Middle East Twitter bots and the covid-19 infodemic

Alexei Abrahams and Noura Aljizawi

Citizen Lab, Munk School of Global Affairs and Public Policy, University of Toronto
alexei@citizenlab.ca and noura@citizenlab.ca

Abstract

Misinformation is often found to be propagated by coordinated networks of accounts ('bots', 'sockpuppets'). We looked for such networks on Middle East Twitter by following 149 search filters related to covid-19 during March-April, 2020. Contrary to the hype, the share of plausible bots ran on the order of 4-7% per filter, comparable to pre-covid-19 rates. Leveraging contextual knowledge, we inspected two Arabic and one Persian networks more closely. We found reasons to doubt their inauthenticity, and scant evidence of misinformation. We conjecture that fears of Twitter being overrun by misinformative bot armies may be overblown, at least on Middle East filters during the covid-19 'infodemic'.

Introduction

As the covid-19 pandemic has rampaged across the globe in early 2020, it has been accompanied by what the World Health Organization (WHO) alleges to be an 'infodemic' of misinformation (World Health Organization 2020). Consistent with this allegation, journalists and researchers have surfaced numerous examples across social media of false medical advice; conspiracy theories insinuating that the pandemic originated in the United States; and networks of coordinated inauthentic accounts ('bots') advancing pro-Chinese, anti-American propaganda (Molter and Webster 2020, Kao and Shuang Li 2020).

So far, however, comparatively little research has focused on pandemic-related manipulation in the Middle East (an important and prominent exception has been the numerous, insightful Twitter investigations conducted by researcher Marc Owen Jones, often self-published to his Twitter timeline https://twitter.com/marcowenjones). This is surprising given that, in recent years, Middle Eastern Twitter has gained considerable notoriety as a hotbed for social media misinformation and manipulation, with the Saudi Arabian regime going so far as to infiltrate Twitter with spies (Hubbard 2019). Indeed, it is widely believed that much of the region's social media misinformation and manipulation is state-backed, as authoritarian regimes like Iran, Saudi Arabia, and Egypt, attempt to quell domestic political dissent and influence regional politics (Abrahams 2019). Insofar as the covid-19 pandemic threatens to expose inadequate health infrastructure and lack of political accountability, it arguably gives rise to a volatile political moment where we should expect regimes to ramp up information controls to quell dissent. Indeed, both the Iranian and Egyptian regimes appear to be stepping up censorship activities (Shahbaz and Funk 2019); it makes sense that we should see a concomitant uptick in misinformation.

Focusing on Middle East Twitter during the covid-19 pandemic, we look for social media manipulation on one specific dimension -- coordinated inauthentic activity (bots). In recent years, researchers have surfaced tens of thousands of Twitter bots operating on Middle Eastern hashtags (Jones and Abrahams 2018, Jones 2019). Twitter itself, including during the pandemic, has suspended tens of thousands of inauthentic users pushing Emirati, Egyptian, and Saudi political talking points (AlJazeera 2019, Gadde and Derella 2020).

We build a dataset by following 149 pandemic-related Twitter search terms (mostly hashtags) in Arabic, Turkish, and Persian, using Twitter's REST API from late March to late April, 2020. To each hashtag, we apply the widely used 'birth anomaly' detection method, which looks for bots among statistically anomalous waves of account births (Jones 2019) . Sure enough, we succeed in surfacing thousands of anomalously born accounts, many of them exhibiting further characteristics widely associated with bots.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

We find, however, that across our sample of hundreds of pandemic-related hashtags, anomalously born users represent a relatively small share of all users (roughly 3.6%-7.1%, depending on how the average is taken). Notably, this average share is not statistically different from what was found using a sample of 279 Middle East hashtags corresponding to the period October 2019 to January 2020 -- *before* the start of the 'infodemic' (Abrahams and Leber 2020).

Consistent with that earlier study, we also find an inverse relationship between the 'size' of the hashtag (as proxied for by the number of participating users) and the share of anomalously born users. That is to say, accounts that are potentially bots tend to constitute a larger share of users on fringier, more parochial hashtags that draw fewer total participants, trend less widely and/or over a shorter time.

