

## **Kingdom of Trolls? Influence Operations in the Saudi Twittersphere**

CHRISTOPHER BARRIE

University of Edinburgh, UK

ALEXANDRA A. SIEGEL

University of Colorado at Boulder, USA

Saudi Arabia has one of the highest rates of Twitter penetration in the world. Despite high levels of repression, the platform is frequently used to discuss political topics. Recent disclosures from Twitter have revealed state-backed attempts at distorting the online information environment through influence operations (IOs). A growing body of research has investigated online disinformation and foreign-sponsored IOs in the English-speaking world; but comparatively little is known about online disinformation in other contexts or about the domestic use of IOs. Using public releases of IO tweets, we investigate the extent of such activity in Saudi Arabia. Benchmarking these tweets to four samples of Saudi Twitter users, we find that engagement with IO accounts was lower than engagement with the average user, but equal to engagement with news accounts. Network analysis reveals that engagement with IO accounts was largely driven by a small number of influential accounts.

*Keywords:* *influence operations, disinformation, Twitter, Saudi Arabia*

---

Christopher Barrie: [christopher.barrie@ed.ac.uk](mailto:christopher.barrie@ed.ac.uk)

Alexandra A. Siegel: [alexandra.siegel@colorado.edu](mailto:alexandra.siegel@colorado.edu)

Date submitted: 2020-12-23

From September 2019 to April 2020 Twitter announced three public disclosures of now-suspended Twitter accounts attributed to organized state-backed “information operations” originating in Saudi Arabia (Twitter, 2020).<sup>1</sup> Included in the first of these releases was the official account of Saud al-Qahtani, a former associate of Crown Prince Mohammed bin Salman. Al-Qahtani had been leading an organized effort aimed at distorting the information environment to the benefit of the Saudi regime (Leber and Abrahams, 2019). Nicknamed the “Lord of the Flies” in reference to his army of “electronic flies”—paid trolls who persistently attacked online critics—he was also implicated in the murder of journalist Jamal Khashoggi in 2018 (Benner et al., 2018; b33lz3bub, 2019; Ignatius, 2018). In November 2019, the US Justice Department then charged two former Twitter employees with spying for the Saudi state after they shared the personal data of online critics.

This series of revelations led some to speculate that Twitter has become “more useful to the repressors than the repressed” in countries like Saudi Arabia (Post, 2019). Given the brutality of the Saudi regime offline, its history of repressing online critics (Pan and Siegel, 2020), and the apparent extent of its reach online, this conclusion is understandable. However, we lack empirical evidence regarding the scale of these domestically deployed influence operations. In this article, we quantify domestic engagement with online influence operations in Saudi Arabia. Our goal is to contribute to recent efforts aimed at providing descriptive analysis of the scale of online disinformation (Allen et al., 2020; Allcott et al., 2019).

We employ multiple descriptive strategies to gauge the extent of these operations and any traces of their influence on individual users. We draw on five datasets of tweets including tweets from IO accounts, tweets produced by a random sample of Saudi users, tweets from geolocated Saudi users, tweets from fifteen major Saudi news organizations, and tweets produced by the fifty most followed users in Saudi Arabia. We benchmark observed engagement with IO tweets against engagement with our samples of ordinary Saudi Twitter users and high-profile Saudi accounts. We then estimate network positions of IO users and estimate their influence by analyzing mentions of IO accounts over time.

Overall, we find that although domestic engagement with IO tweets was approximately

---

<sup>1</sup>Twitter refers to these as “information operations”; we refer in what follows to “influence operations,” which is the more common designation in the literature.

equal to engagement with the top fifteen Saudi news outlets, it was markedly lower than engagement with tweets in our representative samples, and dramatically lower than engagement with the fifty most influential accounts in the Saudi Twittersphere. Engagement with IO accounts was also relatively lower on controversial topics such as the murder of Khashoggi, the Qatar crisis, and other foreign and domestic political topics. In our representative samples of Saudi Twitter users, IO accounts were mentioned in  $\sim.05$ – $.10\%$  of all tweets, and in  $\sim.25$ – $.50\%$  of all mentions. Our network analysis suggests that engagement with IO accounts was largely driven by engagement with a few prominent accounts. The vast majority of IO accounts saw very little engagement during the period under study.

### Quantifying Online Influence Operations

Online disinformation encompasses a broad range of activities including fake news websites, hired and independent trolls, hyper-partisan news media, and online influence efforts (Siegel, 2018). We focus on state-backed influence operations as a distinct sub-type of online disinformation. We adopt the definition of Martin et al. (2020) for domestic online influence and influence operations as “coordinated campaigns by a state to impact one or more specific aspects of domestic politics, through media channels, including social media, by producing content designed to appear as though it is produced by normal users.” The term “influence operation” might imply a specific political goal while “information operation” better captures the broader goal of distorting the online information environment. Both terms have been used interchangeably in the literature, we opt for the catch-all “IO” when referring to this activity.

Existing research into online IO campaigns takes the IO as the unit of analysis, finding that since at least 2011, online “influence operations,” “information operations” or forms of “computational propaganda” have been steadily increasing (Bradshaw and Howard, 2018, 2019; Martin et al., 2019, 2020). Here, Russia, stands out as a key sponsor of IO activity, alongside China, Iran, Saudi Arabia, and the UAE (Martin et al., 2020). Twitter and Facebook remain the principal platforms for the dissemination of IO content, and there is evidence that states are increasingly reliant on out-sourcing IO campaigns to ostensibly independent media companies (Martin et al., 2020; Bradshaw and Howard, 2019). These cross-national reviews provide valuable information on the extent and nature of IO activity as a phenomenon but are silent on the scale of IO activity relative to total content on a given platform in each country.

Other research investigates the effect of IO activity or “fake news” *outside* the country of origin.<sup>2</sup> Some have argued that consumption of false or malicious information affected voting outcomes in both the US and UK (Gorodnichenko et al., 2018; Gunther et al., 2019). However, a growing body of research has found reason to be skeptical of these conclusions. A first cause for skepticism is that engagement with false or unreliable information promoted by coordinated IO campaigns may not actually affect behaviour or attitudes (Guess et al., 2020; Bail et al., 2020). The second cause for skepticism is that fake news often represents only a tiny proportion of overall information consumption (Guess et al., 2019; Allen et al., 2020). To date, the overwhelming majority of empirical work investigating these questions has focused on the US and foreign IO campaigns—a limitation acknowledged by these authors (Bail et al., 2020). Whether or not IO campaigns work differently when deployed domestically therefore remains an open empirical question.

A final body of work that does investigate domestically deployed IO campaigns focuses specifically on China. A key finding is that the objective of IOs inside China’s “Great Firewall” is, more than anything, to distort the online information environment through sheer volume (King et al., 2017; Roberts, 2018). It is worth noting, however, that findings from China may be unlikely to generalize to other authoritarian contexts given that China’s ability to shape domestic social media is—at least in part—a function of the dominance of domestic firms in China’s market for Internet content (Pan, 2017). Limited information is available on the scale of IOs in other settings.<sup>3</sup> As such, we also do not know whether IOs outside of such high-capacity regimes can effectively distort the information environment—particularly on mainstream, rather than state owned, social media platforms.

In sum, we know that IO activity has been increasing worldwide. Foreign-backed operations may have had limited success in affecting political outcomes or attitudes in target countries. When deployed by high-capacity regimes domestically, these information campaigns have, however, been successful in distorting the online information environment. Descriptive information on the extent and potential influence of these operations outside these contexts remains lacking. A move towards providing descriptive information on these forms of commu-

---

<sup>2</sup>Fake news is a phenomenon distinct from IO activity, but IO campaigns have also participated in the dissemination of fake news (Badawy et al., 2019; Martin et al., 2020).

<sup>3</sup>For an exception see Alizadeh et al. 2020. For descriptive details of IO tweet content attributed to Iran see Elswah et al. (2019).



nication is therefore critical if we are to understand their likely influence.

To that end, we take inspiration from recent research quantifying the likelihood of exposure to false, misleading or actively malicious forms of information (Allen et al., 2020; Guess et al., 2019). These initiatives are crucial: without descriptive information on the actual scale of particular forms of online information, we can only conjecture about their potential effects. In what follows, we quantify the extent of engagement with influence operations in Saudi Arabia—a country that has made headlines for its use of computational propaganda in the pursuit of domestic political aims.

