


Sustained Online Amplification of COVID-19 Elites in the United States

Ryan J. Gallagher¹ , Larissa Doroshenko¹ , Sarah Shugars², David Lazer¹, and Brooke Foucault Welles¹

Social Media + Society
April-June 2021: 1–16
© The Author(s) 2021
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/20563051211024957
journals.sagepub.com/home/sms


Abstract

In the absence of clear, consistent guidelines about the COVID-19 pandemic in the United States, many people use social media to learn about the virus, public health directives, vaccine distribution, and other health information. As people individually sift through a flood of information online, they collectively curate a small set of accounts, known as crowdsourced elites, that receive disproportionate attention for their COVID-19 content. However, these elites are not all created equal: not all accounts have received the same attention during the pandemic, and various demographic and ideological groups have crowdsourced their own elites. Using a mixed-methods approach with a panel of Twitter users in the United States over the first year of the COVID-19 pandemic, we identify COVID-19 crowdsourced elites. We distinguish sustained amplification from episodic amplification and demonstrate that crowdsourced elites vary across demographics with respect to race, geography, and political alignment. Specifically, we show that different subpopulations preferentially amplify elites that are demographically similar to them, and that they crowdsource different types of elite accounts, such as journalists, elected officials, and medical professionals, in different proportions. In light of this variation, we discuss the potential for using the disproportionate online voice of crowdsourced COVID-19 elites to equitably promote public health information and mitigate misinformation across networked publics.

Keywords

crowdsourced elites, amplification, networked publics, communication networks, COVID-19

Introduction

Local COVID-19 caseloads, community guidelines, national directives, international travel bans, testing locations, and vaccine availability all change rapidly, sometimes on a daily or even hourly basis, as we endure the pandemic. The fluidity of the ongoing public health crisis makes it difficult for individuals to find reliable information online, forcing them to sift through both high-quality expert information and life-threatening misinformation, particularly on social media platforms. As people individually engage with an imperfect assortment of local news media, epidemiological specialists, political pundits, public health officials, and financially motivated opportunists online, they collectively amplify a small set of disproportionately influential COVID-19 information sources, or *crowdsourced elites* (Papacharissi & de Fatima Oliveira, 2012). Crowdsourced elites were originally theorized by Papacharissi and de Fatima Oliveira within the context of Twitter and the 2011 Egyptian revolution to describe the process of how bloggers, activists, and ordinary citizens gained visibility and “elite” status by being amplified by

other ordinary Twitter users. Subsequently, the concept of crowdsourced elites has been used widely to explain organic grassroots amplification and the emergence of non-traditional elites on social media during other mass protests, elections, and social justice movements (Jackson & Foucault Welles, 2016; Jungherr, 2016; Papacharissi & de Fatima Oliveira, 2012). With respect to COVID-19, crowdsourced elites are key pillars in the online information ecosystem because they are the most capable of disseminating both critical health information and divisive anti-scientific opinions.

However, different publics amplify different information sources, meaning that different publics crowdsource different elites. While prior work has demonstrated the multiplicity of

¹Northeastern University, USA

²New York University, USA

Corresponding Author:

Ryan J. Gallagher, Network Science Institute, Northeastern University, 177 Huntington Ave., Boston, MA 02115-5005, USA.
Email: ryanjgallag@gmail.com



networked publics on social media platforms (Jackson et al., 2020; Shugars et al., 2021), quantitative approaches still typically identify crowdsourced elites in aggregate without distinguishing how they vary in prominence and visibility across different subnetworks. Furthermore, we also often only study crowdsourced elites around particular focal events, and operationalize amplification accordingly by counting reshapes—or retweets on Twitter. While counting reshapes measures the *volume* of amplification, it does not measure its *consistency*. This distinction is critical when we consider crowdsourced elites with respect to sustained events, like the pandemic, rather than more episodic conversations that quickly emerge and dissipate online. If we do not account for the sustained nature of amplification or the diversity of crowdsourced elites across different publics, we will be unable to produce reliable maps of public dialogue around ongoing events like the pandemic, presidential elections, or social justice movements.

Here, within the context of COVID-19, we identify elites that have been crowdsourced through sustained amplification across different publics on Twitter. We first introduce a measure of sustained amplification. Using a panel of Twitter users matched to public US voter registration files over the first year of the pandemic, we apply our measure to enumerate nearly 1,900 crowdsourced COVID-19 elites—both nationally and across different demographic, geographic, and ideological subgroups—and show how they differ from those produced by episodic amplification. Nationally, we find that journalists, media outlets, and political accounts are most often crowdsourced as COVID-19 information sources, while epidemiologists, public health, and medical professionals make up only a small portion of all COVID-19 elites. There is also considerable variation in COVID-19 elites across demographics though, which accentuates racial, political, and geographic homophily among elites and their corresponding demographic publics. We conclude with how our findings can be leveraged to disseminate high-quality public health information among marginalized populations most devastated by the pandemic, and how our methods can be used to measure crowdsourced elites across sustained conversations and networked publics more generally.

Related Work

Networked Gatekeeping and Crowdsourced Elites

The networked mix of journalists, media outlets, politicians, celebrities, and everyday people on Twitter makes it a critical conduit of both credible information and misinformation. The platform plays a particularly important role during breaking news events and rapidly evolving situations when mainstream media are not able to quickly adapt and provide on-the-ground, authoritative information (Hu et al., 2012), or when those outlets are blocked, restricted, or mistrusted (Howard, 2010; Papacharissi, 2009). In these information gaps, ordinary users and alternative sources are able to rise in

prominence, acting as on-the-ground citizen journalists (Jackson & Foucault Welles, 2016; Lotan et al., 2011), resource coordinators (Pourebrahim et al., 2019), and afar commentators (Jackson et al., 2020; Tufekci, 2017). As traditional broadcast media outlets catch up to these events, they often rely upon these emergent sources when framing their news stories (Chadwick, 2011; Molyneux & McGregor, 2021), treating them as representatives of the ensuing online and offline conversations (McGregor, 2019). The synergy between social media and news media can help direct critical information (Tufekci, 2017) and shine light on important social issues (Jackson & Foucault Welles, 2016). However, while journalists can be receptive to crowdsourced stories, potentially widening the array of voices represented in mainstream media, they are also vulnerable to media manipulators who use that receptiveness to inject misinformation into still fluid and fragile situations (Lukito et al., 2020; W. Phillips, 2018). So although Twitter can be an important source of credible information, there is a delicate balance involved in properly curating and amplifying that information, particularly during emerging and evolving events.

Journalists play an important role in striking that balance, as gatekeepers that select which information is broadcasted through traditional print and news media outlets (Shoemaker & Vos, 2009). Despite their disproportionate influence though, they do not assert unique control over what gains attention on social media. They are positioned among many other regular users who curate information by choosing what to share or ignore on their own feeds. Because social media affords the ability to amplify particular pieces of content and not others (Meraz, 2009), they act together as *networked gatekeepers*, determining what others—including journalists—see on a platform (Barzilai-Nahon, 2008; Meraz & Papacharissi, 2013). Through their individual acts of curation and filtering, users collectively amplify *crowdsourced elites*, that is, accounts that gain the most visibility through the networked gatekeeping process (Jackson & Foucault Welles, 2015; Papacharissi & de Fatima Olivera, 2012). At times, these crowdsourced elites may be traditional political and media elites or those who already have established followings (Goel et al., 2016), but Twitter's particular affordances allow for more permeability in who is deemed "elite." As such, domain experts, social justice activists, fringe antagonists, and ordinary users can also emerge as crowdsourced elites during acute and rapidly evolving events (Jackson & Foucault Welles, 2016; Lotan et al., 2011; Pourebrahim et al., 2019), and have their local, technical, and situational knowledge placed alongside traditional elites (Chadwick, 2011; Molyneux & McGregor, 2021). Overall, the complex interactions of networked gatekeeping with mainstream media produce social media conversations that are an agglomeration of both established elites and emergent crowdsourced elites.

As an ongoing public health crisis consisting of a constant stream of evolving news, the COVID-19 pandemic sits at the

delicate intersection of mainstream media, scientific experts, locally affected citizens, and financially motivated opportunists. In the United States, there has been a lack of unanimous, authoritative guidance from federal and local health officials regarding the pandemic. In its place, elected congressional officials have broadcasted polarized messages about the virus (Green et al., 2020). Together with the historically low trust in mainstream news (Fink, 2019), this has prompted many to turn to social media—and Twitter in particular—for information about COVID-19 (Fisher & Weiner, 2020). Although it is possible to find reliable health information and updates about the pandemic online, it is generally variable in quality and can lead people to make erroneous conclusions that are not grounded in scientific fact (Fisher & Weiner, 2020). So in the absence of top-down, authoritative messaging informed by epidemiologists, medical professionals, and public health officials, the uncertainty around COVID-19 creates space for politically and financially motivated actors to rise through the networked gatekeeping process and disseminate conspiracy theories and misinformation (Bridgman et al., 2020; Donovan, 2020) that further politicize pandemic responses (Allcott et al., 2020; Enders et al., 2020). To measure the extent to which various benevolent and malicious actors have gained visibility during the pandemic, we need effective methods for identifying crowdsourced elites.