We complement this panoramic sweep with a more in-depth investigation of two Arabic networks and one Persian network surfaced by birth anomaly detection. Randomly sampling users from each network, we confirm that user profiles tend to exhibit a set of common characteristics associated with bot networks (Digital Forensic Research Lab (DFRL) 2017). One of the Arabic networks, however, appears to be a coordinated campaign by Egyptian secondary school students hoping to postpone end-of-year exams until after the pandemic has subsided. Thus, while sharing the characteristics of inauthentic coordinated networks, a bit of contextual knowledge suggests the network may in fact be an *authentic* coordinated network of protestors associating with each other to pressure the Egyptian ministry of education. The second Arabic network, by contrast, appears to be state-aligned; but it too seems fairly innocuous, largely retweeting the Saudi Ministry of Health's decree to stay home and practice social distancing, interspersed with praise for the King and Crown Prince for their leadership during the crisis. Finally, the Persian network, which retweets Iran's supreme leader, Ayatollah Khamenei, engages in amplification of hashtags supportive of Bahraini human rights prisoners, and other hashtags expressing solidarity with the Palestinian liberation struggle. None of these three networks appears to be advancing any notable misinformation campaigns.

Taken together, these findings run contrary to the hype and hysteria around misinformation -- both the 'infodemic' narrative around covid-19, and the more general trend over the past few years to view social media in the Middle East as pervasively inauthentic and manipulated (Abrahams 2019). The findings of our broad sweep across 149 covid-19-related hashtags suggests marginal bot incidence, comparable to pre-pandemic rates. While researchers and journalists around the world have in recent months done an excellent job of drawing attention to various bot networks and misinformation campaigns around the globe, For example, see (Kao and Shuang Li 2020). it is more rare to see how those findings scale against overall activity. The failure to discover a bot network is a null finding, and null findings are well known to be an under-published category at scientific outlets, and moreover not newsworthy.[1] By only reporting positive findings, however, we collectively contribute to an alarmist narrative of the social media landscape as teeming with bots and misinformation, whereas these phenomena may in fact constitute only a small percentage of overall activity.

Our in-depth investigation of three networks casts further doubt. The statistical anomalies of some Twitter networks may turn out, as in the case of the Egypt network, to be artefacts of authentic coordination by citizens engaging in associational activities. And even when state-aligned networks emerge, they are not necessarily pro-authoritarian and reactionary, but may be promoting public health or even expressing solidarity with social justice causes. Users, therefore, may often not be part of coordinated networks. When they are, coordinated networks may be authentic, may not necessarily be reactionary, may not be misinformative, and might even convey a positive message. The findings of our closer inspection of three networks suggests that algorithmic bot detection efforts that lack contextual nuance may lead to many false positives and exaggerate bots as a problem (Rauchfleisch and Kaiser 2020, Allyn 2020).

Findings

Over late March to late April, 2020, we tracked 149 search filters (mostly hashtags) in Arabic, Turkish, Persian, and English, trending in the Middle East and relevant to covid-19. Across all 149 filters, we find that on average 7.1% (s.d. 4.3%) of users were born during anomalous spikes (an important though imperfect predictor of bots, as discussed below). Notably, this average share is not statistically different from what was found by Abrahams and Leber (2020) using a sample of 279 Middle East hashtags corresponding to the period October 2019 to January 2020 -- *before* the start of the 'infodemic'.[2] This suggests that, at least using this widely adopted bot detection technique, the share of bots on pandemic-related hashtags does not statistically exceed the share of bots on pre-pandemic hashtags; there is no 'uptick' in bot incidence.

Consistent with that earlier study, we also find here that the shares of anomalously born accounts are

[1] See the editorial in Nature, "The importance of no evidence" (2019)
[2] Across 279 hashtags, they found an average share of 6.1% (s.d. 3.1%) accounts were anomalously born. (Abrahams and Leber 2020)

higher on smaller hashtags (hashtags that attracted fewer participating users). Given that smaller hashtags have a larger share of birth anomalies, it makes sense to recalculate the average share across all hashtags weighting by user count. When we do this, the initial estimate of 7.1% falls to 3.6% -- exactly the same weighted mean as found in Abrahams and Leber (2020).

