

**Pandemics, Protests, and Publics:  
Demographic Activity and Engagement on Twitter in 2020**

SARAH SHUGARS

New York University, USA  
George Washington University, USA

ADINA GITOMER

Northeastern University, USA

STEFAN D. MCCABE

Northeastern University, USA

RYAN J. GALLAGHER

Northeastern University, USA

KENNETH JOSEPH

University at Buffalo, USA

NIR GRINBERG

Ben-Gurion University, Israel

LARISSA DOROSHENKO

Northeastern University, USA

BROOKE FOUCAULT WELLES

Northeastern University, USA

DAVID LAZER

Northeastern University, USA

Harvard University, USA

---

Sarah Shugars: [sarah.shugars@nyu.edu](mailto:sarah.shugars@nyu.edu)

Date submitted: 2020-12-31

When researchers collect and aggregate social media data, they are making explicit decisions about the populations and behaviors under study. However, there is little available guidance to ensure that these methodological choices are conceptually and empirically grounded. For example, how should researchers conceptualize a topical sample of social media content? Can it be understood as a self-contained world? Can we interpret individual accounts as participating in the *same* discourse? Should we disaggregate specific mechanisms of user activity and engagement? In short: when do researchers need to consider variation in user experience and behavior, and when can they meaningfully aggregate over such behavior? Leveraging a panel of 1.6 million Twitter accounts matched to U.S. voting records we provide empirical guidance on these questions through the conceptual lens of public sphere theory. We focus on the first nine months of 2020, giving particular attention to the Black Lives Matter movement and the COVID-19 pandemic. Examining the demographics, activity, and engagement of 800,000 American adults who collectively posted nearly 300 million tweets, our findings help establish meaningful bounds around populations and behaviors to study. We find that topics are imperfect but useful bounds, though topically selected tweets must be understood to be capturing segments of numerous, overlapping, and disconnected conversations. We further find that researchers should always conduct a disaggregated analysis of tweet activity, separately examining behavior around authored tweets, retweets, quote tweets, and replies. Additionally, we find retweets and quote tweets appear to be used in distinctly different ways, potentially reflecting that retweets amplify content while quote tweets modify that content. Finally, we find that while temporal bias is inherent to social media data, its effects are manageable within our period of study. Overall, this work paints a picture of Twitter as a fluid, contextual environment best conceptualized as networked publics and characterized by enormous variety in user identity, activity, and engagement. While there are no self-contained “Twitter publics” around which perfect boundaries can be drawn, our findings provide valuable empirical guidance to researchers grappling with the conceptual implications of their methodological choices.

*Keywords:* *public sphere, COVID-19, Black Lives Matter, Twitter, demographics, networked publics*

## Introduction

Social media is integral to public discourse, and Twitter, in particular, has come to serve as a key source for scholarship examining that discourse (Tufekci, 2014; Fiesler and Proferes, 2018). Every year, numerous papers leverage Twitter data in order to better understand public opinion, elite messaging, media engagement, social justice activism, political polarization, disinformation campaigns, and broader dynamics of information consumption and amplification (Olteanu et al., 2019). However, despite the prevalence of Twitter as a venue of scholarship, there are notable gaps in our conceptual and empirical understanding of engagement on the platform. What are we really capturing when we aim to examine “the discourse” on Twitter?

This question highlights an inherent tension in Twitter research: there is a great deal of variation in user identity, activity, and engagement, yet this complexity is often disregarded in large-scale, aggregate analysis. While such aggregation is valuable and necessary to scientific inquiry, it must be conceptually and empirically grounded. Proceeding without a clear understanding of what populations and behaviors are being aggregated over risks erroneous inference about the nature of that activity and engagement.

For example, a reasonable first step in Twitter analysis is to use keyword lists in order to identify tweets related to a given topic. Even assuming that we are able to retrieve all tweets connected to our keywords, how should we conceptualize the resulting corpus? Can it be understood as a self-contained world, fully capturing discourse around our topic of interest? Can we interpret the individual accounts as participating in the *same* discourse or do we need to consider the ways in which different subpopulations engage? Should we disaggregate the specific mechanisms of activity and engagement or can we consider, for example, a tweet to be the same as a retweet? In other words, when do we need to consider the breadth of variation in user experience and behavior, and when can we meaningfully aggregate over such behavior?

### *Conceptualizing Twitter Public(s)*

Public sphere theory provides a valuable conceptual framework for interrogating these questions. Previous work has suggested that Twitter can be understood as a *public*, in the sense

that it captures more than isolated, individual expressions and promotes a range of back-and-forth discourse (Shugars and Beauchamp, 2019). While some research questions lend themselves to treating the platform as a single, unified public (Habermas, 1984, 1991), another line of work has demonstrated that social media can be more richly understood as containing multiple *publics* (Warner, 2002; Fraser, 1990; Squires, 2002; boyd, 2010; Jackson et al., 2020)—diverse and at times divergent populations which form meaningful communities in their own right. Importantly, *both* conceptualizations are valuable. Each play distinctive roles in supporting our collective understanding of public discourse and are appropriate for different types of research questions.

When researchers collect and aggregate social media data, they are making explicit decisions about the populations and behaviors under study. Conceptually, these decisions can be understood as reflecting implicit assumptions about what populations and behaviors meaningfully capture “a public.” Research which aims to situate Twitter in relation to broader public discourse, for example, appropriately treats the platform as a singular public—aggregating over the internal complexity in order to better understand its role within the media environment (Hu et al., 2012; Green et al., 2020). On the other hand, research which aims to understand discourse *within* Twitter should more appropriately consider the multiple publics which coexist across the the platform (Jackson et al., 2020; Gallagher et al., 2020; Brock, 2012; boyd, 2010). In the digitally connected world, these can further be understood as *networked publics*, collections of individuals whose norms and behaviors are shaped through networked technologies (boyd, 2010). Under this conception, the boundaries of “a public” are permeable and intersecting, as individuals engage across multiple, overlapping publics.

While “networked publics” is arguably the most accurate conceptualization, it often not practical to operationalize in large scale research. Such research necessarily focuses on subsamples of Twitter data and aggregates over various populations and behaviors. However, a great deal of research overlooks the implicit assumptions underlying collection and aggregation strategies, and does not grapple with what these methodological choices mean conceptually. For example, research which focuses on topically coherent tweets implicitly assumes that this is a meaningful sample to study; that topical discourse constitutes, to some degree, “a public” which can be extracted as a discrete, categorical construct. This

disregards the context in which those tweets occurred—losing information about concurrent topics, as well as users’ perspectives and discursive histories. Without this context, researchers risk making implicit and potentially erroneous assumptions about the identity, perspective, and motivations of users engaged with that topic. For example, a single on-topic tweet from a highly active Twitter user may be given disproportionate weight, or that user may be misinterpreted as being inactive on the platform, rather than quiet on the specific topic. Furthermore, without access to users’ demographic data, researchers may further assume that those who do engage with a given topic share certain demographic features or are drawn from distinctive demographic clusters. Similarly, though, research which begins by subsetting on users’ demographic characteristics may also lose important context around the broader discourse in which those populations’ conversations take place.

This leaves researchers with a significant challenge: it is not reasonable to leverage Twitter’s full contextual history and user demographics are typically unavailable. This means that sampling and aggregation—while a simplification of Twitter’s rich networked publics—are both necessary and inherent to Twitter research. Unfortunately, researchers have little empirical guidance as to what types of aggregation are appropriate. If researchers must consider some populations to be “publics” that can be studied in isolation from their context, what are the implications of different definitions and choices? Little work has provided the needed insight into this question by examining behavioral variation across both demographics and topical interests. This leaves open important questions about what populations and behaviors can be meaningfully aggregated over and highlights the need for empirical assessment as to what is lost or gained through those methodological choices.

### *Pandemics and Protests*

This gap in empirical insight is particularly salient within our study period of 2020. During this time, two topics saw significant volumes of discourse on Twitter: the COVID-19 pandemic and the social justice protests associated with the Black Lives Matter movement (Alshaabi et al., 2020). Each of these topics engaged sizable portions of the population and were characterized by significant political polarization and important demographic stratification. This makes questions of demographic activity and engagement around these topics worthy of study in their own right, while further offering an opportunity to examine

how demographic identities and topical interests overlay the conceptualization of Twitter publics. By presenting a disaggregated, multifaceted description of American Twitter users in 2020, we are able to expand scholarship on the formation and function of publics, while also clarifying the discursive landscape of Twitter in particular. Each of these outcomes, taken separately or together, may be used to inform future studies of online conversational behaviors.

Leveraging a panel of 1.6 million Twitter accounts matched to U.S. voting records, we examine the demographics, activity, and engagement of nearly 800,000 American adults who were active on Twitter between January 1, 2020 and September 30, 2020. We collect over 284 million tweets from members of our panel, examining variation in the content produced and engagement received by different demographic segments of the population. Furthermore, we delve into variations in the specific mechanisms of activity and engagement, examining differences in the volume of types of tweets produced, and different types of engagements received from others.

By examining user behavior for both Twitter as a whole and for topical and demographic subpopulations, we capture what may be lost or gained through different types of aggregation and provide substantive insight for the topics of COVID-19 and the Black Lives Matter movement in particular. Which populations are over or under represented in terms of the volume of content they produce relative to their share of the population? How frequently are different kinds of tweets (e.g., authored tweets and retweets) made by different users in different contexts? How much engagement (e.g. favorites and follows) do different users receive for the content they produce? The answers to these questions provide the empirical context needed for researchers to make informed choices about the collection, aggregation, and interpretation of Twitter content. This includes insight into which populations, in which contexts, can meaningfully be considered a public, which behaviors can be reasonably treated as equivalent, and what it means for someone to be a “high” or “low” activity user, or for them to get “a lot” or “a little” attention.

We close by examining some of the temporal challenges in studying a rapidly changing platform (Munger, 2019), characterizing the data loss and temporal biases found in our sample. While these effects are minimal within our dataset, an explicit acknowledge-

ment and analysis of these challenges adds further context to aid in the interpretation and application of these findings.

Our findings suggest that, in line with the networked publics conceptualization, Twitter is a richly varied platform with notable diversity in user activity and engagement. There is no single dimension along which these publics can be defined—neither topical nor demographic variation perfectly encapsulates “a Twitter public.” However, while Twitter publics can be best understood as fluid, contextual communities, our findings provide practical and empirical guidance for researchers aiming to establish meaningful bounds around populations and behaviors to study. We stress, however, that these bounds will always be arbitrary, imperfect, and researcher-imposed—there simply are not self-contained Twitter publics around which clear lines can be drawn.

Specifically, we find that topics are imperfect but useful bounds for delineating populations of study, though topically selected tweets should not be assumed to represent a unified “discourse” and must be considered to capture segments of numerous, overlapping, and disconnected conversations. We further find that researchers should always conduct a disaggregated analysis of tweet activity, separately examining behavior around authored tweets, retweets, quote tweets, and replies. While in some contexts these behaviors may be safely aggregated over, examining their variation has the potential to reveal subpopulations of substantive interest and serves as an important robustness check for aggregated findings. Additionally, we find several notable dynamics around received engagement which researchers should be mindful of and which warrant further investigation. Specifically, retweets and quote tweets appear to be used in distinctly different ways, reflecting the fact that retweets simply amplify content while quote tweets modify that amplified content. Furthermore, the popularity of a tweet should not be conflated with the popularity of that tweet’s author, as even smaller accounts can have tweets that “go viral.” Finally, we find that while temporal bias is inherent to the rapidly changing context of social media, its effects are manageable within at least a year-long period of time. We find that rehydrating tweets even ten months after their post date results in only minimal data loss and does not appear to artificially inflate engagement with tweets which have had more time in which to accrue interactions. We caution, however, that content from some topics and populations may be more difficult to retrieve, and that this retrieval relies upon uninterrupted API

access from the platform (Freelon, 2018). Overall, this work paints a picture of Twitter as a fluid, contextual environment best conceptualized as networked publics and characterized by enormous variety in user identity, activity, and engagement. While there are no self-contained “Twitter publics” around which perfect boundaries can be drawn, our findings provide empirical guidance to researchers grappling with the conceptual implications of their methodological choices.

### Related Work

Since its launch in 2006, a significant volume of work has examined various aspects of Twitter. The “model organism” of social media research (Tufekci, 2014), Twitter is studied both out of substantive interest and because the platform makes it easy to acquire data and make preliminary observations about more general online phenomena.

An important venue for “everyday” political conversations (Shugars and Beauchamp, 2019; Jaidka et al., 2019), Twitter allows individuals to connect with and voice their opinions to numerous elected officials (Barberá et al., 2019; Farina et al., 2013; Kavanaugh et al., 2012), journalists (McGregor and Molyneux, 2020), and other public figures. Racial justice movements like Black Lives Matter and hashtag campaigns like #MeToo have used this to their advantage (Jackson et al., 2020; Freelon et al., 2018), leveraging social media to garner the attention of journalists and boost otherwise marginalized narratives (Chadwick, 2017; Jackson and Foucault Welles, 2015). Similarly though, right-wing extremists and disinformation campaigns have also targeted journalists and political elites to amplify false and distorted claims (Phillips, 2018; Marwick and Lewis, 2017; Lukito et al., 2020). While misinformation is estimated to make up only a small fraction of content on Twitter (Grinberg et al., 2019), the ease with which it can spread and be brought offline has made the platform a critical conduit in the broader information ecosystem (Phillips and Milner, 2020; Allen et al., 2020). Furthermore, Twitter’s ability to rapidly disseminate information has made the platform an important alternative to traditional broadcast media during a range of breaking news events (Hu et al., 2012; Grusin, 2010) from social protests (Jackson et al., 2020) to natural disasters (Pourebrahim et al., 2019). This has put Twitter at the center of discourse and information-seeking efforts around two key topics of 2020: COVID-19 and the Black Lives Matter movement.



Despite the prevalence of Twitter data in research examining social media phenomena, relatively little work has aimed to describe the demographics, activities, and engagement of users on this platform. While notable strands of work has delved into each of these topics individually, the existing literature has not explored the intersection of these topics—e.g., the volume of tweets produced or the number of engagements received by different demographics groups.

Twitter regularly self-reports some demographic information to its shareholders (Twitter, 2020), and marketing platforms (Hootsuite, 2020) similarly release summary reports using proprietary data and methods. Some of the most methodologically transparent work aimed at characterizing the demographics and news consumption of users on this platform comes from Pew Research Center (2018, 2019a,b). Their survey data relies on self-reported measures and representative samples of U.S. adults, both of which may introduce bias when it comes to actual behavior and representation of the target population on Twitter. By relying on a non-probability sample, our panel offers a complementary approach. Pew Research Center (2019b) suggests that 22% of American adults, approximately 46 million individuals, have ever used Twitter. This places it among the smaller social media sites, with fewer adult users in the U.S. than YouTube (73%), Facebook (69%), Instagram (37%), LinkedIn (27%) and Snapchat (24%). Some sources, however, estimate a much larger portion of U.S.-based Twitter users, with one social media management company finding that up to 40% of American internet users between the ages of 16 to 64 reported using Twitter (Hootsuite, 2020). Furthermore, Twitter is one of the most news-centric social media sites, with 71% of the platform's users saying in 2018 that they got news from Twitter (Pew Research Center, 2018). That makes Twitter the third-ranking social media site for news, with 12% of U.S. adults saying they get news from Twitter, compared to 43% of U.S. adults who get news from Facebook, and 21% who get news from YouTube (Pew Research Center, 2018).

Additionally, existing work examining the demographic characteristics of Twitter users has found that this population is generally younger and more highly educated than the general U.S. population (Pew Research Center, 2019b). The median age of a Twitter user is 40, while the median U.S. adult is 47. Twitter users are also among the most highly educated social media users. Approximately 31% of the U.S. adult population has a

college or advanced degree (Pew Research Center, 2019b), a share that is overrepresented on Twitter, where 41% of users have a college degree or more. Only LinkedIn (61%) has a more highly educated user base. Previous work has further found that Twitter users are also more likely to be Democratic than the general population, with 31% of U.S. adults and 36% of adult Twitter users identifying as Democratic, and 26% of adults and 21% of Twitter users identifying as Republican (Pew Research Center, 2019b). Twitter users are also slightly more likely to be Independents, making up 27% of the general population and 29% of adult Twitter users (Pew Research Center, 2019b). Another line of work relevant to understanding the demographic make up of the Twitter population has aimed to infer the demographic characteristics of Twitter users. This inference has been done using Twitter profile pictures (An and Weber, 2016), full profile data (Wang et al., 2016), and users' tweet history (Mislove et al., 2011), as well as the accounts they follow (Culotta et al., 2015).

Beyond these demographic descriptions of the platform, additional work has aimed to characterize the behavior of Twitter users. As with many social media platforms, activity on Twitter follows a heavy-tailed distribution. Previous work has found that while the median user only tweets twice a month, the top 10% of active users tweet approximately 138 times per month and are responsible for nearly 80% of posted content (Pew Research Center, 2019b). Furthermore, a smaller line of work has aimed to understand the differing meanings and uses behind different types of tweets. For example, retweets amplify a message and therefore, whether intended or not, are often interpreted as endorsements (Kim and Yoo, 2012). The act of retweeting has also been found to integrate users into the broader conversation at play, despite a lack of original authorship (boyd et al., 2010). Replies, on the other hand, indicate direct engagement with another user that may be negative or positive (Kim and Yoo, 2012); the likelihood of either sentiment has been linked to the following relationship between the communicating users (Liu and Weber, 2014). Additionally, the introduction of the quote tweet has been found to increase the volume and reach of political discourse (Garimella et al., 2016). While functionally similar to retweets, quote tweets have been found to be more comparable to replies than retweets, often manifesting as one of either public opinion, public reply, or public forwarding of a message (Garimella et al., 2016). While any single tweet type may not always have a universal meaning (Tufekci, 2014), in aggregate, these platform interactions have been found to reflect survey-based measures of political sentiment (Joseph et al., 2019).