### **Influence Operations in Saudi Arabia**

Over the past several years, the Saudi regime has poured increasing resources into influence operations aimed at promoting the Kingdom’s strategic goals domestically and abroad. Researchers have identified evidence of Saudi Arabia using computational propaganda to spread pro-government messages and discredit opposition since 2013 (Woolley, 2016). Domestically, Saudi influence operations have engaged in a variety of tactics from spreading sectarian hate speech to promoting content praising Crown Prince Mohammad bin Salman and his social reform plan, Vision 2030 (Jones, 2016; Pfeiffer, 2019).

These operations include the use of automated bots and paid trolls alike. Research on bot use in Saudi Arabia finds that automated content creation of this kind was used to promote anti-Qatar hashtags during and after the 2017 Gulf Crisis (Jones, 2019), even if the volume of automated online content creation was likely only limited (Abrahams and Leber, 2020). Some such accounts have been linked to Smaat, a social media and marketing firm based in Saudi Arabia, that has managed the accounts of many high-profile Saudis as well as official accounts for some government departments (Grossman et al., 2020b). Inauthentic accounts have been found to engage in a variety of activities from harassing activists or opposition figures to flooding the Saudi Twittersphere with diversionary trending hashtags (al Sharif, 2019; DiResta et al., 2019). Saudi influence campaigns abroad, in particular their efforts to shape global opinion after the Khashoggi murder and the Qatar Crisis, have received a reasonable share of journalistic and scholarly attention (Grossman et al., 2020b; Jones, 2019). In comparison, we know relatively little about the size of domestic influence operations in the Saudi Twittersphere.

## Data & Method

Our principal data source is a set of IO tweets originating in Saudi Arabia over 2019. These data were released publicly by the Twitter Transparency Center, and provided “unhashed”—in a format that does not redact the names of users with fewer than 5,000 followers—to academic researchers. We obtained the unhashed versions of the releases for December 2019 and April 2020.<sup>4</sup> Twitter attributes the accounts released in December 2019 to Smaat, the Saudi marketing firm, stating “Rigorous investigations by our Site Integrity team have allowed us to attribute these accounts to a significant state-backed information operation on Twitter originating in Saudi Arabia...Our investigations have traced the source of the coordinated activity to Smaat, a social media marketing and management company based in Saudi Arabia. Our in-house technical indicators show that Smaat appears to have created, purchased, and/or managed these accounts...Smaat managed a range of Twitter accounts for high-profile individuals, as well as many government departments in Saudi Arabia.” The second of these releases is described by Twitter as “A network of accounts associated with Saudi Arabia and operating out of multiple countries including KSA, Egypt and UAE.”

While it can often be difficult to ascertain the origins of IO activities, Twitter’s investigations suggests that these accounts were indeed coordinated from within Saudi Arabia. Twitter’s descriptions of the activities of these accounts suggest that they were engaged in a combination of activities amplifying messages in line with Saudi domestic and geopolitical interests—including a small amount of English language content. This suggests that these were likely aimed at Saudi citizens, regional audiences, and—to a much lesser degree—Western audiences.

Because we are interested in tracking the domestic activity of these accounts, we manually filtered the second release to retain only users for whom “`user_reported_location`” referenced a location within Saudi Arabia. After restricting our tweets to the year 2019, we end up with a combined sample of ~9.8m IO tweets and ~4.5k unique accounts. In the Appendix, we provide descriptive statistics on each release, including information on the proportion of tweets that included user reported locations.

---

<sup>4</sup>See <https://tinyurl.com/twitIODec19> and <https://tinyurl.com/twitIOApr20> for the blog post write-ups supplied by Twitter for each release. Last accessed: 09 Nov, 2020.

In order to gauge the scale and influence of IO content in the Saudi Twittersphere, we compare these tweets with four other datasets of tweets originating from Saudi Arabia. The first two are random samples of Saudi users' tweets obtained using two different sampling techniques. The second two are 2019 tweets from the top fifty most followed accounts in Saudi Arabia and the top fifteen most-followed news sources.<sup>5</sup> For ease of communication we refer to each of our datasets using a shortened title: **IO** = Influence Operations tweets; **SA** = 1% Saudi user sample; **GEO** = Geo-located Saudi user sample; **TOP** = top fifty Saudi user tweets; **NEWS** = top fifteen Saudi news organization tweets. We describe how each of these was collected below.

Our first sampling method was designed to get a random sample of Saudi Twitter users. We streamed tweets from the public Twitter POST statuses/filter API endpoint. For this, we used the “tweepy” Python library (Roesslein, 2020) and filtered tweets with bounding box coordinates for Saudi Arabia. After ingesting an initial set of tweets, we filtered users by country code, restricting our user set to users located to Saudi Arabia. We then used the GET statuses/user\_timeline API endpoint via the R package “rtweet” (Kearney et al., 2020) to obtain the last 3,200 tweets of each of the 10,400 Twitter users in our sample. The resulting SA sample contains ~4.7m tweets. While the vast majority of tweets in this dataset are produced by everyday Saudi Twitter users, manually coding the top 100 most followed users in this sample suggests that this sample also includes social media influencers, football accounts, government accounts, entrepreneurs, and clerics. The top 100 most followed accounts in this sample are displayed in Table A7.

Our second approach used Twitter's Streaming API to collect tweets that had been geolocated to within the borders of Saudi Arabia in early 2019.<sup>6</sup> The advantage of this second technique is that tweets were collected in real time. This means that we also capture any tweets engaging with IO accounts that users might have later chosen to delete after the public disclosure of their IO origin. The Streaming API captures tweets within a large bounding-box that crosses into non-Saudi territory. We therefore first filtered our sample to users located within the borders of Saudi Arabia. We plot the location of our GEO users in Figure 1. Like

---

<sup>5</sup>Information on the top fifty most followed accounts was obtained from social media analytics company “Socialbakers.” See: <https://www.socialbakers.com/statistics/twitter/profiles/saudi-arabia>.

<sup>6</sup>This data was collected at NYU's Social Media and Political Participation (SMaPP) lab.

the random sample of Saudi Twitter users, while this dataset mostly consists of tweets produced by regular Saudi Twitter users, the most followed accounts include influencers, football accounts, government accounts, and entrepreneurs. This dataset also includes a number of entertainment accounts (artists, musicians, comedians etc.), some of whom visited Saudi Arabia for performances. The top 100 most followed accounts in this sample are displayed in Table A8. After collecting these tweets, we obtained the 3200 most recent tweets for the unique users that appeared in our sample. This gave us a final dataset of  $\sim 11.3$ m GEO tweets over the year 2019 (of which  $\sim 367$ k from live stream).


In addition to these sampling techniques, we also used the Historical Power Track API to collect all tweets over the year 2019 for the fifty most followed accounts generally and fifteen most-followed accounts of Saudi news organizations. The first of these includes  $\sim 145$ k tweets; the second  $\sim 449$ k tweets. This gives us two samples of high-profile accounts—TOP and NEWS—against which to further benchmark engagement with IO accounts. These users are displayed in Tables A9 and A10.

A note on sampling techniques is worthwhile at this stage. Recent research on the Twitter sampling API reveals that it is vulnerable to artificial distortion and may not represent an actual 1% random sample (Morstatter et al., 2013; Pfeffer et al., 2018). While this type of distortion may be a vulnerability during periods of intense use—e.g., during elections—we do not expect that it threatens the representativeness of our sample. By using two separate sampling techniques we are also able to compare our two samples of Saudi user tweets to each other to check for any obvious dissimilarities.<sup>7</sup>

The account language provided by Twitter for the IO tweets was 92% Arabic. For our other datasets, we did not capture the Twitter account language upon first ingest. To overcome this, we manually coded a random sample of 200 tweets from each sample for whether they contained Arabic text (see Table A6). The majority language in each dataset is Arabic. Our GEO tweet sample contains the smallest percentage of Arabic tweets at 73%. The SA tweets are 89% Arabic; NEWS tweets are 99.5% Arabic; TOP tweets are 86% Arabic. The relatively

---


<sup>7</sup>Given that our GEO random sample was collected after IO accounts had been removed, we can also be sure that there is no crossover in users. We also checked whether any IO accounts appeared in our tweets collected with the streaming API (before the IO account removal) and we found no crossover.