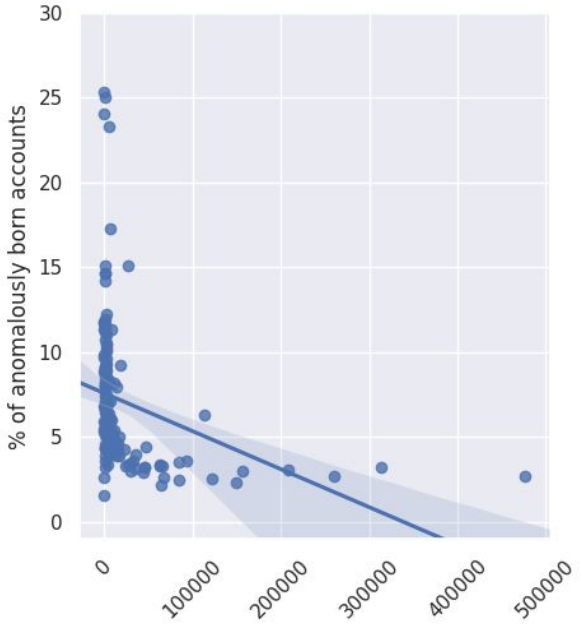

Figure 1: Percentage of anomalously born accounts plotted against number of users participating in each hashtag. The inverse relationship indicates smaller hashtags exhibit higher rates of anomalous births (indicative of potentially inauthentic activity).

Figure 1 plots the percentage of anomalously born accounts against the count of users per hashtag, which exhibits a decidedly downward slope. Taking anomalously born accounts as a proxy for bots, this inverse relationship seems to suggest that hashtags with larger reach and salience tend to be less penetrated by bots, whereas smaller, fringier hashtags, with fewer participants and likely narrower audiences, tend to suffer higher rates of infiltration by bots. Incidentally, this finding appears to support the empirical approach of many journalists and researchers we have interacted with, who tend to hunt for bots and misinformation on more 'peripheral' hashtags.

Is anomalous birth a good indicator of being inauthentic?

The reader may understandably be skeptical that bots are only or even mostly anomalously born. To allay this concern, we use account attrition rates as a second proxy for bots. The logic here is that Twitter itself has far better data by which to identify abusive behavior, and has been at pains to curb inauthentic coordinated activity on its platform. Accordingly, Twitter accounts that participated in a hashtag and later get suspended are more likely to be bots. We can 'piggy-back' on Twitter's decision to suspend, and calculate the suspension rates for a random sample of (<=1000) users from each of our 149 hashtags.

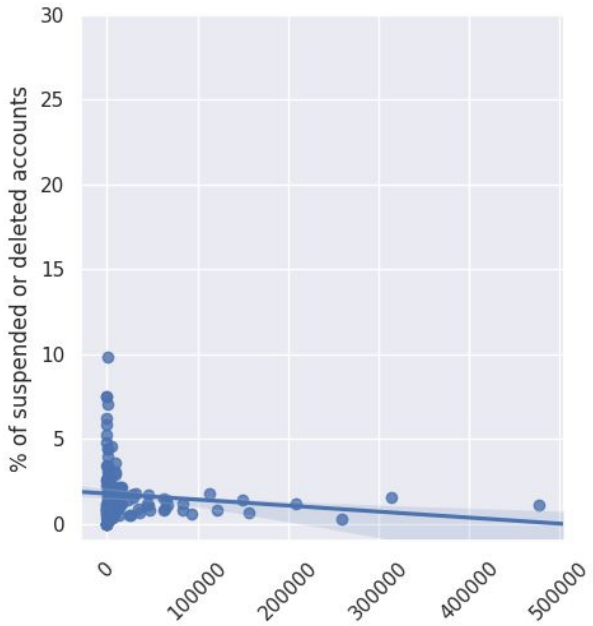

Figure 2: Percentage of suspended and deleted users plotted against number of users participating in each hashtag. The inverse relationship matches Figure 1, suggesting once again that smaller, fringier hashtags exhibit more suspicious activity and provoke Twitter to intervene.

Figure 2 depicts the resulting scatter plot of percentage suspensions against number of users. On average, 1.72% (s.d. 1.55%) of users on any given hashtag were suspended or deleted. Consistent with Figure 1, we find that deletions and suspensions are more frequent among smaller, fringier hashtags, though the slope is notably flatter. Both Figure 1 and 2 are consistent with similar findings of Abrahams and Leber (2020) on pre-pandemic hashtags.

To get a better sense of whether the anomalously born accounts truly exhibit bot-like characteristics, and to see what kind of content they amplify, we investigate three networks surfaced across three of our 149 search filters.

Egyptian Arabic network:

We begin by taking a closer look at the pandemic-related Arabic hashtag, #حياتنا_أهم_من_التعليم, which translates roughly as "our lives are more important than [going to school]". We tracked this hashtag between March 27th and April 25th, 2020, finding two statistically anomalous birth spikes on April 18th and 19th corresponding to 802 account births. Some or perhaps many of these accounts

may belong to real people; the fact that they were born anomalously merely piques our interest and invites further investigation.