Finally, previous work has also pointed to the importance of considering the temporal nature of social media (Munger, 2019). This temporality has numerous implications when studying a platform. For example, work on Twitter, in particular, has found that tweets typically receive the bulk of their interactions within 24 hours of posting (Shugars and Beauchamp, 2019; Wang et al., 2016; Kwak et al., 2010; Starbird and Palen, 2012; Karpf, 2019, 2020). This gives behavioral insight into how quickly new content turns old, and further illustrates how much time must pass in order for a researcher to accurately measure the interactions received. Furthermore, data loss is a very real concern on these platforms, as posted content may be deleted, taken down over time, or rendered inaccessible via a platform’s API (Lazer et al., 2020; Freelon, 2018).

While the existing literature separately captures information about demographic composition and user behavior on Twitter, it is missing an understanding of demographic variation *in* user behavior. Such an analysis is particularly important given what we know about the heavy-tailed nature of Twitter activity. That is, knowing that a small segment of the population is responsible for the majority of tweets suggests that it not enough to know the demographic distribution of Twitter *accounts*, we must further study the demographic distribution of produced tweets themselves, along with the interactions they receive. Examining demographics, activity, and engagement together provides insight into subgroup representation and appropriate levels of aggregation. Furthermore, decomposing activity and engagement into specific modalities—such as comparing authored tweets to retweets—helps distinguish between the types of discursive moves currently in use on the platform. Finally, by comparing the activity and engagement of two salient topics, this paper examines how topical interests and demographic identities intersect with the concept of Twitter publics, and provides valuable context for understanding the behavior of other publics which may be studied.

## Materials and Methods

### *Panel of U.S. voters on Twitter*

Our primary dataset consists of a panel of 1.6 million Twitter users whose accounts have been linked to public U.S. voter records. As described in prior work (Grinberg et al., 2019;

Hughes et al., 2020), the matching was done by first using Twitter’s 10% Decahose sample to collect data from 290 million profiles, representing a near-complete set of accounts which were active between January 2014 and March 2017. These profiles were then compared against public voter records compiled by the data vendor TargetSmart in October 2017. Each individual in the voter file was then compared to a list of Twitter profiles whose name—extracted from either the Twitter handle or display name—was an exact match for the name in the voter file. If fewer than 10 Twitter profiles had a matching full name, the location of those accounts were then examined. Final matches were accepted if a single Twitter profile from these candidates successfully matched both the city and state of the individual voter record. For users who only list a state in their profile, we allowed a match if the name was unique at the state level.

The resulting panel of 1.6 million profiles covers all 50 states in the U.S. as well as the District of Columbia, and accounts for about 3% of all adult U.S. Twitter users (Hughes et al., 2020). This panel has been found to be notably more white and slightly more female than a comparable sample constructed by Pew Research Center using addressed-based sampling and volunteered handles (Hughes et al., 2020). While African-Americans are appropriately represented in our panel, Hispanics are slightly underrepresented and Asians are significantly underrepresented. These differences between our panel and the survey-constructed sample could reflect demographic biases in the unique name and location restrictions of the matching process or uneven response rates in the survey data. However, we cannot be sure whether this underrepresentation arises from restricting our sample to unique name and city matches—in which case alternative matching strategies could be used, or whether individuals in these populations are less likely to use their real names, in which case matching to administrative data would not be possible. Furthermore, it should be noted that non-voters are completely unrepresented, and these individuals may differ from voters in key ways. However, there are a number of advantages of using our panel approach. Golder and Macy (2011) have previously demonstrated the value of baseline panels as an important strategy for overcoming several inference pitfalls of social media analysis, and other scholars have argued in support of such a panel approach as well (Tufekci, 2014). As such, this panel provides valuable insight into the activities and engagement of a large sample of known users. The sheer size of the panel offers us the ability to study specific subpopulations at a resolution that would not be possible with a sample drawn from a

national survey. Furthermore, matching Twitter accounts to U.S. voter records allows us to be largely confident that we are studying the behavior of real humans—not bots, organizations, or other non-human entities (Ferrara et al., 2016; Gorwa and Guilbeault, 2020). Finally, because the voter file provided by the data vendor is also matched to demographic information, this approach allows us to study variation in age, race, gender, and political affiliation. Age and gender are directly included as part of the voter file for every state. States affected by the Voting Rights Act (VRA) also include self-reported race as part of the voter file. For states where this data is not included, race is inferred by TargetSmart. As discussed in Appendix B, these estimates of race have high (85% - 97%) agreement with other methods of inferring race. We calculate this agreement for both self-reports from a subsample of 182 panelists (85% agreement) and using the `wru` package from Imai and Khanna (2016) (82%-97% agreement, varying by state). Additionally, while some states include party registration as part of the voter file, there is high variation in party status across states (Ansolabehere and Hersh, 2017), so we instead use TargetSmart’s inferred party score to estimate panelists’ political affiliation, as further discussed in Appendix B. Our matching approach has been approved by the Northeastern University IRB, and the analysis solely relied on publicly available data. The analysis consists of de-identified data, analyzed in large, demographic bins of the population. No member of the research team investigated individual panel members’ profiles in the course of this research.

For this study, we collect all tweets made between January 1, 2020 through September 30, 2020 from all 1.6 million Twitter users in our panel. We find that 783,697 panel members—about 47.7% of our entire panel—were active during that time window, meaning they posted at least one tweet of any type. Collectively, these active panelists produced a total of 284,581,223 tweets. In October–December of 2020, we used the Twitter API to rehydrate tweets made by these users and to retrieve user metadata. This yielded the most up-to-date counts of retweets, favorites, and followers. At the time of this retrieval, 4,165 (0.53%) accounts had been suspended, and an additional 14,802 (1.89%) accounts had been deleted by their owners.

### *Identifying Tweets Related to COVID-19 and Racial Justice*

In addition to collecting all tweets from panel members in the sample time period, we also specifically identify which of their tweets relate to two of the most prominent topics of discussion in the United States during 2020: the COVID-19 pandemic and racial justice protests centered around the Black Lives Matter movement. For both topics, we take a similar approach to tweet identification by using extensive keyword lists. As one of the most common approaches for topical tweet identification, our use of keyword lists, overlaid with demographic and behavioral analysis, supports our goal of providing researchers with insight into the hidden complexities of their own keyword-generated datasets. These keywords, described for each topic below, are intentionally broad. For COVID-19, the keywords identify tweets that explicitly reference the pandemic, as well those that mention related topics, like “quarantine life” or the social and economic fallout caused by COVID-19. For racial justice, the keywords identify tweets that explicitly reference the murder of and subsequent protests for George Floyd, as well as those about other topics, such as Black Lives Matter, the murder of Breonna Taylor, and police abolition. While this broad approach requires us to carefully curate our keyword lists to balance precision and recall, it also allows us to examine the far-reaching scope of these topics beyond the narrow view of limited and rigid keywords. Both keyword lists can be found in Appendix A and online<sup>1</sup>.

For each topic of COVID-19 and racial justice, a tweet is considered to be related to that topic if at least one keyword from its respective list appears in the tweet’s text. We define the tweet’s text to include its main text, as well as any text from a quote retweet, any hashtags, and any URL substrings from a link in the tweet. We then go one step further, and classify URLs that are related to each topic, even if those URLs did not explicitly use one of the keywords. We perform this classification by deeming a URL as related to a topic if it is used with keywords from a respective list at least 100 times and at least 20% of the URL’s use is with those keywords. We found that this heuristic expands the recall of topic-related URLs while maintaining precision.

In order to evaluate the validity of our identification method, we randomly sampled 1500 tweets evenly distributed across the three categories: COVID-19, Black Lives Matter,

---

<sup>1</sup><https://sarahshugars.com/twitter-publics/>

and tweets related to neither topic. Each tweet was then coded by two undergraduate Research Assistants, who were shown the tweets in a random order. For each tweet, RAs were asked to label the tweet as being related to “racial justice,” “COVID”, “neither,” or “both.” Krippendorff’s alpha on the nominal four-category task (COVID, racial justice, neither, both) was .81. On the task of identifying only tweets that were COVID-related or not, agreement was .85, and for racial justice, .84. Disagreements, which only occurred in 6.6% of the sample, were resolved with a third annotation by an author of the paper. Finally, we used the agreed-upon annotation to evaluate the accuracy, precision, and recall of our COVID-related and racial justice-related keyword classifiers. Accuracy for the racial justice and COVID-related keyword classifiers was (by chance) 90.7% for both classifiers. Precision for the racial justice and COVID-related classifiers was 90.7% and 88.4%, respectively. And recall for the racial justice and COVID-related classifiers was 88.6% and 92.8%, respectively. Overall, these results give us confidence that our model both identifies content that is largely related to the relevant topics, and does not miss a large proportion of related content.

### **Tweets Related to COVID-19**

To curate a list of COVID-19 related keywords, we started with three sources: Chen, Lerman, and Ferrara’s keyword list for their COVID-19 Twitter Dataset (2020), Green et al.’s keyword list for their study on COVID-19 elite polarization (2020), and Twitter’s official keyword list for their COVID-19 streaming endpoint.<sup>2</sup> We then manually added additional keywords as necessary to increase our coverage across COVID-19 related topics, and removed a small number of keywords that were too broad for the entire nine month sampling period (e.g. “china”). Our final multi-lingual keyword list consists of 974 keywords, including specific references to the virus like “covid-19” and “coronavirus”; phrases about the pandemic’s social impact (e.g. “six feet apart” and “reopening”); terms about responses from public officials (e.g. “contact tracing” and “ventilators”); and terms associated with misinformation (e.g. “plandemic” and “dr. immanuel”).

While the emergence of COVID-19 was recognized as early as November 2019 in China, discussion about the virus in the United States was not widespread until 2020. To identify COVID-19 related tweets, we therefore consider the entire time period from January

---

<sup>2</sup><https://developer.twitter.com/en/docs/labs/covid19-stream/filtering-rules>

1, 2020 to September 30, 2020. Using our keyword identification strategy, we find a total of 27,922,328 tweets which are related to COVID-19, where over 97% of the tweets from our American panel members were posted after February. These tweets represent 9.8% of all tweets made by our panel during this time window. In total, the nearly 28 million COVID-related tweets come from 480,773 unique users, who represent 61% of all panelists that were active from January to September.

### **Tweets Related to Racial Justice**

To create a list of keywords related to racial justice, we began with two prior pieces of work: the foundational Black Lives Matter report from Freelon et al. (2016), and the comprehensive hashtag activism research by Jackson et al. (2020). We then manually added a number of terms specific to the 2020 protests and conversations that arose after the murder of George Floyd on May 25th. For example, these include direct references to the Minnesota protests (e.g. “justiceforgeorgefloyd” and “minneapolis police department”), names of various bail funds (e.g. “bail fund” and “minnesota freedom fund”), phrases concerning police abolition (e.g. “defund the police” and “abolitionist”), and keywords regarding right-wing responses (e.g. “armed militia” and “alllivessplatter”). In total, we curated 345 keywords related to racial justice and the corresponding protests.

Racial injustice in the U.S. has a much longer history than Twitter, and even “Black Lives Matter” as a named movement is nearly a decade old, having been founded by Alicia Garza, Patrisse Cullors, and Opal Tometi in 2013 following the murder of 17-year-old Trayvon Martin. For the purposes of this study, however, we are particularly interested in examining discourse within the 2020 calendar year, and intentionally choose to focus on the unprecedented level of discourse and awareness which took place following the murder of George Floyd on May 25, 2020 (Anderson et al., 2020). Further, we expect the validity of keywords as an identification strategy to be more time-dependent. While some keywords, such as “BlackLivesMatter” always appropriately identify tweets about racial justice, others, such as “no-knock warrant,” may only indicate topical relevance within a specific time window. For these reasons, we conduct our keyword search for tweets related to racial justice between the dates of May 1, 2020 and September 30, 2020.



During this time window, we find 13,148,448 tweets related to racial justice. This unprecedented surge in attention for Black Lives Matter has been noted by other sources (Anderson et al., 2020), but here we will dig deeper into the demographics and dynamics of this conversation. Tweets related to racial justice make up 8% of all tweets between May 1 and September 30, 2020, and 4.6% of tweets since January 1, 2020, even though we did not start accounting for these tweets until four months into the year. These tweets come from 320,589 unique users, representing 40.9% of panelists who were active during the first nine months of 2020.

### *Examining Activity & Engagement*

A core focus of this work is in examining topical and demographic differences in activity—what panelists *do*—and engagement—what panelists *receive*. We examine these differences for the full panel of active users as well as for those who tweet about COVID-19 or the Black Lives Matter movement. Furthermore, we examine variation within activity and engagement, studying the different mechanisms through which panelists may create content or have their content engaged with.

### **Measuring Activity & Amplification**

There are four distinct ways in which individuals may create content: they may author tweets themselves, retweet others' tweets, add commentary by quote retweeting, or directly reply to another tweet. All four of these behaviors appear within the Twitter API as unique tweets generated by a given user. The id of any tweets which are retweeted, quote retweeted or replied to are also included in the returned tweet object, making it possible to determine which types of activities were involved. Note that these activities are not mutually exclusive—a person may reply to a tweet with a quote of a separate tweet, for example.

Previous work suggests that these different tweet types may signal distinct relationships with the content being shared, as users engage with each other to argue, persuade, commiserate, and amplify each other's ideas (Garimella et al., 2016; Kim and Yoo, 2012). We therefore not only examine the total volume of tweets produced by panelists, but further examine topical and demographic variation in the specific types of tweets used. An

authored tweet consists of new content which a person has intentionally chosen to put into the world. While this content is frequently assumed to be written solely by a Twitter user, authored tweets also include automatically triggered content, such as posts generated from an RSS feed. Retweets, on the other hand, contain no new content and simply amplify the content of others. Quote tweets fall between these two extremes, amplifying existing content while also containing additional commentary from the quoting user. Finally, replies capture more directed interactions as a user posts new content in response to another user.

Since these categories are not mutually exclusive, we define retweets, quote tweets and replies as any tweet which includes the given activity, regardless of whether or not it falls into another category as well. Authored tweets are restricted to those which do not include any of the other activities.

Examining demographic and topical variation in use of these tweet modalities highlights the behavioral similarities and differences between these populations. In turn, this behavioral coherence or divergence lends insight into the degree to which topical context or identity constitute an appropriate level of aggregation.

### **Measuring Engagement & Attention**

In addition to producing different types of content, users on Twitter may receive engagement from others who retweet, quote tweet, or favorite their produced content. The interactions that Twitter content receives is further tied to the attention enjoyed by the content creator themselves. The spread of a tweet is a cascading process, as each new engagement may garner additional engagements. We typically think of this as being a follower-driven process (Goel et al., 2016), with a user's tweet first shared with an account's followers and then continuing to spread as their followers and followers' followers further engage with the content. However, the Twitter algorithm's use of hashtags, trending topics, and other features may expose non-followers to a user's content as well. Furthermore, while we may imagine a Twitter author garnering an increased number of followers as their content spreads to new users, it is not clear that those who interact with a viral tweet will generally also take the step of following that tweet's author.

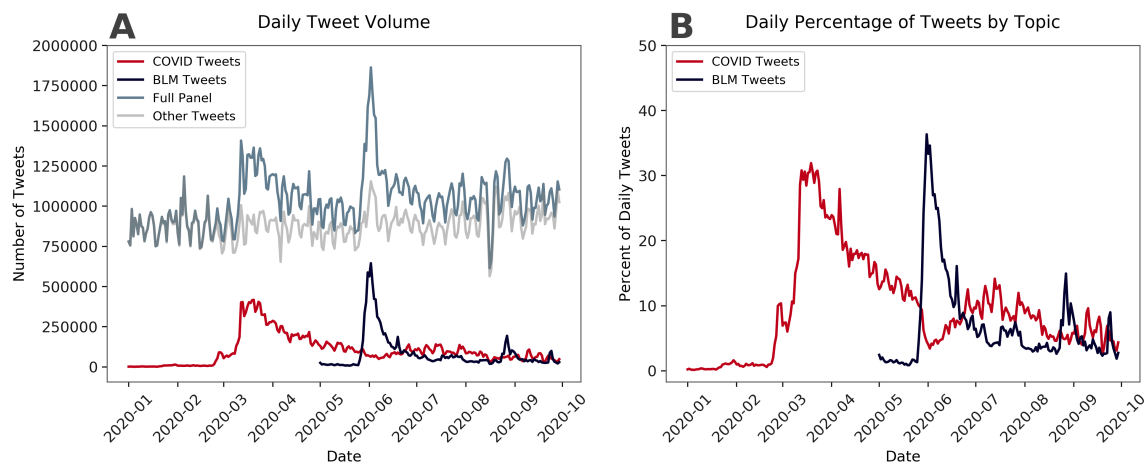
We focus our analysis of engagement on the number of followers the posting account has as well as on four types of tweet interactions available through the API—retweets, quote tweets, favorites, and favorites attributed to quoted tweets. Note that these four interaction types occur at the tweet level, while follower count is best understood at the user level. While Twitter content may also receive interactions in the form of replies, this count is not currently available through the Twitter API without enterprise-level access, and so we do not include it in our analysis. It’s also important to note that the behavioral trace data available through the API are imperfect measures of the full breadth of attention and engagement made by users (Lomborg and Bechmann, 2014; Freelon, 2018; Bruns, 2019; Puschmann, 2019; Tromble, 2021). For example, attention is only visible to the API in the form of platform actions; it cannot detect “lurkers” who read and engage with content but take no on-platform actions (Lomborg and Bechmann, 2014; Bernstein et al., 2013).

The amount and types of engagement received by panelists can provide insight into the attention afforded to different publics. A particularly interesting question here is the degree to which received engagement is reflective of user’s on-platform behavior or their personal characteristics. For example, we might imagine that those who tweet the most receive the most engagement, regardless of their demographics or topical interests. Alternatively, we may find interactions between these elements, with some demographics, or some topics, generally receiving more attention for their content. Again, topical and demographic variation in this regard can lend insight into which populations can be meaningfully considered as publics and hold implications for aggregation over different dimensions of data.

