

The Hague, 17/04/2018

Who disseminates *Rumiyah*?

Examining the relative influence of sympathiser and non-sympathiser Twitter users

This paper was presented at the 2nd European Counter Terrorism Centre (ECTC) Advisory Group conference, 17-18 April 2018, at Europol Headquarters, The Hague.

The views expressed are the authors' own and do not necessarily represent those of Europol.

Authors:

Daniel Grinnell (Cardiff University), Stuart Macdonald (Swansea University),
David Mair (Swansea University) & Nuria Lorenzo-Dus (Swansea University)



1 Introduction

In a speech delivered at the United Nations General Assembly in September 2017, the U.K. Prime Minister Theresa May called on social media companies to do more to remove and block terrorist content from their platforms [1]. In the speech she stated that the average lifespan of online propaganda from the so-called Islamic State (IS) was 36 hours. For such content to be disrupted effectively, she claimed that this figure needed to be reduced to one to two hours. This has since come to be known as the ‘golden window’: if terrorist material can be detected and removed within one to two hours, its spread will be prevented.¹

Recent research by Conway et al. found that IS and its supporters are already being significantly disrupted by Twitter suspension activity, and now struggle 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’ [2]. This is consistent with the findings of our own pilot study, which we presented to this conference last year [3]. In our pilot study we examined the release of issue 15 of IS’s online magazine *Dabiq* on Twitter. The accounts within our dataset that had been suspended were largely very young accounts with relatively low numbers of followers and were suspended before they gained much of a following. In other words, they seemed to be throwaway accounts, whose purpose was to disseminate that particular issue of *Dabiq* magazine. We concluded our pilot study by pointing instead to a different challenge. The number of throwaway accounts that were created to disseminate issue 15 of *Dabiq* was relatively small. Far greater was the number of other accounts, not sympathetic to IS, that posted links to the new issue or to discussion of it. Moreover, these accounts were not being suspended. It was these accounts, we felt, that were maintaining the presence and discoverability of *Dabiq* issue 15 on Twitter. We thus concluded that ‘the IS sympathisers who disseminated issue 15 of *Dabiq* caused a fairly small splash; it was others that caused the ripples to travel a long way’ [3].

The objective of this paper is to test whether the findings of our pilot study hold true for a larger dataset – or whether the dynamic by which issue 15 of *Dabiq* was shared on Twitter was attributable more to other factors, such as the theme of that particular issue. In this study we, therefore, examine the release on Twitter of a total of nine issues

¹ Notice that this is subtly different to the guidelines produced by the European Commission. The Commission’s guidelines call on social media companies to remove terrorist content within an hour of it being referred (‘A Europe that protects: Commission reinforces EU response to illegal content online’ *European Commission Press Release*, 1 March 2018). Prime Minister May’s statement refers to one to two hours after the content is posted.

of IS's online magazine *Rumiyah*, the successor to the now-abandoned *Dabiq*. In pursuit of this objective, we have two sets of research questions. First, our pilot study suggested that pro-IS throwaway accounts were only creating a small splash on Twitter. Has this been the case for *Rumiyah*? Are these pro-IS accounts being disrupted effectively, before they manage to exert much influence? Second, what is the relative influence of the pro-IS throwaway accounts in comparison to other accounts that are not sympathetic to IS, but nonetheless disseminate its propaganda – whether that be for research purposes, personal interest or even to provide an oppositional voice or engage in debate? In other words, how great is the ripple effect generated by these non-IS sympathiser accounts?

2 Methodology

In this section we outline how we collected data, provide an overview of our dataset and explain the focus of this study, and then finally describe how we approached the task of data analysis.