Identifying Sustained Crowdsourced Elites Across Publics

Amplification is often operationalized by counting the number of reshares—or retweets in the context of Twitter—that an account receives. This is not without reason: it is easily calculated, interpretable, and effective at identifying crowdsourced elites across contexts. Its simplicity, though, hides richer contextual and temporal dynamics.

First, measuring amplification by counting reshares does not account for diversity among the amplifiers. While reshares are often counted in aggregate across entire datasets, many online communication networks are polarized (Conover et al., 2011; Yardi & boyd, 2010) or otherwise fragmented. While such fragmentation is not necessarily bad, it does have implications for the networked gatekeeping process. Because different networked publics tend to amplify different sources, they can also have different crowdsourced elites. So, methodologically, if we only produce a single ranked list of crowdsourced elites across an entire dataset, we may overestimate the importance of general elites for some publics, or underestimate the importance of specific elites for other publics. This has downstream consequences: behavioral change campaigns—including those designed to spread accurate information about things like health risks, treatment, and vaccination—are considerably more effective when trusted opinion leaders are enlisted to encourage their adoption and compliance (Valente & Davis, 1999; Xu et al.,

2014). Community-specific crowdsourced elites can trigger larger informational cascades (Cha et al., 2010) and reach more people with public health information. When identifying crowdsourced elites, we must account for how different publics amplify different accounts.

Second, measuring amplification by only counting reshares collapses several possible outcomes of amplification into one number. This conflates three types of visibility that may be possible through amplification. One, some accounts receive many shares because they have a single post “go viral,” momentarily elevating them on the platform. Two, other accounts, particularly spam accounts, may garner many shares because they produce an abundant amount of low-quality posts that each happen to get a little amplification. Third and finally, an account may sit between these two extremes: it may not always receive abundant amplification at the magnitude of virality, but it consistently posts about a topic and receives attention. When we are studying the conversations around a *particular* focal event, like the announcement of a vaccine, a presidential campaign gaffe, or an extrajudicial police murder, the distinction between amplification outcomes is less important: an account is seen widely at a particular moment whether it had one viral post or many impactful ones. However, when we consider *ongoing* conversations online, like those around the pandemic, US elections, or Black Lives Matter more broadly, the differences between *sustained* and *episodic* amplification become more acute. Health practitioners would not want to seed a campaign with an account that was only viral at the beginning of the pandemic, and platforms would not want to expend significant effort suspending spam accounts that receive little amplification. By measuring sustained amplification of crowdsourced elites across different publics, we help ensure that relevant crowdsourced elites can be employed for health campaigns or targeted for platform interventions.

We focus primarily on three research questions.

1. In the United States, who have people crowdsourced as COVID-19 elites on Twitter over the course of the pandemic?
2. How does our understanding of who those crowdsourced elites are change depending on whether we measure them according to sustained or episodic amplification?
3. How do COVID-19 crowdsourced elites vary across different demographics, including race, gender, age, state of residence, and political affiliation?

Data and Methods

We use a mixed-methods approach for identifying the sustained amplification of COVID-19 crowdsourced elites. Drawing on a panel of Twitter users matched to public US voting registration records, we use a set of keywords to identify tweets about COVID-19 posted during the first year of

the pandemic in the United States. We introduce a measure of sustained amplification and use it to enumerate *sustained COVID-19 crowdsourced elites* nationally and across various demographically induced publics. Having identified these elites, we undertake a qualitative hand-coding process to annotate them according to their Twitter account type, race, gender, political affiliation, and geographic locality. We use these annotations to demonstrate how sustained amplification differs from episodic amplification, and how elites differ across demographic publics.

Twitter Panel of Registered Voters

To focus our analysis on real US citizens, we use a panel of Twitter users linked with public US voter registration records. Matching Twitter users to these records has several benefits over other methodological approaches. First, we can be confident that we are studying the online behavior of real people, and not that of media outlets, organizations, social bots, and other non-human entities (Ferrara et al., 2016; Gorwa & Guilbeault, 2020). Second, using a voter registration file gives us demographic information about the Twitter users—which is not otherwise available using the Twitter application programming interface (API)—and so we are able to study demographic variation in how users crowdsource elites. Finally, the size of the panel allows us to focus our analyses on demographic minorities and other marginalized populations while retaining a statistically reliable sample size (Foucault Welles, 2014).

We constructed the panel following the approach of Grinberg et al. (2019). We started with a 10% sample of tweets collected from the Twitter Decahose. From this sample, accounts were extracted if they reported a US location in their profile and included a full name in their account or handle. These accounts were then joined with public voter registration files—sembled by the data vendor TargetSmart—by matching unique entries according to both name and location. In total, the matching process resulted in a Twitter panel of about 1.6 million registered voters across all 50 states and Washington, D.C. We estimate this to be about 3% of all adult Twitter users in the United States (Perrin & Anderson, 2019).

We make use of several demographic attributes from the panel data: state, race, gender, age, and political affiliation. The state of residence, age, and gender of our panel members come directly from the voter registration records. For race, only the nine states required by the Voting Rights Act consistently report race through voter registration records. For the remaining states, we use race as provided by TargetSmart, which is reported by the company to be determined through a combination of statistical inference and matching to other commercial data files. For party affiliation, we use TargetSmart's estimate of party score, which is a statistical estimate based on registered political party and a number of other variables, including number of votes in Republican and

Democrat primaries and Federal Election Commission (FEC) contributions. Party status can vary notably across states, and so this score helps maintain more consistency across states. It falls within an interval from 0 to 100, indicating the probability that a registered voter supports the Democratic party. Following TargetSmart's recommended guidance, we classify any individual with a score less than 35 as a Republican and more than 65 as a Democrat. The remaining panel members are classified as Independent. In separate work, we have validated that TargetSmart's proprietary inferences of both race and political affiliation are reliable (Shugars et al., 2021).

There are a number of factors in the construction of the panel that affect its final composition. Grinberg et al. (2019) compared the demographics of their panel with a survey of voters on Twitter by Pew Research and found that the composition of their panel largely agreed with the survey in terms of gender, political affiliation, and race, though it contained slightly more Democrats—which was explained by the number of users reporting no party. Hughes et al. (2021) conducted a more detailed comparison of the voter file approach to several survey-based approaches that constructed panels through voluntary responses. Compared with survey-constructed Twitter panels, the voter file panel contained younger users, more women, more White users, and fewer Hispanic and Asian users. Given these studies and the construction of our panel, we confidently report results for panelists with respect to state, age, and gender because there is no overwhelming over or underrepresentation. We are also confident of the results with respect to White and Black panelists, but results related to COVID-19 elites of Hispanic and Asian panelists should be interpreted more tentatively as they are notably underrepresented in our panel.

Ethical Considerations

The construction and use of the Twitter panel was approved by our university's Institutional Review Board. To create the panel, we used names and locations that were clearly and publicly disclosed on Twitter, as well as publicly available voter registration records. However, many Twitter users do not anticipate having their information extracted from Twitter and linked across datasets (Fiesler & Proferes, 2018), and their explicit consent was not obtained to do so. We balance this against the other methodological approaches that we could have possibly taken. One option would have been to work generally with a sample of Twitter data, based just on keywords. With this approach though, we would not be able to reliably identify human users and, among those, different demographics. Given that different demographic groups—particularly racial ones—have disproportionately borne the burden of the pandemic (Azar et al., 2020), it is important to identify elites that are trusted and amplified among those communities. A demographic-agnostic approach would not proportionately benefit those most harmed by the pandemic,

which would be an unjust research outcome. Alternatively, we could have constructed a panel through a survey allowing for informed consent, but it would have resulted in a substantially smaller panel (Hughes et al., 2021). While we would have access to demographic data, our identification of elites would be less reliable because we have fewer users for any particular demographic subset. In this case, we would not be able to produce research that adequately benefits those most affected by COVID-19.

In light of these considerations, we believe that our voter file approach is most capable of identifying COVID-19 crowdsourced elites in such a way that we can justly distribute the outcomes of our research. We defend against misuse of the data by storing and analyzing it on secure servers with multiple checkpoints of restricted access. We do not include private accounts in our analyses, and only report aggregate trends across the panel. These precautions minimize the likelihood that harm will come to our panel members because we linked their Twitter profiles with their voter registration records.