In order to get the most recent counts of attention received by panelists, we rehydrate all tweets related to COVID-19 and Black Lives Matter, as well as the additional 243,584,517 tweets posted by members of our panel. Completed in October–December of 2020, we retrieve 90% of all COVID-19-related tweets, 88% of Black Lives Matter tweets, and 89% of the remaining panelist tweets. Additionally, we get follower counts by retrieving current metadata for for active members of our panel. Completed in October 2020, we were able to retrieve metadata for 98% of our active panelists.

It is important to note the effect of time on these analyses. While we were able to retrieve updated data for the vast majority of tweets and users, some of this content was no longer available at the time of retrieval. On the other hand, variation in the time between initial posting and initial data collection could have induced a different type of temporal bias, as the number of engagements recorded in a retrieved tweet object may depend heavily on whether retrieval is done five minutes, five hours, or five days after initial posting (Shugars and Beauchamp, 2019; Wang et al., 2016; Kwak et al., 2010; Starbird and Palen, 2012; Karpf, 2019, 2020). We therefore end our analysis by looking at these temporal biases and recommending best practices to future researchers.

## Results



**Figure 1. Daily tweet totals from January 1 - September 30, 2020.**

*Note.* **A)** All tweets and topics. “Other tweets” captures the volume of tweets with neither a “COVID” or “BLM” tag. **B)** Percentage of daily tweets categorized as related to COVID-19 or Black Lives Matter

In this paper, we provide empirical insight into the behaviors and identities of Twitter publics, which contributes essential guidance on how researchers should approach potential aggregation and interpretation of Twitter data. Specifically, we describe the demographic identity, activity, and engagement of U.S. voters on Twitter, both as a whole and as realized through two highly relevant topics. These distinct but overlapping populations focus on two major events of 2020: the COVID-19 pandemic and racial justice protests

centered on the Black Lives Matter movement. For simplicity, we refer to the latter topic in our analyses and figures as “Black Lives Matter,” even though it is broader than just the movement itself.

### *Topical Populations*

As seen in Figure 1A, these topics each dominated the conversation at different points of the year, with COVID-19 reaching a single-day peak of 417,476 tweets on March 20, 2020, and the Black Lives Matter movement seeing an astounding 645,066 tweets on June 2, 2020 alone. Interestingly, we further see that these topical peaks do not appear to coincide with a decrease in other Twitter content, which remains relatively stable at a median of 887,304 tweets per day, or 1.13 daily tweets per panelist. We define “other” content as being all remaining tweets, e.g., those which are not related to either COVID-19 or Black Lives Matter. This suggests that increased attention to one topic does not necessarily imply a trade-off of decreased attention to other topics. Figure 1B further illustrates just how monumental these peaks were, showing the percent of each day’s tweets attributed to each topic. At its peak, COVID-19 was discussed in 30.6% of all tweets made on a single day. In early June, Black Lives Matter was the subject of over one third (34.6%) of all posted tweets.

Furthermore, both of our studied topics engaged distinct subpopulations of panelists. Of the 783,967 panelists who were active in 2020, only 511,590 (65%) engaged with at least one of the topics. Of those, a majority (56.6%, 289,772 accounts) tweeted about both topics at least once during our nine-month time window. An additional 30,817 panelists (6.0%) only tweeted about Black Lives Matter while the remaining 191,001 (37.3%) tweeted only about COVID-19.

While there is substantial overlap in panelists who tweeted about both topics, the difference in these populations illustrates that, far from being monolithic, Twitter is composed of users with a multitude of interests and behaviors. Even the most popular, salient topics do not engage all users as a single public, and many users engage across multiple topics. These populations are neither entirely separate nor quite the same; they exist within a shared platform context, but may engage with and experience the platform in differing

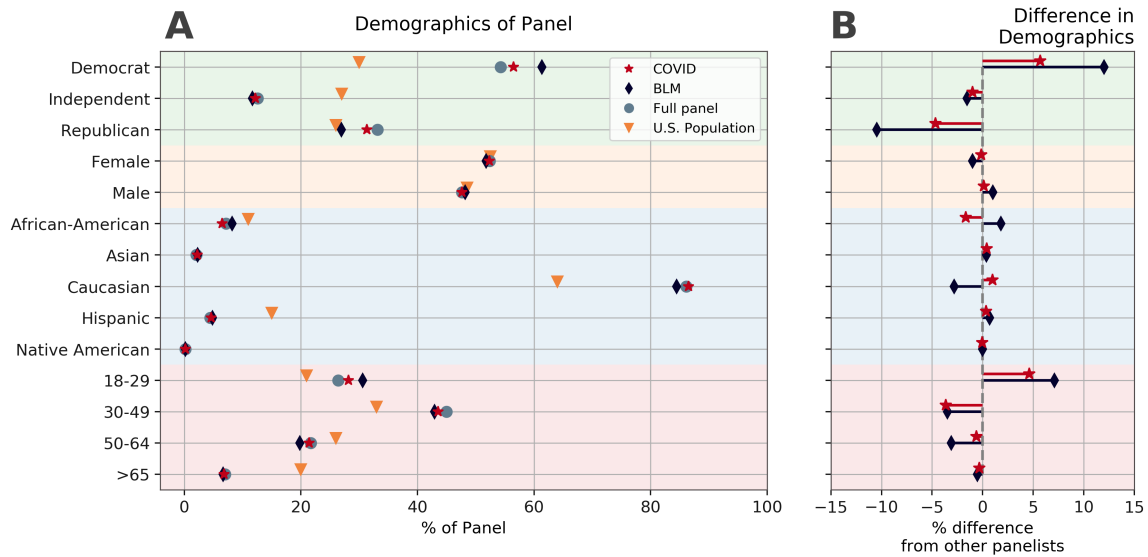
ways. This finding has important implications for researchers studying topically-induced subsamples of tweets, suggesting that topic alone is not a meaningful indicator of a coherent public. Topical inquiries can still be extremely valuable, but researchers should be aware that the individuals generating this data may not be acting within clearly delineated, topic-oriented bounds. Rather, a sample of topically induced tweets is likely to contain users who seriously and selectively engage on that issue alongside a potentially larger population for whom that topic is just one of many things they are talking about.

In the following sections, we examine the similarities and differences between these topical populations as they are expressed through demographics, activity and engagement. We end with analysis of the temporal bias inherent in our data.

### *Demographic Populations*

We find broad demographic similarities between Twitter users as a whole and the populations who engage in discussions around COVID-19 and Black Lives Matter. In Figure 2A, we see the demographic distributions of the full panel, along with the demographics of panelists who tweeted about each of these topics. For reference, we also include estimates of the demographic distribution of all U.S. adults (Pew Research Center, 2019b), which is notably different from the population of adult Twitter users. Despite the difference from the general population, previous work has found that the demographic composition of our panel is broadly reflective of the Twitter population (Hughes et al., 2020), though our panel is slightly more female and notably more white. Specifically, our panel slightly underrepresents Hispanics and notably underrepresents Asians compared to the overall Twitter population (Hughes et al., 2020). The portion of African-American users in our panel is reflective of this demographic's share of the overall Twitter population (Hughes et al., 2020). While the underrepresentation of Hispanic and Asian users is an important limitation of our dataset, this previous validation of our matching strategy leads us to believe that our results are broadly reflective of demographic trends in user activity on Twitter in the United States during the time under study.

Figure 2B gives more detailed insight into the differences of these populations, showing the percent difference between the share of panelists who tweet about a given topic and



**Figure 2. Demographic characteristics of panelists.**

*Note.* Full panel covers all 783,967 active panelists, COVID covers 480,773 panelists who tweet about COVID-19, and BLM covers 320,589 panelists who tweet about the Black Lives Matter movement. Additionally U.S. Population reflects the demographic proportions among all U.S. adults as reported by Pew Research Center (2019a) **A**) The portion of panelists within each demographic. **B**) Percent difference in demographic proportion between panelists who tweet about a given topic and those who do not.

those who do not. This calculation is made independently at the topic level. For all demographic analyses, we calculate proportions using only tweets from panelists who have a known label within each demographic category.

We find that Democrats are overrepresented in the set of users who tweeted about COVID-19 (+5.7 percentage points), and Republicans are underrepresented (-4.6 percentage points). This pattern also arises in the set of users who tweeted about Black Lives Matter, but with much greater disparities: Democrats are overrepresented by 12 percentage points, and Republicans are underrepresented among this population by 10.5 points. In the case of gender, the breakdowns of users who tweet about either topic almost exactly matches the demographics of those who do not tweet about these topics. The same holds true for race, although the set of users who tweeted about Black Lives Matter is slightly

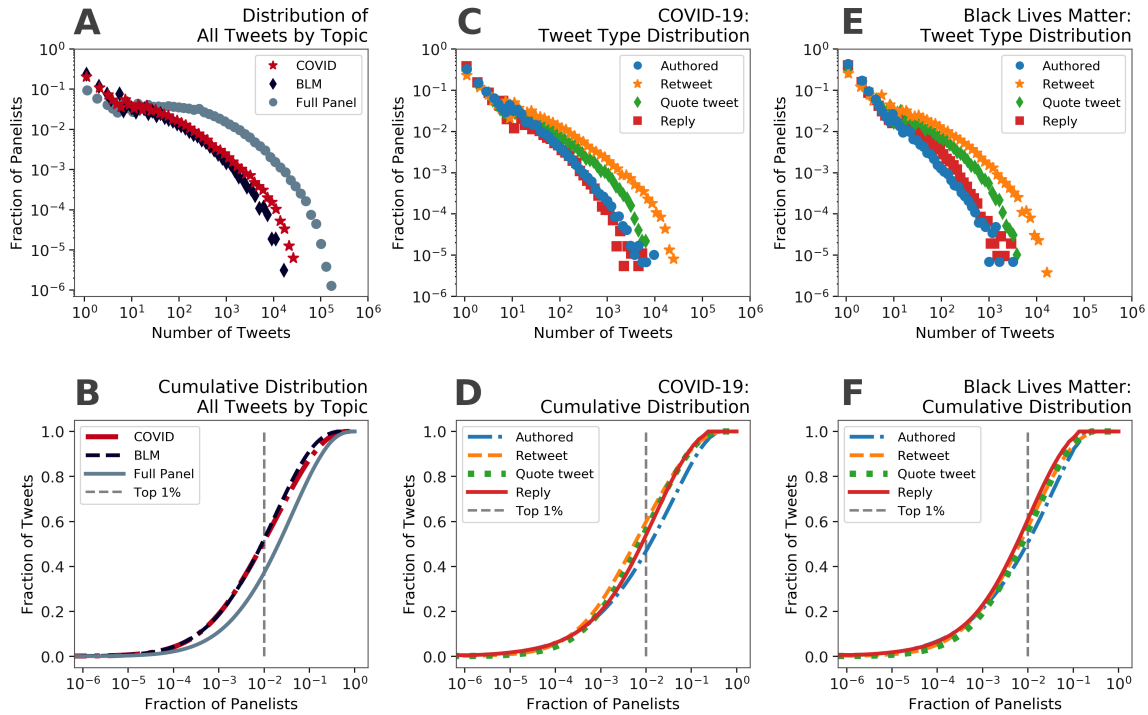
more African-American (+1.8 points) and slightly less Caucasian (-2.8 points) than the full panel. In the case of age, a disproportionate number of 18-29 year-olds tweeted about Black Lives Matter (+8.0 points) and, to a lesser extent, COVID-19 (+4.6 points). Using a t-test to compare our observed demographics to a comparable sample bootstrapped from the full panel, we find all these differences to be significant at the  $p < 0.01$  level.

In many respects, the demographic similarity across these topics is striking. This suggests that, at least for highly salient topics, researchers should avoid assumptions about the demographic composition of Twitter samples. For example, while we do see that the corpus of Black Lives Matter tweets tends to include slightly more Democrats and African-Americans, it would be would not be accurate to conceptualize this corpus as *representative* of Democratic or African-American voices. Rather each topic appears to engage a cross-section of users who are broadly reflective of the Twitter population as a whole. This further emphasizes our finding that “a topic” should not be conflated with “a public,” as topical tweets are likely to contain a wide variety of users who may not be meaningfully participating in the same conversation. While researchers may not always be able to obtain demographic information for users, they should be be mindful that a given sample of tweets likely captures a diversity of perspectives and motivations. We therefore urge researchers to carefully consider approaches for disaggregating behavioral indicators of this diversity. In the following sections we examine two such behavioral features easily measured through the Twitter API.

### *Activity & Amplification*

Next, we examine the implications of data aggregation in the context of content created and amplified by members of our panel. In addition to considering topical and demographic aggregation, we give particular attention to the implications of aggregating over different types of tweet activities, namely authoring tweets themselves, retweeting others’ tweets, quote tweeting, and replying to others’ tweets. We first consider aggregation at the user-level and examine distributions of tweet activities by topic. We then take a closer look at how these activities breakdown across both demographic and topical categories, illustrating the types of insights which may be lost through higher levels of data aggregation.





**Figure 3.** **A)** Distribution of total number of tweets made by panelists. Distributions are visualized using log binning. **B)** Cumulative distribution of fraction of tweets made by fraction of panelists. **C)** Distribution of tweets of each type made about COVID-19. **D)** Cumulative distribution of fraction of tweets about COVID-19 made by fraction of panelists. **E)** Distribution of tweets of each type made about Black Lives Matter. **F)** Cumulative distribution of fraction of tweets about Black Lives Matter made by fraction of panelists.

### Topical Activity

Figure 3 illustrates the distribution of tweet activities and the topical distribution aggregated over all tweet types. In Panel A, we see that the distributions for COVID-19 tweets and Black Lives Matter tweets are both heavy tailed and nearly identical. We further see in the cumulative distribution shown in Panel B, that for both topics, the top 1% most active tweeters account for nearly half of all tweets on that topic. Interestingly, we further see that while tweets from the full panel similarly follows a heavy-tailed distribution, this broader content is produced by a larger fraction of panelists. The top 1% of most active users in this category account for only 37% of the content. This is likely due to the sheer difference in the volume of tweets; however, this may also suggest that there are notable

distributional differences between content intentionally generated around a specific topic versus that which is not. While we are fairly confident that the users in our panel are real people, many of these accounts do display semi-automated behavior, such as the daily posting of horoscopes or content from RSS feeds. We therefore expect that the full panel data reflects a mix of intentionally generated user content as well as an additional volume of automatically generated tweets, which may account for the slight difference in distribution.

This highlights an important *benefit* of topical aggregation—while these subsamples should not be considered as unified publics, they do appear to represent something more coherent than the full corpus of Twitter as a whole. This suggests that keyword-based, topical analysis can indeed place meaningful bounds on a construct of interest, though, as our earlier analysis indicates, these bounds should always be understood to be arbitrary, imperfect, and researcher-imposed.

### **Disaggregating Tweet Activity**

In Panels C–F of Figure 3, we see how this overall tweet distribution breaks down across the different types of tweet modalities. As we might expect, all four types of tweets follow heavy-tailed distributions, with some users producing a lot of tweets, and others producing hardly any at all. However, we further see that these tweet types do not follow the exact same distribution, suggesting that they may be leveraged in distinctly different ways. Authored tweets overall account for the smallest portion of panelist tweets, and retweets are by far the most common type of tweet. Across the nine months (274 days) of the full dataset, the average panelist produces a total of 363 tweets, consisting of 56.3 (15.5%) authored tweets, 188.1 (51.8%) retweets, 55.8 (15.4%) quote tweets and 95.4 (26.3%) replies. Note that non-authored tweets may fall into multiple categories. This pattern could point to retweets having a lower cost of engagement, since the retweeting user does not need to expressly articulate their own position. Alternatively, it could be a topic-specific effect, or reflect larger platform norms about the degree to which one should exercise their own voice versus amplifying the voice of others.

