recorded a median of 33 ‘friends’ per Pro-IS account – a decrease of 87%.¹⁶ In an analysis of 20,000 IS supporter accounts over five months (September 2014 to January 2015), Berger and Morgan observed suspension of just 678 accounts, a total loss of 3.4%.¹⁷ In our dataset, the total loss of Pro-IS accounts in just four months (between January and April 2017) was 63%. It is worth noting that the total loss of pro-IS accounts over the period studied would have been dramatically higher had we included not just accounts with a minimum of one follower, but all the throwaway accounts generated in the same period. Considering also those accounts we were unable to capture due to their suspension within minutes of creation, the total loss of IS-supportive accounts over the period was probably greater than 90%.

15 Berger and Morgan, *The ISIS Twitter Census*, 2015, p.30.

16 Berger and Morgan, *The ISIS Twitter Census*, 2015, p.32.

17 Berger and Morgan, *The ISIS Twitter Census*, 2015, p.33.

In the IS Twitter ‘Golden Age’ in 2013 and 2014, a variety of official IS ‘fighter’ and an assortment of other IS ‘fan’ accounts could be accessed with relative ease. For the uninitiated user, once one IS-related account was located, the automated Twitter recommendations on ‘who to follow’ accurately supplied others. For those ‘in the know’, pro-IS users were easily and quickly identifiable via their choice of carousel and avatar images, along with their user handles and screen names. Therefore, if one wished, it was quick and easy to become connected to a large number of like-minded other Twitter users. If sufficient time and effort was invested, it was also relatively straightforward to become a trusted – even prominent – member of the IS ‘Twittersphere.’¹⁸ Not only was there a vibrant overarching pro-IS Twitter community in existence at this time, but also a whole series of strong and supportive language (e.g. Arabic, English, French, Russian, Turkish) and/or ethnicity-based (e.g. Chechens or ‘al-Shishanis’) and other special interest (e.g. females or ‘sisters’¹⁹) Twitter sub-communities. Most of these special interest groups were a mix of: a small number of users actually on the ground in Syria; a larger number of users seeking to travel (or with a stated preference to do so); and an even larger number of so-called ‘jihobbyists’²⁰ with no formal affiliation to any jihadist group, but who spent their time lauding fighters, celebrating suicide attackers and other ‘martyrs’ and networking around and disseminating IS content.

In 2014, pro-IS users were already under some pressure from Twitter; for example, official IS accounts were some of the first to be suspended in summer 2014. Twitter’s disruption activity increased

18 See, for example, the extensive media coverage of the Twitter user @ShamiWitness who was revealed in December 2014 to be Mehdi Biswas, a 24-year-old Bangalore-based business executive, who prior to his arrest was one of the most prominent IS supporters on social media. Interestingly, his Twitter account was only suspended in early 2017, despite being dormant since his arrest. Biswas is awaiting trial in India.

19 Elizabeth Pearson, ‘Wilayat Twitter and the Battle Against Islamic State’s Twitter Jihad.’ *VOX-Pol Blog*, 11 November, 2015, www.voxpol.eu/wilayat-twitter-and-the-battle-against-islamic-states-twitter-jihad.

20 Jarret M. Brachman, *Global Jihadism: Theory and Practice*, Abingdon: Routledge, 2009, p.19.

significantly over time, forcing pro-IS users to develop and institute a host of tactics to allow them to maintain their Twitter presences, remain active and preserve their communities of support on the platform.²¹ For example, the group used particular hashtags, such as #baqiyyafamily ('baqiyya' means 'remain' in Arabic) to announce the return of suspended users to the platform, in an attempt to regroup after their suspension. Twitter eventually responded by including these hashtags in their disruption strategies. Interestingly, this increased disruption strengthened some IS supporters' resolve and

Twitter's disruption activity increased significantly over time, forcing pro-IS users to develop and institute a host of tactics to allow them to maintain their Twitter presences, remain active and preserve their communities of support on the platform.

they became even more determined to re-establish their accounts, even after repeated suspensions. This may have resulted in decreased numbers of pro-IS users, but also more close-knit and unified communities, because those who remained needed a high level of commitment and virtual community support to do so.²²

Eventually, however, the costs of remaining began to outweigh the benefits. Research from 2016 shows that "the depressive effects of suspension often continued even after an account returned and was not immediately re-suspended. Returning accounts rarely reached their previous heights,"²³ in terms of numbers of followers and friends. This was probably due to the eventual discouragement of many IS supporters subjected to rapid and repeated suspension.