COVID-19 Tweet Identification

Using the Twitter panel as our source of tweets, we identify COVID-19-related content through a keywords-based approach. We compiled a multilingual keyword list from three sources: the keyword list for the COVID-19 Twitter dataset assembled by Chen et al. (2020), the keyword list from Green et al.'s (2020) study on elite polarization around COVID-19, and the keyword filters for Twitter's official COVID-19 streaming endpoint.¹ As the pandemic evolved, we continued to augment those lists with additional terms to expand our coverage across COVID-19-related topics and misinformation. We also removed select words from the lists that were likely to produce false positives as the pandemic continued (e.g., "china"). After several iterations, we further removed keywords that rarely uniquely identified COVID-related tweets: specifically, we dropped all keywords that appeared by themselves and no other keywords less than 30 times. In total, our list consisted of 570 keywords, including terms explicitly mentioning the pandemic ("covid-19" and "coronavirus"), phrases referencing societal responses ("social distancing" and "flatten the curve"), names for different types of masks ("n95," "surgical mask," and "face covering"), and hashtags associated with misinformation ("plandemic," "faucifraud," and "arrestbillgates2020"). The full keyword list is available as an online supplement (covid_keywords.tsv).

We labeled a tweet as COVID-related if at least one keyword matched the tweet text, including any quoted text, any hashtags, or any substring of any URL included in the tweet. We also identified "COVID-related" URLs and included any tweet using at least one of those URLs. A URL was labeled as "COVID-related" if it was used with our keywords at least 100 times and at least 20% of its use was with COVID

keywords—even if the URL itself did not contain any COVID keywords. These thresholds were chosen by manually inspecting the most frequently occurring URLs and domains used by our panel and iteratively altering the thresholds until they appeared to filter non-COVID URLs while still returning relevant links.

We searched for the keywords and COVID-related URLs in all tweets posted by our panel members from January 1, 2020, to December 31, 2020. This resulted in 34,149,869 tweets posted by 504,096 panel members. Of those 9,840,552 were originally authored tweets and 24,309,317 were retweets, where 5,402,637 were retweets of quote tweets. The vast majority of the tweets (98%) were posted after March 1, 2020, around which time COVID-19 first became a major point of conversation in the United States.

Sustained Amplification

To identify *sustained* amplification around COVID-19, we define a retweet *h*-index. The *h*-index, originally proposed for measuring academic citations (Hirsch, 2005), is the maximum value of *h* such that an account has *h* COVID-related tweets that each have at least *h* retweets. This metric balances the volume and consistency of amplification: in order for an account to have a high retweet *h*-index they must produce many tweets about COVID-19 *and* receive notable attention for each of those tweets. If an account did not consistently post about the pandemic and only had a single viral tweet, then they will only have an *h*-index of 1. If an account made many posts about COVID-19, but did not receive more than one retweet on any of them, then they will also only have an *h*-index of 1. However, if an account had 20 COVID-related tweets that each had at least 20 retweets, then they will have an *h*-index of 20. The *h*-index discerns sustained amplification from episodic amplification by prioritizing volume *and* consistency, mitigating the noisy effects of virality and spam when measuring amplification for ongoing events like the pandemic.

Identifying COVID-19 Crowdsourced Elites

We want to demonstrate the utility of measuring sustained amplification through the *h*-index compared with measuring episodic amplification by counting retweets. So for every account retweeted by one of our Twitter panel members, we calculated both the total number of COVID-19-related retweets received by the panel and their COVID-19 *h*-index. We then ranked all the accounts by each of these measures, yielding two ranked lists from which we extract the top 0.1% of accounts. We did this nationally across the entire panel.

We also want to demonstrate the utility of identifying crowdsourced elites specific to different publics. While there are a number of ways to identify (networked) publics, we choose to use demographically induced publics, given the demographic variation in COVID-19 health disparities and

Table 1. The Typology of Account Types of Crowdsourced Elites.

Account type	Examples	Political affiliation?	Geographic locality?
Public health official	Surgeon general, director of CDC		
Medical professional	Doctors, nurses, hospitals		
Epidemiologist	Epidemiological researcher		
Elected official	Mayors, governors, senators	X	X
Public service	City offices, police departments		
Media outlet	News outlets, TV news channels	If hyperpartisan	X
Journalist (hard news)	Investigative journalists, public health and economics reporters	If hyperpartisan	X
Journalist (soft news)	Sports and entertainment reporters	If hyperpartisan	X
Journalist (broadcast)	TV anchors, radio news hosts	If hyperpartisan	X
Journalist (new media)	Twitch streamers, podcast hosts	If hyperpartisan	X
Public intellectual	Academic researchers, mainstream opinion columnists		
Political (other)	Activists, campaign staffers, political personalities, anonymous hyperpartisan accounts, political pundits	X	
Entertainer	Musicians, comedians, social media personalities		
Athlete	Baseball players, sports team managers		
Religious leader	Priests, rabbis, churches		
Organization	Organizations not classified above		
Other	Accounts not classified above, primarily personal accounts of non-public figures		
Deactivated/suspended	Accounts deactivated or suspended by the time of annotation		

CDC: Centers for Disease Control and Prevention. Only certain types of accounts were annotated with a political affiliation or geographic locality, as indicated in the table.

vaccine hesitancy. As such, we also constructed the two ranked lists—one ranked by retweet counts and another by *h*-indices—for each demographic of age, race, gender, state of residence, and political affiliation. For example, we identified the COVID-19 elites crowdsourced by residents of the state Massachusetts by calculating the *h*-index using retweets from only panel members in Massachusetts. For demographics where the top 0.1% resulted in less than 50 accounts, we included all 50 of the top accounts, beyond just the top 0.1%. We did this so that we had enough elites to reliably disaggregate them across account types, which we explain below. Overall, across both amplification metrics and all demographics, this process yielded 1,896 unique accounts: 1,580 accounts are crowdsourced elites according to total retweets, and 1,455 are sustained crowdsourced elites according to the *h*-index.

Classifying COVID-19 Crowdsourced Elites

To provide a meaningful differentiation of COVID-19 crowdsourced elites, we annotated them according to account type, race, gender, political affiliation, and geographic locality. Two authors used their prior general knowledge of Twitter, along with a related report on “COVID-19 “expertise” on Twitter (Gligorić et al., 2020), to devise a classification scheme of different types of accounts that can act as COVID-19 opinion or information sources. Upon multiple iterations, the typology was finalized such that an account could be categorized as 1 of 18 account types (see Table 1 for

details). Journalists could be labeled according to whether they report hard news, report soft news, broadcast news, or publish through new media. Both journalists and media outlets could be identified as hyperpartisan. Hyperpartisan journalists and media outlets, elected officials, and other political accounts were annotated with political affiliations. While other accounts can also express political views, we focused on identifying accounts that are primarily political actors. We also labeled other types of accounts with a geographic locality if relevant: elected officials, media outlets, and journalists. While other accounts may have geographic information, we focused on extracting it for people’s local officials and news sources. We list the account types, summarize the annotations, and provide examples in Table 1.

The two authors formalized this typology into a codebook. Three authors and seven research assistants were trained on it and instructed to use publicly available information to research the COVID-19 crowdsourced elites, identify their demographic attributes, and classify them according to the typology. The codebook was updated iteratively based on feedback from the coding team, and all data were re-coded following codebook updates. Any ambiguous cases, as well as all cases marked “Other” were reviewed by at least two coders prior to finalizing a code. Krippendorff’s alpha was calculated for a 5% random sample of the data that were coded by at least one author and one research assistant. For account types, Krippendorff’s alpha was 0.78, which is a reasonable indicator of intercoder reliability given the number of account types in the typology.

Results

Sustained Amplification of COVID-19 Elites

For our comparison of how counting retweets and the retweet h -index differ in how they measure amplification, we focus on the national-level COVID-19 elites. Each measure separately ranked the top 0.1% of accounts and produced 1,110 national elites. Together between the two metrics, there are 1,298 unique elites, where 922 of them (71%) are shared between the h -index and retweet count rankings. President Joe Biden is ranked first according to both his h -index and total number of retweets (217,967 retweets and h -index=273), while former president Donald Trump ranks second according to h -index (h -index=251) and third according to total retweets (147,417 retweets). While we do not otherwise discuss individual accounts here and instead focus on aggregate trends, we make the national COVID-19 elites available as an online supplement (US_national_COVID-19_elites.tsv).