**Table 1: Proportions of tweet activity by topic and demographic.**

**(a) Tweeting activity of full panel**

	All Tweets		Authored		Retweets		Quote Retweets		Replies	
	#	%	#	%	#	%	#	%	#	%
<b>Democrat</b>	176,194,814	61.9%	27,270,794	61.7%	90,599,663	61.4%	27,926,739	63.8%	46,756,648	62.5%
<b>Independent</b>	30,223,054	10.6%	4,925,818	11.2%	15,673,071	10.6%	4,405,772	10.1%	7,798,141	10.4%
<b>Republican</b>	78,163,355	27.5%	11,974,384	27.1%	41,219,840	27.9%	11,420,546	26.1%	20,270,268	27.1%
<b>Female</b>	139,063,038	50.4%	20,917,274	48.9%	78,234,613	54.8%	21,581,665	50.9%	31,676,417	43.5%
<b>Male</b>	136,754,581	49.6%	21,877,163	51.1%	64,432,209	45.2%	20,785,468	49.1%	41,142,293	56.5%
<b>African-American</b>	25,171,461	9.1%	4,707,920	11.0%	13,135,088	9.2%	4,215,882	10.0%	5,480,558	7.5%
<b>Asian</b>	5,653,689	2.1%	869,658	2.0%	3,025,064	2.1%	867,455	2.1%	1,402,187	1.9%
<b>Caucasian</b>	230,731,435	83.7%	35,111,482	82.0%	118,233,391	82.8%	35,042,020	82.8%	62,904,191	86.6%
<b>Hispanic</b>	13,721,492	5.0%	2,059,569	4.8%	8,119,376	5.7%	2,124,756	5.0%	2,764,239	3.8%
<b>Native American</b>	424,831	0.2%	75,114	0.2%	217,050	0.2%	57,080	0.1%	106,031	0.1%
<b>18-29</b>	69,085,777	24.4%	10,886,176	24.8%	38,982,334	26.6%	11,251,437	25.9%	14,680,558	19.8%
<b>30-49</b>	101,505,120	35.9%	17,689,534	40.3%	43,001,508	29.4%	14,310,501	33.0%	33,776,542	45.5%
<b>50-64</b>	73,860,740	26.1%	10,454,270	23.8%	40,297,317	27.5%	11,440,470	26.3%	18,869,670	25.4%
<b>65+</b>	38,148,722	13.5%	4,836,953	11.0%	24,205,522	16.5%	6,422,678	14.8%	6,971,875	9.4%

**(b) Tweet activity related to COVID-19**

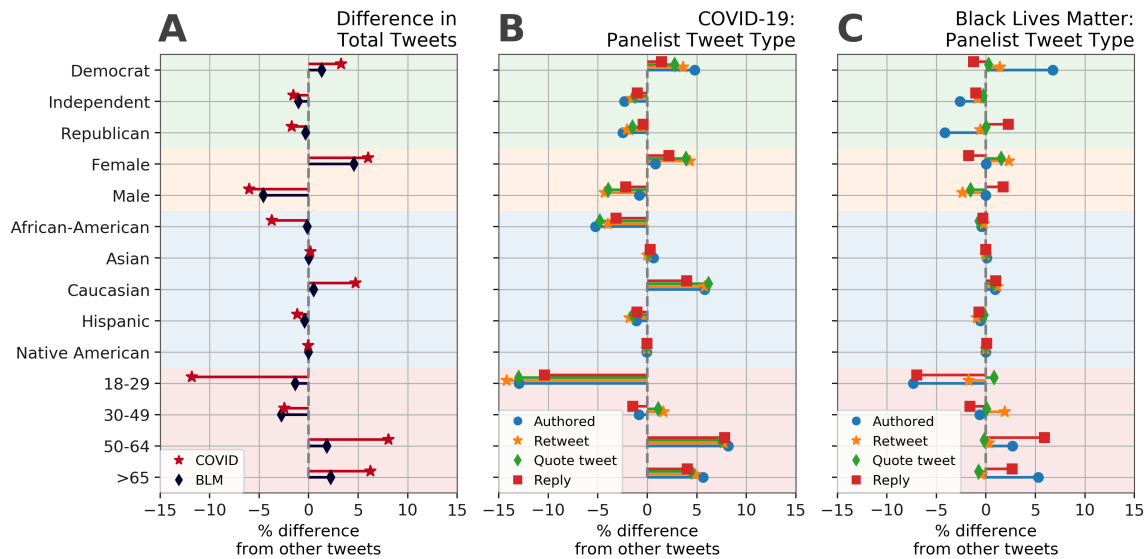
	All Tweets		Authored		Retweets		Quote Retweets		Replies	
	#	%	#	%	#	%	#	%	#	%
<b>Democrat</b>	18,109,804	64.9%	2,321,458	66.1%	12,905,264	64.6%	4,660,106	66.2%	1,272,991	63.9%
<b>Independent</b>	2,572,878	9.2%	316,889	9.0%	1,852,275	9.3%	634,171	9.0%	187,996	9.4%
<b>Republican</b>	7,239,645	25.9%	871,931	24.8%	5,232,304	26.2%	1,749,654	24.8%	531,860	26.7%
<b>Female</b>	15,252,547	55.8%	1,706,280	49.6%	11,446,831	58.6%	3,736,275	54.2%	892,745	45.6%
<b>Male</b>	12,070,695	44.2%	1,731,036	50.4%	8,103,425	41.4%	3,154,177	45.8%	1,064,590	54.4%
<b>African-American</b>	1,562,609	5.8%	209,881	6.2%	1,121,117	5.8%	406,043	6.0%	86,334	4.5%
<b>Asian</b>	599,430	2.2%	87,909	2.6%	413,578	2.1%	147,270	2.2%	42,921	2.2%
<b>Caucasian</b>	23,808,534	88.0%	2,966,201	87.3%	17,028,535	87.8%	6,002,718	88.0%	1,751,475	90.4%
<b>Hispanic</b>	1,064,813	3.9%	129,049	3.8%	801,186	4.1%	258,588	3.8%	53,757	2.8%
<b>Native American</b>	28,157	0.1%	4,033	0.1%	19,378	0.1%	6,293	0.1%	2,039	0.1%
<b>18-29</b>	3,828,926	13.8%	448,881	12.9%	2,853,195	14.4%	1,050,876	15.0%	191,204	9.7%
<b>30-49</b>	9,350,161	33.7%	1,378,184	39.5%	6,112,150	30.8%	2,370,709	33.9%	872,300	44.0%
<b>50-64</b>	9,268,110	33.4%	1,093,451	31.4%	6,794,953	34.2%	2,280,685	32.6%	653,933	33.0%
<b>65+</b>	5,306,040	19.1%	565,470	16.2%	4,114,831	20.7%	1,294,666	18.5%	263,434	13.3%

**(c) Tweet activity related to Black Lives Matter**

	All Tweets		Authored		Retweets		Quote Retweets		Replies	
	#	%	#	%	#	%	#	%	#	%
<b>Democrat</b>	8,260,991	63.2%	570,689	68.4%	6,488,495	62.7%	2,477,138	64.1%	564,776	61.2%
<b>Independent</b>	1,262,308	9.7%	71,703	8.6%	1,016,454	9.8%	379,432	9.8%	86,915	9.4%
<b>Republican</b>	3,551,080	27.2%	192,364	23.0%	2,835,440	27.4%	1,009,606	26.1%	270,570	29.3%
<b>Female</b>	6,932,495	54.8%	397,139	48.9%	5,696,560	57.0%	1,954,853	52.4%	376,636	41.8%
<b>Male</b>	5,723,927	45.2%	414,944	51.1%	4,294,535	43.0%	1,778,745	47.6%	524,832	58.2%
<b>African-American</b>	1,134,080	9.0%	85,050	10.6%	898,600	9.0%	347,015	9.3%	64,610	7.2%
<b>Asian</b>	262,499	2.1%	16,950	2.1%	209,354	2.1%	79,514	2.1%	17,068	1.9%
<b>Caucasian</b>	10,615,208	84.2%	666,976	82.9%	8,365,852	83.9%	3,109,617	83.5%	782,677	87.6%
<b>Hispanic</b>	578,890	4.6%	34,213	4.3%	482,354	4.8%	181,272	4.9%	27,704	3.1%
<b>Native American</b>	18,951	0.2%	1,244	0.2%	14,293	0.1%	5,222	0.1%	1,893	0.2%
<b>18-29</b>	3,008,523	23.2%	145,975	17.6%	2,567,393	25.0%	1,023,304	26.7%	117,733	12.9%
<b>30-49</b>	4,323,514	33.3%	329,091	39.7%	3,195,668	31.1%	1,267,428	33.0%	401,548	43.9%
<b>50-64</b>	3,622,884	27.9%	219,381	26.5%	2,854,895	27.8%	1,004,125	26.2%	285,772	31.2%
<b>65+</b>	2,027,432	15.6%	134,294	16.2%	1,650,936	16.1%	542,238	14.1%	109,995	12.0%

The divergence in use of different tweet types suggests that these modalities may indeed carry different meaning and implications. This, in turn, suggests that researchers must be cautious when conflating different types of tweet activities. There may be some research questions for which using the full volume of tweets, regardless of type, is appropriate, and others for which the distinctiveness of these modalities must be considered. We recommend that researchers examine the robustness of their findings across different aggregations of tweet activity in order to determine if there are particular modalities driving results.

### Demographic & Topical Activity



**Figure 4. Difference in proportion of tweets from each demographic.**

*Note.* Differences are measured by comparing demographic distribution within a topic to the distribution among all other tweets. **A)** Differences across panelist tweets for both COVID-19 and Black Lives Matter. **B)** Difference in demographic proportions for each tweet type for COVID-19. **C)** Differences in demographic proportions for each tweet type for Black Lives Matter.

We further examine demographic differences in panelists’ tweet activity and their amplification of others’ content, highlighting the types of results which may be lost in aggregations that treat users as identical. Table 1 details the volume and proportion of

tweets of each type made by different demographic segments of our panel, with Table 1A capturing the full panel, and Tables B and C capturing tweets related to COVID-19 and Black Lives Matter respectively.

These differences in panel activity are further highlighted in Figure 4. While previously, Figure 2B showed the differences in the proportion of *panelists* who engage with each topic, Figure 4A shows demographic differences in the proportion of *tweets* posted on each topic. This highlights variation in the volume of tweets produced by panelists from different demographic groups. For example, we see that women, white people, and those age 50-64 as well as over the age of 65 all tweet more about COVID-19 than about other topics. This is particularly notable since all four of these demographic groups make up roughly equal proportions of the population of panelists who tweet about COVID-19 compared to those who do not. We also see that the youngest members of our panel, those 18-29, tweet significantly less about COVID-19 than they do about other topics.

Figures 4B and C show how these differences breakdown by tweet type for COVID-19 (Panel B) and Black Lives Matter (Panel C). Here we see that the four types of tweets follow broadly similar patterns, though each modality often represents a different share of the total tweets made. Furthermore, there are some notable exceptions, where different types of tweets seem to signify substantively different discursive behaviors. Interestingly, these differences are particularly prominent around Black Lives Matter and are not seen as strongly on the topic of COVID-19. In Figure 4B, we see that the different types of tweets almost always follow the same trend as the total number of tweets. For example, white panelists are, in general, more likely to tweet about COVID-19, and that increased likelihood extends across authored tweets, retweets, quote tweets, and replies.

However, this consistency in behavior changes around the topic of Black Lives Matter. For example, we see that tweets from Republicans are nearly equally represented in the corpus of tweets about Black Lives Matter compared to the corpus of all other tweets (Figure 4A). However, this equality in the proportion of total tweets is achieved through an underrepresentation of authored tweets and an overrepresentation in replies (Figure 4B). In other words, compared to Republicans tweeting on other topics, these panelists are less likely to author stand-alone tweets about Black Lives Matter, but more likely to post replies

on this topic. The exact opposite is true for Democrats, who are more likely to author tweets about Black Lives Matter, but less likely to respond to tweets on this topic.

The differences in these patterns highlight the complexity of Twitter publics. While some user behavior can be attributed to a particular topic or a given demographic, neither of these dimensions appear to be wholly sufficient for categorizing “a Twitter public.” Rather, these publics must be more fully understood as contextual communities which emerge around salient topics and are informed by demographic identities. These findings also highlight why disaggregating specific tweet types can be important to interpreting the behavior of a Twitter public. If we were to consider all tweet types to be equivalent, than we might erroneously conclude that some populations have little engagement with a topic when, in fact, different populations are engaging with the same topic in different ways. Even in datasets where demographic information is unavailable, examining behavioral differences in tweet activity has the potential to reveal distinctive sub-populations which could then be further studied through hand coding and content analysis.

Overall in this section we have seen that, while topically-induced subsamples should not be assumed to capture a unified public, topical filtering can introduce meaningful bounds for creating a corpus in which the diversity of user activity can then be interrogated. We have further seen that the specific types of activities users engage in—authoring tweets, retweeting, quote tweeting, and replying—may sometimes capture meaningful variation in user behavior. While researchers should take advantage of disaggregating by demographic data when it is available, they should always examine whether different aggregations of tweet activity effect their results.

### *Engagement & Attention*

Now that we have some insight into user activity on Twitter, we turn our attention to examining the implications of aggregating over different types of engagement users might receive. Specifically, users may receive engagement themselves in the form of follows from other users, or their content may receive engagements in the form of interactions from other users. In this paper, we consider follower counts along with four distinct types of interactions a user’s content may receive: retweets, quote tweets, favorites, and quote favorites. This

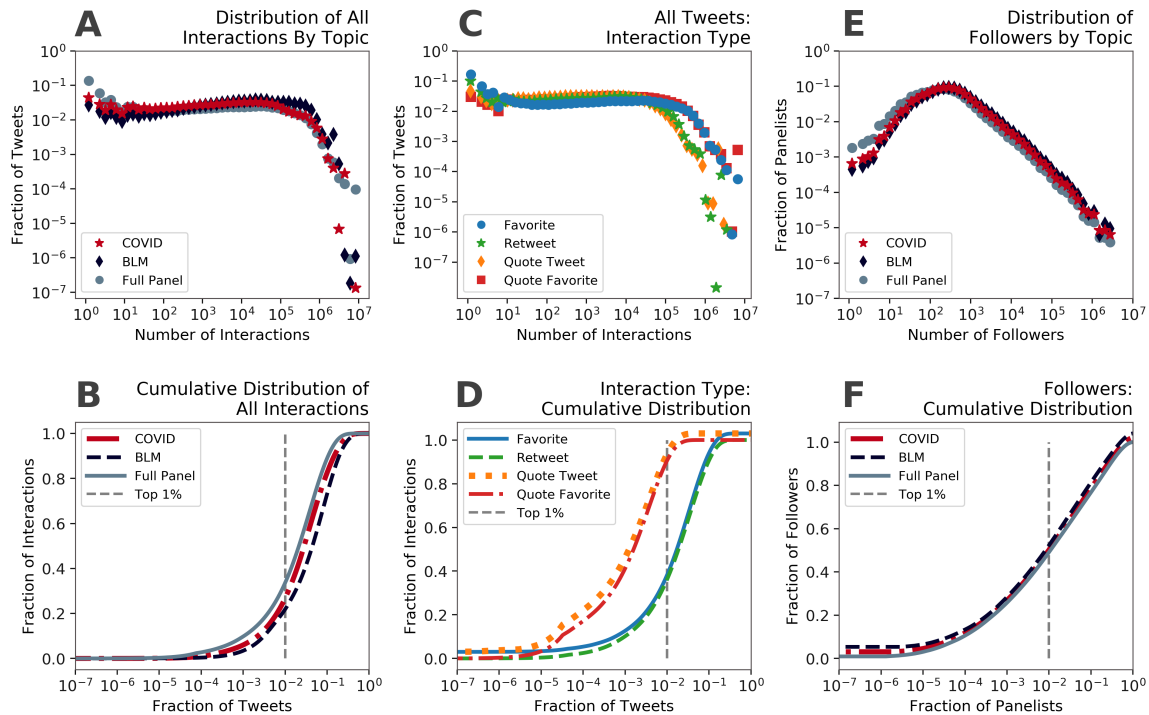
last form of interaction is the most indirect, capturing favorites attributed to a quote tweet of a user's content. While users may also receive attention in the form of replies, that data is not readily available through the Twitter API and therefore is not included as part of this analysis. Again, it is important to note that the behavioral trace data available through the API are imperfect measures of the full breadth of attention and engagement made by users on the platform (Lomborg and Bechmann, 2014; Bernstein et al., 2013).

As we did in Section 4.3, we again proceed by first examining data which has been aggregated over all users in our sample. This allows us to interrogate topic-level analysis as well as aggregation of different engagement types. We then add demographic detail to this analysis, examining how engagement varies over both topics and demographic groups. Finally, we add to this section an analysis of the interplay between activity and engagement, showcasing how these behavioral features can be considered together in both aggregate and disaggregate ways.

### **Engagement by Topic and Type**

Figure 5 illustrates the distribution of interactions received by panelist tweets. Panel A shows the topic-level distribution for an engagement measure which aggregates the total number of retweets, favorites, quote tweets, and quote favorites. Here we see that, for all topics, these distributions are relatively stable over several orders of magnitude. While there are a small number of tweets which receive an outsized proportion of interactions, roughly equal proportions of tweets are engaged with anywhere from 10 to  $10^5$  times. Within our corpus, the top 1% of COVID-19 tweets account for 26.8% of interactions, the top 1% of Black Lives Matter tweets account for 21.9% of interactions, and the top 1% of all tweets capture 33.6% of interactions.

This again suggests that while topical bounds can serve as a meaningful starting point for Twitter analysis, researchers should be cautious to not assume uniformity of behavior or experience within the resulting sample. Just because users are tweeting about a popular topic does not imply that individual tweets will be popular.



**Figure 5. Distribution of engagement received by panelists.**

*Note.* **A)** Distribution of total interactions received on panelist tweets. **B)** Cumulative distribution of fraction of interactions received by fraction of tweets. **C)** Distribution of interactions of each type made across the full corpus of rehydrated tweets. **D)** Cumulative distribution of fraction of interaction received by fraction of tweets, disaggregated by engagement type. **E)** Distribution of the number of followers across all panelist accounts. **F)** Cumulative distribution of fraction of panelists with a given fraction of followers.

We further see in Figure 5A that, while tweets on both topics cover a similar range of interaction counts, tweets about Black Lives Matter receive slightly more interactions than tweets about COVID-19. While both topics contain tweets which receive no interactions as well as tweets garnering over 9.8 million interactions, the tweets about Black Lives Matter receive a median of 3,687 and an average of 78,098 interactions (92nd percentile), while tweets about COVID-19 receive a median of 646 and an average of 39,519 interactions (88th percentile). This is particularly interesting because the COVID-19 tweets in our sample are generally older and therefore have had more time in which to accrue interactions. The fact



that the Black Lives Matter tweets received more interactions, then, points to the salience of this topic and the immense work activists have done to build the prominence, network, and shared language of this movement (Jackson et al., 2020). We will return to the temporal nature of these data in Section 4.5.

Panel C shows how this engagement breaks down across the four interaction types we examine. We only show these distributions for the full corpus of all tweets because, notably, we find substantively similar distributions among the subsets of topical tweets. We see that all four types of engagement follow roughly similar distributions though favorites are by far the most common. In total, 72% of all interactions in our dataset are favorites, followed by 18.7% of interactions which are retweets. Quote tweets are significantly less common, making up only 1.8% of interactions in our dataset, though this still represents a total of more than  $10^{11}$  interactions. Finally, favorites attributed to those quote tweets make up 7.3% of all interactions across the full dataset. The disparity between favorites and retweets is notable in part because retweeting and favoriting are mechanically similar, requiring approximately equally amounts of exertion. However, if these actions serve different social roles, these engagement types may come with different “costs.” Favoriting a tweet signals affirmation to that tweet’s author and generally has lower-visibility to other users. On the other hand, retweeting or quote tweeting intentionally amplifies and makes visible the content to one’s followers. This may cause engaged users to make strategic choices in terms of the specific content they choose to amplify.

Panel D shows the cumulative distribution of interactions received by fraction of tweets. Here we see the distinct difference between quote tweets and quote favorites on the one hand and retweets and direct favorites on the other. As quote favorites are interactions made with a quote tweet, it is not surprising that the cumulative distribution of quote favorites so closely matches that of quote tweets. However, it is notable that quote tweets themselves vary so distinctly from favorites and retweets, with the top 1% of tweets receiving 90% of quote tweet interactions. Retweets, on the other hand, are much more spread out, with the top 1% of tweets accounting for 36% of all retweet interactions. Note that this does not appear to be a temporal effect as the volume of quote tweets and retweets are consistent across the nine months of dataset. The fact that quote tweets seem to be more concentrated than retweets could be reflective of the difference in volume between retweets and quote

tweets, or could suggest that some content is more “quotable.” This could indicate content which is more controversial—in which people feel compelled to quote tweet with their own opinion—or could reflect tweets which intentionally prompt others to quote with their own story or joke.