**Figure 1. Coordinate locations of GEO-identified Saudi users.**

*Note. N= 367,000.*

lower percentage of Arabic tweets in our GEO sample suggests that our SA sample may be more reliably capturing users normally resident in Saudi Arabia. In contrast, our GEO sample may be capturing some individuals in transit through Saudi Arabia or only temporarily in the country.



**Figure 2. IO tweets attributed to Saudi state-backed operations.**

*Note.*  $N = 33,740,460$ , includes only accounts active in 2019. Shaded area indicates observation period for the analyses below.

We focus our analysis on 2019, which was the year that saw the highest overall incidence of IO tweets in the public Twitter release (see Figure 2). It is also the year for which we were able to obtain a reliable live-collected sample of Saudi users. Using a more proximate time period also meant that we could retrospectively obtain a second random sample of Twitter users without having to bypass the limit of 3,200 historic tweets permitted by Twitter ToS.

We first measure engagement by calculating the observed number of retweets for IO tweets.<sup>8</sup> We repeat this for a subsample of tweets mentioning a range of political topics, including those tweets that mention "Jamal Khashoggi" or "Qatar"—two polarizing topics that were particularly prevalent in the more actively political content disseminated by IO accounts (DiResta et al., 2019). The motivation for carrying out this analysis is that, while overall engagement with IO accounts may be low, they may receive outsized attention for their dissemination of false or malicious information. This would be consistent with previous published findings that false news stories diffuse more rapidly than verifiable news stories (Vosoughi et al., 2018). We then benchmark observed engagement against our two random samples and two samples of news and high-profile users.

Using mentions of IO accounts within our random samples as a measure of influence, we then estimate for our random samples of Saudi users the percentage of total mentions, and percentage of total tweets, that include an IO account. Finally, we locate influential IO users within overall mention networks for all of our samples combined.

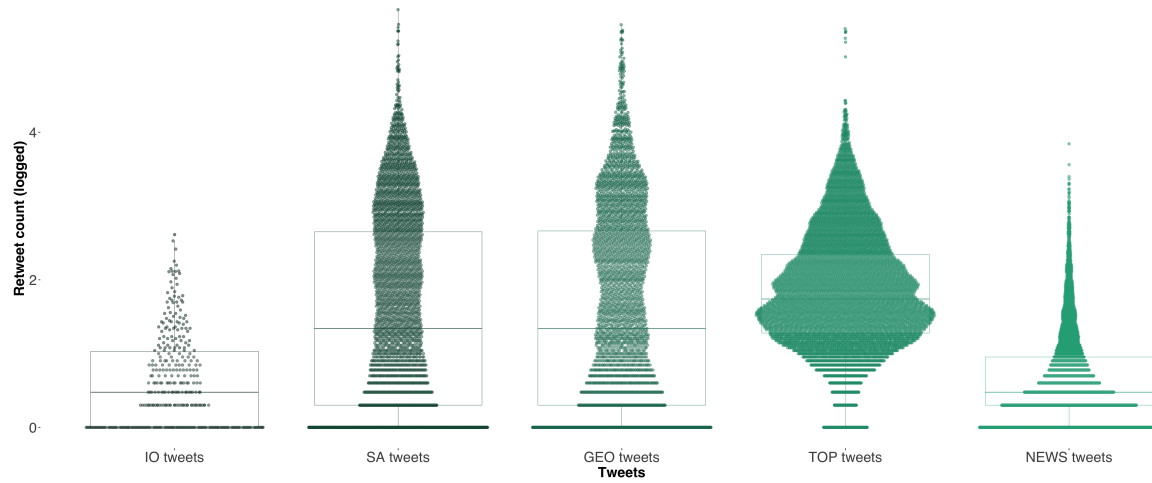
## Results

We display density plots of observed engagement (retweets) in our IO accounts alongside our reference samples in Figure 3. The width of the density curve corresponds to the relative frequency of observations at given regions on the y-axis. Given the size of these datasets, we choose to display a random subsample of 10,000 observations for each tweet sample.

We see that while retweets of IO accounts are approximately equivalent to retweets of major Saudi news outlets, retweet volume is markedly higher in our SA and GEO samples when compared to the IO accounts. As expected, our TOP accounts receive substantially more retweets. We then subset our samples, filtering tweets to include only those that mention Khashoggi or Qatar. These represent more contentious topics, and so we may expect to see heightened engagement.

---

<sup>8</sup>We were unable to recover accurate estimates of favorites for our SA and GEO benchmark samples. This was due to a change in the Twitter endpoint, which meant the "rtweet" package failed to recover favorite counts. In any case, retweets and favorites tend to correlate relatively highly, meaning we would unlikely reach different conclusions due to this omission. Retweeting is also, arguably, a surer measure of engagement than liking.




**Figure 3. Retweets of IO accounts and four benchmark samples.**

*Note. Kernel density scatter plots with overlaid boxplots showing median, first, and third quartiles of distribution.*


We see in Figures 4a and 4b that tweets mentioning these subjects do see substantially more engagement, but the increase in engagement is far more pronounced in the SA and GEO samples than it is in the IO accounts. That is, when it comes to contentious political subjects, IO accounts invited far less engagement than the average Saudi user. In the Appendix, we verify that this finding is not attributable to our choice of topics. One concern is that these topics might be more aimed at regional than domestic audiences, which could explain the lower engagement we observe. In place of Khashoggi and Qatar topics, we instead analyze subsets of tweets based on whether they include a politically-relevant hashtag. We do so by taking the top hashtags used by the NEWS tweet accounts, and manually code those hashtags that pertain to domestic or foreign political themes.<sup>9</sup> We then group these into two categories that relate to domestic political and foreign political topics, as well as three further categories relating to Iran, Yemen, and the Saudi Royal family specifically. We list the hashtags and their assigned categories in Table A11. Across all five groups of tweets, we see that SA and GEO tweets see far higher engagement than IO tweets, and as much, if not more, engagement, than TOP accounts. Again, NEWS accounts see relatively muted engagement irrespective of topic. We

<sup>9</sup>We defined top hashtags as those that appeared with frequency >500 in the NEWS tweets. In total, we coded 145 hashtags for relevance to different political topics.





(a) Tweets mentioning “Khashoggi” in IO accounts and four benchmark samples.



(b) Tweets mentioning “Qatar” in IO accounts and four benchmark samples.

Figure 4. Tweets mentioning political topics.

*Note.* Kernel density scatter plots with overlaid boxplots showing median, first, and third quartiles of distribution.

display these findings in Figure A1.

Providing suggestive evidence beyond Twitter’s investigations that the Saudi IO accounts are likely intentionally targeting a domestic audience—in addition to a foreign one—we see that IO accounts tweet about domestic topics (including discussions of specific royal family members, the economy, the education system and the justice system) at about the same rate that they tweet about the Qatar crisis—the most discussed geopolitical event in our dataset. We display the distributions of these topics in our data in Figure A3.


Mentions of IO accounts provide us with another measure of their influence. If a user within our other random samples mentions an IO user, this is indication that IO activity has managed to penetrate the online discussion. We calculate for both our SA and GEO samples of Saudi Twitter users the percentage of overall mentions and overall tweets that include mention of an IO user. We display these as daily means, seven-day rolling means, and overall means in Figures 5a and 5b.<sup>10</sup>

We see that, for our SA sample, no more than .09% on average of overall tweets contain mention of an IO accounts, and for our GEO sample, no more than .05%. As a percentage of overall mentions, this quantity is .4% and .27% respectively for the two samples. We also see from our two samples of Saudi Twitter users, however, that a relatively high percentage of users interacted with (mentioned) an IO account at least once. In the SA sample, 17.9% of users mentioned an IO account; in the GEO sample 9.6% of users mentioned an IO account. These proportions are comparable to available estimates of US users’ interactions with Russian Internet Research Agency accounts, where 11.3% of users directly interacted with these troll accounts (Bail et al., 2020).