Bots are well known to retweet or amplify official or organic accounts in order to amplify a particular idea or position favored by the bot-herder. Obtaining as of April 29th all of the available tweets of each of these 802 accounts[3], we build a graph object using their retweets and, applying the page-rank algorithm, append to our list of 802 accounts the top 10 page-rank percentiles of accounts they choose to retweet.[4] We plot the resulting 3,093 accounts in Gephi, color-coding the anomalously born accounts as green and their preferred 'retweetees' as pink (node size corresponds to in-degree, i.e. number of retweets received).

Figure 3: the retweet network (Force-Atlas-2 layout, Gephi) of the 802 anomalously-born accounts (green) and the top 10% of accounts they have retweeted since January 1st, 2020 (as scored using page-rank). The graph shows remarkable density: 133,344 retweets to 3,093 nodes, implying a ratio of ~43.1.

Remarkably, the edge-to-node ratio of the graph is over 43, constituting one of the most dense retweet networks we have seen across hundreds of hashtags. The graph also confirms that not only were the green accounts born on April 18th and 19th, but they largely choose to retweet (quite intensively) the same 'pink' accounts.

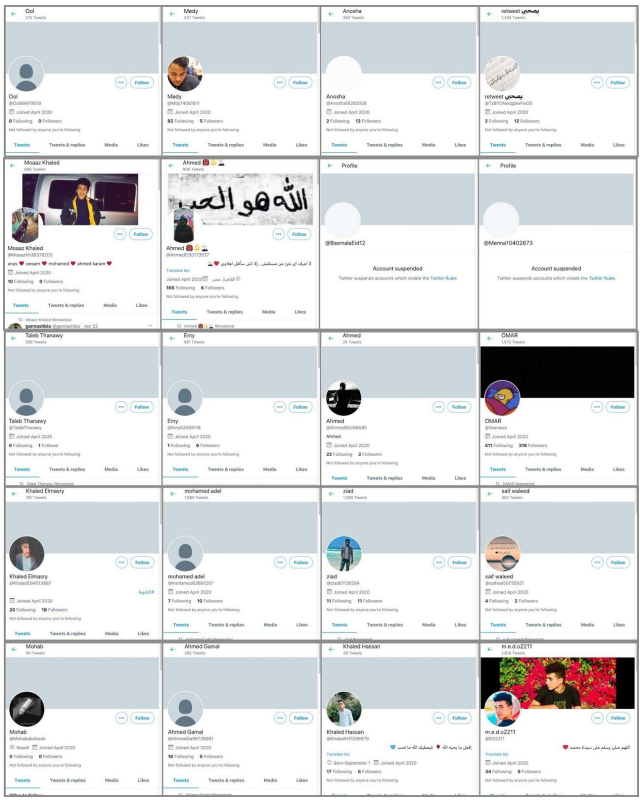

Figure 4: a random sample of 20 of the 802 anomalously born accounts on #حياتنا_أهم_من_التعليم. Accounts all exhibit low friend and follower counts, generally lack images or depict male users.

On April 29th, we checked 20 accounts at random (Figure 4) from among the sample of 802 anomalously born accounts. Already, 2 of the 20 had been suspended. Across the remaining 18 accounts, we found that user screen names followed a recurring pattern of some alphabetic characters followed by a series of digits. 8/20 accounts had no profile picture; another 8 depicted male teenagers, while one had a cartoon profile picture and another presented a verse from the Qur'an. All but one had a male user name.

The reader should note that, although these commonalities would seem to imply that these accounts are coordinating with each other, or can be thought of as one 'network', we cannot jump to the conclusion that they are *inauthentic*. Indeed, the network does not at first blush appear to be deployed by a state or political actor. Rather, it appears to be a campaign launched by secondary school students demanding exams cancellations in view of the covid-19 pandemic. This raises the possibility that this is simply a network of high-schoolers trying to avoid examinations under trying circumstances, and coordinating together on Twitter to make their demands felt. If they agreed amongst each other to campaign against exams, then that would explain the sudden surge of account births over April 18th-19th. The recentness of those births

[3] Twitter REST API limits us to collecting up to the most recent 3,200 tweets of any account.

[4] This technique is developed in (Abrahams, and van der Weide 2020)

(relative to the time of writing, April 29th), moreover, goes a long way to explaining why they generally have low follower and friend counts, why their user bios are short, and why they sometimes lack profile or cover images. Evidently, proving coordination is not the same as proving inauthenticity.