2.1 Data collection

Data was collected using Cardiff University's 'Sentinel' research tool [4]. The data collection period was 1 November 2016 to 31 October 2017. During this period, all Twitter posts were collected that: (1) mentioned the term 'Rumiyah'; (2) were posted within 21 days of the release of a new issue; and, (3) were posted from an account that used the English language interface (U.S. or U.K.). The last of these reflected our decision to focus specifically on users that posted about the English language version of *Rumiyah* (which is also published in multiple other languages). Given the study's focus on the characteristics of those who post about *Rumiyah*, we also collected the public openly available user data of these posters, the details of the first post by each user following the release of a new issue, the onward distribution counts of those posts, and the account status (at the end of the data collection period).²

As Table 1 shows, during the 12-month data collection period a total of 11 issues of *Rumiyah* were published (issues 3 to 13). As a result of collection drop outs, the data collection for issues six and eight was incomplete and so these issues have been excluded from the study. Our dataset thus encompasses a total of nine issues.

² For the purposes of this study, Sentinel functioned only as a repository of structured data supplied by the Twitter Streaming API.

Table 1: Original linkers by account type

Issue	Date and time of first tweet collected	Date and time of last tweet collected	Notes
3	11/11/2016 17:18	02/12/2016 17:18	
4	07/12/2016 19:33	28/12/2016 19:33	
5	06/01/2017 16:57	27/01/2017 16:57	
6	04/02/2017 18:53	25/02/2017 18:53	Incomplete collection so not included in the dataset
7	07/03/2017 16:04	28/03/2017 16:04	
8	05/04/2017 15:23	26/04/2017 15:23	Incomplete collection so not included in the dataset
9	04/05/2017 14:40	25/05/2017 14:40	
10	08/06/2017 21:32	29/06/2017 21:32	Timing moderately uncertain due to similarly timed presence of a “fake issue”
11	13/07/2017 15:00	03/08/2017 15:00	Timing highly uncertain due to hashtag flooding and a “fake issue”
12	06/08/2017 14:47	27/08/2017 14:47	
13	09/09/2017 18:58	30/09/2017 18:58	

2.2 Overview of dataset

There was a total of 9 968 distinct users that posted about one (or more) of the nine issues of *Rumiyah* that we examined. This is an average of 1 108 users per issue, which is noteworthy given that in our previous study a total of 11 586 distinct users posted about a single issue of *Dabiq* [3].

Table 2 breaks the 9 968 users down according to: (1) whether the user account was extant, suspended or self-deleted by the end of our data collection period; and, (2) the number of issues that the user posted about. It shows, first, that over a quarter of the accounts (26.2%) were suspended and, unsurprisingly, that almost all of the accounts that were ultimately suspended only posted about one issue of *Rumiyah* (2 543 out of the 2 608 suspended accounts, i.e., 97.5%). (Though it should be acknowledged that accounts could have been suspended as a result of other Twitter activity and not because of the posts mentioning *Rumiyah*). More generally, Table 2 shows that the vast majority (90.3%) of users only posted about a single issue, and less than 2% posted about four issues or more.

Table 2: Number of issues each user posted about (within the 21-day window)

Number of issues	Extant accounts	Suspended accounts	Deleted accounts	All users
1	6209 (62.3%)	2543 (25.5%)	252 (2.5%)	9004 (90.3%)
2	538 (5.4%)	45 (0.5%)	21 (0.2%)	604 (6.1%)
3	154 (1.5%)	14 (0.1%)	2 (0.0%)	170 (1.7%)
4	76 (0.8%)	3 (0.0%)	1 (0.0%)	80 (0.8%)
5	44 (0.4%)	2 (0.0%)	0	46 (0.5%)
6	24 (0.2%)	0	0	24 (0.2%)
7	17 (0.2%)	0	0	17 (0.2%)
8	13 (0.1%)	1 (0.0%)	1 (0.0%)	15 (0.2%)
9	8 (0.1%)	0	0	8 (0.1%)
Total	7083 (71.1%)	2608 (26.2%)	277 (2.8%)	9968 (100%)

2.3 Focus of this study

Given the nature of Twitter as a gateway platform [5], coupled with IS's resort to the use of throwaway accounts, we chose to focus our analysis on out-linking posts, i.e., posts containing a link to an external site (be that to the magazine itself, extracts or excerpts from it, or news or commentary about the magazine). We also chose to focus on original posts, i.e., posts that were not a repost, link to or quote of any other post. In short, our analysis focuses on 'original linkers': users who independently sought to direct other users to content outside of the Twitter platform (including, but not limited to, copies of *Rumiyah* itself).