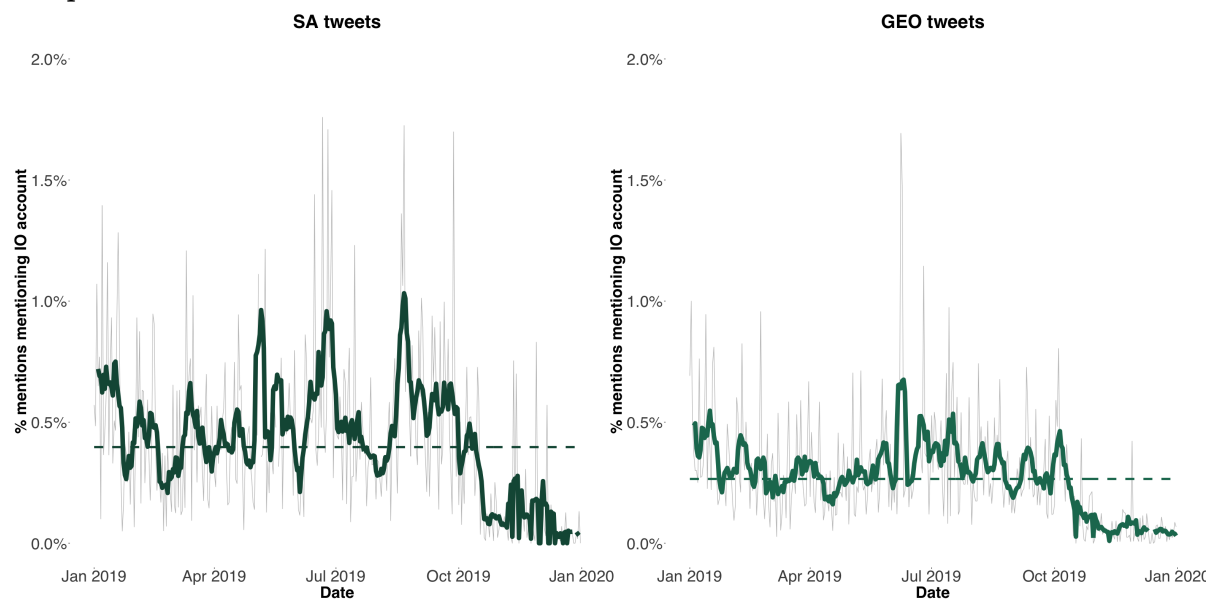
In the article by Bail et al. (2020), the sample of US users included only politically engaged users. Conversely, our sample is blind to users’ political engagement. The elevated levels of interaction we see in our Saudi users sample may therefore be surprising. However, it is worth bearing in mind that the IO accounts did not just tweet political content. In fact, many tweeted an abundance of commercial and religious content (DiResta et al., 2019; Grossman et al., 2020a). As a check on this, we subset our SA and GEO sample users by

---

<sup>10</sup>See Table A12 for a list of the top mentioned IO accounts across all of our samples.



(a) Percentage of tweets that include mention of an IO account, SA and GEO samples.



(b) Percentage of mentions that include mention of an IO account, SA and GEO samples.

**Figure 5.** Percentage of tweets and mentions mentioning an IO account.

*Note.* Grey lines represent daily mean percentages; colored lines seven-day rolling means, and dashed horizontal lines overall means.

their volume of political posting. We identify the top 25% most “political” users in our SA and GEO samples by subsetting out those users who post most frequently with politically relevant hashtags taken from our NEWS sample. Here, we use all hashtags that coded as 1 for either domestic or foreign political relevance. We see that engagement with IO accounts among these users is slightly more pronounced: .1% of tweets mention an IO account among these users in the SA sample, and .08% of tweets in the GEO sample; while .55% and .50% of all mentions include an IO account in the SA and GEO samples respectively.<sup>11</sup> Nonetheless, overall levels of interaction with IO accounts remains muted. It is worth noting, too, that GEO accounts posted about political topics (used politically-relevant hashtags) with less frequency than our SA accounts (.60% versus .97%). One interpretation of this is that those individuals who do not take the precaution of disabling geolocation are also likely less disposed to talking about potentially sensitive political topics in public. This provides suggestive lessons about appropriate sampling techniques when analyzing repressive contexts. A second interpretation is that our GEO sample captures different types of users, some of whom are not normally resident in Saudi Arabia, as suggested by the relatively lower percentage of Arabic-language tweets in this sample.<sup>12</sup>

These elevated engagement levels also seem to be driven largely by a small number of influential accounts. We therefore determine which IO users received most mentions in each of our samples. Combining information on IO mentions across all of our samples, we then locate the position of influential IO accounts within an overall mention network. The network visualization in Figure 6 displays the mention network of the combined SA, GEO, and TOP samples.<sup>13</sup> Here, nodes are sized by in-degree and represent IO users that are mentioned in any of our three samples of Saudi Twitter users. We also color sets of nodes and edges by modularity score, following the community-detection technique proposed by Blondel et al. (2008). We see that the mention network is dominated by a small number of IO users. Manually coding the users who most frequently mention IO accounts, we see that most of the engagement comes from regular Saudi Twitter users with relatively low numbers of followers, rather than influencers or

---

<sup>11</sup>For the bottom 25% of users by frequency of posting with political hashtags, the percentages were .063% and .034% of all tweets, and .26% and .17% of all mentions for the SA and GEO samples respectively.

<sup>12</sup>Our GEO and SA samples are nonetheless essentially identical in terms of follower count and daily tweet frequency (see Figure A2).

<sup>13</sup>We do not include NEWS accounts here as these accounts mentioned zero IO accounts.

popular figures.<sup>14</sup>



**Figure 6. Network visualization of mentioned IO accounts.**

*Note. Nodes sized by in-degree and colored by cluster.*

---

<sup>14</sup>Grossman et al. (2020a) note that one tactic used by Smaat—the organization subcontracted to conduct the Saudi IO campaign—was to target online influencers in an attempt to amplify their influence. Our analysis here suggests that this strategy was largely unsuccessful.

## Discussion and Conclusion

The absence of descriptive information on the scale of different forms of information exposure is a central problem in the communications literature. There are several dimensions to this issue. First, information is consumed across multiple platforms, meaning that even if inaccurate or malicious forms of information proliferate on one medium, it may yet constitute a small proportion of overall information diets when assessed across multiple media—at the level of the “information ecosystem” (Allen et al., 2020). Second, we lack good measures of human attention on the Internet (Lazer, 2020), meaning that we do not know whether and how individuals actually engage with particular forms of information. Third, much of the information shared online remains proprietary, inaccessible to the public, and subject to unpredictable changes in terms of use (Freelon, 2018; Pasquetto et al., 2020). The ostensibly straightforward task of estimating the share of online misinformation as a proportion of overall content therefore remains difficult.

In this article, we solve some of these issues, while others remain open. Our findings are—first and foremost—limited to domestic engagement with known IO content on Twitter. We are constrained to analyzing engagement with *known* IO accounts that have been made publicly available. But it is of course possible that there is other IO activity in the Saudi Twittersphere during this period that we are not able to observe. Moreover, while understanding domestic engagement with these accounts is undoubtedly important, Twitter’s investigations suggest these accounts likely target individuals both within and beyond the borders of Saudi Arabia.

More generally, our results are also limited to Twitter. Given that Twitter is so widely used in Saudi Arabia, conclusions that pertain solely to this platform are still relevant. But we know that other state-backed media remains more popular overall, in particular among older generations (Everette et al., 2019), even if younger Saudis are consuming the large majority of their news on social media (Shikaki, 2020). Given the popularity of social media for news consumption, the relatively muted engagement we see with news organizations might be taken as indication that online mainstream news media is not particularly influential.

However, there is reason to question such an interpretation. When measured as likes or retweets, engagement may not be immediately comparable between different account types. This issue is most stark for news accounts, which tweet frequently, often with links to articles hosted on the relevant organization's website. These accounts nonetheless receive often only limited engagement in the form of likes and retweets (Ackland et al., 2019). Part of the reason for this may be that news organizations are unlikely to engage in conversation or reciprocate another user's engagement. But given the absence of alternative measures of human attention, we remain reliant on standard engagement criteria.

Despite these limitations, by utilizing multiple sampling strategies and benchmarking IO content against relevant country-level samples, we are nonetheless able to provide systematic measures of domestic engagement with known influence campaigns in Saudi Arabia. The Twitter API remains relatively open, and by using the Streaming API to collect data in real time we are able to get a clearer picture of the overall platform-level information environment and catch any IO influence on Saudi users before content removal.