In Figure 1, we show the distribution of account types among both the retweet and h -index rankings. Regardless of the ranking, we see clear trends in what types of accounts have been amplified around COVID-19. Journalists and media outlets command a significant portion of the COVID-19 discussion, followed by political accounts on both the left and right. Notably across both rankings, medical professionals, public health officials, and epidemiologists make up only a small portion of all crowdsourced elites (4%–5%). Relative to the h -index though, counting retweets overestimates the prevalence of media outlets among COVID-19 elites, while underestimating the number of elites that are hard news journalists. To a lesser extent, counting retweets also underestimates the number of right-wing political accounts, and overestimates the prevalence of “other” accounts among the elites.

While the two amplification metrics may look similar if we only look at the distributions of account types, they rank those accounts very differently. For Figure 2, we calculate the median difference between each account’s h -index rank

and retweet count rank. Accounts with a higher difference are ranked higher when counting retweets, while those with a lower difference are ranked higher by the h -index. When ranking by total retweets, media outlets rank 341 positions higher on average compared with the h -index rankings. Medical professionals and public health officials are also overranked by counting retweets as well: they rank 128 and 62 positions higher on average. On the contrary, the total retweet rankings underrank the sustained amplification of deactivated and suspended accounts, Republican elected officials, and other right-wing political accounts.

The median rank difference appears to show that the retweet counts and h -index rank some types of accounts similarly, like left-wing accounts, entertainers, and other accounts. We see a more complex story in Figure 3, which shows how each individual account is ranked by the two metrics. For media outlets, we can clearly see how they are frequently overranked by counting retweets. There are a number of accounts, including left-wing political accounts, medical professionals, public intellectuals, entertainers, and other accounts, that have much lower h -index ranks than their total retweets ranks. These are the accounts that were momentarily amplified, but have not received consistent amplification over the course of the pandemic. These plots provide a clear indicator that counting retweets is not appropriate for measuring amplification over an ongoing event.

Together, Figures 1 to 3 demonstrate that the choice of how we measure amplification notably affects what kinds of accounts are determined to be important. In practice, we often use such rankings for further content analysis, information campaigns, or platform interventions. Failing to account for the sustained nature of the pandemic results in a muddled view of who has been deemed a COVID-19 elite.

Demographic Crowdsourcing of COVID-19 Elites

Having demonstrated the need for measuring the sustained amplification of COVID-19 elites, we now turn to show how




Figure 1. The distribution of COVID-19 crowdsourced elite account types, according to the retweet h -index (blue, left) and their total number of retweets (red, right).

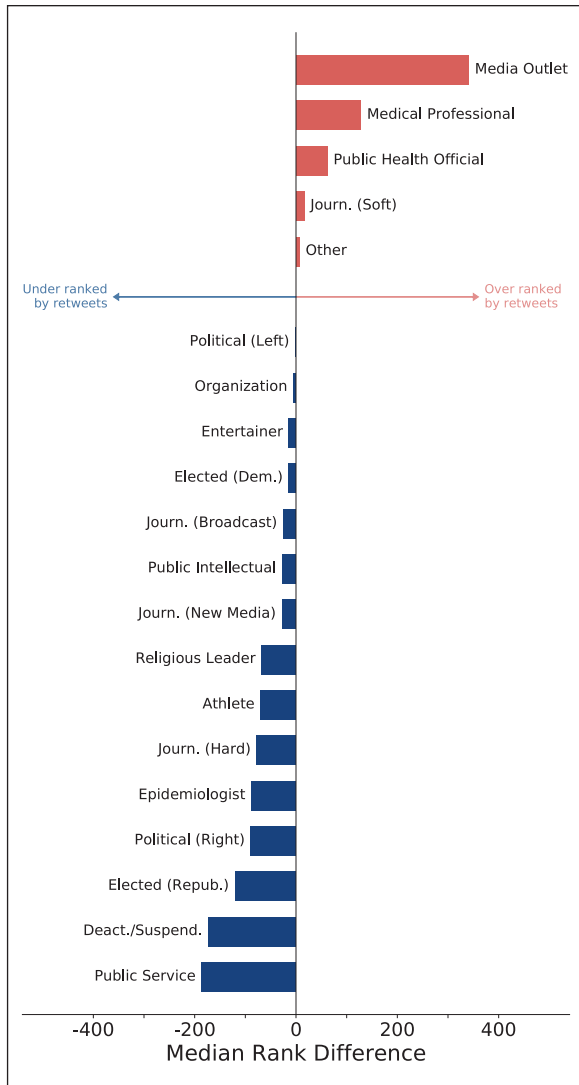


Figure 2. The median difference between *h*-index rank and total retweet rank across all accounts of each type. Positive values indicate that the account was ranked more highly according to its total number of retweets, indicating it is overranked relative to the *h*-index. Negative values indicate that an account was ranked more highly according to its *h*-index, indicating the total number of retweets underranked the account according to the *h*-index.

different demographically induced publics crowdsourced different elites. From here forward, we identify sustained elites by using only the retweet *h*-index. According to the *h*-index, there are 1,401 elites across all demographics. In Figure 4, we map the different types of COVID-19 elites crowdsourced by different demographics.

In line with the national trends, all demographics have consistently amplified many journalists, and few medical professionals, public health officials, or epidemiologists. There are notable demographic-specific trends though. Democrats and Republicans amplify more left- and right-wing accounts as their elites, respectively, which we discuss

in further detail below. By the account type breakdown though, we see that both political demographics crowdsource relatively more general, but not elected political accounts than elected officials themselves. In contrast, non-White panel members amplify more elected Democrats than left-wing accounts, and fewer right-wing accounts overall. Political accounts of all types are crowdsourced more as elites by older users: 30.1% of elites for panel members ages 18–29 are political, 39.3% are for those ages 30–49, 44.4% are for those ages 50–64, and 49.5% are for those ages 65 and older. Opposite this trend, 13.9% of elites for those ages 18–29 are media outlets, while that number falls for those ages 30–49 (8.5%), 50–64 (6.8%), and older (5.5%). There are also more subtle differences in elites. For example, those who are 18–29 years old amplify less public intellectuals and more “other” assorted accounts, while Republicans amplify more accounts that are now suspended or deactivated.

As we see then, there is a rich diversity in the types of accounts that are amplified across different demographic publics. That diversity would be lost if we were to only consider the national elites in aggregate. As we show further in Figure 5, these publics exhibit homophilic amplification, amplifying elites that are like themselves.

In Figure 5a, we look at the COVID-19 elites of Republicans and Democrats in our panel. For Republicans, we measure the percent of their elites that are right-wing accounts (both elected officials and others) and compare that with the percentage nationally. Similarly, for Democrats we measure the percent of their elites that are left-wing accounts and compare with the national elites. Figure 5a shows that both sets of panel members exhibit homophilic amplification with respect to their political parties. When looking specifically at political COVID-19 elites—rather than all elites—82.1% of Democrat’s political elites are left wing (58.5% nationally), while 74.6% of Republicans political elites are right wing (39.7% nationally).

While homophilic amplification with respect to political ideology may be expected, we also find similar patterns across the lines of race and geography (we did not find any preferential amplification according to gender with respect to the national baseline). In Figure 5b, we show that people of color are more likely to have COVID-19 elites that align with their racial demographic, compared with the national elites. Black panel members in particular have significantly more Black COVID-19 elites than the panel as a whole: while only about 8.8% of COVID-19 elites nationally are Black, 25.7% of COVID-19 accounts amplified by Black panelists are themselves Black. While the absolute numbers of Latinx/Hispanic and Asian COVID-19 elites are lower both nationally and for each demographic, Latinx/Hispanic panel members are about three times more likely to amplify Latinx/Hispanic people as elites, and Asian panel members are about two times more likely to crowdsource an Asian elite. Generally, people of color are more likely to have COVID-19 elites that are of color: when looking at accounts