Table 3 contains the same breakdown as Table 2, but is limited to just original linkers. As with the dataset as a whole, the vast majority (91.4%) of original linkers only posted about a single issue of *Rumiyah*. What is striking, however, is that a far higher proportion (71.4%) of original linkers had been suspended by the end of our data collection period. Only very rarely did an account that went on to be suspended post about more than one issue (11 out of 1006 suspended accounts, i.e., 1.1%).

Table 3: Number of issues each original linker posted about (within the 21-day window)

Number of issues	Extant accounts	Suspended accounts	Deleted accounts	All users
1	279 (19.8%)	995 (70.6%)	14 (1.0%)	1288 (91.4%)
2	51 (3.6%)	6 (0.4%)	2 (0.1%)	59 (4.2%)
3	12 (0.9%)	2 (0.1%)	0	14 (1.0%)
4	20 (1.4%)	1 (0.1%)	0	21 (1.5%)
5	7 (0.5%)	1 (0.1%)	0	8 (0.6%)
6	6 (0.4%)	0	0	6 (0.4%)
7	4 (0.3%)	0	0	4 (0.3%)
8	5 (0.4%)	1 (0.1%)	1 (0.1%)	7 (0.5%)
9	2 (0.1%)	0	0	2 (0.1%)
Total	386 (27.4%)	1006 (71.4%)	17 (1.2%)	1409 (100%)

Given the low number of original linkers that subsequently chose to close their account, in this paper we focus on those whose account was still in existence at the end of our data collection period ('extant original linkers', n=386) and those whose account was suspended by this date ('suspended original linkers', n=1006).

2.4 Data analysis

Using the data collected by Sentinel, we sought to identify and compare the features of 'extant original linkers' and 'suspended original linkers'. This quantitative analysis focused on: the size of each sub-group's social network; the age of the accounts; the number of times each sub-group was retweeted; the number of repeat posts from each sub-group; and, the external sites to which each sub-group out-linked. We supplemented this by identifying the language in which each tweet the two sub-groups posted was written. The findings follow below.

In addition to this quantitative analysis, we also conducted some qualitative analysis. The qualitative analysis focused on two features. First, the account type: whether it was the account of someone posting in a personal capacity, the account of someone posting in their capacity as an intelligence analyst or practitioner, the account of a news or media organisation, the account of an identified group, the account of someone posting in their capacity as an academic or researcher, the account of someone posting in their capacity as a journalist, or some other type of account (including where it was not

possible to classify it). Second, the combined tone of the account and post: whether it was sympathetic to the message of *Rumiyah*, overtly critical of it, or only sought to communicate factual material that was neither sympathetic nor critical. To classify users for each of these two features, we examined the account username, the account profile and the content of the post. Google Translate was used to translate non-English language account profiles and posts into English. In some cases, Google Translate was unable to detect the original language, meaning that translation was not possible and these accounts could not be included in this part of the analysis.

3 Findings

In this section we present the findings of our analysis. We begin with the quantitative findings before moving on to the qualitative ones.

3.1 Follower/following network and account age

First, we examine the follower/following count of the original linkers. Figure 1 shows the percentile distribution of the total number of users that the original linkers follow.