21 For examples, see Berger and Perez, *The Islamic State's Diminishing Returns on Twitter*, 2016, pp.15–18.

22 Pearson, 'Wilayat Twitter and the Battle Against Islamic State's Twitter Jihad'. *VOX-Pol Blog*, 11 November, 2015, www.voxpol.eu/wilayat-twitter-and-the-battle-against-islamic-states-twitter-jihad. See also Elizabeth Pearson, 'Online as the New Frontline: Affect, Gender, and ISIS-takedown on Social Media.' *Studies in Conflict & Terrorism* (forthcoming).

23 Berger and Perez, *The Islamic State's Diminishing Returns on Twitter*, 2016, p.9.

Even those who persisted had to take counter-measures such as locking their accounts so they were no longer publicly accessible, or diluting the content of their tweets so their commitment to IS was no longer as readily apparent. By April 2017, these measures had taken such hold that the vast majority of Pro-IS account avatar images were default ‘eggs’ or other innocuous images, and many of the account user handles and screen names were meaningless combinations of letters and numbers (see Table 4). A conscious, supportive and influential virtual community is almost impossible to maintain in the face of the loss of access to such group or ideological symbols and the resultant breakdown in commitment. Therefore, IS supporters have re-located their social media community-building activity elsewhere, primarily to Telegram,²⁴ which is no longer just a back-up for Twitter.²⁵

As a result [of disruption by Twitter], pro-IS Twitter activity has largely been reduced to tactical use of throwaway accounts for distributing links to pro-IS content on other platforms, rather than as a space for public IS support and influencing activity.

From a quantitative perspective, the data discussed in this section demonstrates three key findings. First, IS and their supporters were being significantly disrupted by Twitter, where the rate of disruption depended on the content of tweets and out-links. Second, although all accounts experienced some type of suspension over a period of time, Pro-IS accounts expe-

rienced this at a much higher rate compared to the Other Jihadist accounts in the dataset. Third, this has severely affected IS’s ability to develop and maintain robust and influential communities on Twitter. As a result, pro-IS Twitter activity has largely been reduced to tactical use of throwaway accounts for distributing links to pro-IS content on other platforms, rather than as a space for public IS support and influencing activity.

24 Prucha, ‘IS and the Jihadist Information Highway’, 2016.

25 Berger and Perez, *The Islamic State’s Diminishing Returns on Twitter*, 2016, p.15.

Table 4. Changes in account name types due to disruption activity*

| TYPICAL USER HANDLES 2014–2015 | TYPICAL USER HANDLES 2017 |
|---|--------------------------------------|
| Mujahid1985 | 4iM7EjZphT3OXYG |
| BintSham | 5Asdf68 |
| ukhtialalmani | Omar_08 |
| Khilafah78 | t7dYqgYMaSB4Ecl |
| ShamGreenbird | GilUllul |

** These are not real account screen names but composite examples constructed for illustration purposes.*



**5. BEYOND TWITTER:
THE WIDER
JIHADIST SOCIAL
MEDIA ECOLOGY**

RESEARCH ON THE intersections of violent extremism and terrorism and the Internet have, for some time, been largely concerned with social media. They have often had a singular focus on Twitter because of its particular affordances – e.g. ease of data collection due to its publicness, and the nature of its application programming interface (API) – which is problematic.²⁶ For example, EUROPOL's Internet Referral Unit reported that, by mid-2016, they had identified “70 platforms used by terrorist groups to spread their propaganda materials”.²⁷ Therefore, this section of the report is concerned with the wider social media ecology where IS supporters and other non-IS jihadist users operate, with a particular focus on out-links from Twitter.