Taken together, our analysis suggests that although domestic engagement with IO tweets was approximately equal to engagement with the top fifteen Saudi news outlets, it was significantly lower than engagement with tweets in our representative samples, and dramatically lower than engagement with the fifty most influential accounts in the Saudi Twittersphere. These findings complement existing research into misinformation originating from Saudi Arabia. Abrahams and Leber (2020), for example, find evidence of bot activity in Saudi Arabia promoting pro-regime messaging, but conclude that even for contentious political topics, the volume of such content was generally low. Engagement with IO account tweets in our analyses was also relatively low. Compared to engagement with general Saudi user tweets, engagement with IO content on controversial topics such as the murder of Khashoggi or the Qatar crisis was muted. And this finding is consistent across a diverse range of political topics. Our network analysis indicates that engagement with IO accounts was largely driven by engagement with a few prominent accounts, with the vast majority of IO accounts receiving very little engagement during the period under study.

We hope that future research will continue to descriptively map engagement with influence operations in diverse contexts and across multiple platforms. Quantifying basic levels of

engagement is a first step in understanding the importance of rapidly evolving computational propaganda campaigns world-wide.

### Acknowledgments

The authors thank the editors and reviewers for their helpful suggestions and gratefully acknowledge the financial support of the NYU Social Media and Political Participation (SMaPP) lab from the INSPIRE program of the National Science Foundation (Award SES-1248077), the William and Flora Hewlett Foundation, the Rita Allen Foundation, the Knight Foundation, the Bill and Melinda Gates Foundation, Craig Newmark Philanthropies, the Democracy Fund, the Intel Corporation, the New York University Global Institute for Advanced Study, and the Faculty of Arts and Sciences Research Investment Fund at New York University. Replication materials are available at: <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/T97FMH>.

### References

- Abrahams, A. and Leber, A. (2020). Electronic Armies or Cyber Knights? The Sources of Pro-Authoritarian Discourse on Middle East Twitter. *International Journal of Communication*, 14:1–29.
- Ackland, R., O’Neil, M., and Park, S. (2019). Engagement with news on Twitter: Insights from Australia and Korea. *Asian Journal of Communication*, 29(3):235–251.
- al Sharif, M. (2019). Why I Quit Twitter and Facebook. <https://web.archive.org/web/20190624211055/https://www.manal-alsharif.com/blog-entries/2019/1/31/why-i-quit-twitter-and-facebook>.
- Alizadeh, M., Shapiro, J. N., Buntain, C., and Tucker, J. A. (2020). Content-based features predict social media influence operations. *Science Advances*, 6(30):eabb5824.
- Allcott, H., Gentzkow, M., and Yu, C. (2019). Trends in the diffusion of misinformation on social media. *Research & Politics*, pages 1–8.
- Allen, J., Howland, B., Mobius, M., Rothschild, D., and Watts, D. J. (2020). Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances*, 6(14):eaay3539.



- b33lz3bub (2019). Lord Of The Flies: An Open-Source Investigation Into Saud Al-Qahtani. <https://www.bellingcat.com/news/mena/2019/06/26/lord-of-the-flies-an-open-source-investigation-into-saud-al-qahtani/>.
- Badawy, A., Addawood, A., Lerman, K., and Ferrara, E. (2019). Characterizing the 2016 Russian IRA influence campaign. *Social Network Analysis and Mining*, 9(1):31.
- Bail, C. A., Guay, B., Maloney, E., Combs, A., Hillygus, D. S., Merhout, F., Freelon, D., and Volfovsky, A. (2020). Assessing the Russian Internet Research Agency’s impact on the political attitudes and behaviors of American Twitter users in late 2017. *Proceedings of the National Academy of Sciences*, 117(1):243–250.
- Benner, K., Mazzetti, M., Hubbard, B., and Isaac, M. (2018). Saudis’ Image Makers: A Troll Army and a Twitter Insider (Published 2018). *The New York Times*.
- Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *arXiv:0803.0476 [cond-mat, physics:physics]*. Comment: 6 pages, 5 figures, 1 table; new version with new figures in order to clarify our method, where we look more carefully at the role played by the ordering of the nodes and where we compare our method with that of Wakita and Tsurumi.
- Bradshaw, S. and Howard, P. N. (2018). Challenging Truth and Trust: A Global Inventory of Organized Social Media Manipulation. Technical report, Oxford Internet Institute, Oxford, UK.
- Bradshaw, S. and Howard, P. N. (2019). The Global Disinformation Order 2019 Global Inventory of Organised Social Media Manipulation. Technical report, Oxford Internet Institute.
- DiResta, R., Grossman, S., and Miller, C. (2019). Analysis of Twitter Takedown of State-Backed Operation Attributed to Saudi Arabian Digital Marketing Firm Smaat. Technical report, Stanford Internet Observatory.
- Elswah, M., Howard, P. N., and Narayanan, V. (2019). Iranian Digital Interference in the Arab World. COMPROP DATA MEMO, Oxford Internet Institute, Oxford, UK.
- Everette, D. E., Martin, J. D., and Hassan, F. (2019). Media Use in the Middle East: A Seven Nation Survey. Technical report, Northwestern University in Qatar.

- Freelon, D. (2018). Computational Research in the Post-API Age. *Political Communication*, 35(4):665–668.
- Gorodnichenko, Y., Pham, T., and Talavera, O. (2018). Social Media, Sentiment and Public Opinions: Evidence from #Brexit and #USElection. Technical Report w24631, National Bureau of Economic Research, Cambridge, MA.
- Grossman, S., H., K., DiResta, R., Kheradpir, T., and Miller, C. (2020a). Blame it on Iran, Qatar, and Turkey: An analysis of a Twitter and Facebook operation linked to Egypt, the UAE, and Saudi Arabia. Technical report, Stanford Internet Observatory, Stanford.
- Grossman, S., H., K., and Ross, E. (2020b). Royal Sockpuppets and Handle Switching: How a Saudi Arabia-Linked Twitter Network Stoked Rumors of a Coup in Qatar. Technical report, Stanford Internet Observatory.
- Guess, A., Nagler, J., and Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances*, 5(1).
- Guess, A. M., Lockett, D., Lyons, B., Montgomery, J. M., Nyhan, B., and Reifler, J. (2020). “Fake news” may have limited effects on political participation beyond increasing beliefs in false claims. *Harvard Kennedy School Misinformation Review*.
- Gunther, R., Beck, P. A., and Nisbet, E. C. (2019). “Fake news” and the defection of 2012 Obama voters in the 2016 presidential election. *Electoral Studies*, 61:102030.
- Ignatius, D. (2018). Opinion | How a chilling Saudi cyberwar ensnared Jamal Khashoggi. *Washington Post*.
- Jones, M. O. (2016). Automated sectarianism and pro-saudi propaganda on twitter. *Exposing the Invisible (Tactical Technology Collective)*.
- Jones, M. O. (2019). Propaganda, Fake News, and Fake Trends: The Weaponization of Twitter Bots in the Gulf Crisis. *International Journal of Communication*, 13:1389–1415.
- Kearney, M., Heiss, A., and Briatte, F. (2020). Rtweet.