Pro-Saudi and pro-Iranian networks

Although the Egypt network is plausibly a grassroots campaign led by students, other birth-anomaly networks surfacing in our datasets align with state-level agendas. In particular, following the hashtag #كلنا_مسؤول ("we are all responsible [for public health]"), we investigate 2,333 accounts born during an anomalous spate of births over a four-day period, 20th-23rd of March, 2020. We also track mentions of the Iranian Supreme Leader's English-language Twitter account (@khamenei_ir), identifying a spike of 492 account births on April 7th and 8th, 2020.[5]

Examining a random sample of 50 accounts from the 492 anomalously born accounts mentioning @khamenei_ir, we find that 45/50 express pro-regime views. Across pro-regime accounts, several commonalities emerged, including similar screen names (characters followed by digits); religious bio descriptions; and religious cover photos. Many of the accounts adopted a profile photo of Qassem Suleimani, the Iranian Republic Guards general assassinated by the United States in early January, 2020. Many accounts appeared to tweet on #ThePromisedSaviour. Nevertheless, Persian was the most common language adopted across all accounts.

Of the 50 accounts, 20 appeared to be inactive since April 20th, 2020, while another 20 seemed active on a constellation of liberation-related hashtags, namely #أطلقوا_سجناء_البحرين ("free the Bahraini [human rights] prisoners"), and various hashtags associated with Palestinian nationalism: #Covid1949, #QudsDay, #شهيد_القدس ("Martyr of Jerusalem") #القدس_درب_الشهداء ("Jerusalem is the path of the martyrs") #وفاء_لشهيد_القدس (loyalty to the martyr of Jerusalem).

We also investigated a random sample of 50 accounts from among 2,333 anomalously born accounts on #كلنا_مسؤول ("we are all responsible"). The hashtag appeared to be a part of the Saudi Ministry of Health's (MOH) campaign to encourage Saudi citizens and residents to stay home. While 6/50 accounts were suspended, 37 of the remaining 44 accounts exhibited a clear affiliation with Saudi Arabia, whether retweeting the tweets of Saudi ministries, or content panegyric of Crown Prince Muhammad bin Salman.

Of the 44 accounts still alive, 32 appeared active while 12 appeared inactive. Browsing their timelines, 33 of the accounts seemed only to retweet, never contributing original content of their own. As with the Khamenei network, accounts shared several commonalities, including screen names composed of letters followed by digits; religious or nationalistic bio descriptions (often including the Saudi flag); and profile photos that were either religious in nature (even sometimes containing a verse from the Qur'an) or nationalistic, containing a photo of Crown Prince MBS or King Salman.

Across the three networks that we investigated, then, we find pro-Saudi, pro-Iranian, and grassroots mobilizations. This mixture is consistent with the findings of research in recent years, which has documented the rise of state-backed mobilizations after an initially grassroots-dominated social media landscape circa 2011 during the Arab Spring. That said, neither of the regime-aligned networks exhibited a reactionary agenda. The seemingly Iran-aligned network tweeted supportively of the Palestinian liberation struggle, and in solidarity with Bahraini human rights defenders imprisoned by the Saudi-aligned Bahraini regime. The seemingly Saudi-aligned network, meanwhile, merely encouraged citizens to obey the regime's covid-19 curfew and stay home. Thus, even if these accounts could be definitively verified as inauthentic (which seems hard to conclude from the information at our disposal), it seems difficult to conclude that they behaved in a misinformative manner within the time period we observed them.

Methods

Previous research on Twitter bots in general, and Middle East Twitter in particular, has suggested that bots are often deployed to make hashtags trend, or to 'hijack' already-trending hashtags (Jones 2019). Accordingly, it makes sense to build a 'dataset of datasets', where we query Twitter's REST API for all tweets associated with each of a number of hashtags associated with the pandemic. The resulting database of hashtag tables can then be analyzed either *within* each table (to understand the dynamics of each particular hashtag) or *across* tables (to understand the broader ecosystem of hashtag-delimited conversations).

We build a dataset by periodically downloading all tweets pertaining to Arabic, Persian, and Turkish hashtags related to the covid-19 pandemic, using Twitter's

[5] During this period of time, Khamenei (among other Middle Eastern leaders) has commented frequently on the covid-19 pandemic, which motivated our decision to track mentions of his Twitter handle.

REST API.[6] These hashtags were chosen manually, by monitoring trending topics on Twitter. As such, we acknowledge that this dataset inherently suffers from selection bias: it is not a random sample of hashtags trending during this time period. Rather, it is a compilation of all the covid-19-related hashtags we spotted trending in the Middle East region on Twitter during late March to late April, 2020.