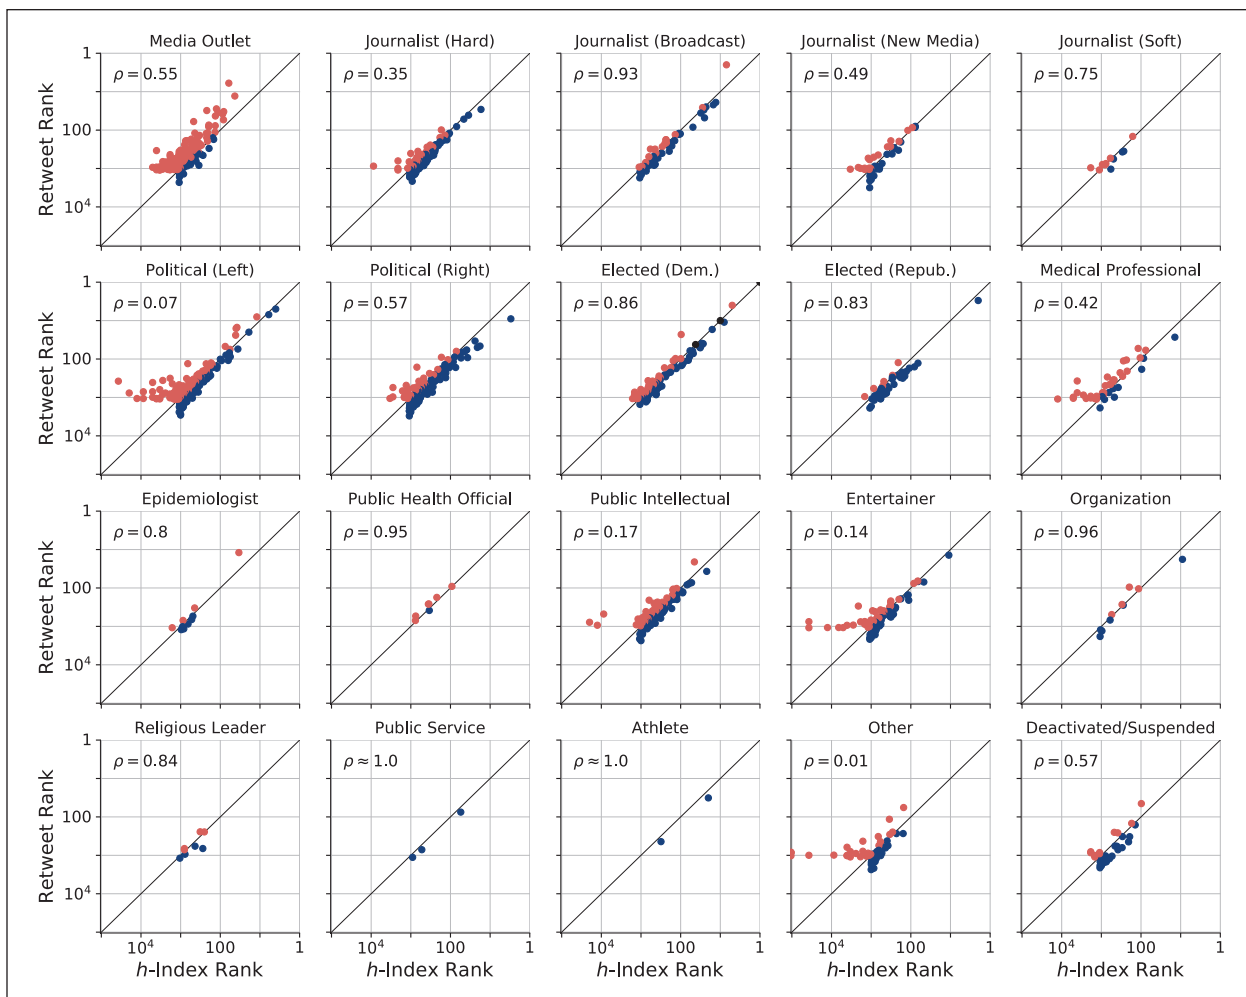


Figure 3. Comparisons of accounts when ranking by their total number of retweets and their h -index across COVID-19 account types. The diagonal line in each plot indicates the line of equality. Accounts above the line of equality in red are ranked higher by their total number of retweets. Accounts below the line of equality in blue are ranked higher by their h -index. The Pearson correlation of the ranks is indicated in the top left of each subplot. Note the log-scale axes, which mean that small visual deviations from the line of equality can be large in magnitude.

that have a race (e.g., as opposed to organizations), it is 2.4 times more likely for the elites of Black panel members to be non-White compared with those of White panel members, and 1.1 and 1.4 times more likely for the elites of Latinx/Hispanic and Asian panel members to be non-White, respectively.

With regard to geography, we find that there are relatively more local COVID-19 elites among each state's elites than there are nationally. Figure 5c shows the difference in percentage points between the percent of a state's elites that are local journalists, media outlets, and elected officials and the percent of all national elites that are from that same state. The average state has 7.1 percentage points more local representation than is seen nationally, where 17 states have no local representation at all among national COVID-19 elites. Even states that are already represented relatively more among national elites, including California (2.3% of all

national elites) and New York (1.2% of all national elites), crowdsource more local elites: 6.1% of all Californian elites and 3.0% of all New York elites are affiliated with California and New York, respectively.

Given the variety of accounts that different demographic publics amplify, and the homophily they exhibit in doing so, it is clear that it is important to account for how different publics amplify different elites. Doing so allows us to identify COVID-19 elites that are relevant and local to particular political, racial, geographic, and age-based populations. Measures which only look at nationally elevated elites disproportionately represent the most dominant populations and overlook the contributions and discourse happening in marginalized populations. Such a misplaced focus may fundamentally misrepresent dynamics of discourse and information spreading and could lead to poorly targeted intervention strategies.




Figure 4. The distribution of COVID-19 elite account types nationally and across different demographics. Darker shades of color indicate that relatively more accounts of that type are crowdsourced as elites by that demographic.

URL Sharing by Sustained COVID-19 Elites

We finish by briefly examining the extent to which the sustained COVID-19 elites used URLs to share information and the amount of amplification they received for doing so. Looking again only at the national COVID-19 sustained elites, we take the tweets of the COVID-19 elites amplified by our panel (i.e., tweets retweeted at least once by our panel), and count the number of them that contained at least one URL. We then calculate what proportion of all their retweets came from those tweets with URLs. If elites received an equal number of retweets on URL and non-URL tweets, then we would expect the number of retweets to be proportional to the number of URL and non-URL tweets they sent, respectively. The results are shown in Figure 6. Perhaps unsurprisingly, media outlets shared the most links relatively: over 88% of their tweets contained a URL. In accord with this, about 80% of all the retweets that media outlets received came from those tweets with URLs. Unlike media outlets though, while most other account types included URLs approximately 30%–40% of the time, they often did not receive a proportional number of retweets from those tweets. In particular, irrespective of political party or status as an elected official, all political accounts linked to external websites in about 30% of their tweets, but only received 15% of all their retweets from those tweets, which is half of what we would expect based on how many URL tweets they sent. So although political accounts have been consistently amplified in COVID-19 discussions on Twitter as we showed earlier, they are more often amplified for non-URL tweets than

tweets directed toward off-platform information. Generally, most other account types received much more amplification for their non-URL tweets relative to the number of URL tweets they sent.

There are some accounts other than media outlets that were amplified at the same rate as the number of URLs they shared. In addition to media outlets, journalists (except soft news journalists) received a proportional number of retweets on their tweets containing URLs. As professional information gatekeepers, this proportionality reflects a balance in the amount of amplification they receive with respect to the amount of off-platform information that they share. The other notable exceptions to the URL amplification imbalance are epidemiologists and public health officials. Both types of accounts included URLs in 36%–39% of their tweets, and both received 39% of their retweets from those tweets with URLs. This suggests that epidemiologists and public health officials receive balanced amplification with respect to URLs, which is only on par with media outlets and journalists. Still though, while epidemiologists and public health officials—those who have direct experience managing the pandemic—may receive amplification proportional to their information sharing, they make up only a small portion of all COVID-19 elites, as we showed at the start of our analysis.

Discussion

We have taken a census of who Americans have given disproportionate attention to during the first year of the pandemic on Twitter around COVID-19 by measuring sustained