- King, G., Pan, J., and Roberts, M. E. (2017). How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, Not Engaged Argument. *American Political Science Review*, 111(3):484–501.
- Lazer, D. (2020). Studying human attention on the Internet. *Proceedings of the National Academy of Sciences*, 117(1):21–22.
- Leber, A. and Abrahams, A. (2019). A Storm of Tweets: Social Media Manipulation During the Gulf Crisis. *Review of Middle East Studies*, 53(2):241–258.
- Martin, D. A., Shapiro, J. N., and Ilhardt, J. G. (2020). Trends in Online Influence Efforts. *Working paper*, pages 1–107.
- Martin, D. A., Shapiro, J. N., and Nedashkovskaya, M. (2019). Recent Trends in Online Foreign Influence Efforts. *Journal of Information Warfare*, 18(3):15–48.
- Morstatter, F., Pfeffer, J., Liu, H., and Carley, K. M. (2013). Is the Sample Good Enough? Comparing Data from Twitter’s Streaming API with Twitter’s Firehose. *arXiv:1306.5204 [physics]*. Comment: Published in ICWSM 2013.
- Pan, J. (2017). How market dynamics of domestic and foreign social media firms shape strategies of internet censorship. *Problems of Post-Communism*, 64(3-4):167–188.
- Pan, J. and Siegel, A. A. (2020). How Saudi Crackdowns Fail to Silence Online Dissent. *American Political Science Review*, 114(1):109–125.
- Pasquetto, I., Swire-Thompson, B., and Amazeen, M. A. (2020). Tackling misinformation: What researchers could do with social media data. *Harvard Kennedy School Misinformation Review*.
- Pfeffer, J., Mayer, K., and Morstatter, F. (2018). Tampering with Twitter’s Sample API. *EPJ Data Science*, 7(1):50.
- Pfeffer, S. (2019). Inside saudi arabia’s disinformation campaign. *National Public Radio*.
- Post, W. (2019). Opinion | In Saudi Arabia, Twitter becomes more useful to the repressors than the repressed. *Washington Post*.

- Roberts, M. E. (2018). *Censored: Distraction and Diversion Inside China's Great Firewall*. Princeton University Press, Princeton.
- Roesslein, J. (2020). Tweepy.
- Shikaki, M. (2020). Arab Youth Survey 2020. Technical report, ASDA'A BCW.
- Siegel, A. A. (2018). Producers of Disinformation. Technical report, Hewlett Foundation.
- Twitter (2020). Information Operations: Twitter Transparency Center. <https://transparency.twitter.com/en/reports/information-operations.html>.
- Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380):1146–1151.
- Woolley, S. C. (2016). Automating power: Social bot interference in global politics. *First Monday*.

## Appendix

### *Sample statistics*

#### *IO tweets*

We use two releases of IO tweets from the unhashed datasets provided by Twitter Transparency Center in December, 2019 and April, 2020. The first of these datasets, detailed in Table A.1, contained only tweets identified as originating from Saudi state-backed IO activity. Because all of these tweets were attributed to Saudi Arabia, we do not filter these tweets by "user\_reported\_location". Unfortunately, Twitter does not provide details on how these tweets were ultimately attributed to Saudi state-backed IO activity. In the main analyses, we decide to retain all tweets from this release, irrespective of user-reported location, given that they were attributed to Saudi state-backed activity by Twitter. In Table A.2 we provide descriptive statistics for the sample used in the main analyses: tweets in 2019. Here, we see that  $\sim 6.5$ m tweets had some location information (61%). The N of tweets from this release included in the final IO analysis sample is bolded in red.

**Table A.1: IO tweets Release 1: full date-range provided in data release.**

IO dataset	# tweets	# tweets with location	% with location
saudi_arabia_112019	32054257	18840595	58.78

**Table A.2: IO tweets Release 1: tweets in 2019.**

IO dataset	# tweets	# tweets with location	% with location
saudi_arabia_112019 (2019)	<b>6456312</b>	3947510	61.14

The second IO release was in April, 2020. These tweets came bundled with tweets attributed to state-backed Information Operations in Egypt and the UAE. We therefore filtered these tweets by user-reported location to obtain a sample of tweets for Saudi Arabia. We

present aggregate statistics for this sample in Table A.3 below. Of the  $\sim 36.5\text{m}$  tweets,  $\sim 19.3\text{m}$  included user-reported locations (53%), and of the reported locations,  $\sim 9.9\text{m}$  were in Saudi Arabia (27%). In Table A.4 we provide descriptive statistics for the sample used in the main analyses: tweets in 2019. Here, we see that  $\sim 7.6\text{m}$  tweets had some location information (48%), and of the reported locations  $\sim 3.4\text{m}$  were in Saudi Arabia. The N of tweets from this release included in the final IO analysis sample is bolded in red.

**Table A.3: IO tweets Release 2: full date-range provided in data release.**

IO dataset	# tweets	# tweets with location	tweets in Saudi	% with location	% with location in Saudi Arabia
sa_eg_ae_022020	36523980	19344502	9924380	52.96	51.3

**Table A.4: IO tweets Release 2: tweets in 2019.**

IO dataset	# tweets	# tweets with location	tweets in Saudi	% with location	% with location in Saudi Arabia
sa_eg_ae_022020 (2019)	15769043	7616583	<b>3370759</b>	48.3	21.38

Finally, in Table A.5 we the total N of tweets after combining the two sets of tweets for 2019 from both releases. Note: the total N of tweets after combining the two datasets is slightly less than the sum of the two sets of tweets for 2019 in Saudi Arabia displayed in Tables A.2 and A.4 (i.e.,  $6456312 + 3370759$ ) because there was some tweet IDs that appeared in both datasets. These duplicates were therefore removed for the final analysis.

**Table A.5: IO tweets analysis sample: all tweets in 2019 combined.**

IO dataset	# tweets	# accounts
Combined sa_eg_ae_022020 (2019) & saudi_arabia_112019 (2019)	9826132	4536

*Sampled account details*

**Table A.6: Percentage Arabic language tweets by tweet dataset.**

Tweet dataset	% Arabic tweets
IO tweets	92.3
GEO tweets	73.0
NEWS tweets	99.5
SA tweets	89.0
TOP tweets	86.0

Table A.7: 100 Most Followed Users in SA Sample

rank	screen_name	followers_count	type
1	Talhabeeb	6668550.00	influencer
2	_m4nb_	547382.00	influencer
3	Fahad_Alotaibi_	508620.00	football
4	AhmedJallalah	452684.00	business
5	Saif_Hamoh	409020.00	media
6	dr_alimalki	387835.00	cleric
7	lbinkasem	276961.00	influencer
8	moh_Altwaijri	241763.00	sports
9	imede_	162196.00	influencer
10	tourismpictures	123285.00	entertainment
11	lWether	121325.00	weather
12	Abdulhalsalemm	104618.00	journalist
13	Turki_binFaisal	104508.00	government
14	m_albryk	103209.00	entertainment
15	azoz6670	98474.00	football
16	Riyadh_mmm	91489.00	business
17	mrj112	86007.00	influencer
18	AL_HJARI	83228.00	influencer
19	B_swed	82968.00	influencer
20	Y_senan	80997.00	entertainment
21	azizooghali	75920.00	sports
22	bahaforest	75627.00	influencer
23	swsynbd	74909.00	football
24	ahmed_alabdallh	74294.00	government
25	meteb_alnofay	73834.00	influencer
26	OmarExplains	72473.00	entertainment
27	snap_tabuk	71631.00	media
28	abdoh4magic	70772.00	entertainment
29	IlabraJ	64637.00	entertainment
30	BahaUniversity	63168.00	university
31	hamadGR	62302.00	influencer
32	bader2992	61698.00	influencer
33	Hamad_L77	60261.00	influencer
34	Moh1Rz2H3	59734.00	influencer
35	su_subit	58460.00	cleric
36	alqassimwater	58365.00	business
37	iiSultt	55103.00	influencer
38	AsirTourism	53991.00	business
39	i_____300	53938.00	influencer
40	hash_salman__	53045.00	government
41	Arafatbinhassan	52507.00	cleric
42	majidr99	49896.00	football
43	AliAlmeshaal	48794.00	influencer
44	Mkhnews	47970.00	media
45	DrKhalidAloudah	46990.00	business
46	OmarAlwhibi	46495.00	influencer
47	Manafa_co	44117.00	business
48	aiydhaltqhtani	42557.00	sports
49	Brand_Gifts1	41512.00	business
50	MOE_ONZ	40413.00	government
51	AlObayyd	39989.00	influencer
52	WaleedMuath	38540.00	influencer
53	OMAR_____911	38235.00	influencer
54	ahmadanaji	37328.00	influencer
55	gunfdhnet	34461.00	media
56	ayalyami	34450.00	media
57	AliAlmohsen1	33795.00	influencer
58	majed_otayf	33182.00	media
59	s3dzxx	33143.00	influencer