Percentile distribution graphs indicate the proportion of the sample which have a certain value. So, if two curves rise to the same value (X) but at different percentiles (say 10% and 20%), the one that rises to value X at the lower percentile has a greater proportion of users at that value or higher (in our example, the first curve shows that 90% of users were at the value X or higher, compared to 80% for the second curve). Note also that: (1) in order to represent the data in its entirety, a logarithmic scale is used on the vertical axis in all figures; and, (2) where no value is present for a percentile it should be considered zero.

In addition to the blue line (which shows the distribution for the extant original linkers) and the red line (which shows the distribution for the suspended original linkers), for the sake of comparison Figure 1 also shows the distribution for the dataset as a whole (green line). As Figure 1 shows, up until the 99th percentile the number of users that extant original linkers were following was greater than for the dataset as a whole. What is most striking, however, is the fact that over half (53.8%) of the suspended original linkers were following no other users at all. Only a small proportion of suspended original linkers (~5%) had a large following count (100+).




Figure 1: Percentile distribution of the number of users the original linkers follow

Figure 2 shows the percentile distribution of the total number of followers that the original linkers had.




Figure 2: Percentile distribution of the number of users that follow the original linkers

Figure 2 follows a similar pattern to Figure 1. Up until the 99th percentile the extant original linkers had a greater number of followers than for the dataset as a whole. Most striking, however, is the fact that over half (55.5%) of the suspended original linkers had no followers at all, whilst an additional 11.3% had just one follower. And, as with the following count, only a small proportion of suspended original linkers (~5%) had a large follower count (100+).

In addition to the follower/following network, we also examined the age of the user accounts. Figure 3 shows the percentile distribution.