To detect bots, we search for birth anomalies -- an approach pioneered in the Middle East region by researcher Marc Owen Jones, and favored by other researchers beyond the region (DFRL 2017). The broader logic behind this approach is that bot-herders, in order to substantially impact a hashtag, have to marshall many bots to tweet in concert on the same hashtag. If these bot brigades exhibit *group* characteristics (whether in their metadata or their behavior) that distinguish them from organic users, then we can identify them. This approach is distinct from, say, the approach of Botometer (Rauchfleisch and Kaiser 2020), which seeks to assess a user's authenticity based on the user's *own* characteristics.

Within this category of group-based bot detection, the birth anomaly approach is yet further distinguished by implicitly piggy-backing on Twitter's internal detection schemes. The idea here is that Twitter has far more data on users than outside researchers have access to -- for example, Twitter knows the IP addresses from which accounts are being operated, so it can detect if hundreds or thousands of accounts are being operated from the same IP. Based on this privileged data access, Twitter often suspends or provokes the deletion of suspicious accounts. Bot-herders, whether they be employees of regimes, or private firms who offer to manipulate social media discourse on behalf of clients, must therefore re-build their brigades in the wake of Twitter's periodic purges. This leads, so the thinking goes, to notable 'spikes' in account creations, as hundreds of accounts are born on the same day. When these accounts are subsequently deployed together to tweet on a particular hashtag, they create discontinuities in the hashtag's birth histogram, drawing our attention as researchers.

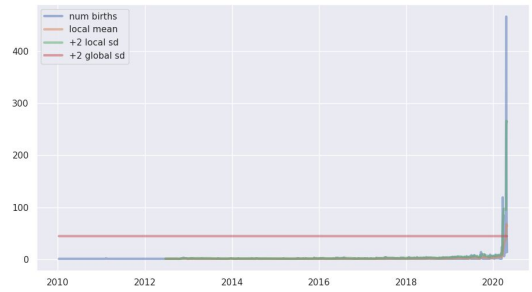

Figure 5: Example of a birth chart, in this case for the Arabic, pandemic-related hashtag #حياتنا_أهم_من_التعليم ("our lives are more important than [going to school]"), collected March 27 - April 25, 2020. The tall blue spikes on the 18th and 19th of April, 2020, correspond to the births of some 336 and 466 accounts, respectively, which are well outside two standard deviations of the global and local (30-day rolling window) means.

Figure 5 depicts an example of a birth chart, in this case corresponding to the Arabic, pandemic-related hashtag #حياتنا_أهم_من_التعليم ("our lives are more important than [going to school]"). The chart tallies the total number of account births per day (colored blue), across all accounts that participated in the hashtag. Any day that exceeds two standard deviations over the global mean number of births (the red horizontal line) is globally statistically anomalous. The reader will note, however, that Twitter's popularity may trend over time or vary seasonally in a manner which renders whole months or years above the global anomaly threshold. To account for this, we define a rolling window of approximately 1 month (previous 30 days), and calculate a rolling threshold of 2 rolling standard deviations above the rolling mean (colored green). If the number of births on a given day (blue) exceeds this rolling threshold (green), then we classify that spate of births as locally statistically anomalous. We conclude that accounts are anomalously born if they were born during spate of births that is *both* globally and locally anomalous within the birth chart of the hashtag in question.

While there may be extenuating circumstances that can explain away anomalous birth spikes, it seems clear from our dataset that anomalously born accounts exhibit distinct characteristics from non-anomalously born accounts in each hashtag of our sample. Across our 149 hashtags, we find that non-anomalously born accounts are generally older, have larger counts of followers and friends, have posted more statuses, and are more often verified, than their anomalous counterparts (Figure 6).

[6] For more information on the options and limitations of Twitter's REST API, we recommend (Steinert-Threlkeld 2018)