Partly because of its 140-character limit, Twitter functions as a ‘gateway’ platform²⁸ to other social networking sites and a diversity of other online spaces. In 2014, it was estimated that one in every 2.5 pro-IS tweets contained a URL. It was acknowledged at the time that it would be useful to analyse these links, but this was not undertaken due to complications around Twitter's URL-shortening practices.²⁹ The roll-out of auto-expanding link previews by Twitter in July 2015 remedied this difficulty. In terms of link activity in our data, most links were not out-links, but rather in-links (i.e. within Twitter): 8,086 or 14% for Pro-IS and 4,650 or 7.5% for Other Jihadist tweets. Of the Pro-IS and Other Jihadist Twitter accounts we identified, 1 in 8 (around 13%) contained non-Twitter URLs or out-links. This is a considerable reduction from the 40% of tweets reportedly containing URLs in 2014. Analysis of Twitter out-links nonetheless provides an interesting snapshot of the

26 Maura Conway, ‘Determining the Role of the Internet in Violent Extremism and Terrorism: Six Suggestions for Progressing Research.’ *Studies in Conflict and Terrorism*, 40(1), 2017, p.9 and p.12.

27 EUROPOL, *EU Internet Referral Unit: Year One Report*, The Hague: EUROPOL, 2016, p.11, <https://www.europol.europa.eu/publications-documents/eu-internet-referral-unit-year-one-report-highlights>.

28 Derek O’Callaghan, Derek Greene, Maura Conway, Joe Carthy and Pádraig Cunningham, ‘Uncovering the Wider Structure of Extreme Right Communities Spanning Popular Online Networks.’ In *WebSci ’13: Proceedings of the 5th Annual ACM Web Science Conference*, New York: ACM Digital Library, 2013, pp.276–285.

29 Berger and Morgan, *The ISIS Twitter Census*, 2015, p.21.

Top 10 platforms linked to by Pro-IS and Other Jihadist accounts in our data-collection period (see Table 5).

Table 5. Top 10 other platforms (based on out-links from Twitter)

| PRO-IS | | | OTHER JIHADIST | | |
|------------------------------------|--------|------------------------|--|--------|--------------------------------|
| PLATFORM | NUMBER | % OF ALL PRO-IS TWEETS | PLATFORM | NUMBER | % OF ALL OTHER JIHADIST TWEETS |
| 1. YouTube | 1,330 | 2.3% | 1. YouTube | 2,488 | 4.0% |
| 2. Google Drive | 792 | 1.4% | 2. Facebook | 1,294 | 2.1% |
| 3. justpaste.it | 472 | 0.82% | 3. justpaste.it | 479 | 0.77% |
| 4. Google Photos | 431 | 0.75% | 4. Islamic prayers website | 316 | 0.51% |
| 5. sendvid.com | 410 | 0.71% | 5. Taliban news website | 244 | 0.39% |
| 6. archive.org | 353 | 0.61% | 6. Official Taliban website | 228 | 0.37% |
| 7. archive.is | 243 | 0.42% | 7. Taliban's official Urdu website | 208 | 0.33% |
| 8. Bahasa IS fan site | 198 | 0.34% | 8. Hizb ut-Tahrir website | 189 | 0.30% |
| 9. medium.com | 155 | 0.27% | 9. Telegram | 111 | 0.18% |
| 10. Unofficial Arabic IS news site | 139 | 0.24% | 10. Taliban's official English website | 103 | 0.17% |

Interestingly, YouTube was the top linked-to platform for both Pro-IS and Other Jihadist accounts. This points to the overall importance of YouTube and of video to Web 2.0 in the jihadist online scene. Facebook does not appear in the Top 10 out-links for Pro-IS accounts. This indicates that, like Twitter, Facebook is also engaged in differential disruption as it is the second most preferred platform for out-linking by Other Jihadists. The somewhat obscure justpaste.it content upload site has been known for some time as a core node in the 'jihadisphere' and so it is

YouTube was the top linked-to platform for both Pro-IS and Other Jihadist accounts. This points to the overall importance of YouTube and of video to Web 2.0 in the jihadist online scene. Facebook does not appear in the Top 10 out-links for Pro-IS accounts.

unsurprising that it should appear as the third most linked-to site for both Pro-IS and Other Jihadist accounts.