60	zalfadah1408	32643.00	influencer
61	hamood_9h	32526.00	influencer
62	_5LiD_	32438.00	influencer
63	A_Alrasheed5	32200.00	influencer
64	aa5sdfg802	32085.00	influencer
65	jubailnow	31597.00	media
66	WaIeedGH	31052.00	sports
67	TasiPro	30887.00	business
68	FaisalALa7mari	30511.00	entertainment
69	aalsanie	30509.00	business
70	yanbu___news	30249.00	media
71	ALHarbi_seif	30060.00	influencer
72	alkhamisziad	29717.00	academic
73	alshnaif	29293.00	media
74	khobar_history	28659.00	media
75	RBahais	28509.00	business
76	na7da7n	28021.00	influencer
77	mycruiseksa	27908.00	business
78	DrAlashgar	27814.00	business
79	azmrsh	26808.00	entertainment
80	MajidSociety	26527.00	business
81	MajedAlFahad	26069.00	influencer
82	3z_w1	26009.00	influencer
83	Azizooo1	25982.00	cleric
84	Enfahadd	25872.00	influencer
85	Artisana_Decor	25791.00	business
86	saleh_br	25586.00	influencer
87	k1k21	25021.00	influencer
88	AloniniYb	24824.00	influencer
89	alknfri8	24799.00	entertainment
90	jameel_hazzazi	24745.00	suspended
91	rahla22	24653.00	influencer
92	fu___82	24214.00	sports
93	asg1391h	23955.00	influencer
94	fawwazalhrbi	23818.00	influencer
95	yutm_511	23569.00	influencer
96	alashiri747	23340.00	influencer
97	___hm78	23182.00	business
98	malaysiatravel1	23031.00	business
99	alshetri_dr	22760.00	business
100	burydahprojects	22628.00	business

Table A.8: 100 Most Followed Users in GEO Sample

rank	screen_name	followers_count	type
1	FrencHMonTanA	2956783.00	entertainment
2	melly_goeslaw	2326600.00	entertainment
3	PeterPsquare	1847002.00	entertainment
4	AnisaRahma_Adi	1272228.00	entertainment
5	Almoj_alazra8	966210.00	football
6	faizalfbi76	671525.00	entertainment
7	emma_maembong	630200.00	entertainment
8	HamadAlMuneef	536749.00	government
9	bourashed	534029.00	media
10	salimafillah	528097.00	entertainment
11	NAmrabat53	517448.00	football
12	khalid88990	513100.00	football
13	mjeedalfawzan	471353.00	entertainment
14	mq8y	446310.00	influencer
15	dataacmizal	417949.00	entertainment
16	nahlah_aljammal	388198.00	journalist
17	ka3am_al3ayel	376709.00	entertainment
18	HaqueMarissa	375918.00	entertainment
19	jeddahnews_	369562.00	media
20	rohara_group	332886.00	business
21	_tu00	320112.00	suspended
22	MoaSalem	318990.00	entertainment
23	mainohustlehard	317475.00	entertainment
24	5o__51	305668.00	influencer
25	DA2_Ridho	297461.00	entertainment
26	DA2_Rizki	293574.00	entertainment
27	yasser_madkhli	291085.00	entertainment
28	Hafez_AlMedlej	280293.00	football
29	O__vip	277818.00	influencer
30	Mhzh_	261793.00	football
31	moh_Altwaijri	258195.00	sports
32	Ai9i9i	251851.00	influencer
33	fi9_z	246577.00	media
34	qap_1	245207.00	influencer
35	99t	243429.00	influencer
36	ab_raddadi	229705.00	entertainment
37	Ruleyork	229155.00	entertainment
38	beautyouamal	229004.00	influencer
39	Ahmadbinnaqi	229002.00	suspended
40	farooi	213172.00	business
41	mltqa_al3nzh	209435.00	suspended
42	Sami_Alhomood	207970.00	cleric
43	alzhirri	203013.00	entertainment
44	TMFaisaIx	196044.00	influencer
45	md_almousa	188310.00	entertainment
46	MoustafaELFarra	186430.00	influencer
47	alromansiahKSA	182856.00	business
48	hemsh8	179713.00	entertainment
49	laylaskandar	175717.00	entertainment
50	3zizez	171817.00	influencer
51	ArwaHomaidan	169893.00	entertainment
52	tdrebnews	160606.00	media
53	Sharqiya_event	160379.00	entertainment
54	Ss_m_f	159563.00	suspended
55	KSA_A52	152717.00	suspended
56	WonhoChung	150202.00	entertainment
57	HatoonKadi	149854.00	entertainment
58	AliKuemkh	148277.00	football
59	_x19_	147959.00	influencer

60	H_____60	143811.00	suspended
61	__467__	139397.00	suspended
62	joshua00966	137874.00	suspended
63	Abualfawarss123	137866.00	media
64	souza5	137373.00	football
65	kupinang	136010.00	entertainment
66	ShougAA6	133265.00	influencer
67	Nada__Fadel	132922.00	business
68	FHAD__M__305	127811.00	entertainment
69	MFMiraFilzah	124719.00	entertainment
70	BassemCHRISTO	122404.00	entertainment
71	JeddahChamber	121225.00	government
72	MajeedAlrhedi	120798.00	entertainment
73	Jfurlanich	120400.00	journalist
74	s7sn__	117439.00	influencer
75	aljubailtoday1	116337.00	media
76	sz__v27	112689.00	influencer
77	triggar6	111393.00	influencer
78	mkki1982	107393.00	suspended
79	snake__toe	106736.00	influencer
80	rahneputri	106082.00	influencer
81	nabilhdad	104708.00	suspended
82	SultanQhtani	102822.00	media
83	abb__hf	102616.00	influencer
84	NIZARAALGHANNAM	100776.00	suspended
85	3__li	100278.00	influencer
86	TurkiAlmohsen	100080.00	influencer
87	F__abarh	99910.00	suspended
88	ZulAbrarHakim	98867.00	influencer
89	abualmshx	98087.00	suspended
90	abdullahacbadem	97533.00	government
91	sredusa	97523.00	university
92	ABUALMSH	97067.00	suspended
93	Cydiagold	96437.00	suspended
94	Ms dossary7	95900.00	football
95	Ahoud40009	95564.00	business
96	abualmsh	94776.00	suspended
97	DrFahad55	92732.00	business
98	alaqsaclinics	92227.00	business
99	ALM3TA9M	91390.00	influencer
100	Mabkhot__b13	91306.00	suspended

Table A.9: Top 50 Most Followed Saudi Accounts

rank	screen_name	type
1	MohamadAlarefe	religious
2	Dr_alqarnee	religious
3	shugairi	religious
4	SaudiNews50	news
5	sabqorg	news
6	salman_alodah	religious
7	Alwaleed_Talal	royal family
8	HashKSA	aggregator
9	Alhilal_FC	football
10	KingSalman	royal family
11	abdulrahman	football
12	SamiAlJaber	football
13	fayez_malki	celebrity
14	waleedalfarraj	news
15	SalehAlmoghamisy	religious
16	Talhabeeb	celebrity
17	Y20	football
18	battalalgoos	football
19	alhabibali	religious
20	naseralomar	religious
21	AzzamAlDakhil	government
22	spagov	government
23	k_alshenaif	football
24	PTT_Riyadh	religious
25	akaworldwide	aggregator
26	Nawal_Al3eed__	celebrity
27	AlRiyadh	news
28	SaudiMOH	government
29	MOISaudiArabia	government
30	KSA24	news
31	stc_ksa	company
32	TurkiAldakhil	government
33	Turki_alalshikh	government
34	ittihad	football
35	tfrabiah	government
36	Mobily	company
37	BidzSaleh	celebrity
38	TatweirT	aggregator
39	faisalbinturki1	royal family
40	afaaa73	royal family
41	DaleSteyn62	football
42	KSASociety	aggregator
43	moe_gov_sa	government
44	AdelAljubeir	government
45	altemyat	football
46	aleqtisadiah	news
47	AlNassrFC	football
48	KSAMOFA	government
49	ActionYaDawry	football
50	Dr_almosleh	religious

**Table A.10:** Top 15 Saudi News Accounts

rank	screen_name
1	AlwatanSA
2	AlRiyadh
3	OKAZ_online
4	al_jazirah
5	aawsat_News
6	Almadinanews
7	News_Brk24
8	aleqtisadiah
9	alweeamnews
10	SPAregions
11	SaudiNews50
12	sabqorg
13	Akhbaar24
14	alShrqNews
15	twasulnews