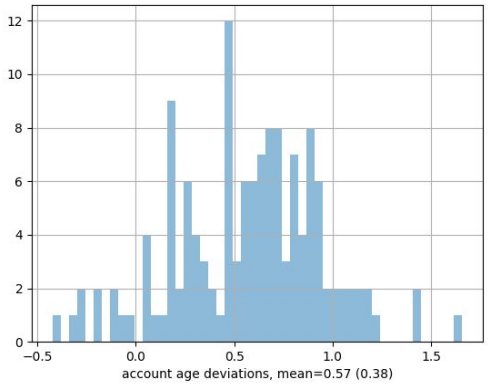

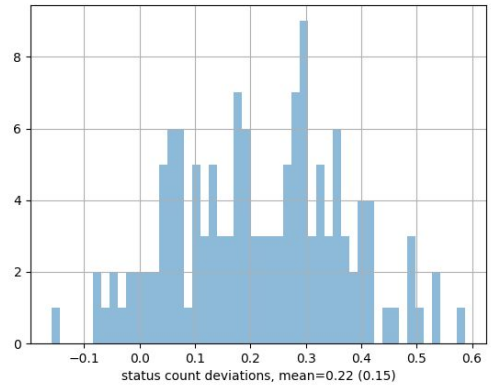

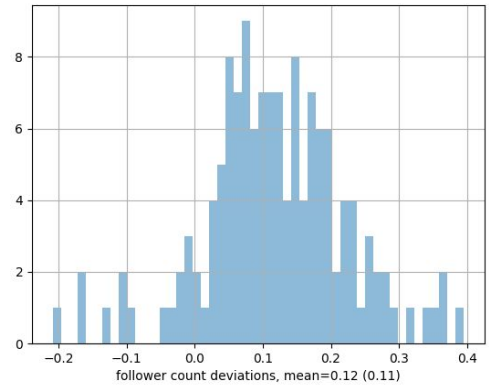

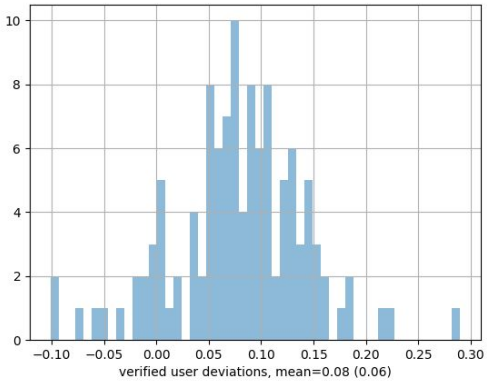

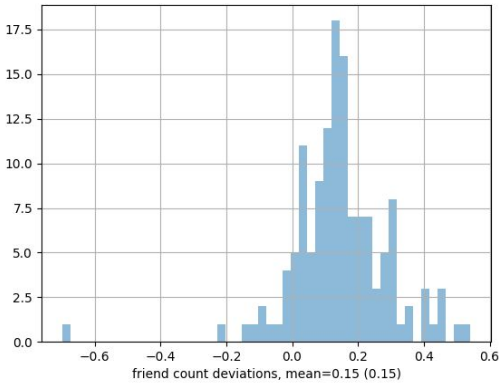

Figure 6: Across our sample of 149 hashtags, non-anomalously-born accounts tend to be older, and (perhaps consequently) exhibit higher follower counts, friend counts, status counts, and are more often 'verified' by Twitter.

Limitations

The broader implications of the study must of course be tempered by the fact that we only focused on one platform (Twitter), one region (the Middle East), a short time period (late March to late April, 2020), and only investigated one dimension of social media manipulation (bots). Even within this narrow category, the reader should note caveats to our research design. Our sample of 149 search filters was hand-picked based on what we saw trending; not randomly, representatively drawn. Thus, while we have no particular reason to suppose that selection bias is driving our results in one direction or other, the sample is clearly selected, not random. Future studies ought to improve upon this by developing a credible random sampling scheme.

Secondly, in order to track so many search terms in parallel, we devised an approach that relied on Twitter's REST API, which we would periodically query for tweets mentioning each of our search terms. Since the REST API is inherently retrospective, we cannot rule out the possibility that Twitter may have been vigilantly eliminating bot networks during the interim. Indeed, a

conclusion consistent with our findings is that Twitter has largely mitigated the problem of misinformative bot networks with its in-house filtering techniques.

A related caveat to using the REST API is that it can only recall tweets from the previous 10 days. Especially on high-volume search terms, this occasionally means that we end up with gaps in our data, where we missed tweets between our previous and current tranches. These lapses are fairly small, however, rarely exceeding 10% of the time period of investigation (roughly 3 days over the month-long period we studied here). There is no reason to believe that these lapses were during critical moments when misinformation campaigns were most likely to be observed.

Finally, from among these 149 filters, we chose to investigate three networks in-depth. We chose only three due to our own personal resource limitations; the possibility remains that had we investigated more networks in our dataset, we may have discovered more nefarious activity. And while we were blind *a priori* to what our findings might be, we surely did not choose these three at random.

Given all of these caveats, we encourage researchers to view our results with caution. Nevertheless, we hope these results stimulate further investigations along these lines of inquiry.