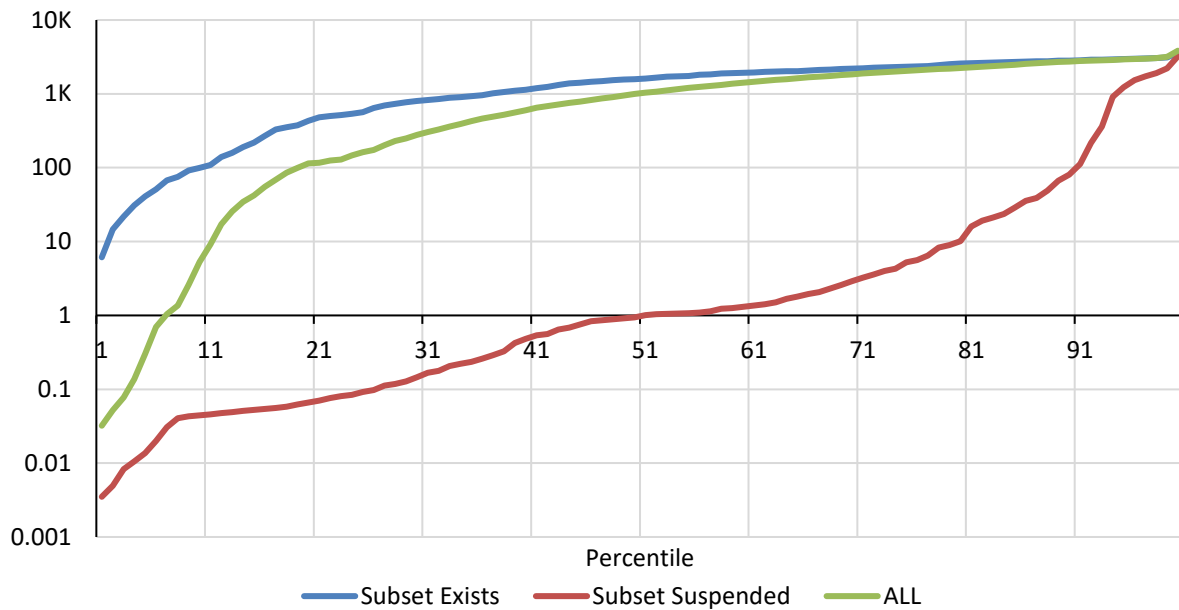


Figure 3: Percentile distribution of account age (in days) of original linkers

As with the follower/following count, Figure 3 shows a stark difference between extant original linkers and suspended original linkers. The extant original linkers were typically older accounts. By contrast, the vast majority (86.0%) of the suspended original linkers' accounts were less than 28 days old, and roughly half (50.8%) were less than one day old. Of the accounts that were less than one day old, 60% were less than three hours old and 20% were less than one hour old. Our findings are thus consistent with Conway et al.'s statement that IS Twitter activity has largely been reduced to the tactical use of throwaway accounts. The fact that most suspended original linker accounts were very young and had only a small – or non-existent – social network suggests that they were set up specifically to advertise the release of the new issue, particularly when coupled with the additional fact that 14.5% of suspended original linkers had never posted before (and a further 32.4% had posted only four or fewer tweets previously).

3.2 Retweets and repeat postings

Our next set of findings examine the number of times the original linkers' first post was retweeted and the number of tweets the original linkers posted that mentioned *Rumiyah*. Figure 4 shows the percentile distribution of the number of times the original linkers' first post was retweeted.




Figure 4: Percentile distribution of the number of times the first post of an original linker was retweeted

For both extant and suspended original linkers, the majority of first posts were not retweeted (76.4% and 89.6% respectively). As Figure 4 shows, when retweeting did occur the frequency was greater for extant original linkers. For this sub-group, the 99th percentile was 16 retweets and the highest (outlier) value was 220 retweets. By contrast, the 99th percentile for suspended original linkers was just four retweets, and the highest (outlier) value was ten. With few or no followers and lack of retweeting, the overall visibility of the suspended original linker accounts was low. These accounts thus seem to have relied upon being directly searched for, before they were suspended.

Turning from retweets to repeat posting, Figure 5 examines the number of times original linkers posted tweets that mentioned a new issue of *Rumiyah* within 21 days of its release.




Figure 5: Percentile distribution of the number of times original linkers posted a tweet mentioning *Rumiyah* within 21 days of its release

Beginning with the extant original linkers, only a small proportion posted *Rumiyah* mentioning tweets in high volumes within the 21 days following the release of a new issue. As Figure 5 shows, 5.2% of these users posted over 100 times, with the highest figure being 150 times. This may be contrasted with the figures for suspended original linkers. Roughly 70% of these users posted more than once, 40% posted more than five times, and 10% posted more than ten times prior to suspension. This repeat posting, we suggest, reflects an attempt to increase the likelihood of being seen, particularly given the lack of retweets.

In terms of the findings for the dataset as a whole, it is worth noting that the highest number of posts mentioning *Rumiyah* by one account within the 21 day window was 9 547. This appears to us to have been something like a chaff account. It began posting at around the time that issue ten of *Rumiyah* was released. Its tweets contained the hashtag “#Rumiyah”, followed by a random string of characters, then a link to another of its own posts. It continued posting until the day on which issue eleven of *Rumiyah* was released, when it was suspended. Whilst the intention behind the creation of the account cannot be stated with any certainty, the effect of the tweets it posted was to make it more difficult for someone using the hashtag “#Rumiyah” to search for a link to an e-copy of the magazine to find one.

3.3 Out-links

Across the data for all suspended and extant original linkers, we identified a total of 164 unique URL hostnames. Table 4 below shows the websites that were out-linked to most frequently by each group of original linker. (Note that the table only includes those websites that were out-linked to ten times or more).