Finally, in Panel E we see the distribution of followers by topic for every active member of our panel, accompanied in Panel F by the cumulative distribution for these follower counts. While this distribution is heavy-tailed, there also appears to be a minimum threshold, with the highest density of active panelists having around 100 followers. This finding may be an artifact of the age of our panel—since panelists were originally identified in 2017, they have all had at least three years in which to acquire followers. Additionally, we see that the the follower distribution for panelists who tweet about COVID-19, panelists who tweet about Black Lives Matter, and the full panel are nearly identical. However, it is notable that users who do not engage in our two core topics tend to have much fewer followers. In our dataset, users who tweet about COVID-19 have a median of 230 and an average of 1,165 followers while users who tweet about Black Lives Matter have a median of 262 and an average of 1,373 followers. However, users who do not engage with either of these topics have a median of 79 and an average of only 284 followers. Because we measure follower counts in October 2020, we do not know how a user’s topical tweets influenced their follower count. But, the fact that users who engage with either topic have similar distributions of followers suggests that the shape of this distribution may be fairly stable when examining discourse around a salient topic. This further suggests that there may be reinforcement dynamics between topical engagement and follower count. Perhaps people who tweet about the most pressing matters garner more followers for their engagement, or perhaps accounts with more followers are more likely to engage in the salient topics of the day.

**Table 2: Proportions of received engagement and user follower counts.**

(a) Engagement received across all rehydrated tweets and active panelists.

	All Engagement		Retweets		Favorites		Quote Retweets		Quote Favorites		Followers	
	tweet mean	% of total	tweet mean	% of total	tweet mean	% of total	tweet mean	% of total	tweet mean	% of total	user mean	% of total
Democrat	34,003	64.0%	6,387	63.6%	24,366	63.8%	659	66.7%	2,590	66.6%	1,035	64.8%
Independent	40,843	12.7%	7,825	12.8%	29,715	12.8%	667	11.1%	2,636	11.2%	669	9.5%
Republican	28,935	23.3%	5,536	23.6%	20,870	23.4%	512	22.2%	2,017	22.2%	693	25.7%
Female	37,287	58.3%	7,156	59.3%	27,028	58.9%	628	52.1%	2,474	52.2%	781	46.0%
Male	27,571	41.7%	5,085	40.7%	19,544	41.1%	598	47.9%	2,345	47.8%	996	54.0%
African-American	42,525	11.4%	8,007	11.3%	30,230	11.2%	857	12.3%	3,430	12.5%	869	7.2%
Asian	43,753	2.7%	8,442	2.7%	31,681	2.7%	749	2.5%	2,881	2.4%	1,043	2.5%
Caucasian	30,186	77.1%	5,677	76.7%	21,665	76.9%	577	79.0%	2,266	78.8%	885	86.9%
Hispanic	61,578	8.7%	12,139	9.1%	45,462	8.9%	803	6.1%	3,174	6.1%	657	3.3%
Native American	45,449	0.2%	8,720	0.2%	33,238	0.2%	704	0.2%	2,786	0.2%	391	0.1%
18-29	75,397	50.3%	14,638	51.6%	55,719	51.7%	985	35.2%	4,055	36.9%	523	15.2%
30-49	23,853	26.1%	4,268	24.7%	16,391	24.9%	654	38.4%	2,539	37.9%	973	50.3%
50-64	19,313	15.8%	3,632	15.7%	13,659	15.6%	426	18.7%	1,597	17.8%	1,003	25.2%
65+	18,373	7.9%	3,543	8.0%	13,232	7.9%	332	7.6%	1,266	7.4%	1,164	9.4%

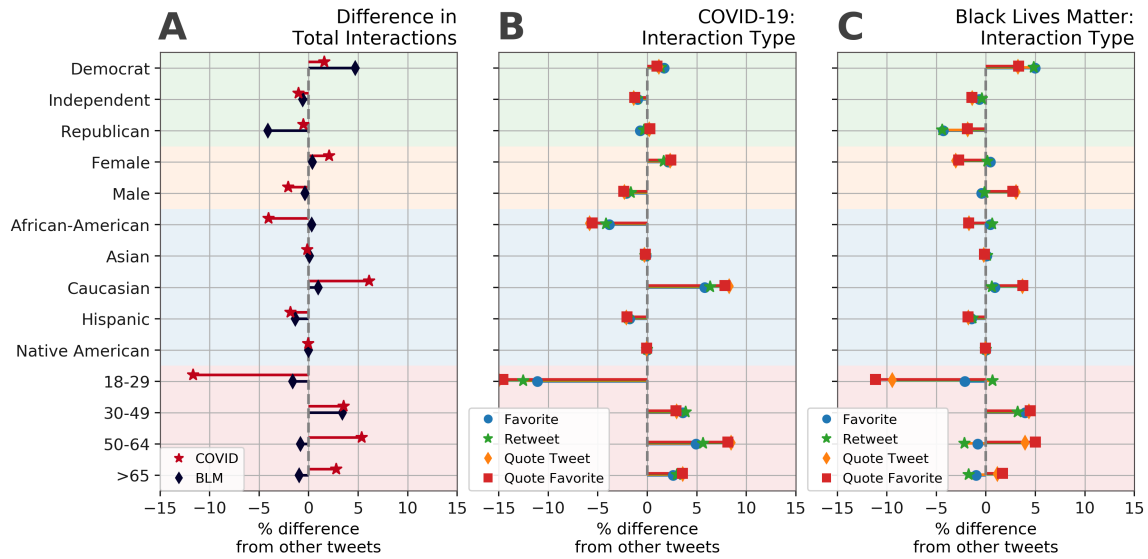
(b) Engagement received by tweets related to COVID-19 and follower counts

	All Engagement		Retweets		Favorites		Quote Retweets		Quote Favorites		Followers	
	tweet mean	% of total	tweet mean	% of total	tweet mean	% of total	tweet mean	% of total	tweet mean	% of total	user mean	% of total
Democrat	38,956	65.0%	7,390	64.3%	27,891	64.8%	779	67.4%	2,896	67.3%	1,377	66.4%
Independent	52,191	11.9%	10,131	12.1%	37,997	12.1%	853	10.1%	3,211	10.2%	893	9.1%
Republican	36,433	23.1%	7,129	23.6%	26,076	23.1%	681	22.4%	2,546	22.5%	935	24.5%
Female	41,312	59.9%	7,994	60.6%	29,861	60.5%	731	54.3%	2,726	54.4%	1,050	46.0%
Male	35,181	40.1%	6,626	39.4%	24,856	39.5%	785	45.7%	2,914	45.6%	1,344	54.0%
African-American	55,364	8.2%	10,427	8.0%	39,955	8.2%	1,023	7.8%	3,958	8.1%	1,259	7.0%
Asian	46,106	2.6%	8,910	2.6%	33,482	2.6%	782	2.3%	2,932	2.3%	1,372	2.6%
Caucasian	36,402	81.7%	6,958	81.5%	25,986	81.4%	734	85.3%	2,724	85.0%	1,188	87.0%
Hispanic	75,416	7.4%	15,065	7.7%	56,123	7.7%	891	4.5%	3,338	4.5%	864	3.3%
Native American	64,617	0.2%	12,864	0.2%	47,243	0.2%	977	0.1%	3,533	0.1%	465	0.1%
18-29	127,284	41.4%	24,772	42.0%	95,414	43.3%	1,423	24.0%	5,675	25.7%	618	13.9%
30-49	33,112	28.6%	6,089	27.4%	22,803	27.4%	902	40.3%	3,319	39.8%	1,354	50.8%
50-64	23,287	20.0%	4,497	20.1%	16,193	19.4%	567	25.2%	2,030	24.2%	1,377	25.6%
65+	20,493	10.1%	4,062	10.4%	14,521	9.9%	411	10.5%	1,500	10.2%	1,634	9.7%

(c) Engagement received by tweets related to Black Lives Matter and follower counts

	All Engagement		Retweets		Favorites		Quote Retweets		Quote Favorites		Followers	
	tweet mean	% of total	tweet mean	% of total	tweet mean	% of total	tweet mean	% of total	tweet mean	% of total	user mean	% of total
Democrat	82,700	67.7%	18,565	67.3%	58,690	67.7%	1,268	69.4%	4,177	69.4%	1,593	70.1%
Independent	100,799	12.3%	23,427	12.6%	72,091	12.4%	1,233	10.0%	4,048	10.0%	1,053	8.6%
Republican	58,838	20.0%	13,401	20.1%	41,544	19.9%	909	20.6%	2,984	20.6%	1,134	21.3%
Female	81,119	58.5%	18,534	59.3%	58,058	59.0%	1,050	49.5%	3,476	49.8%	1,266	46.0%
Male	70,036	41.5%	15,489	40.7%	48,984	41.0%	1,302	50.5%	4,261	50.2%	1,592	54.0%
African-American	102,532	12.0%	23,598	12.2%	72,656	11.9%	1,455	11.3%	4,822	11.4%	1,309	7.8%
Asian	103,902	2.8%	24,133	2.9%	74,174	2.8%	1,311	2.3%	4,284	2.3%	1,673	2.7%
Caucasian	71,107	77.3%	15,909	76.5%	50,357	77.1%	1,128	81.5%	3,713	81.5%	1,439	86.2%
Hispanic	134,979	7.8%	32,184	8.2%	97,591	7.9%	1,230	4.7%	3,975	4.6%	945	3.2%
Native American	109,666	0.2%	25,111	0.2%	79,500	0.2%	1,169	0.1%	3,886	0.1%	529	0.1%
18-29	176,973	50.1%	42,672	53.5%	127,892	51.0%	1,504	28.4%	4,905	28.2%	661	13.8%
30-49	65,675	28.5%	13,977	26.8%	45,541	27.8%	1,441	41.7%	4,716	41.4%	1,684	52.0%
50-64	40,192	14.6%	8,341	13.4%	28,016	14.4%	887	21.5%	2,948	21.7%	1,755	25.2%
65+	33,449	6.8%	7,034	6.3%	23,688	6.8%	620	8.4%	2,107	8.7%	1,874	9.1%

### Engagement by Demographic, Topic & Type



**Figure 6. Difference in proportion of interactions received by each demographic.**

*Note.* Differences measured by comparing against the proportion of interactions received by that demographic among all other tweets. **A)** Differences in total interactions received for both COVID-19 and Black Lives Matter. **B)** Difference in demographic proportions for each type of interaction for COVID-19. **C)** Differences in demographic proportions for each type of interaction for Black Lives Matter.

We further explore how engagement varies by demographic across these topics. Table 2 outlines this variation in detail, examining differences in the average interactions received by different demographics, how that engagement breaks down across retweets, favorites, quote tweets, and quote favorites, as well as demographic differences in the number of followers accounts have. Table 2A shows these numbers across the full panel, while Tables B and C capture interactions received by tweets related to COVID-19 and Black Lives Matter respectively. All four interaction measures as well as their aggregate are measured at the tweet level, while follower counts are measured at the user level. Since these different demographic groups make up different shares of the population and are responsible for different shares of the tweet corpus, we report the mean number of interactions per tweet and the mean number of followers per user. While these full distributions span several orders

of magnitude, considering the mean provides per-unit context to these counts. Finally, the reported percentages are derived from each demographic's share of the total number of engagements or followers.

The topical and demographic differences in received engagement are visually represented in Figure 6. In Panel A, we see the difference in the aggregated count of interactions received by each demographic within a given topic. Here we see that tweets from Democrats relating to Black Lives Matter receive slightly more engagement than tweets from Democrats on other topics. This contrasts with Republican tweets about Black Lives Matter which generally receive fewer interactions than other Republican tweets. However, we previously saw, in Figure 4C, that Republican tweets on this topic are more likely to be replies, rather than other kinds of tweets. This decrease in attention, then, could be primarily topical—reflecting a disinterest in engaging with this demographic on this topic—or could be logistical—reflecting a broader trend of replies receiving fewer engagements. Additionally, we see that tweets from white people about COVID-19 receive more engagement than tweets from white people on other topics. The same holds true for panelists 50-64 as well as those over 65. Tweets related to COVID-19 from African-Americans and those 18-29 also receive fewer engagements than other tweets from these populations, though both groups also tweet less about COVID-19 in general.

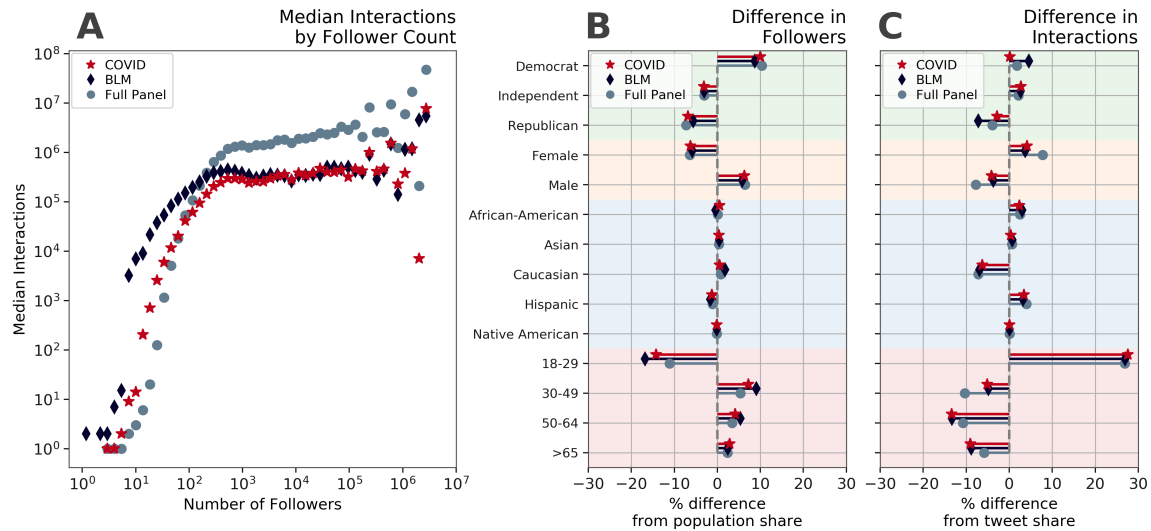
Figures 6B and C show how this received engagement varies across interaction types for tweets related to COVID-19 and Black Lives Matter. Here we see that the aggregate measure of interactions is generally reflective of the individual engagement measures. As was the case with the disaggregation of tweet type in Figure 4, this similarity is particularly true around the topic of COVID-19, where all engagement types generally share the same degree of over or under representation. In Panel C we start to see some difference in how tweets are engaged with, but only in terms of age. Tweets from younger panelists about Black Lives Matter are retweeted more often than tweets from young people on other topics, but these same tweets generally receive fewer quote tweets. On the other hand, panelists 50-64 as well those 65 and up are more likely to be quote tweeted, but less likely to be directly retweeted. This suggests that while retweets, quote tweets, and favorites may broadly reflect a shared type of engagement, subtle nuance between these behaviors may point to substantive differences in regards to different demographics and topics.

In particular, differences in retweets and quote tweets warrant further study as these differences may point to different beliefs about whose voice should be amplified (through a retweet) versus whose voice should be modified (through a quote tweet). This also highlights the importance of considering these differing modalities and what they may signal within the context of a particular topic. For example, young people's tweets about Black Lives Matter garnering more retweets than quote tweets could signal that other users are intentionally aiming to increase the visibility of this group without inserting their own voice on top of it.

Broadly, these findings suggest that the need to disaggregate interactions may vary significantly by research topic. Both retweets and quote tweet have the effect of amplifying content, so it would be appropriate to aggregate these measures for questions which are focused purely on amplification or attention. Favorites are by far the most common form of interaction and can be used alone if a single measure of "popularity" is desired. However, retweets and quote tweets do appear to be used differently, and there is interesting work to be done in further examining what content is most "quotable" and how different populations use these different mechanisms in different contexts.

### **Relationships Between Interactions and Followers**

Finally, we examine the intersection of our engagement measures along with how this engagement reflects user activity. While we might generally expect accounts with more followers to receive more interactions, we find this relationship to be only weakly true, with a correlation of only 0.015. Some of the most highly-followed accounts receive virtually no interactions, while some of the most modestly-followed accounts receive a large number of interactions. Figure 7A shows this relationship by topic, highlighting the median number of interactions received by tweets made from accounts with a given number of followers. While there is a weak relationship between these factors, there is a sharp inflection point between 1 and 100 followers, after which accounts receive roughly equal amounts of interaction regardless of their follower count. While this finding is confounded by the lack of temporal data relating to the order in which an account acquired followers or a tweet received interactions, it broadly points to a high level of content visibility on the platform. Even accounts with relatively few followers may "go viral" and ultimately garner a high number of interactions for some of their content.



**Figure 7. Relationships between follower counts and received engagements.**

*Note.* **A)** Median engagements received on tweets produced by accounts with a given number of followers. Tweet interaction types are aggregated into a single measure for the purpose of this analysis. **B)** Demographic differences in follower count, compared to each demographic’s population share. **C)** Demographic difference in proportion of total interactions received, compared to each demographics share of tweets.

In Figures 7B and C, we break this relationship down by demographic. Panel B shows the percent difference in followers compared to each demographic’s population share, while Panel C shows differences in interactions compared to that demographic’s tweet share. Comparing these suggests some striking differences between voice and attention afforded to different demographic groups. Across all topics, we see that Democrats in our panel typically have more followers and receive more interactions with their tweets. The reverse is true for Republicans, who have fewer followers than their population share and receive fewer interactions per tweet. The differences in gender are particularly striking—women typically have fewer followers than we might guess based on their population share, but generally receive more engagement with their content. Men, on the other hand, typically have more followers but receive less engagement on their content. Notably here, we have only examined the volume of this engagement, not its tenor or meaning. For example, increased interactions on women’s content could indicate greater interest in this content, a

desire among other Twitter users to elevate women's voices, or a tendency to use Twitter's quote tweet feature to explicitly rebut women's contributions. We see similar differences in followers and interactions by age. Panelists 18-29 generally have fewer followers but receive notably more interactions for their content, particularly around Black Lives Matter. On the other hand, panelists age 30-49 as well as 50-64 typically have more followers, but receive fewer interactions. Finally, while we see little difference in terms of followers by race, we do find that African-Americans and Hispanics receive slightly more interactions for their content related to Black Lives Matter.