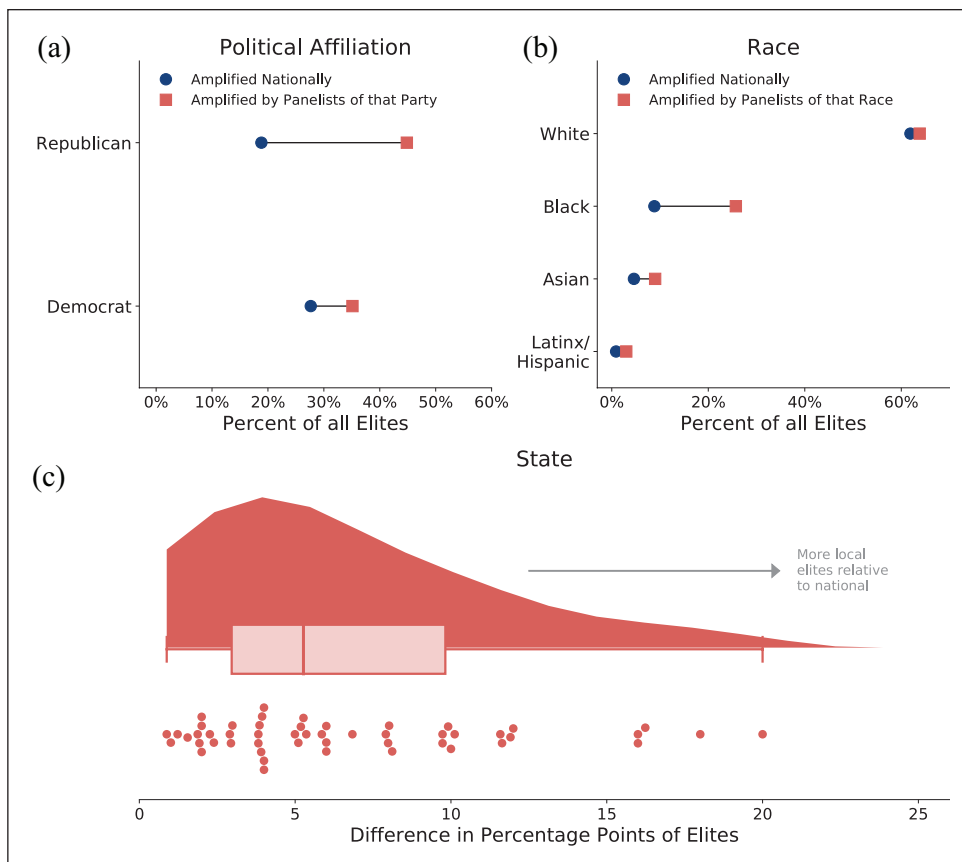


Figure 5. (a) The political affiliation of elites amplified nationally and panelists of that political party. For example, the percent of right-wing accounts amplified by Republican panel members is compared with the percent of right-wing accounts amplified across the panel nationally; the percent of left-wing accounts amplified by Democrats is compared with the percent nationally. (b) The racial composition of elites amplified nationally and by particular racial demographics. For example, the percent of Black COVID-19 elites amplified by Black panel members is compared to the number of Black elites amplified nationally. (c) The distribution of differences between the percent of local elites in each state and the percent of those same local elites among the national elites. Larger values indicate that more accounts from that state are amplified locally than nationally.




Figure 6. URL information sharing by different types of sustained COVID-19 elites. Blue bars to the left indicate the percent of tweets that contained at least one URL. Red bars to the right indicate the percent of received retweets that resulted from tweets with URLs. If a type of COVID-19 elite received equal amplification of their URL and non-URL tweets, then we would expect the bars to be level. Accounts where the right (red) bar is lower than the left (blue) bar are those that received less retweets for their URL tweets than their non-URL tweets, relative to the number of URL tweets they sent.

amplification across various demographic populations. By constructing a measure of sustained amplification and applying this measure to a panel of Twitter users in the United States, we have shown that people have primarily amplified journalists and media outlets as COVID-19 elites, followed by assorted political accounts across the political spectrum. Furthermore, we have also shown that epidemiologists, public health officials, and medical professionals only make up a small fraction of all COVID-19 elites crowdsourced nationally and across demographics. We discuss the immediate implications of these results for public health communication regarding COVID-19, and the more general consequences of our methods for identifying sustained crowdsourced elites across networked publics beyond the pandemic.

COVID-19 Elites and Opportunities for Fighting “Infodemics”

Previous research suggests that the US public, across demographic and ideological lines, trusts medical professionals and scientists more than elected officials and journalists, and this trust is on the rise since the start of the pandemic (Funk et al., 2020). We see evidence of this in our own results, as epidemiologists and public health officials (though not medical professionals) are the only types of accounts other than journalists and media outlets to receive amplification proportional to the amount of information they share via URLs. Yet, our results also suggest that scientists and medical professionals have not consistently played a central role in driving conversations about COVID-19 on Twitter, while political accounts and journalists—who also may be perceived as partisan (Iyengar & Hahn, 2009)—have. This research cannot say exactly why this is the case; it may be that epidemiologists, public health officials, and medical professionals do not have established followings to amplify their messages, their messaging itself is ineffective, or they do not post about the pandemic as much as other types of accounts. All of these explanations suggest different communication strategies: the information of scientists and public authorities may need to be amplified further, their public communication needs to be improved, or they generally need to engage with the Twitter public more. Regardless, at the start of the pandemic in March 2020, all of the major social media platforms made a statement that they were “elevating authoritative content” on their platforms in response to the pandemic (Shu & Shieber, 2020). Despite this, Americans appear to be receiving a sizable portion of their COVID-19 information from political sources—likely resulting from and contributing to political polarization around the pandemic (Green et al., 2020). Our work suggests that social media platforms may not have gone far enough over the first year of the pandemic to ensure that experts are *consistently* centered in COVID-19 discussions with their messages *widely* visible.

If relevant scientific and public authorities are not being regularly consulted for COVID-19 information, then effective public health messaging campaigns on these platforms should focus on crowdsourced COVID-19 elites who already share a sustained portion of COVID-19 discussions online. Importantly, we have shown that these elites vary across demographic groups along the lines of race, geography, and political affiliation: people are more likely to amplify accounts that share their beliefs and background. In particular, we have shown that people of color have been more likely to share COVID-19 information from elites of color. People of color are disproportionately affected by the pandemic (Azar et al., 2020), and our results suggest that there is a need for more non-White experts and authorities to communicate with these marginalized populations. Such figures could be algorithmically amplified and prioritized on social media platforms, so that they may transmit reliable information about COVID-19 to their respective communities online. Another route might involve institutional intervention, following the example of the United Kingdom where the government paid social media celebrities to encourage participation in COVID-19 testing and contact tracing (Bolat, 2020). Prior research suggests that paid celebrity promotions may be more effective than the same information provided by the World Health Organization (WHO) and Centers for Disease Control and Prevention (CDC), partially because politicians’ and celebrities’ tweets tend to have more positive undertones (Kamiński et al., 2021). If crowdsourced elites can advocate among people of color, then they may be able to combat particular types of COVID-19 misinformation that exploit the understandable mistrust that portions of these communities have toward scientists and the United States government because of unethical experimentation on them (Collins-Dexter, 2020). Leveraging COVID-19 crowdsourced elites to communicate public health messaging to these marginalized communities may be a critical component in combating COVID-19 misinformation over the coming years, as the virus likely becomes endemic in parts of the world (N. Phillips, 2021) and vaccines continue to be distributed.

Not all COVID-19 crowdsourced elites are appropriate for disseminating information though: certainly, a number of them are distributors of COVID-related misinformation. In itself, it is not necessarily concerning that content from journalists and media outlets are elevated significantly more than content from epidemiologists, public health officials, and medical professionals. Indeed, if the latter do not already have an established social media following, they may need media platforms and journalists to amplify their expertise in the news. However, media outlets vary widely in their commitment and ability to share high-quality, accurate information, and the WHO and other experts have warned of an online “infodemic” coinciding with the offline pandemic (Donovan, 2020; World Health Organization [WHO], 2020). For instance, disinformation agents, such as the Russian Internet Research Agency, acted as local news aggregators

on Twitter in 2016 and disproportionately focused on crime rates in their respective areas (Bastos & Farkas, 2019). Thus, there is a potential that similar local strategies have been used during the pandemic. Previous research has also highlighted the prevalence of religious misinformation in the Middle East and North Africa online sphere (Alimardani & Elswah, 2020), suggesting that such beliefs might be also amplified in other countries and communities affiliated with these regions. Future research needs to evaluate the quality of information shared by crowdsourced COVID-19 elites so that we can better monitor and intervene upon misinformation in future health crises (Gallotti et al., 2020; Shams et al., 2021; WHO, 2021).

Sustained Amplification Across Publics Beyond COVID-19

Our work speaks more broadly to how crowdsourced elites are identified in practice. First, we have distinguished *sustained* amplification from *episodic* amplification, and have argued that a reshare *h*-index should be used in place of simply counting reshares when studying crowdsourced elites over the course of longer, ongoing events. While we have focused on COVID-19 elites here and which accounts have received sustained attention over the course of the pandemic, sustained amplification can be measured for other discourses that operate at longer time scales. For example, identifying crowdsourced elites by counting reshares is sufficient when studying any particular Black Lives Matter protest. In that case, it is less necessary to concern ourselves with the consistency of amplification because it is all relatively brief in that moment. However, if we want to string together the conversations around multiple protests that take place over the course of several weeks, months, or even years, it becomes increasingly important to distinguish between episodic and sustained amplification. We have demonstrated how to operationalize each type of amplification and encourage other researchers to carefully consider how the timescale of their data and how they measure amplification affect who is considered a crowdsourced elite.

We have also emphasized the need to acknowledge how different publics crowdsource different elites. Here, we have used demographically induced publics, as they are particularly relevant to the health disparities regarding COVID-19. However, crowdsourced elites can be identified for any networked public, not just demographic ones. For example, any politically contentious topic in the United States will very likely result in a polarized interaction network online. Rather than extracting crowdsourced elites in aggregate across all of the users in that network, one could run a community detection algorithm and identify crowdsourced elites that correspond to each of the different network components. While it may seem like an obvious point that different publics, especially political ones, will have different crowdsourced elites,

researchers need to explicitly account for that in their methods going forward.