Table 4: External website hostnames out-linked to by original linkers >10 times			
Suspended original linkers	Total	Extant original linkers	Total
drive.google.com	235	fb.me	39
archive.org	185	dlvr.it	35
cldup.com	166	bit.ly	32
pc.cd	96	ref.gl	30
cloud.mail.ru	81	counterjihadreport.com	24
yadi.sk	64	icct.nl	21
justpaste.it	30	memri.org	16
dropbox.com	27	ift.tt	16
mediafire.com	26	heavy.com	15
1drv.ms	24	clarionproject.org	15
goo.gl	12	terrortrendsbulletin.com	14
		express.co.uk	12

As Table 4 shows, suspended original linkers generally out-linked to websites where content can be uploaded, viewed and downloaded. This would suggest that this group of original linkers was seeking to disseminate e-copies of *Rumiyah* itself, or associated content. As for extant original linkers, whilst we cannot say with any certainty what was shared via the out-links to dlvr.it, bit.ly, ref.gl or ift.tt, it is nonetheless clear that these original linkers out-linked to a more diverse set of websites that included other social media (Facebook), sites offering commentary and analysis (International Centre for Counter Terrorism, MEMRI, Clarion Project), blogs and activist movements (Counter Jihad Report, Terror Trends Bulletin) and news (Heavy, Daily Express).

3.4 Languages in which posts were written

Table 5 shows the language in which the tweet posted by the original linker was written.³

Table 5: Language in which post was written			
Language	Extant original linkers	Suspended original linkers	Total
English	410 (27.0%)	741 (48.8%)	1151 (75.9%)
Arabic	5 (0.3%)	185 (12.2%)	190 (12.5%)
Italian	36 (2.4%)	1 (0.1%)	37 (2.4%)
German	7 (0.5%)	18 (1.2%)	25 (1.6%)
French	16 (1.1%)	2 (0.1%)	18 (1.2%)
Turkish	1 (0.1%)	14 (0.9%)	15 (1.0%)
Croatian	0 (0.0%)	12 (0.8%)	12 (0.8%)
Sindhi	0 (0.0%)	11 (0.7%)	11 (0.7%)
Urdu	0 (0.0%)	10 (0.7%)	10 (0.7%)
Other (<10) / Multiple	24 (1.6%)	24 (1.6%)	48 (3.2%)
Total	499 (32.9%)	1018 (67.1%)	1517 (100%)

By far the most common language (for both extant and suspended original linkers) was English, with three-quarters (75.9%) of the tweets posted by these accounts written in the English language. This is unsurprising, given our decision to focus on posts from accounts that used the English language interface. What is noteworthy, however, is the number of tweets posted by suspended original linkers that were written in the Arabic language (185; 12.2% of all original linker posts). This is particularly striking given that the accounts that posted the tweets were all ones that used the English language

³ For the sake of completeness, two further points should be noted: (1) some user accounts may feature in Tables 5, 6 and 7 more than once. For example, if user X's first *Rumiyah* mentioning post in the collection window for both issues 3 and 13 was an original content tweet containing an out-link, then each of these posts from user X will be included in the count in Tables 5, 6 and 7; and, (2) the totals for Tables 5, 6 and 7 are different to the total for Table 3, because a different counting method was employed for Table 3. Unlike Tables 5, 6 and 7, Table 3 includes some posts that did not contain original content and/or did not contain an out-link. For example, if user Y's first post mentioning *Rumiyah* in the collection window for issue 3 was an original out-linking post, then the count for Table 3 will include Y's first *Rumiyah* mentioning post in any of the other collection windows even if the post contained no original content and/or no out-link.

interface. Our understanding is that, during the creation of a Twitter account, the language setting is set to a default based on a number of factors including locale and machine configuration. This suggests that the accounts that posted the Arabic-language tweets were set up by users who are able to understand or use written English. More speculatively, it may also suggest that the accounts were either created in a location where English is the first language or that the account language was purposefully set to English by the user (perhaps either in an attempt to make detection more difficult or because the account was created automatically as part of a botnet).

3.5 Account types

Table 6 shows the breakdown of extant and suspended original linkers by account type.