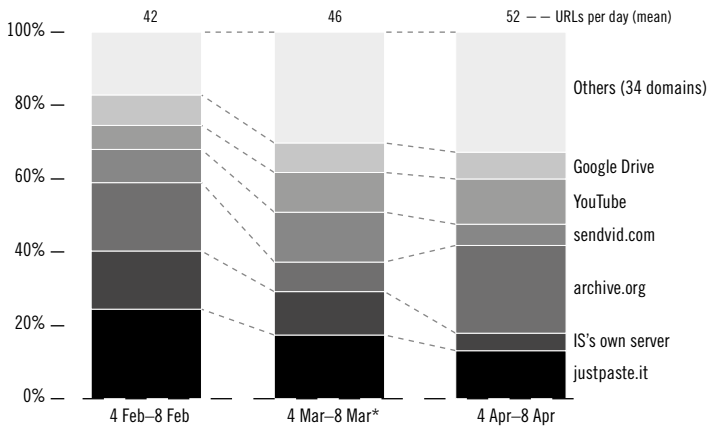
Other content upload destinations preferred by Pro-IS users, including Google Drive, Sendvid, Google Photos and the Web Archive, do not appear in the Other Jihadist Top 10. One reason for this is probably the focus of Other Jihadists on

linking to traditional proprietary websites, such as the Taliban's suite of sites. It is worth mentioning that, while Telegram slips into the Top 10 for Other Jihadists, only 20 (or 0.04%) of all tweets from Pro-IS accounts contained a telegram.me link. The paucity of such links caused us to explore further; we were surprised to find that just two of 722 Pro-IS users' biographies and two of 451 Other Jihadist users' biographies contained Telegram links. Neither group of accounts was using Twitter to advertise ways into Telegram.

5.1 CASE STUDY: DESTINATIONS OF OFFICIAL IS PROPAGANDA

As mentioned, when we undertook our research, IS was operating a 24-hour 'news cycle,' disseminating a daily batch of new official propaganda via social media channels, including Twitter. Links to the propaganda were circulated via tweets and other means. These links point to a wide variety of other social media platforms and content hosts that contained uploaded propaganda daily. We analysed a sample of these propaganda destinations at three time points: 4–8 February, 4–8 March (excluding 7 March, see below), and 4–8 April 2017. We obtained the full daily roster of IS propaganda and the sites where it appeared for each of these time periods. This allowed us to identify the most frequently linked-to platforms, along with how many pieces of propaganda were posted by host domains, and what proportion of these URLs were subsequently taken down (see Figure 6).

Figure 6. Destinations of official IS propaganda: Number of URLs and URL destinations February to April 2017



* Excludes 7 March which had 240 URLs (Rumiyah release)

Overall, over these three time periods, Pro-IS users linked to 39 different third-party platforms or sites hosting its propaganda material, as well as running its own server to host material.³⁰ It is important to note that the former were exclusively (we believe) ‘leaf’ destinations. That is, they contained content but no links to other sites, so did not have a networking or community-building aspect. Someone visiting such a page would learn nothing about the network of other sites. Important exceptions to this were YouTube and a small number of other sites which algorithmically ‘recommend’ similar content in their inventory, which may have resulted in their pointing to other available IS propaganda.³¹ During our analysis, the average number of URLs populated rose from 42 per day in February to 52 per day in April. This hints at increasing fragmentation and