Limitations

We acknowledge that there are limitations to our study. With respect to our Twitter panel itself, we have had to make decisions that affect how representative it is of the US Twitter population, and the US offline population more broadly. We have also used inferences of political affiliation and race provided by the data vendor TargetSmart. Although prior work suggests our results are not overwhelmingly biased (Grinberg et al., 2019; Hughes et al., 2021), it is important for future work to continue mapping COVID-19 elites across demographics. With respect to COVID-19 elites, we only labeled a small portion of all users retweeted by our panel members—which may partly explain why epidemiologists, public health officials, and medical professionals do not appear highly among the COVID-19 elites—and we did not measure the quality of the information being shared by these accounts. With respect to information seeking behavior, we cannot say to what extent COVID-19 crowdsourced elites were amplified algorithmically. Furthermore, our panel members certainly acquired information beyond just retweets, and sought information directly from external sources. We also do not know to what extent our panel members have taken advantage of COVID-19 “information portals” that many social media platforms have introduced since the start of the pandemic. More broadly, Twitter is just one platform in the larger information ecosystem, and it is valuable to understand how amplification on Twitter compares to information crowdsourcing on other platforms like Facebook, Reddit, YouTube, Google, Instagram, WhatsApp, and TikTok, and with traditional news media. For example, emerging evidence suggests that users of more niche social media platforms, such as Gab, are more susceptible to posts with questionable sources of information about COVID-19 than users of other mainstream platforms (Cinelli et al., 2020).

Conclusion

COVID-19 will have a lasting impact on our global community, and people will continue to need relevant and trusted information about the virus to safeguard themselves and adapt to its social consequences. Our work begins to map the information ecosystem that has emerged in the wake of the pandemic, and identifies sustained COVID-19 crowdsourced elites: those who have been consistently amplified around the pandemic across different demographic groups. These elites are focal points in online communication networks because they have been given an authoritative platform with respect to COVID-19, and their voices are more likely to resonate with the audiences that crowdsource them. By working with COVID-19 elites to promote scientifically

informed guidelines and counter misinformation, particularly among already devastated populations, we may be able to leverage the natural crowdsourcing potential of social media to promote the public's health more equitably.

Acknowledgements

The authors thank Alex Madaras, Eric Gu, Andrea Barrios, Kristen Flaherty, Sagar Kumar, V Lange, and Adina Gitomer for their assistance labeling the COVID-19 elites. They also thank Alan Mislove, Christo Wilson, Kenneth Joseph, Nir Grinberg, and Stefan McCabe for their help maintaining the panel data. They thank the anonymous reviewers for helping them improve this work.


Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was partially funded by grants from Northeastern University, the US Army Research Office (award W911NF-18-1-0421), and the National Science Foundation (award 2026631). We are grateful to our sponsors for their support. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of our sponsors.

ORCID iDs

Ryan J. Gallagher  <https://orcid.org/0000-0003-3040-4500>

Larissa Doroshenko  <https://orcid.org/0000-0001-6763-628X>

Supplemental Material

Supplemental material for this article is available online.

Note

1. <https://developer.twitter.com/en/docs/labs/covid19-stream/filtering-rules>. As of writing, the filtering rules have not been publicly updated since May 13, 2020.

References

- Alimardani, M., & Elswah, M. (2020). Online temptations: COVID-19 and religious misinformation in the MENA region. *Social Media + Society*, 6(3), 2056305120948251.
- Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., & Yang, D. (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics*, 191, 104254.
- Azar, K. M., Shen, Z., Romanelli, R. J., Lockhart, S. H., Smits, K., Robinson, S., . . . Pressman, A. R. (2020). Disparities in outcomes among COVID-19 patients in a large health care system in California: Study estimates the COVID-19 infection fatality rate at the US county level. *Health Affairs*, 39(7), 1253–1262.
- Barzilai-Nahon, K. (2008). Toward a theory of network gatekeeping: A framework for exploring information control. *Journal of the American Society for Information Science and Technology*, 59(9), 1493–1512.
- Bastos, M., & Farkas, J. (2019). “Donald Trump is my president!” The internet research agency propaganda machine. *Social Media + Society*, 5(3), 2056305119865466.
- Bolat, E. (2020, September 9). Why the UK government is paying social media influencers to post about coronavirus. <https://the-conversation.com/why-the-uk-government-is-paying-social-media-influencers-to-post-about-coronavirus-145478>
- Bridgman, A., Merkley, E., Loewen, P. J., Owen, T., Ruths, D., Teichmann, L., & Zhilin, O. (2020). The causes and consequences of COVID-19 misperceptions: Understanding the role of news and social media. *Harvard Kennedy School Misinformation Review*, 1(3). <https://misinfoeview.hks.harvard.edu/article/the-causes-and-consequences-of-covid-19-misperceptions-understanding-the-role-of-news-and-social-media/>
- Cha, M., Haddadi, H., Benevenuto, F., & Gummadi, K. (2010, May 23–26). *Measuring user influence in Twitter: The million follower fallacy* (Vol. 4, No. 1) [Conference session]. Proceedings of the International AAAI Conference on Web and Social Media, Washington, DC, United States.
- Chadwick, A. (2011). The political information cycle in a hybrid news system: The British prime minister and the “bullygate” affair. *The International Journal of Press/Politics*, 16(1), 3–29.
- Chen, E., Lerman, K., & Ferrara, E. (2020). Tracking social media discourse about the COVID-19 pandemic: Development of a public coronavirus Twitter data set. *JMIR Public Health and Surveillance*, 6(2), Article e19273.
- Cinelli, M., Quattrociocchi, W., Galeazzi, A., Valensise, C. M., Brugnoli, E., Schmidt, A. L., Zola, P., Zollo, F., & Scala, A. (2020). The COVID-19 social media infodemic. *Scientific Reports*, 10(1), 1–10.
- Collins-Dexter, B. (2020). *Canaries in the coal mine: COVID-19 misinformation and Black communities*. Harvard Kennedy School Shorenstein Center on Media, Politics and Public Policy.
- Conover, M., Ratkiewicz, J., Francisco, M., Gonçalves, B., Menczer, F., & Flammini, A. (2011). Political polarization on Twitter. *Proceedings of the International AAAI Conference on Web and Social Media*, 5(1). <https://ojs.aaai.org/index.php/ICWSM/article/view/14126>
- Donovan, J. (2020, April 14). Social-media companies must flatten the curve of misinformation. *Nature*. <https://www.nature.com/articles/d41586-020-01107-z>
- Enders, A. M., Uscinski, J. E., Klofstad, C., & Stoler, J. (2020, November 16). The different forms of COVID-19 misinformation and their consequences. *The Harvard Kennedy School Misinformation Review*. <https://misinfoeview.hks.harvard.edu/article/the-different-forms-of-covid-19-misinformation-and-their-consequences/>
- Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2016). The rise of social bots. *Communications of the ACM*, 59(7), 96–104.
- Fiesler, C., & Proferes, N. (2018). “Participant” perceptions of Twitter research ethics. *Social Media + Society*, 4(1), 2056305118763366.
- Fink, K. (2019). The biggest challenge facing journalism: A lack of trust. *Journalism*, 20(1), 40–43.