References

Abrahams, A. 2019. Regional Authoritarians Target the Twittersphere. *MERIP: Middle East Research and Information Project.* Fall/Winter 2019. https://merip.org/2019/12/regional-authoritarians-target-the-twittersphere/

Abrahams, A., and Leber, A. 2020. Electronic Armies or Cyber Knights? The origins of pro-authoritarian discourse on Middle East Twitter. Unpublished manuscript. https://sites.google.com/site/alexeiabrahams/alexeis-files-for-download/abrahams_leber_ijoc_draft_feb29_2020.pdf

Abrahams, A., and van der Weide, R. 2020. Ten thousand whispering: inequality of voice on Twitter. Unpublished manuscript.

AlJazeera. 20 Sept 2019. Twitter suspends thousands of fake accounts from UAE. *AlJazeera.* https://www.aljazeera.com/news/2019/09/twitter-suspends-thousands-fake-accounts-uae-190920083258996.html

Allyn, B. May 20, 2020. Researchers: Nearly Half Of Accounts Tweeting About Coronavirus Are Likely Bots. *NPR*. https://www.npr.org/sections/coronavirus-live-updates/2020/05/20/859814085/researchers-nearly-half-of-accounts-tweeting-about-coronavirus-are-likely-bots

Digital Forensic Research Lab (DFRL). August 28, 2017. #BotSpot: Twelve Ways to Spot a Bot. *Medium*. https://medium.com/dfrlab/botspot-twelve-ways-to-spot-a-bot-aedc7d9c110c

Gadde, V. and Derella, M. 16 March, 2020. An update on our continuity strategy during COVID-19. *Twitter blog*. https://blog.twitter.com/en_us/topics/company/2020/An-update-on-our-continuity-strategy-during-COVID-19.html

Hubbard, B. Nov. 7, 2019. Why Spy on Twitter? For Saudi Arabia, It's the Town Square. *The New York Times*. https://www.nytimes.com/2019/11/07/world/middleeast/saudi-arabia-twitter-arrests.html

Jones, M. O. 2019. The gulf information war| propaganda, fake news, and fake trends: The weaponization of twitter bots in the gulf crisis. *International Journal of Communication*, 13, 27. https://ijoc.org/index.php/ijoc/article/view/8994.

Jones, M. O. September 15, 2019. Saudi, UAE Twitter takedowns won't curb rampant disinformation on Arab Twitter. *The Washington Post*. https://www.washingtonpost.com/politics/2019/09/25/saudi-uae-twitter-takedowns-wont-curb-rampant-disinformation-arab-twitter

Jones, M.O., and Abrahams, A. June 5, 2018. A plague of Twitter bots is roiling the Middle East. *The Washington Post*. https://www.washingtonpost.com/news/monkey-cage/wp/2018/06/05/fighting-the-weaponization-of-social-media-in-the-middle-east/

Kao, J. and Shuang Li, M. March 25, 2020. How China Built a Twitter Propaganda Machine Then Let It Loose on Coronavirus. *ProPublica*.

https://www.propublica.org/article/how-china-built-a-twitter-propaganda-machine-then-let-it-loose-on-coronavirus

Molter, V. and Webster, G. March 31, 2020. Coronavirus Conspiracy Claims: What's Behind a Chinese Diplomat's COVID-19 Misdirection. Internet Observatory, Cyber Policy Center. *Stanford*. https://cyber.fsi.stanford.edu/io/news/china-covid19-origin-narrative

Rauchfleisch, A. and Kaiser, J. March 2020. The False Positive Problem of Automatic Bot Detection in Social Science Research. *Berkman Klein Center Research Publication* No. 2020-3. SSRN preprint. https://ssrn.com/abstract=3565233 or http://dx.doi.org/10.2139/ssrn.3565233

Shahbaz, A. and Funk, A. 2019. Freedom on the Net 2019. The Crisis of Social Media. *Freedom House*. https://freedomhouse.org/sites/default/files/2019-11/11042019_Report_FH_FOTN_2019_final_Public_Download.pdf

Steinert-Threlkeld, Z. C. 2018. Twitter as data. Cambridge: *Cambridge University Press*.

The importance of no evidence. *Nat Hum Behav* 3, 197 (2019). https://doi.org/10.1038/s41562-019-0569-7

World Health Organization. 2020. Novel Coronavirus(2019-nCoV) Situation Report-13. *World Health Organization*. https://www.who.int/docs/default-source/coronaviruse/situation-reports/20200202-sitrep-13-ncov-v3.pdf