30 This server had five names over the three periods studied because each domain name was rapidly taken down.

31 Derek O’Callaghan, Derek Greene, Maura Conway, Joe Carthy and Pádraig Cunningham, ‘Down the (White) Rabbit Hole: The Extreme Right and Online Recommender Systems.’ *Social Science Computer Review*, 2015, 33(4), pp.459-478.

dispersal, possibly in response to takedown activity by a variety of platforms and sites. However, there was a large inter-day variation (20 to 65) and we excluded one outlier day on 7 March, the publication date of issue 7 of *Rumiyah* magazine. On this day, IS pushed 240 separate URLs, a quarter of which contained direct reference to *Rumiyah* in the link, and many more which probably linked to the new issue of the magazine.

Out of the 40 domains used (39 external, one internal server) there was a consistent 'big 6' across the three time periods: justpaste.it; IS's own server; archive.org; sendvid.com; YouTube; and Google Drive. These six domains accounted for 83%, 70% and 67% of the URLs in the February, March and April sampling periods respectively. However, there was a noticeable declining trend in the use of justpaste.it and IS's own servers. Between them, this accounted for 40% of URLs in February declining to only 18% by April. Recently the Amaq News Agency website has come under repeated attack, which may be responsible for its relative downgrading.³² Use of sendvid.com and archive.org varied across the time periods, while Google Drive and YouTube were consistently heavily used; YouTube use showed an increasing trend (7%, 11% and 12%, respectively). The remaining URLs (17% in February rising to 33% of URLs by April) were spread across a wide variety of mainly, though not exclusively, content upload sites: 34 in total.

We also analysed what proportion of IS propaganda content had been taken down successfully. We found that the takedown rate (as of 12 April) was 72%, 66% and 72% for the February, March and April samples respectively. Overall, 30% of links were still live on 12 April. This suggests that takedown activity is relatively rapid (occurring over a matter of days after propaganda is posted) and widespread (across a multiplicity of sites and platforms).

32 Lizzie Dearden, 'ISIS Losing Ground in Online War Against Hackers After Westminster Attack Turns Focus on Internet Propaganda', *The Independent*, 1 April 2017.



6. RECOMMENDATIONS

MODERN SOCIAL MEDIA monitoring systems have the ability to dramatically increase the speed and effectiveness of data gathering, analysis and (potentially) intervention. To work effectively, however, they must deploy a combination of suitable technology solutions, including analytical systems, with trained human analysts who are versed in the domain deployed and preferably also the relevant languages. This is particularly the case where an adversary is actively trying to evade tracking efforts. Technology such as Method52 helps by allowing the analyst to rapidly develop new analytical pipelines that take into account day-to-day changes in modes of operation. However, technology cannot detect such changes; these can generally only be spotted by a human well-versed in the particular domain of interest.

Some IS supporters remain active on Twitter. Content disseminators using throwaway accounts could probably be degraded further – though this may have both pros (e.g. detrimental impact on last remaining significant IS supporter Twitter activity) and cons (e.g. further degradation of Twitter as a source of data or open source intelligence on IS). Like all disruption activity, whether this is viewed positively or negatively depends on one's perspective and institutional interests. For example, law enforcement tends to favour this approach, whereas free-speech advocates warn against corporate policing of political speech, even if that speech is deeply objectionable. Some intelligence professionals, on the other hand, advocate for greater attention to social media intelligence.³³

Our focus in this report has not just been on Twitter, but we also point to the importance of the wider jihadist social media ecology. Also, our analysis was not restricted to IS users and content; we underline, too, the presence and often uninterrupted online activity of non-IS jihadists. In recent years, many counter-terrorism professionals tasked with examining the role of the Internet in violent extremism and terrorism have narrowed their focus to IS.

33 David Omand, Jamie Bartlett and Carl Miller, 'Introducing Social Media Intelligence (SOCMINT)'. *Intelligence and National Security*, 27(6), 2012, pp. 801–823.