Table A.11: Politically-relevant hashtags and codings from NEWS accounts.

rank	frequency	translated hashtag	domestic	political	foreign	political	iran	yemen	royal
1	6813	#iran	0		1		1	0	0
2	6424	#yemen	0		1		0	1	0
3	4769	#crown_prince	1		0		0	0	1
4	3462	#cabinet	1		0		0	0	0
5	3297	#sudan	0		1		0	0	0
6	2846	#Custodian_of_the_Two_Holy_Mosques	1		0		0	0	1
7	2696	#Custodian_of_the_Two_Mosques	1		0		0	0	1
8	2581	#iraq	0		1		0	0	0
9	2514	#turkey	0		1		0	0	0
10	2496	#the_alliance	0		1		0	1	0
11	2405	#syria	0		1		0	0	0
12	2287	#education	1		0		0	0	0
13	2264	#aramco	1		0		0	0	0
14	2231	#lebanon	0		1		0	0	0
15	2197	#economy	1		0		0	0	0
16	2030	#egypt	0		1		0	0	0
17	2018	#Trump	0		1		0	0	0
18	1820	#Washington	0		1		0	0	0
19	1818	#UAE	0		1		0	0	0
20	1705	#america	0		1		0	0	0
21	1412	#britain	0		1		0	0	0
22	1396	#kuwait	0		1		0	0	0
23	1340	#houthi	0		1		0	1	0
24	1243	#libya	0		1		0	0	0
25	1242	#qatar	0		1		0	0	0
26	1213	#algeria	0		1		0	0	0
27	1160	#tunisia	0		1		0	0	0
28	1150	#bahrain	0		1		0	0	0
29	1140	#russia	0		1		0	0	0
30	1138	#commerce	1		0		0	0	0
31	1105	#pakistan	0		1		0	0	0
32	1084	#oil	1		0		0	0	0
32	1084	#jobs	1		0		0	0	0
33	1070	#king_salman	1		0		0	0	1
34	945	#economic_reports	1		0		0	0	0
35	935	#india	0		1		0	0	0
36	889	#civil_defence	1		0		0	0	0
37	867	#public_prosecution	1		0		0	0	0
38	844	#france	0		1		0	0	0
39	826	#mohammed_bin_salman	1		0		0	0	1
39	826	#houthis	0		1		0	1	0
40	813	#palestine	0		1		0	0	0
41	759	#Tehran	0		1		1	0	0
42	752	#economy	1		0		0	0	0
43	735	#citizens_account	1		0		0	0	0
44	728	#shura_council	1		0		0	0	0
45	714	#arab_summit	0		1		0	0	0
46	687	#employment	1		0		0	0	0
47	643	#UN	0		1		0	0	0
48	630	#civil_service	1		0		0	0	0
49	628	#new_zealand	0		1		0	0	0
50	627	#ministry_of_education	1		0		0	0	0
51	625	#justice	1		0		0	0	0
52	585	#sanaa	0		1		0	0	0
53	581	#ameerca	0		1		0	0	0
54	568	#jordan	0		1		0	0	0
55	557	#Trump	0		1		0	0	0
56	549	#crownprinceinindia	0		1		0	0	1
57	548	#ISIS	0		1		0	0	0

58	546	#khalid_bin_faisal	1	0	0	0	0
59	544	#interior_ministry	1	0	0	0	0
60	543	#police	1	0	0	0	0
61	538	#riyadh_agreement	0	1	0	0	0
62	505	#justice_ministry	1	0	0	0	0

---


**Table A.12:** Most mentioned IO accounts in SA, GEO, TOP, and NEWS samples combined.

rank	screen_name	# times mentioned
1	zV1mp	641
2	506MBS	627
3	Plolive	577
4	ETDose	402
5	ThbtY	316
6	giile7	291
7	ajmidm	286
8	iiSnB	284
9	phzw50	221
10	ij202	154
11	TurkibinTalal	108
12	iAb3al	93
12	RMCFarab	93
13	MubarakAldosari	89
14	AlderbyN	83
15	as7as14	81
16	a7mrL	74
17	i2i2b	71
18	iihaqbani	56
19	xxxt44	54
20	suuny2010	43
20	vbqq	43
21	tweeqksa	40
22	sillicon77	39
23	360motions	34
23	afaqn2030	34
24	beINAlahli	32
24	brtriangle	32
25	albarakati0	30
25	OnsidearV	30
26	KingsAFC	28
27	naifbinsalamah1	26
28	rowef88	25
29	MRakah	23
29	Raeialsalfah	23
30	AA1Omaim	22
30	M7sniii	22
31	S7MiO	21
32	ii2iz	20
32	kljklpopo	20
33	DexSAIF	19
34	abdullahalorfj	18
34	M3LmKsa	18
35	alhejazclub	16
35	anjalbishah	16
36	Heev97	15
36	YahyaSAItaleedi	15
37	2moresa	14
37	AbeerAlsamt	14
37	Kuco33inho	14
38	12345mmggoo	13
38	EngAbuJulia	13
38	farraj966	13
38	panki2015	13
39	ahmyam1	12
39	almdaryah	12
39	iitzkon	12
39	llwdne	12




40	darthwolv	11
40	lanaslittlegurl	11
40	vandettals	11
41	0ii	10
41	Abusaud188Am	10
41	Bander232	10
41	hobaasha	10
41	JazanOffice1	10


---

*Additional analyses*

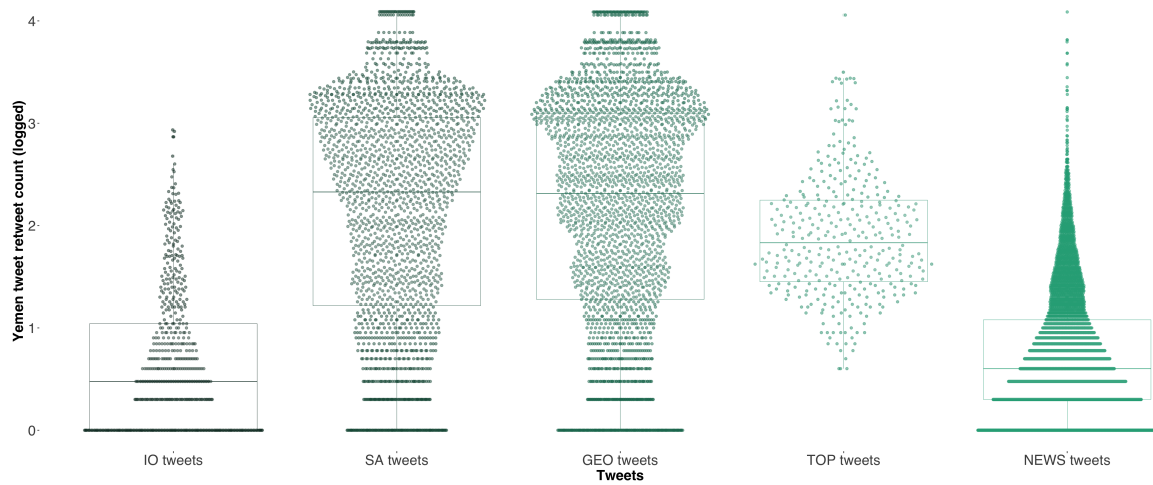
(a) Kernel density scatter plots of retweets for tweets mentioning domestic political hashtags in IO accounts and four benchmark samples.




(b) Kernel density scatter plots of retweets for tweets mentioning foreign political hashtags in IO accounts and four benchmark samples.



(c) Kernel density scatter plots of retweets for tweets mentioning Iran-related hashtags in IO accounts and four benchmark samples.



(d) Kernel density scatter plots of retweets for tweets mentioning Yemen-related hashtags in IO accounts and four benchmark samples.



(e) Kernel density scatter plots of retweets for tweets mentioning Saudi royal family-related hashtags in IO accounts and four benchmark samples.

Figure A.1. Kernel density scatter plots of retweets for tweets mentioning different political hashtags in IO accounts and four benchmark samples. Overlaid boxplots show median, first, and third quartiles of distribution.




Figure A.2. Comparison of GEO and SA sample characteristics.




Figure A.3. Distribution of news topics in Saudi IO Tweets.