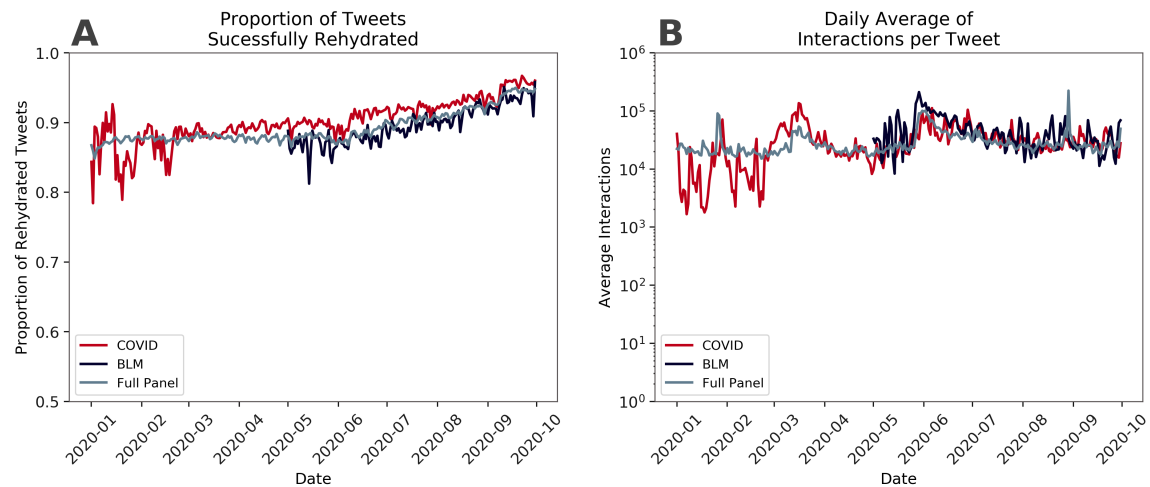
Again, we see that researchers must be cautious when interpreting Twitter data without its full context. A user who has popular tweets is not necessarily popular and receiving interactions is not necessarily positive. When conducting their analyses, researchers should be aware that while interaction with content and following a user are both highly-used forms of engagement, they may serve notably different roles. Furthermore, these findings point to a substantive need for further study around the role of differing types of engagement in which content and publics rise to prominence.

### *Temporal Bias*

Social media is constantly changing, raising important questions about the temporal nature and temporal validity of these data (Munger, 2019). We therefore close our analysis with an explicit acknowledgement of this challenge, examining some of the temporal biases at play in our data. We have already seen several of these temporal effects in our earlier analysis. Of the 1.6 million Twitter accounts which were identified as being active between 2014-2017, only 783,967 (47%) were active during the nine months of 2020 under study. Of those accounts, 18,967 (2.4%) were no longer available through the Twitter API as of October 2020. Of the 284,581,223 tweets we rehydrated in late 2020, 30,682,671 (10.8%) were no longer publicly available. As with any social media platform the population which comprises "the users" and the corpus which comprises "the content" are not static, who is active and participating is constantly changing.

Here, we examine the temporal implications of choosing to rehydrate, in late 2020, metadata for tweets which were posted between January 1 and September 30 of that year.





**Figure 8. Examination of temporal bias within our data.**

*Note.* **A)** Proportion of daily tweets which could not be retrieved by winter 2020. Note the the vertical axis begins at half (.5) of tweets retrieved. A value of 1 indicates that all tweets were retrieved. **B)** Daily average of engagements recorded for tweets retrieved in winter 2020.

We chose to rehydrate these data because we could not be certain that the engagement statistics recorded at the time of initial retrieval were an accurate representation of the level of engagement a tweet had received. Tweets, of course, are not posted with their maximum engagements achieved—these interactions come over time after the tweet has been posted. Our initial retrieval of panelists tweets is done through scheduled archiving of user content—meaning that there may be great variation in the number of hours, minutes, or even days that pass between a tweet’s posting and our initial retrieval. Choosing to recollect this metadata at a later date, however, risks two potential forms of temporal bias. First, we would expect some amount of data loss, as some tweets would no longer be available due to the passage of time. Second, there could be a temporal effect in which older tweets appear to receive more engagements simply because they have had more time in which to accrue those engagements.

Figure 8 illustrates the potential for both effects within our data. Panel A shows the proportion of daily tweets which we were unable to retrieve in late 2020. A value of 1 here indicates that every tweet posted on a given day was successfully retrieved as part of

our rehydration. While overall we were able to retrieve the vast majority of this content—89% of tweets across the full dataset, including 90% of COVID-19 tweets and 88% of Black Lives Matter tweets—we do see a small but notable temporal bias in this retrieval, with older tweets slightly less likely to be returned. By regressing the number of days passed on the proportion of tweets returned, we find that for the full dataset, each day further back in time corresponds to a 0.02% drop in the proportion of tweets received. This effect is significant at the  $< 0.001$  level. This data loss is small enough that we do not expect it to have a notable effect on the findings we have presented here. However, we encourage researchers to keep this data loss in mind and to examine its potential effects on substantive inference, if, for example, certain types of users or certain types of content are less likely to be retrieved.

In Figure 8B, we see the average number of interactions received by tweets posted on a given day. If a temporal bias favored older tweets, we would see a continual decrease in the number of engagements received, with older tweets generally receiving more interactions and newer tweets generally receiving less. However, we do not see evidence for such a temporal effect, finding instead that topical salience seems to drive engagement numbers. In our data, a longer window between tweet posting and retrieval does not favor older content, and indeed, newer content appears to have received more interactions.

Together these findings suggest, in general, researchers may reasonably retrieve the vast majority of tweets without worrying about data loss or favoring older content. However, the potential for data loss over time is likely to vary by the specific context under study, as we see, for example, that Black Lives Matter tweets are slightly less likely to be retrieved than COVID-19 tweets. Furthermore, researchers should be aware that the ability to retrieve these tweets at a later point in time relies heavily on consistent and uninterrupted API access, which social media platforms have the ability to change or restrict without warning (Freelon, 2018; Puschmann, 2019).

## Discussion

Scholars have argued that Twitter is best conceptualized as a collection of intersecting, networked publics (boyd, 2010; Jackson et al., 2020), but this full system cannot possibly

be studied in all its detailed complexity. This means that researchers must necessarily make choices about the collection and aggregation of Twitter content. In principle, these methodological choices should establish meaningful bounds for an imperfect but sufficiently coherent Twitter public. However, researchers have little empirical guidance in this task: there are notable gaps in our understanding of the empirical and conceptual implications of topical filtering or aggregation over users' activity, engagement, or demographic composition. Furthermore, discourse around the COVID-19 pandemic and the Black Lives Matter movement of 2020 is likely to be the subject of increased research focus for years to come, accentuating the need for descriptive statistics capable of providing contextual understanding to these publics and this time period. In taking a disaggregated, descriptive approach to capturing differences in demographics, activity, and engagement among U.S. Twitter users, this paper aims to provide researchers with insight on how to conceptualize and appropriately aggregate Twitter data.

We have found that while these salient topical events engage a substantial fraction of our panelists and generate a large portion of content, both events appear to take place amid a steady background of other discourse. This underscores the networked publics conceptualization of the platform, as users simultaneously inhabit a multiplicity of contexts across different discussions and publics. While there is no single dimension along which Twitter publics can be defined, we do find that topics provide meaningful bounds on populations of interest. However, topically selected tweets should not be assumed to represent a unified "discourse" and should instead be considered to capture segments of numerous, overlapping, and disconnected conversations.

We have also further found notable variation across our dataset, with distinctive and varied patterns in how panelists use different actions—such as retweets, quote tweets, and replies—across different demographics and topics. Specifically, we found notable differences in how these tweet types are used across demographic groups on the topic of Black Lives Matter. Republicans who engage with the topic are more likely to post replies rather than authored tweets, while Democrats are more likely to author tweets than to reply. This suggests that researchers should always conduct a disaggregated analysis of tweet activity, separately examining behavior around authored tweets, retweets, quote tweets, and replies. Variation in this behavior has the potential to reveal subpopulations of substantive

interest—such as those who are more likely to reply. While in some cases, the behaviors can be reasonably aggregated over, a disaggregated analysis serves as an important robustness check to examine any variation in behavior or to identify specific actions which are driving results.

We see similar variation in attention and engagement, suggesting that researchers should be particularly cautious not to conflate tweet interactions with author popularity and that retweets and quote tweets may serve distinctly different roles. Specifically, on the topic of Black Lives Matter, we see differences in the types of interactions tweets receive, but only along a demographic axis of age. Younger users are more likely to have their content on this topic retweeted, but are notably less likely to have their content quote tweeted. The opposite is true for older users. On the topic of COVID-19, we see demographic variation in the volume of tweets produced and interactions those tweets received, but do not see the same fine-grained variation in different modalities—demographic groups which produce more tweets on this topic produce more tweets of all types, while those which produce less, produce fewer tweets of all types. We further see interesting demographic variation which transcends topical categories and apply across the whole panel. Tweets posted by women and young people often receive more interactions, but these same populations tend to have fewer followers than we would expect given their population share.

Finally, we closed by examining some of the temporal challenges in studying a rapidly changing platform, characterizing the data loss and temporal biases found in our sample. We found only a small loss in content availability over time, though there is a notable temporal trend with older content less likely to be retrieved. Conversely, for data which can be retrieved, a prolonged time between posting and retrieving does not appear to produce biased estimates of engagement received by older content. Ability to retrieve data, however, may vary by context and is highly dependent on the platform offering continued API access (Freelon, 2018; Bruns, 2019).

Broadly, this work shows how critical it is for researchers to be mindful of context, as different conversations are happening simultaneously across permeable, dynamic, networked publics. While researchers have practical cause to be concerned about potential restrictions to API access (Freelon, 2018; Bruns, 2019), the technical specifics are just one of many

complications that social media researchers must consider (Tromble, 2021; Puschmann, 2019). Social media research requires conceptual grounding, ethical principles, and critical interpretation of the trace data we can see (Tromble, 2021). Researchers must necessarily make methodological choices to bound the populations and behaviors they study, but they should never forget that such bounds will always be arbitrary, imperfect, and researcher-imposed. Twitter publics are not self-contained entities which can be so cleanly extracted from their context. Any research question which aims to understand the dynamics of dialogue *within* Twitter must take this into account, considering the demographic and topical dimensions which may influence activity and engagement as well as the modalities of that engagement. Such considerations are essential to understanding how different publics act, interact, and react to each other within the platform.

In total, this paper provides valuable context for interpreting observational data generated by Twitter publics, and provides empirical insight into what populations can or should be considered as publics. By taking a disaggregated, descriptive approach, we have demonstrated how demographics, activity, and engagement are interconnected on such a platform, and how the interpretation of observational data is inexorably linked to an understanding of these publics. The importance of social media as a venue for public discourse will only continue to grow, and Twitter is likely to serve as the model organism for such research for years to come. It is therefore essential for scholars to publish and compare their descriptive analyses, and to develop frameworks for interpreting their observations of publics and behavior on those platforms.

### **Acknowledgments**

This work was partially supported by NSF grant #2026631, US Army Research Office grant W911NF-18-1-0421, and a Social Science Research Council Social Data Research Fellowship, with funds provided by Omidyar Network. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation, the US Army, or the Social Science Research Council, or the Omidyar Network.

## References

- Allen, J., Howland, B., Mobius, M., Rothschild, D., and Watts, D. J. (2020). Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances*, 6(14):eaay3539.
- Alshaabi, T., Adams, J. L., Arnold, M. V., Minot, J. R., Dewhurst, D. R., Reagan, A. J., Danforth, C. M., and Dodds, P. S. (2020). Storywrangler: A massive exploratorium for sociolinguistic, cultural, socioeconomic, and political timelines using Twitter.
- An, J. and Weber, I. (2016). #greysanatomy vs. #yankees: Demographics and hashtag use on Twitter. In *Proceedings of the Tenth International AAAI Conference on Web and Social Media (ICWSM)*, pages 523–526.
- Anderson, M., Barthel, M., Perrin, A., and Vogels, E. A. (2020). #BlackLivesMatter surges on Twitter after George Floyd’s death. *Pew Research Center*. Available online at <https://www.pewresearch.org/fact-tank/2020/06/10/blacklivesmatter-surges-on-twitter-after-george-floyds-death/>.
- Ansolabehere, S. and Hersh, E. (2017). Validation: What big data reveal about survey misreporting and the real electorate. *Political Analysis*, 20(4):437–459.
- Barberá, P., Casas, A., Nagler, J., Egan, P. J., Bonneau, R., Jost, J. T., and Tucker, J. A. (2019). Who leads? Who follows? Measuring issue attention and agenda setting by legislators and the mass public using social media data. *American Political Science Review*, 113(4):883–901.
- Bernstein, M. S., Bakshy, E., Burke, M., and Karrer, B. (2013). Quantifying the invisible audience in social networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '13*, page 21–30, New York, NY, USA. Association for Computing Machinery.
- boyd, d. (2010). Social network sites as networked publics: Affordances, dynamics, and implications. In *A networked self*, pages 47–66. Routledge.

- boyd, d., Golder, S., and Lotan, G. (2010). Tweet, tweet, retweet: Conversational aspects of retweeting on Twitter. In *2010 43rd Hawaii International Conference on System Sciences*, pages 1–10.
- Brock, A. (2012). From the blackhand side: Twitter as a cultural conversation. *Journal of Broadcasting & Electronic Media*, 56(4):529–549.
- Bruns, A. (2019). After the ‘APIcalypse’: Social media platforms and their fight against critical scholarly research. *Information, Communication & Society*, 22(11):1544–1566.
- Chadwick, A. (2017). *The hybrid media system: Politics and power*. Oxford University Press.
- Chen, E., Lerman, K., and 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):e19273.
- Culotta, A., Kumar, N. R., and Cutler, J. (2015). Predicting the demographics of Twitter users from website traffic data. In *AAAI*, volume 15, pages 72–8. Austin, TX.
- Farina, C., Epstein, D., B. Heidt, J., and J. Newhart, M. (2013). Regulation room: Getting “more, better” civic participation in complex government policymaking. *Transforming Government: People, Process and Policy*, 7(4):501–516.
- Ferrara, E., Varol, O., Davis, C., Menczer, F., and Flammini, A. (2016). The rise of social bots. *Communications of the ACM*, 59(7):96–104.
- Fiesler, C. and Proferes, N. (2018). “Participant” perceptions of Twitter research ethics. *Social Media + Society*, 4(1):2056305118763366.
- Fraser, N. (1990). Rethinking the public sphere: A contribution to the critique of actually existing democracy. *Social Text*, 25/26:56–80.
- Freelon, D. (2018). Computational research in the post-API age. *Political Communication*, 35(4):665–668.

- Freelon, D., McIlwain, C., and Clark, M. (2018). Quantifying the power and consequences of social media protest. *New Media & Society*, 20(3):990–1011.
- Freelon, D., McIlwain, C. D., and Clark, M. (2016). Beyond the hashtags: #Ferguson, #BlackLivesMatter, and the online struggle for offline justice. *Center for Media & Social Impact, American University*.
- Gallagher, R. J., Doroshenko, L., Shugars, S., Lazer, D., and Welles, B. F. (2020). Sustained online amplification of COVID-19 elites in the United States.
- Garimella, K., Weber, I., and De Choudhury, M. (2016). Quote RTs on Twitter: Usage of the new feature for political discourse. In *Proceedings of the 8th ACM Conference on Web Science, WebSci '16*, page 200–204, New York, NY, USA. Association for Computing Machinery.
- Goel, S., Anderson, A., Hofman, J., and Watts, D. J. (2016). The structural virality of online diffusion. *Management Science*, 62(1):180–196.
- Golder, S. A. and Macy, M. W. (2011). Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051):1878–1881.
- Gorwa, R. and 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., and Cranmer, S. J. (2020). Elusive consensus: Polarization in elite communication on the COVID-19 pandemic. *Science Advances*, 6(28):eabc2717.
- Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., and Lazer, D. (2019). Fake news on Twitter during the 2016 U.S. presidential election. *Science*, 363(6425):374–378.
- Grusin, R. (2010). *Premediation: Affect and mediality after 9/11*. Springer.
- Habermas, J. (1984). *The theory of communicative action*. Beacon Press, Boston.
- Habermas, J. (1991). *The structural transformation of the public sphere: An inquiry into a category of bourgeois society*. MIT Press, Cambridge, MA.



Hootsuite (2020). Digital 2020: The United States of America. <https://hootsuite.com/pages/digital-2020>.

Hu, M., Liu, S., Wei, F., Wu, Y., Stasko, J., and Ma, K.-L. (2012). Breaking news on Twitter. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2751–2754.

Hughes, A. G., McCabe, S. D., Hobbs, W. R., Remy, E., Shah, S., and Lazer, D. M. J. (2020). Using administrative records and survey data to construct samples of tweeters and tweets. *Public Opinion Quarterly*. Forthcoming.

Imai, K. and Khanna, K. (2016). Improving ecological inference by predicting individual ethnicity from voter registration records. *Political Analysis*, 24(2):263–272.

Jackson, S. J., Bailey, M., and Foucault Welles, B. (2020). *#HashtagActivism: Networks of race and gender justice*. MIT Press.

Jackson, S. J. and Foucault Welles, B. (2015). Hijacking #myNYPD: Social media dissent and networked counterpublics. *Journal of Communication*, 65(6):932–952.

Jaidka, K., Zhou, A., and Lelkes, Y. (2019). Brevity is the soul of Twitter: The constraint affordance and political discussion. *Journal of Communication*, 69(4):345–372.