We recommend against continued analytical contraction. Instead, we point to the need to maintain a wide-angle view of online activity by diverse other jihadists across a variety of social media and other online platforms.

Scholarly researchers have acted similarly, many narrowing their focus further to IS Twitter activity. We recommend against continued analytical contraction. Instead, we point to the need to maintain a wide-angle view of online activity by diverse other jihadists across a variety of social media and other online plat-

forms. This is particularly important due to the shifting fortunes of IS and HTS on the ground in Iraq and Syria. In the face of increasing loss of physical territory, the continued – and potentially increasing – importance of online ‘territory’ should not be underestimated. We are not suggesting that a focus on IS should be dispensed with, but the significantly less-impeded online activity of HTS is surely an important asset for them and worth monitoring.

Because data collection and analysis of other terrorist groups and their online platforms has been neglected, very few historical metrics are available for comparative analyses. We should guard against this in future too.



7. FUTURE RESEARCH

AS MENTIONED ABOVE, our Other Jihadist category was a convenience sample of non-IS jihadist accounts. For future research, we therefore propose replicating the present research, but with a larger and more equal sample of HTS, Ahrar al-Sham and Taliban accounts. This would allow for a more systematic and comparative analysis of the disruption levels for a range of non-IS jihadists, including those with a significant international terrorism footprint (i.e. HTS), groups with a significant national and regional terrorism profile (i.e. Taliban), and a party to the Syria conflict (i.e. Ahrar al-Sham).³⁴ Such an analysis could help to ascertain the vibrancy of their contemporary Twitter communities and Twitter out-linking practices, and allow us to identify their preferred other online platforms.

Additional research is clearly warranted into the wider violent jihadist social media ecology. We therefore recommend wider and more in-depth research into:

1. patterns of use, including community-building and influencing activity; and
2. levels of disruption on other platforms besides Twitter, including other major platforms such as YouTube, but also other smaller or more obscure platforms, such as justpaste.it and others.

We also suggest analysing pro-IS and other jihadist activity on Telegram, which is almost certainly where the IS online community has reconstituted, and comparing this with our present findings. It would also be worthwhile analysing out-linking trends on Telegram to see if different platforms have an impact on the effectiveness of linking practice.

34 Nationally, Syria, Russia, Iran, Egypt, and the UAE have designated Ahrar al-Sham as a terrorist organisation. Internationally, the US, Britain, France, and Ukraine blocked a May 2016 Russian proposal to the United Nations to take a similar step.



8. CONCLUSION

OUR DATA SHOWED that the costs for most pro-IS users of engaging on Twitter (in terms of deflated morale, diffused messages and persistent effort needed to maintain a public presence) now largely outweigh the benefits. This means that the IS Twitter community is now almost non-existent. In turn, this means that radicalisation, recruitment and attack planning opportunities on this platform have probably also decreased. However, a hard core of users remain persistent. In particular, a subset of established throwaway disseminator accounts pushed out 'official' IS content in a daily cycle during our data-collection period. These accounts were generally suspended within 24 hours, but not before they promoted links to content hosted on other platforms. This included major new content, such as a new issue of the monthly IS *Rumiyah* magazine.

This report is mainly concerned with pro-IS Twitter accounts and their disruption. However, IS are not the only jihadists active on Twitter, and a host of other violent jihadists were shown to be subject to much lower levels of disruption by Twitter. Also, IS and other jihadist groups remain active on a wide range of other social media platforms, content hosting sites and other cyberspaces, including blogs, forums, and dedicated websites. While it appears that official IS content is being disrupted in many of these online spaces, the extent is yet to be fully determined.

The VOX-Pol Network of Excellence (NoE) is a European Union Framework Programme 7 (FP7)-funded academic research network focused on researching the prevalence, contours, functions, and impacts of Violent Online Political Extremism and responses to it.



This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no. 312827

Email info@voxpath.eu
Twitter [@VOX_Pol](https://twitter.com/VOX_Pol)
www.voxpol.eu