Account type	Extant original linkers	Suspended original linkers	Total
Other	82 (5.4%)	869 (57.3%)	951 (62.7%)
Personal	215 (14.2%)	125 (8.2%)	340 (22.4%)
Intelligence analyst/practitioner	86 (5.7%)	0	86 (5.7%)
News/media organisation	60 (4.0%)	18 (1.2%)	78 (5.1%)
Group	22 (1.5%)	1 (0.1%)	23 (1.5%)
Academic/researcher	20 (1.3%)	0	20 (1.3%)
Journalist	14 (0.9%)	5 (0.3%)	19 (1.3%)
Total	499 (32.9%)	1018 (67.1%)	1517 (100%)

A large proportion of the suspended original linkers fell within the “Other” category (869 of 1018: 85.4%). This category included accounts whose purpose could not be expressly identified, normally because the account had no profile and/or the username was made up of randomised alphanumeric characters. Also contained within this account type were suspected bot accounts whose activities appeared to include initial dissemination of new issue of the magazine and hashtag-hijacking campaigns designed to push the release of a new issue into as many newsfeeds as possible. This is consistent with other studies that have shown that IS relies heavily on the use of bots to disseminate its content outside of its more insular networks [6], [7].

The next most common account type was personal accounts. A total of 215 extant original linkers and 125 suspended original linkers were identified as personal accounts. Within the latter, a number of accounts claimed to be IS supporters returning to Twitter following the suspension of a previous account. Other research has also noted this “ritual” of suspension and return, followed by shout-outs [8].

3.6 Tone of posts

Table 7 shows the breakdown of extant and suspended original linkers by the tone of their post.

Account type	Extant original linkers	Suspended original linkers	Total
Sympathetic	4 (0.3%)	975 (64.3%)	979 (64.5%)
Critical	245 (16.2%)	21 (1.4%)	266 (17.5%)
Factual	190 (12.5%)	9 (0.6%)	199 (13.1%)
Other	60 (4.0%)	13 (0.9%)	73 (4.8%)
Total	499 (32.9%)	1018 (67.1%)	1517 (100%)

As Table 7 shows, the majority of original linkers (979 of 1517: 64.5%) were sympathetic to the message of *Rumiyah*. Of the 979 sympathetic original linkers, only four had not been suspended by the end of our data collection period. This supports the findings of previous research that IS sympathiser accounts now experience significant disruption on the platform [2].

It is also noteworthy that the vast majority of the suspended original linkers (975 of 1018: 95.8%) were found to be sympathetic to the message of *Rumiyah*. In respect of the other 43 suspended original linkers, it should be reiterated that the reason for their suspension may not have been their *Rumiyah* mentioning post; our dataset includes the fact of their suspension but not the reason for it. So, as in our pilot study [3], suspension activity appears to be limited to those users that express sympathy or support for IS.

4 Conclusion

To return to the two sets of research questions that we identified at the outset, we found that most of the suspended original linkers (i.e. user accounts that posted original content, including an out-link, and were suspended by the end of our data collection period) were sympathetic to the message of *Rumiyah*. The usernames of these accounts were commonly randomised collections of letters and numbers, the out-links they shared were generally to content-sharing websites (presumably containing e-copies of the magazine or excerpts from it), and whilst most posted in the English language there was also a significant proportion that tweeted in Arabic. Importantly, as in our pilot study, these accounts caused only a small splash: they were normally young accounts (often less than one day old), with very few followers (or none at all) and received few retweets. They thus appeared to be throwaway accounts, that sought to compensate for their lack of visibility by repeat posting.

The other sub-group that we have examined, extant original linkers (i.e. user accounts that posted original content, including an out-link, that were still extant at the end of our data collection period), had quite a different overall profile. They mostly posted in the English language and tended to be the accounts of private individuals, intelligence analysts or practitioners, or news/media organisations. They were generally older accounts, with higher numbers of followers, and were less likely to engage in repeat posting. Their tweets were generally either factual in tone or critical of IS, and they out-linked to a more diverse range of sources, including news reports, commentary and analysis, and activist sites and blogs.