Joseph, K., Swire-Thompson, B., Masuga, H., Baum, M. A., and Lazer, D. (2019). Polarized, together: Comparing partisan support for Trump’s tweets using survey and platform-based measures. *Proceedings of the International AAAI Conference on Web and Social Media*, 13(01):290–301.

Karpf, D. (2019). Something I no longer believe: Is internet time slowing down? *Social Media + Society*, 5(3):2056305119849492.

Karpf, D. (2020). Two provocations for the study of digital politics in time. *Journal of Information Technology & Politics*, 17(2):87–96.

Kavanaugh, A. L., Fox, E. A., Sheetz, S. D., Yang, S., Li, L. T., Shoemaker, D. J., Natsev, A., and Xie, L. (2012). Social media use by government: From the routine to the critical. *Government Information Quarterly*, 29(4):480–491.

- Kim, J. and Yoo, J. (2012). Role of sentiment in message propagation: Reply vs. retweet behavior in political communication. In *2012 International Conference on Social Informatics*, pages 131–136.
- Kwak, H., Lee, C., Park, H., and Moon, S. (2010). What is Twitter, a social network or a news media? In *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, page 591–600, New York, NY, USA. Association for Computing Machinery.
- Lazer, D. M., Pentland, A., Watts, D. J., Aral, S., Athey, S., Contractor, N., Freelon, D., Gonzalez-Bailon, S., King, G., Margetts, H., et al. (2020). Computational social science: Obstacles and opportunities. *Science*, 369(6507):1060–1062.
- Liu, Z. and Weber, I. (2014). Is Twitter a public sphere for online conflicts? A cross-ideological and cross-hierarchical look. In *Lecture Notes in Computer Science*, pages 336–347. Springer International Publishing.
- Lomborg, S. and Bechmann, A. (2014). Using APIs for data collection on social media. *The Information Society*, 30(4):256–265.
- Lukito, J., Suk, J., Zhang, Y., Doroshenko, L., Kim, S. J., Su, M.-H., Xia, Y., Freelon, D., and Wells, C. (2020). The wolves in sheep’s clothing: How Russia’s Internet Research Agency tweets appeared in U.S. news as vox populi. *The International Journal of Press/Politics*, 25(2):196–216.
- Marwick, A. and Lewis, R. (2017). Media manipulation and disinformation online. *Data & Society Research Institute*.
- McGregor, S. C. and Molyneux, L. (2020). Twitter’s influence on news judgment: An experiment among journalists. *Journalism*, 21(5):597–613.
- Mislove, A., Lehmann, S., Ahn, Y.-Y., Onnela, J.-P., and Rosenquist, J. N. (2011). Understanding the demographics of Twitter users. *ICWSM*, 11(5th):25.
- Munger, K. (2019). The limited value of non-replicable field experiments in contexts with low temporal validity. *Social Media + Society*, 5(3):2056305119859294.

Olteanu, A., Castillo, C., Diaz, F., and Kiciman, E. (2019). Social data: Biases, methodological pitfalls, and ethical boundaries. *Frontiers in Big Data*, 2:13.

Pew Research Center (2018). News use across social media platforms 2018. Report, Pew Research Center.

Pew Research Center (2019a). Share of U.S. adults using social media, including Facebook, is mostly unchanged since 2018. Report, Pew Research Center.

Pew Research Center (2019b). Sizing up Twitter users. Report, Pew Research Center.

Phillips, W. (2018). The oxygen of amplification: Better practices for reporting on extremists, antagonists, and manipulators online. *Data & Society Research Institute*.

Phillips, W. and Milner, R. M. (2020). *You are here: A field guide for navigating polarized speech, conspiracy theories, and our polluted media landscape*. MIT Press.

Pourebrahim, N., Sultana, S., Edwards, J., Gochanour, A., and 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.

Puschmann, C. (2019). An end to the wild west of social media research: A response to Axel Bruns. *Information, Communication & Society*, 22(11):1582–1589.

Shugars, S. and Beauchamp, N. (2019). Why keep arguing? Predicting engagement in political conversations online. *SAGE Open*, 9(1):2158244019828850.

Squires, C. R. (2002). Rethinking the black public sphere: An alternative vocabulary for multiple public spheres. *Communication Theory*, 12(4):446–468.

Starbird, K. and Palen, L. (2012). (How) will the revolution be retweeted? Information diffusion and the 2011 Egyptian uprising. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work, CSCW '12*, page 7–16, New York, NY, USA. Association for Computing Machinery.

Tromble, R. (2021). Where have all the data gone? A critical reflection on academic digital research in the post-API age. *Social Media + Society*, 7(1):2056305121988929.

Tufekci, Z. (2014). Big questions for social media big data: Representativeness, validity and other methodological pitfalls. *arXiv preprint arXiv:1403.7400*.

Twitter (2020). Q3 2020 letter to shareholders. <https://investor.twitterinc.com/>.

Wang, L. X., Ramachandran, A., and Chaintreau, A. (2016). Measuring click and share dynamics on social media: a reproducible and validated approach. In *Tenth International AAAI Conference on Web and Social Media*.

Warner, M. (2002). Publics and counterpublics. *Public Culture*, 14(1):49–90.

## Keywords

### *COVID-19 Keywords*

14dayquarantine	avoidcrowds	caronavirusoutbreak
2019_ncov	azar	caronavirususa
2019ncov	bailout	ccp virus
21dayslockdown	bailouts	ccpvirus
5baje5minute	bares cerrados	cdc
6 feet apart	bares fechados	cegahtangkalcorona
6 feet away	bars closed	center for disease control
6 feet between	bat soup	centers for disease control
9minute9baje	bayarealockdown	cerradmadritya
9minutesforindia	beatcoronavirus	chain of transmission
9बजे9मिनट	bersamamelawancorona	chegadequarentena
aerosol transmission	bersatulawancovid19	china flu
airborne transmission	birx	china virus
aislamientoobligatorio	bloody diwali	chinaflu
aksiberantascovid19	bloqueo de ruas	chinahidthevirus
alcool em gel	bloqueo de vias	chinavirus
alcool gel	bloqueo em vias	chinese flu
alcoolemgel	bloqueioderuas	chinese virus
alcoolgel	bloqueou as ruas	chineseflu
alertacovid19sv	bogotaencasa	chinesevirus
americans have died	bolsonarogenocida	chloroquine
ampliarlacuarentenaes	boomer remover	clapforourcarers
anti mask	boomerremove	clases anuladas
anti-lock down	breakcorona	clorox
anti-lockdown	briefing_covid19	closeynpublicschools
anti-mask	c19	cloth mask
antibodies	c.d.c.	cloth masks
antibody	cadeostestes	clubquarantine
antimask	californialockdown	codvid19
armchair epidemiologist	californiaquarantine	codvid_19
arrestbillgates	californiashutdown	community spread
arrestbillgates2020	calockdown	community spreading
arrestfauci	canceleverything	confinementtotal
asymptomatic	capitaocorona	conronaviruspandemic
aula online	carona virus	contact & trace
auxilio emergencial	caronavirus	contact + trace
auxilioemergencial	caronavirusindia	contact tracer

contact tracers	coronaviru	coronavirussouthafrica
contact tracing	coronavirus	coronavirussuisse
contact+trace	coronavirus india	coronavirusswitzerland
contact-and-trace	coronavirusargentina	coronavirustelangana
contact-trace	coronavirusbilluk	coronavirustruth
convid19	coronavirusbrasil	coronavirusu
coron virus	coronavirusbrazil	coronavirusuk
corona	coronaviruscanada	coronavirusup
corona vairus	coronaviruschile	coronavirusupdate
corona virus	coronaviruscolombia	coronavirusupdates
corona voucher	coronavirusde	coronavirusuruguay
coronaa virus	coronavirusdelhi	coronavirususa
coronaalert	coronavirusecuador	coronavoucher
coronaapocalypse	coronavirusencolombia	coronawarriors
coronacation	coronavirusesp	coronazeit
coronachainscare	coronavirusespana	coronga vairus
coronacitizenresponsibility	coronavirusfr	coronga virus
coronacrises	coronavirusfromminorchastisement	corongavirus
coronacrisis	coronavirusindia	coronials
coronaday	coronavirusindonesia	corono virus
coronadeutschland	coronavirusinindia	coronovirus
coronaferien	coronavirusinsa	coronvirus
coronafest	coronavirusireland	corrona virus
coronaflu	coronavirusitalia	corronavirus
coronafreepakistan	coronavirusitaly	corvid19virus
coronahoax	coronaviruserkerala	cotton mask
coronainpakistan	coronaviruslombardia	cotton masks
coronakindness	coronavirusmadrid	covd
coronakrise	coronavirusmaharashtra	covd19
coronalockdown	coronavirusmexico	covid
coronanasperiferias	coronavirusmumbai	covid 19
coronaoutbreak	coronavirusnobrasil	covid-19
coronapandemic	coronavirusnyc	covid-19uk
coronapandemie	coronavirusoubreak	covid-negative
coronapanik	coronavirusoutbreak	covid-positive
coronapocalypse	coronaviruspandemic	covid19
coronaschlager	coronavirusperu	covid19_ch
coronastopkarona	coronavirusplantao	covid19_de
coronaupdate	coronaviruspuertorico	covid19brasil
coronaupdatesindia	coronavirusrd	covid19canada
coronavid19	coronavirussa	covid19ch
coronavillains	coronavirusschweiz	covid19de

covid19deutschland	curona virus	ecoronavirusescuador
covid19ecuador	curonavirus	ecuadorencrisis
covid19espana	curonavirus	ecuadorenenemergencia
covid19france	cv 19	en primera linea
covid19india	cv19	epidemic
covid19indonesia	death rate	epitwitter
covid19insa	defense production act	escolas fechadas
covid19ireland	depoisdaquarentena	escolas fechando
covid19italia	dexamethasone	escolas sem aula
covid19out	dirumahaja	escolas sem aulas
covid19pt	disciplinaparavolver	essencial service
covid19schweiz	disinfect	essencial services
covid19southafrica	disinfectant	essential service
covid19uk	disinfecting	essential services
covid19usa	disiplincegahcorona	euficoemcasa
covid2019	distance learning	eunaquarentena
covid2019pt	diwali in april	event 201
covid_19	dont touch ur face	event201
covid_19ind	dont touch your face	evitar el contagio
covid_19uk	dontbeaspreader	f.d.a.
covid__19	donttouchyourface	face covering
covidactnow	dr immanuel	face coverings
covideos	dr micovits	face mask
covidiidiots	dr micovitz	face masks
covidiot	dr mikovits	face shield
covidiots	dr mikowitz	face shields
covidissnaf	dr. immanuel	facecovering
covidsafe	dr. micovits	facemask
covidy	dr. micovitz	facemasks
covid—19	dr. mikovits	fauci
crozier	dr. mikowitz	faucifraud
cuarentena	drive-through test	faucigate
cuarentenacoronavirus	drive-through testing	fda
cuarentenahastajunioes	drive-through tests	ficaemcasa
cuarentenainformando	drive-thru test	fightcoronawithjokowi
cuarentenametadata	drive-thru testing	fiqueemcasa
cuarentenaya	drive-thru tests	firefauci
cubaporlasalud	duringmy14dayquarantine	flatten the curve
cubasalvidas	earlyhydroxychloroquineworks	flattened the curve
cubriendoeelcoronavirus	earlyhydroxychloroquinezincworks	flattening the curve
curfewinindia	economic disaster	flatteningthecurve
curona vairus	economic stimulus	flattenthecurve

flexibilizarelaislamientos	homemade mask	kungflu
food and drug administration	homemade masks	laprevisiónestáentusmanos
foracoronavirus	homemademask	lava tu manos
free masks	homeschool	lavatumanos
free testing	homeschooling	lave as maos
frenalacurva	hometasking	laveasmaos
front line worker	howtokeeppeoplehome	lawancorona
front line workers	human trials	lawancoronabersama
frontline	hybrid learning	lawancovid19
frontline worker	hydroxychloroquine	layoffs
frontline workers	hydroxychloroquineworks	letsdefeatcovid19together
frontline_warriors_intern	immunity passport	linha de frente
frontlineheroes	immunocompromised	local transmission
frontlineworkers	immunodeficiencias	lock down
frontlineworkersappreciation	imunidade de rebanho	lock downs
furlough	in-person classes	lockdownhustle
furloughed	india for 21	lockdown
furloughs	india hum honge kamyab	lockdown21
gerakansocialdistancing	indiafightcorona	lockdowneffect
getmeppe	indiaprotectdoctors	lockdownend
gripezinha	indonesialawancorona	lockdownextended
had a mask	indonesialawancovid19	lockdownlife
hand out masks	infocoronavirus	lockdownnow
hand sanitisers	inmyquarantinesurvivalkit	lockdownparcial
hand sanitizer	isolamento parcial	lockdowns
handed out masks	isolamentoparcial	lockdowntillmay3
handing out masks	jagadirijagajak	losangeleslockdown
handsanitizer	janatacurfewmarch22	lysol
have a mask	jantacurfewchallenge	mandatory mask
have your mask	jobless claims	mandatory masks
hcq	judy anne mikovits	marchapelocorona
hcws	judy micovits	mascarasalva
health worker	judy micovitz	mascarilla
health workers	judy mikowitz	mascarillas
hepa filter	judy milkovits	mascarillas desechables
hepa filters	kamitidaktakutviruscorona	mascarillassolidarias
herd immunity	kanikacoronarow	mask mandate
holdthevirus	kanikakacoronacrime	mask requirement
home made mask	koronaindonesia	mask was worn
home made masks	koronavirus	mask wearing
home-made mask	koronavirusindonesia	maskachusetts
home-made masks	kung flu	maskingforafriend



maskmandate	maskupma	maskupsouthcarolina
masks required	maskupmaine	maskupsouthdakota
masks save lives	maskupmaryland	maskuptennessee
masks4all	maskupmassachusetts	maskuptexas
maskssavelives	maskupmd	maskuptn
maskup	maskupme	maskuptx
maskupak	maskupmi	maskuput
maskupal	maskupmichigan	maskuputah
maskupalabama	maskupminnesota	maskupva
maskupalaska	maskupmississippi	maskupvermont
maskuparizona	maskupmissouri	maskupvirginia
maskupaz	maskupmn	maskupvt
maskupca	maskupmo	maskupwa
maskupcalifornia	maskupmontana	maskupwashington
maskupco	maskupms	maskupwestvirginia
maskupcolorado	maskupmt	maskupwi
maskupconnecticut	maskupnc	maskupwisconsin
maskupct	maskupnd	maskupwv
maskupdc	maskupne	maskupwy
maskupde	maskupnebraska	maskupwyoming
maskupdelaware	maskupnevada	medical mask
maskupfl	maskupnewhampshire	medical masks
maskupflorida	maskupnewjersey	medical supplies
maskupga	maskupnewyork	medicalmask
maskupgeorgia	maskupnh	micasaesmiplaza
maskuphawaii	maskupnj	michiganshutdown
maskuphi	maskupnorthcarolina	miss rona
maskupia	maskupnorthdakota	missrona
maskupid	maskupnv	modikibaatmano
maskupidaho	maskupny	modly
maskupil	maskupoh	mortonaocompra
maskupillinois	maskupohio	mortonaovota
maskupin	maskupok	mp927
maskupindiana	maskupoklahoma	mpdafome
maskupiowa	maskupor	mpdamorte
maskupkansas	maskuporegon	mpdobolsonaro
maskupkentucky	maskuppa	mumbailockdown
maskupks	maskuppennsylvania	mycovid19summer
maskupky	maskuprhodeisland	n95
maskupla	maskupri	n95 mask
maskuplagos	maskupsc	n.i.h.
maskuplouisiana	maskupsd	national institutes of health

ncov	pcr tests	quarantineandchill
ncov2019	personal protective equipment	quarantined
neuercoronavirus	phase 1 trial	quarantinelifelife
neuercoronavirusschweiz	phase 1 trials	quarantinemoneymakingideas
new cases	phase 1 vaccine	quarantines
nih	phase 1 vaccines	quarantining
no mask	phase 2 trial	quaratinelifelife
noalaislamientointeligente	phase 2 trials	quarentena
notdying4wallstreet	phase 2 vaccine	quarentenabrasil
notessential	phase 2 vaccines	quarentenalgbtq
nouveau coronavirus	phase 3 trial	quarentenou
nouveaucoronavirus	phase 3 trials	quarentine
novel coronavirus	phase 3 vaccine	quarentined
novelcorona	phase 3 vaccines	quarentinelife
novelcoronavirus	physical distancing	quaretenabrasil
nuovocoronavirus	physicaldistancing	quarona virus
obrasilnaopodeparar	plandemic	quaronavirus
obrasilnaovaiparar	plandemic2020	quedateencasa
obrasilvaiparar	pmcares	r0
ohiocoronavirus	polymerase chain reaction	re open
oneteamfromhome	positive cases	re opening
onlineclasses	positivity rate	re-open
outbreak	post-rona	re-opening
oxygen level	ppe	realmenwearmasks
oxygen levels	ppeshortage	receitasdaquarentena
oxygen saturation	ppp	regn-cov2
pandemic	pralernaquarentena	regncov2
pandemie	presidential disaster declaration	relief bill
panic buy	presymptomatic	relief package
panic buying	primera línea de la salud	remdesivir
panic shop	prioridaddineroosalud	remote learning
panic shopping	protective equipment	remotework
panic-buy	provide masks	remoteworking
panic-shop	provided masks	remove a mask
panicbuy	providing masks	remove masks
panicbuying	purell	remove your mask
panickbuing	put a mask on	removed her mask
panicshop	put masks on	removed his mask
paremoselvirus	put your mask on	removed masks
passaportes de imunidade	pánicoporcoranovirus	reopen
paycheck protection program	quarantaine	reopening
pcr test	quarantine	reproduction number