- Fisher, M., & Weiner, R. (2020, July 20). With little clarity on coronavirus, Americans crowdsource how to live in a pandemic. https://www.washingtonpost.com/politics/with-little-clarity-on-coronavirus-americans-crowdsource-how-to-live-in-a-pandemic/2020/07/20/6abcb8ea-c5f6-11ea-b037-f9711f89ee46_story.html
- Foucault Welles, B. (2014). On minorities and outliers: The case for making Big Data small. *Big Data & Society*, 1(1), 2053951714540613.
- Funk, C., Kennedy, B., & Johnson, C. (2020, May 21). Trust in medical scientists has grown in the U.S., but mainly among democrats. <https://www.pewresearch.org/science/2020/05/21/trust-in-medical-scientists-has-grown-in-u-s-but-mainly-among-democrats/>
- Gallotti, R., Valle, F., Castaldo, N., Sacco, P., & De Domenico, M. (2020). Assessing the risks of “infodemics” in response to COVID-19 epidemics. *Nature Human Behaviour*, 4(12), 1285–1293.
- Glorigić, K., Ribeiro, M. H., Müller, M., Altunina, O., Peyrard, M., Salathé, M., Colavizza, G., & West, R. (2020). *Experts and authorities receive disproportionate attention on Twitter during the COVID-19 crisis*. arXiv preprint:2008.08364.
- Goel, S., Anderson, A., Hofman, J., & Watts, D. J. (2016). The structural virality of online diffusion. *Management Science*, 62(1), 180–196.
- Gorwa, R., & Guilbeault, D. (2020). Unpacking the social media bot: A typology to guide research and policy. *Policy & Internet*, 12(2), 225–248.
- Green, J., Edgerton, J., Naftel, D., Shoub, K., & Cranmer, S. J. (2020). Elusive consensus: Polarization in elite communication on the COVID-19 pandemic. *Science Advances*, 6(28), Article eabc2717.
- Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., & Lazer, D. (2019). Fake news on Twitter during the 2016 US presidential election. *Science*, 363(6425), 374–378.
- Hirsch, J. E. (2005). An index to quantify an individual’s scientific research output. *Proceedings of the National Academy of Sciences*, 102(46), 16569–16572.
- Howard, P. N. (2010). *The digital origins of dictatorship and democracy: Information technology and political Islam*. Oxford University Press.
- Hu, M., Liu, S., Wei, F., Wu, Y., Stasko, J., & Ma, K. L. (2012, May 5–10). Breaking news on Twitter. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 2751–2754). ACM.
- Hughes, A., McCabe, S. D., Hobbs, W. R., Remy, E., Shah, S., & Lazer, D. (2021). Using administrative records and survey data to construct samples of tweeters and tweets. *Public Opinion Quarterly*.
- Iyengar, S., & Hahn, K. S. (2009). Red media, blue media: Evidence of ideological selectivity in media use. *Journal of Communication*, 59(1), 19–39.
- Jackson, S. J., Bailey, M., & Foucault Welles, B. (2020). *#HashtagActivism: Networks of race and gender justice*. MIT Press.
- Jackson, S. J., & Foucault Welles, B. (2015). Hijacking #myNYPD: Social media dissent and networked counterpublics. *Journal of Communication*, 65(6), 932–952.
- Jackson, S. J., & Foucault Welles, B. (2016). #Ferguson is everywhere: Initiators in emerging counterpublic networks. *Information, Communication & Society*, 19(3), 397–418.
- Jungherr, A. (2016). Twitter use in election campaigns: A systematic literature review. *Journal of Information Technology & Politics*, 13(1), 72–91.
- Kamiński, M., Szymańska, C., & Nowak, J. K. (2021). Whose tweets on COVID-19 gain the most attention: Celebrities, political, or scientific authorities? *Cyberpsychology, Behavior, and Social Networking*, 24(2), 123–128.
- Lotan, G., Graeff, E., Ananny, M., Gaffney, D., & Pearce, I. (2011). The revolutions were tweeted: Information flows during the 2011 Tunisian and Egyptian revolutions. *International Journal of Communication*, 5, 31.
- Lukito, J., Suk, J., Zhang, Y., Doroshenko, L., Kim, S. J., Su, M.-H., Xia, Y., Freelon, D., & Wells, C. (2020). The wolves in sheep’s clothing: How Russia’s Internet Research Agency tweets appeared in US news as vox populi. *The International Journal of Press/Politics*, 25(2), 196–216.
- McGregor, S. C. (2019). Social media as public opinion: How journalists use social media to represent public opinion. *Journalism*, 20(8), 1070–1086.
- Meraz, S. (2009). Is there an elite hold? Traditional media to social media agenda setting influence in blog networks. *Journal of Computer-Mediated Communication*, 14(3), 682–707.
- Meraz, S., & Papacharissi, Z. (2013). Networked gatekeeping and networked framing on #Egypt. *The International Journal of Press/Politics*, 18(2), 138–166.
- Molyneux, L., & McGregor, S. C. (2021). Legitimizing a platform: Evidence of journalists’ role in transferring authority to Twitter. *Information, Communication & Society*. Advance online publication, 1–19. <https://doi.org/10.1080/1369118X.2021.1874037>
- Papacharissi, Z. (2009). The citizen is the message: Alternative modes of civic engagement. In Z. Papacharissi (Ed.), *Journalism and citizenship: New agendas in communication* (pp. 29–43). Routledge.
- Papacharissi, Z., & de Fatima Oliveira, M. (2012). Affective news and networked publics: The rhythms of news storytelling on #Egypt. *Journal of Communication*, 62(2), 266–282.
- Perrin, A., & Anderson, M. (2019, April 10). Share of U.S. adults using social media, including Facebook, is mostly unchanged since 2018. <https://www.pewresearch.org/fact-tank/2019/04/10/share-of-u-s-adults-using-social-media-including-facebook-is-mostly-unchanged-since-2018/>
- Phillips, N. (2021). The coronavirus is here to stay—here’s what that means. *Nature*, 590(7846), 382–384.
- Phillips, W. (2018). *The oxygen of amplification: Better practices for reporting on extremists, antagonists, and manipulators online*. Data & Society Research Institute.
- Pourebrahim, N., Sultana, S., Edwards, J., Gochanour, A., & Mohanty, S. (2019). Understanding communication dynamics on Twitter during natural disasters: A case study of Hurricane Sandy. *International Journal of Disaster Risk Reduction*, 37, 101176.
- Shams, A. B., Hoque Apu, E., Rahman, A., Sarker Raihan, M., Siddika, N., Preo, R. B., Hussein, M. R., Mostari, S., & Kabir, R. (2021). Web Search Engine Misinformation Notifier Extension (SEMiNExt): A machine learning based approach during COVID-19 pandemic. *Healthcare*, 9, 156.
- Shoemaker, P. J., & Vos, T. (2009). *Gatekeeping theory*. Routledge.
- Shu, C., & Shieber, J. (2020, March 16). Facebook, Reddit, Google, LinkedIn, Microsoft, Twitter, and YouTube issue

- joint statement on misinformation. <https://techcrunch.com/2020/03/16/facebook-reddit-google-linkedin-microsoft-twitter-and-youtube-issue-joint-statement-on-misinformation/>
- Shugars, S., Gitomer, A., McCabe, S., Gallagher, R. J., Joseph, K., Grinberg, N., Doroshenko, L., Foucault Welles, B., & Lazer, D. (2021). Pandemics, protests, and publics: Demographic activity and engagement on Twitter in 2020. *Journal of Quantitative Description: Digital Media*, 1. <https://doi.org/10.51685/jqd.2021.002>
- Tufekci, Z. (2017). *Twitter and tear gas: The power and fragility of networked protest*. Yale University Press.
- Valente, T. W., & Davis, R. L. (1999). Accelerating the diffusion of innovations using opinion leaders. *The Annals of the American Academy of Political and Social Science*, 566(1), 55–67.
- World Health Organization. (2020). *Managing the COVID-19 infodemic: Promoting healthy behaviours and mitigating the harm from misinformation and disinformation*. <https://www.who.int/news/item/23-09-2020-managing-the-covid-19-infodemic-promoting-healthy-behaviours-and-mitigating-the-harm-from-misinformation-and-disinformation>
- World Health Organization. (2021). *WHO launches pilot of AI-powered public-access social listening tool*. <https://www.who.int/news-room/feature-stories/detail/who-launches-pilot-of-ai-powered-public-access-social-listening-tool>
- Xu, W. W., Sang, Y., Blasiola, S., & Park, H. W. (2014). Predicting opinion leaders in Twitter activism networks: The case of the Wisconsin recall election. *American Behavioral Scientist*, 58(10), 1278–1293.
- Yardi, S., & boyd, d. (2010). Dynamic debates: An analysis of group polarization over time on Twitter. *Bulletin of Science, Technology & Society*, 30(5), 316–327.

Author Biographies

Ryan J. Gallagher is a PhD candidate at Northeastern University's Network Science Institute. His research interests include online amplification, social media platforms as ecological landscapes, and network science methodology.

Larissa Doroshenko (PhD, University of Wisconsin) is a postdoctoral teaching associate in the Department of Communication Studies at Northeastern. Her research interests include the effects of new media on political campaigning, populism, nationalism, and disinformation campaigns.

Sarah Shugars (PhD Northeastern University) is a Moore-Sloan Faculty Fellow at New York University's Center for Data Science. Their research interests include how people express their political views, reason about political issues, and engage with others around matters of common concern.

David Lazer (PhD University of Michigan) is a University Distinguished Professor of Political Science and Computer Sciences and a core faculty member of the Network Science Institute at Northeastern University. His research interests include network science, computational social science, and collaborative intelligence.

Brooke Foucault Welles (PhD Northwestern University) is an Associate Professor in the Department of Communication Studies and core faculty member of the Network Science Institute at Northeastern University. Her research interests include how online communication networks enable and constrain behavior, and how those networks enable the pursuit of individual, team, and collective goals.