Perhaps most significant, however, is the fact that the suspended original linkers outnumbered the extant original linkers by more than two-to-one. This takes us back to our initial objective, to examine whether the findings of our pilot study of *Dabiq* issue 15 hold true for a larger dataset. As will be clear by now, there were some significant differences between the findings of our pilot study and the findings we have presented here. For a start, more distinct user accounts tweeted about issue 15 of *Dabiq* than all nine of the issues of *Rumiyah* we studied combined. On top of this, whereas in this study 71.4% of original linkers went on to be suspended by the end of our data collection period, in our pilot study 84.5% of original linkers remained extant at the end of the data collection period. Together, these findings suggest not merely that more users tweeted about *Dabiq* issue 15, but that this greater level of interest was fuelled by non-IS sympathisers. In other words, the nine issues of *Rumiyah* that we studied failed to capture the same level of attention from non-IS sympathisers as *Dabiq* issue 15. Various explanations might be offered for this. Perhaps it was due to IS's territorial losses. Perhaps it was because there was greater brand awareness of *Dabiq* than *Rumiyah*. Or

perhaps it was because issue 15 of *Dabiq* was particularly successful in targeting a particular sub-audience and provoking a response from them. Whilst we can only speculate as to the relative causal role of these different explanations, we conclude by suggesting that further research that investigates the identifying features of terrorist propaganda that is particularly successful in provoking a response from the audiences it targets would make a valuable contribution to efforts to disrupt the spread of such propaganda and promote alternative narratives.

References

- [1] 'Preventing terrorist use of the internet', speech delivered by Prime Minister Theresa May at the United Nations General Assembly, 20 September 2017. <https://www.gov.uk/government/speeches/pm-speech-at-unga-preventing-terrorist-use-of-the-internet>
- [2] Conway, M., Khawaja, M., Lakhani, S., Reffin, J., Roberston, A. & Weir, D. (2017). *Disrupting Daesh: Measuring Takedown of Online Terrorist Material and its Impacts*. VOX-Pol Network of Excellence. http://www.voxpol.eu/download/vox-pol_publication/DCUJ5528-Disrupting-DAESH-1706-WEB-v2.pdf
- [3] Grinnell, D., Macdonald, S. & Mair, D. (2017). *The Response of, and on, Twitter, to the Release of Dabiq Issue 15*. Paper presented at the 1st European Counter Terrorism Centre conference on online terrorist propaganda, 10-11 April 2017, The Hague. <https://www.europol.europa.eu/publications-documents/response-of-and-twitter-to-release-of-dabiq-issue-15>
- [4] Preece, A., Spasic, I., Evans, K., Rogers, D., Webberley, W., Roberts, C., & Innes, M. (2018). 'Sentinel: A Codesigned Platform for Semantic Enrichment of Social Media Streams'. *IEEE Transactions on Computational Social Systems*, 5(1), 118–131. <https://doi.org/10.1109/TCSS.2017.2763684>
- [5] O'Callaghan D., Greene D., Conway M., Carthy J. & Cunningham P. (2013). '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*, 276-285.
- [6] Berger, J.M. & Morgan, J. (2015). *The ISIS Twitter Census: Defining and Describing the Population of ISIS Supporters on Twitter*. The Brookings Institution. https://www.brookings.edu/wp-content/uploads/2016/06/isis_twitter_census_berger_morgan.pdf
- [7] Berger, J.M. & Perez, H. (2016). *The Islamic State's Diminishing Returns on Twitter: How Suspensions are Limiting the Social Networks of English-speaking ISIS Supporters*. George Washington University Program on Extremism. https://cchs.gwu.edu/sites/g/files/zaxdzs2371/f/downloads/Berger_Occasional%20Paper.pdf
- [8] Pearson, E. (2017). 'Online as the New Frontline: Affect, Gender, and ISIS-Take-Down on Social Media'. *Studies in Conflict & Terrorism*. <https://doi.org/10.1080/1057610X.2017.1352280>