require masks	sinophobia	super-spreaders
required masks	six feet apart	super-spreading
requiring masks	six feet away	superspreader
rescue bill	six feet between	superspreaders
rescue package	slow the spread	superspreading
respirator	slowthespread	surgical mask
respirators	social distance	surgical masks
respiratory droplet	social distancing	surgicalmask
respiratory droplets	social_distancing	suspend foreclosures
restaurantes cerrados	socialdistance	suspendanlasclases
restaurantes fechados	socialdistancing	suspendonlineclasses
restaurants closed	socialdistancingnow	swab
rising faster than tests	socialdistnacing	taxarfortunassalvarvidas
saferathome	sosecuador	teletrabajo
sars-cov-2	spcontraocoronavirus	test & trace
sarscov2	srkdonatesforcovid	test + trace
scauquermorrer	stay at home	test and trace
secazar	stay home	test+trace
self isolate	stay home challenge	test-and-trace
self isolated	stay safe stay home	testes em humanos
self isolating	stayathome	testesmasivosja
self isolation	stayathomechallenge	testesmassivosja
self quarantine	stayathomeorder	testing capacity
self quarantined	stayathomereadabook	thalibajao
self quarantining	stayathomesavelives	thankyouwarriors
selfemployedmatteredoo	stayhome	the coronas
selfisolating	stayhomechallenge	the rona
selfquarantine	stayhomesavelives	the virus
sendusbackhome	stayhomestaysafe	thecoronas
serological test	staying at home	this virus
serological tests	staying home	togetherathome
servicio esencial	staysafestayhome	toiletpaperapocalypse
servicios esenciales	staythefhome	toiletpaperpanic
serviços essenciais	staythefuckhome	trabajador sanitario
sflockdown	stella immanuel	trabajadora sanitaria
shamblesstayathome	stimulus bill	trabajadores a la calle
shankh	stimulus package	trabajadores de la salud
shelter in place	stimulus payment	trabajar desde casa
sheltering in place	stimulus payments	trabajar desde casa
shelteringinplace	super spreader	trabalhador essencial
sideeffectsofquarantinelif	super spreaders	trabalhadores essenciais
simecontagioyo	super-spreader	trabalhando de casa

trabalhar de casa	wearmask	कोरोनावायरस
transmission event	wearing a mask	कोविड-१९
trump pandemic	wearing masks	कोविड_19
trumpdemc	wearyourmask	मोदीजी_हम_दीप_नहीं_जलाएंगे
trumppandemic	wfh	वूहान
ukcoronavirusbill	whencoronavirusisover	うちで過ごそう
uklockdown	whopaysforcovid	おうち時間
unexigimosgarantiaslaborales	whyiwearmask	ウイルス付着
unidosvenceremos	wore a mask	コロナ
unitedagainstcoronavirus	wore her mask	コロナに負けるな
unsuspendanclasesoparamos	wore his mask	コロナウィルス
unsuspendanclasesya	wore your mask	コロナウィルス
unsuspendanlasclasesoparamos	work from home	コロリンピック
usemascara	workfromhome	ステイホーム
uss theodore roosevelt	working from home	ソーシャルディスタンス
vaccine distribution	workingfromhome	ロックダウン
vaipassar	workingfromhomelife	加油武汉
vegasshutdown	workingfromhometips	外出禁止
ventilated	world health organization	外出自粛
ventilation	wuhan	定額給付金
ventilator	wuhan flu	新しい生活様式
ventilators	wuhan virus	新冠病毒
virus corona	wuhancoronavirus	新冠肺炎
viruscorona	wuhanflu	新型コロナウイルス
viviremosyvenceremos	wuhanlockdown	新型冠状病毒
vivirencuarentenaes	wuhanvirus	東京コロリンピック
w.h.o.	yomequedoencasa	武汉加油
waragainstvirus	yomequedoencasa	武汉疫情
wash ur hands	فيروس كورونا	武汉肺炎
wash your hands	فيروس_كورونا	武漢肺炎
washurhands	كورونا	特別定額給付金
washyourhands	كورونا_الجديد	疫情
washyourhandsagain	अंधेर_नगरी_चौपट_राजा	緊急事態宣言
wear a mask	कोरोना	코로나
wear masks	कोरोना वायरस	코로나19
wear your mask	कोरोना_वायरस	코로나바이러스

*Black Lives Matter Keywords*

32 bullets	antiracism	capitol hill autonomous zone
32 shots	antiracist	capitol hill occupied protest
8 can't wait	armed militia	capitol hill organized protest
8 cant wait	armed militias	chauvin
8 min 46 secs	arrest the cops	chicago police
8 minutes 46 seconds	arson	chokehold
8 to abolition	atlanta police	chokeholds
8:46	b.l.m.	commitment march
8cantwait	bail fund	confederate statue
8min 46secs	bailfund	confederate statues
8toabolition	baton	copaganda
a.c.a.b.	batons	counter-protest
aaron torgalski	black lives	counter-protesters
abolish prisons	black out tuesday	counterprotest
abolish the cops	black square	counterprotesters
abolish the police	black squares	crowd control
abolishing cops	black surfing association	deadly force
abolishing police	black trans lives matter	defund lmpd
abolishing prisons	black_lives_matter	defund the police
abolishprisons	blacklivesmatter	defunding police
abolishthepolice	blackout tuesday	defundlapd
abolition	blackouttuesday	defundmlmpd
abolitionism	blacktranslivesmatter	defundnypd
abolitionist	blm	defundthepolice
acab	blue lives matter	demonstrations
alexander kueng	bluelivesmatter	derekchauvin
alexanderkueng	body cam	die-in
all cops are bad	body camera	disband the police
all cops are bastards	body cameras	disbandthepolice
all lives matter	body cams	disorderly conduct
all lives splatter	boogaloo	divest from the police
alllivesmatter	boston police	divest the police
alllivesplatter	breonna taylor	eight minutes 46 seconds
alton sterling	breonna's law	eight minutes forty-six seconds
altonsterling	breonnataylor	eight minutes fortysix seconds
anti racism	breonnataylorprotest	eric garner
anti racist	brett hankison	ericgarner
anti-racism	buffalo police	excessive force
anti-racist	campaign zero	federal monuments
antifa	campaignzero	federal statues

ferguson	martin gugino	pepper ball
fired a projectile	michael brown	pepper balls
firing projectiles	michaelbrown	pepper spray
floyd	mikebrown	pepper sprayed
freddie gray	militarization of police	philando castile
freddiegray	militarization of the police	philandocastile
free capitol hill	militia group	police accountability
freedom to assembly	militia groups	police benevolent association
ftp	minneaapolis police department	police brutality
fuck the police	minneapolis city council	police camera
georgefloyd	minneapolis police	police cameras
georgefloydprotest	minnesota freedom fund	police car
get your knee off our necks	mpd	police cars
handsupdontshoot	mutual aid	police department
i can't breathe	myles cosgrove	police departments
i cant breathe	neo confederates	police killing
icantbreathe	neo-confederates	police killings
impact munitions	new jim crow	police kneeling
in solidarity	new york police	police lives matter
insurrection act	nfac	police militarization
jacob blake	no justice no peace	police misconduct
jacobblake	no justice, no peace	police murder
jalexanderkueng	no knock search warrant	police murders
john losi	no knock warrant	police projectile
jonathan mattingly	no-knock search warrant	police projectiles
juneteenth	no-knock warrant	police reform
justice in policing act	nojusticenopeace	police reforms
justiceforbreonnataylor	noracistpolice	police vehicle
justiceforfloyd	not fucking around coalition	police vehicles
justiceforgeorgefloyd	nypd	police violence
justiceforjacobblake	oath keepers	policelivesmatter
kenneth walker	occupy city hall	portland police
kenosha	occupycityhall	projectile fired
lafayette square	officer cosgrove	property damage
lapd	officer hankison	protect federal
lmpd	officer mattingly	protect monuments
looters	officer thao	protect statues
looting	operation diligent valor	protecting america communities
los angeles police	operation legend	task force
louisville metro police	outside agitator	protecting federal
louisville police	outside agitators	protecting monuments
mappingpoliceviolence	paddle out	protecting statues

protest	sandrabland	threw rocks
protesters	sayhername	threw water bottles
protests	sayhisname	throw water bottles
protests2020	seattle police	throwing rocks
proud boys	shutitdown	throwing water bottles
qualified immunity	silent rally	thrown water bottle
racial bias	smoke canister	thrown water bottles
racial injustice	smoke canisters	tou thao
racial justice	so you want to talk about race	touthao
racism	st johns church	unarmed
racist	st johns episcopal church	unlawful assembly
racist police	st paul police	unrest
racists	st. john's church	uprising
reparation	st. john's episcopal church	uprisings
reparations	st. paul police	vigil
riot	strike for black lives	vigilante group
riot control	strikeforblacklives	vigilante groups
rioters	structural racism	vigilante militia
rioting	take a knee	vigilante militias
riots	take the knee	vigils
riots2020	takeaknee	walkwithus
ripbreonnataylor	taketheknee	when the looting starts the
ripgeorgefloyd	taking a knee	shooting starts
ripjacoblake	tamir rice	when the looting starts, the
robert mccabe	tamirrice	shooting starts
rubber bullet	tear gas	white fragility
rubber bullets	tear gassed	white out wednesday
run them over	theshowmustbepaused	white supremacist
runthemover	thomas k lane	white supremacists
safe policing for safe communi-	thomas lane	white supremacy
ties	thomasklane	whiteout wednesday
saint john's church	thomaslane	whiteoutwednesday
saint johns church	three percenters	yougoodsis
saint paul police	threw a rock	
sandra bland	threw a watter bottle	

## Validating Proprietary Measures

Our analyses draw on two variables that are provided by our data vendor and are effectively proprietary: race and partisanship. In the case of the former, TargetSmart reports race drawn from the voter file where available, and from other linked commercial data sources otherwise. To have a measure of partisanship that is consistent across states, we also use their measure of partisanship; in their provided documentation, this is the output of an “ensemble method classifier model was created to predict the likelihood that an individual supports the Democratic Party” when responding to a survey.

### *Race*

Without ground-truth data, it can be difficult to assess the accuracy of the vendor’s measure of race. In an appendix of Hughes et al. (2020), we investigated 182 members of our sample also present in Pew’s surveys of Twitter users. Of these, 147 individuals self-reported race that agreed with TargetSmart. This is 80% agreement; if you disregard the ten cases where TargetSmart records race as unknown or missing, it is 85% agreement. In the much smaller intersection of individuals present in both samples who are also in states historically under VRA preclearance, 21 of 22 self-reports (95%) agree with TargetSmart.

In the absence of vendor-provided data on race, we likely would draw on existing work on inferring race. This provides another natural benchmark for assessing the utility of TargetSmart’s data for our analyses. We compare the TargetSmart data with these imputations created using the `wru` library provided by Imai and Khanna (2016).<sup>3</sup> The library returns probabilities that an individual was a member of each racial group; while propagating this uncertainty would be useful in other analyses, to facilitate comparison with TargetSmart

---

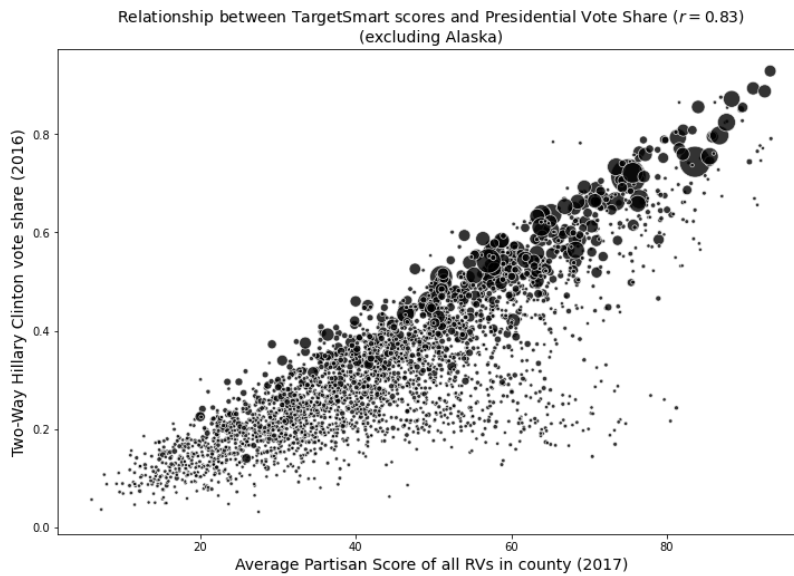
<sup>3</sup>The specific set of arguments used was `predict_race(tsmart_data, census.geo="tract", census.data=census_data, sex=TRUE, age=TRUE)`. The census data was collected using `wru`’s `get_census_data()` function.



**Table 3: State-level patterns of missingness and agreement with imputations for TargetSmart race variable.**

State	NA (TS)	NA (WRU)	Agree	State	NA (TS)	NA (WRU)	Agree
AK	5%	27%	92%	MT	2%	9%	94%
AL	0%	12%	84%	NC	0%	11%	81%
AR	2%	13%	94%	ND	1%	36%	97%
AZ	3%	11%	93%	NE	2%	7%	97%
CA	4%	14%	93%	NH	1%	19%	96%
CO	3%	9%	95%	NJ	3%	10%	94%
CT	3%	12%	95%	NM	2%	10%	89%
DC	6%	27%	86%	NV	4%	13%	93%
DE	3%	9%	91%	NY	4%	11%	94%
FL	0%	12%	81%	OH	2%	10%	95%
GA	1%	10%	79%	OK	3%	13%	93%
HI	9%	36%	79%	OR	3%	7%	94%
IA	2%	6%	96%	PA	2%	11%	95%
ID	3%	14%	96%	RI	3%	8%	95%
IL	3%	8%	95%	SC	0%	10%	82%
IN	2%	9%	96%	SD	1%	8%	94%
KS	2%	8%	96%	TN	1%	13%	92%
KY	1%	8%	95%	TX	3%	12%	93%
LA	0%	11%	82%	UT	3%	11%	95%
MA	3%	9%	95%	VA	3%	10%	91%
MD	3%	9%	92%	VT	2%	12%	96%
ME	1%	8%	95%	WA	4%	9%	94%
MI	2%	8%	95%	WI	2%	23%	97%
MN	2%	8%	96%	WV	1%	9%	94%
MO	2%	9%	96%	WY	3%	28%	95%
MS	1%	24%	85%				

we assign each user an imputed race corresponding to the highest assigned probability. Because there is some missingness in each of the forms of input data (location, age, etc.) to the imputation algorithm, there are a relatively large number of failed imputations. Nevertheless, Table 3 show that, where `wru` and TargetSmart both provide information on race, agreement is generally high.



**Figure 9. Comparison between county-level mean partisanship score and Clinton vote share.**

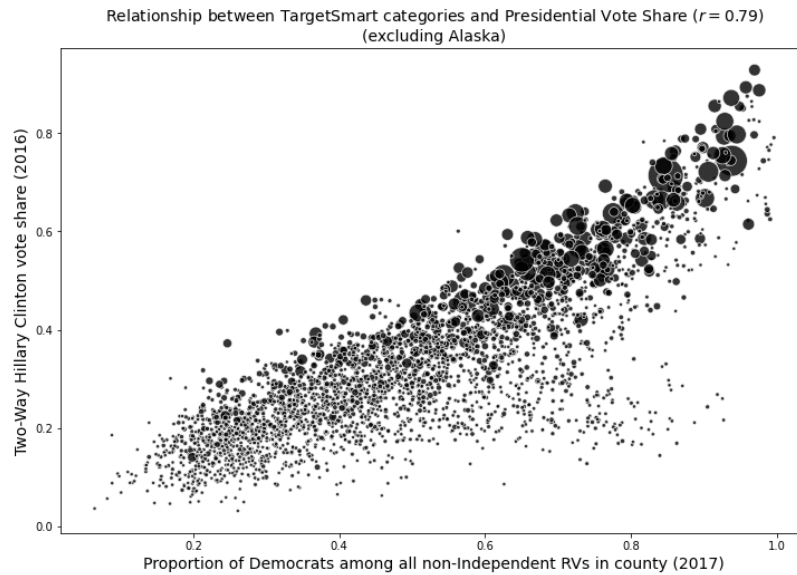
*Note.* Counties are sized by population. Alaska is excluded because it reports results at a different spatial unit (legislative district).

In general, we expect TargetSmart’s data on race to be of the highest quality in states that were historically under preclearance through the Voting Rights Act. To evaluate this, we compare states that were entirely under VRA preclearance, plus North Carolina (where most counties were under preclearance), to all other states.<sup>4</sup> The missingness rate in preclearance

<sup>4</sup>We include states that were under state-level preclearance after 1975, so the full list is: Alabama, Georgia, Louisiana, Mississippi, South Carolina, Virginia, Alaska, Arizona, Texas, plus North Carolina.

states is lower for TargetSmart (2% vs. 8%) and wru (12% vs. 16%), while the agreement between recorded and imputed race is lower (88% vs. 93%).

### *Partisanship*



**Figure 10. Comparison between county-level proportion of TargetSmart-identified Democrats and Clinton vote share.**

*Note.* Counties are sized by population. Alaska is excluded because it reports results at a different spatial unit (legislative district). The proportion is calculated as the number of Democrats relative to the number of Democrats and Republicans (i.e., Independents are excluded).

As mentioned above, the proprietary measure of partisanship we use is a continuous variable representing the probability that an individual self-identifies as a Democrat. In the analyses we present, we have trichotomized this continuous measure using cut points recommended by TargetSmart (0–0.35: Republican, 0.35–0.65: Independent, 0.65–1: Democrat). While we cannot assess TargetSmart’s accuracy as a predictor of survey response directly, here we present various proxies that suggest this is a reliable measure of partisan attitudes.

First, we note that, in aggregate, the continuous form of the partisanship scores match well

to election results. In Figure 9, we aggregate all registered voters in our voter file to the county level; the mean partisanship score correlates well ( $r = 0.83$ ) with Hillary Clinton vote share in 2016. (We use 2016 election results rather than 2020 because the partisan scores are from 2017.) The raw correlation arguably understates the level of agreement between the two measures, because it is not weighted by population. The counties that lie off the main diagonal are primarily Appalachian and rural Midwestern counties that were notable areas of Clinton under-performance in 2016.

To ensure that the conversion of the continuous measure into categories does not significantly affect our understanding of partisanship, in Figure 10 we present a similar aggregate comparison to election results for the categorical party labels; the results are largely unchanged.