

The Private Life of QAnon: A Mixed Methods Investigation of Americans' Exposure to QAnon Content on the Web

RYAN C. MOORE*, Department of Communication, Stanford University, USA

ROSS DAHLKE*, Department of Communication, Stanford University, USA

PETER L. FORBERG, Department of Sociology, University of California, Berkeley, USA

JEFFREY T. HANCOCK, Department of Communication, Stanford University, USA

The QAnon movement has been credited with spreading disinformation and fueling online radicalization in the United States and around the globe. While some research has documented publicly-visible communications and engagements with the QAnon movement, little work has examined individuals' *actual* exposure to QAnon content. In this paper, we investigate the extent to which Americans are exposed to QAnon websites, in what contexts, and to what effect. We employ a mixed methods review of 21 million website visits collected from a nationally representative sample of 1,238 American adults across laptops, smartphones, and tablets during the 2020 U.S. presidential election. Quantitative techniques reveal overall levels of exposure to QAnon and who is more likely to be exposed, and qualitative techniques provide rich information about how participants came to be exposed to QAnon and how it fit within their broader media diets. We find that: (1) exposure to QAnon websites is limited and stratified by political ideology and news consumption; (2) exposure tends to occur within right-wing media ecosystems that align with QAnon beliefs; and (3) mixed methods approaches to analyzing digital trace data can provide rich insights that contextualize quantitative techniques. We discuss the implications of our findings for the design of interventions to lessen exposure to problematic material online and future research on the spread of disinformation and extremist content.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)** → **Empirical studies in HCI**

Additional Key Words and Phrases: QAnon, digital trace data, web browsing, mixed methods, misinformation, disinformation, 2020 US election

ACM Reference format:

Ryan C. Moore, Ross Dahlke, Peter L. Forberg, and Jeffrey T. Hancock. 2024. The Private Life of QAnon: A Mixed Methods Investigation of Americans' Exposure to QAnon Content on the Web. *Proc. ACM Hum.-Comput. Interact.*, 8, CSCW2, Article 518 (November 2024), 33 pages, <https://doi.org/10.1145/3687057>

*Both authors contributed equally to this research.

Authors' addresses: [Ryan C. Moore](mailto:rymoore@stanford.edu), rymoore@stanford.edu, Department of Communication, Stanford University, 450 Jane Stanford Way, Building 120, Stanford, California, USA, 94305; [Ross Dahlke](mailto:rdahlke@stanford.edu), rdahlke@stanford.edu, Department of Communication, Stanford University, 450 Jane Stanford Way, Building 120, Stanford, California, USA, 94305; [Peter L. Forberg](mailto:peterforberg@berkeley.edu), peterforberg@berkeley.edu, Department of Sociology, University of California, Berkeley, 410 Social Sciences Building, Berkeley, California 94720; [Jeffrey T. Hancock](mailto:hancockj@stanford.edu), Department of Communication, Stanford University, 450 Jane Stanford Way, Building 120, Stanford, California, USA, 94305, hancockj@stanford.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

2573-0142/2024/11 - 518

© Copyright is held by the owner/author(s). Publication rights licensed to ACM.

<https://doi.org/10.1145/3687057>

1 INTRODUCTION

There is considerable concern among scholars, politicians, journalists, and the public about the rise of QAnon, an American conspiracy theory and political movement that originated on the online forum 4chan in 2017. While misinformation broadly refers to any content that is false, misleading, or unsubstantiated [70] and has been studied extensively in CSCW alone [2–4, 10, 11, 22, 37, 50, 53, 57, 59, 68, 91], QAnon is a much more specific phenomena. Followers of the conspiracy theory allege that a cabal of Satanists and pedophiles form the “deep state” — a network of politicians, celebrities, and the ultra-wealthy who actually control the U.S. government [104]. More specifically, the theory claims that a group of military operatives codenamed “Q” have been working with former president Donald Trump to overthrow this deep state. Q anonymously posts coded messages (called “Q drops”) to fringe internet forums that their followers interpret to learn more about Trump’s plans [61]. QAnon’s beliefs and online activities have changed following Trump’s loss in the 2020 presidential election and a prolonged disappearance of Q. Still, the conspiracy theory continues to evolve amidst shifting political events, remaining a source of misinformation and political organization [34, 103]. The QAnon movement is particularly engrossing to its members because the QAnon movement is about more than passive consumption, it is distinguished by recruitment into or participation in an online community, and thus understanding QAnon content may require sensitivity to the theory’s shared language and ideas [35, 88].

A key reason for the continued worry about QAnon is the belief that QAnon-related content is pervasive [92]. The concern is that if QAnon content is prevalent in the information environment and reaches many individuals, there is a better chance that it will influence enough people’s beliefs and preferences to undermine the outcomes of democracy. Indeed, recent scholarship has demonstrated that belief in conspiracy theories is associated with participation in far right-wing protests, which raises acute concerns in the U.S. following the storming of the Capitol on January 6, 2021 [17]. One poll suggests that as many as 30 million Americans held beliefs central to the QAnon theory as of March 2021 [77]. Concerned about exposing people to QAnon content, social media platforms have taken several actions to specifically curb QAnon’s proliferation on their platforms [6, 16, 78].

Despite these concerns, little empirical evidence has examined actual exposure to QAnon-related content. QAnon is thought to be popular only among political extremists, but media accounts also suggest that QAnon’s popularity is growing and attracting support among ordinary Americans [92]. In the face of this speculation, several key questions about the consumption of QAnon content remain unanswered: How much are Americans exposed to QAnon websites? What types of people are more likely to be exposed? How do those exposed arrive at QAnon websites online? Answers to these questions can help improve the design of interventions and policies to mitigate exposure to conspiracy theories.

To address these questions, we draw on a dataset of over 21 million website visits made across computers, tablets, and smartphones by a nationally representative sample of American adults ($N = 1,238$) during the 2020 presidential election (August – December 2020), a period of time when QAnon supporters were actively trying to reach broader audiences [34, 51]. We use a database of QAnon-related websites compiled by prior research [49] to identify website visits in our web browsing data, allowing us to quantitatively document Americans’ exposure to QAnon websites. We complement this quantitative analysis with a qualitative investigation into the web browsing behaviors of those exposed to QAnon websites to better understand their instances of exposure and how QAnon fits into their online information diets more broadly.

Our study makes several empirical, methodological and design contributions. First, empirically we find that 3.7% (95% CI = 2.1-5.4%) of American adults were exposed to QAnon websites during the 2020 election. Individuals on the political extremes were more likely to encounter QAnon websites, with most visits being incidental and aligning ideologically with their media consumption. A diverse set of channels, including social networks, email, search engines, and forums, led participants to these QAnon sites, either incidentally or through intentional searches.

Second, our work highlights the utility of web browsing data to study people's exposure to extremist content online. Compared with past work on QAnon, which largely relies on data about people's publicly-visible behaviors on specific platforms (e.g., posts and comments on Reddit) [5, 9, 31, 55, 71, 72, 87, 93, 99, 103], web browsing data cover exposures regardless of whether content is engaged with and tracks users' journeys as people move from platform to platform [63, 96, 97]. We also demonstrate the value of analyzing digital trace data using a mixed-methods approach. Insights from our mixed-methods analysis added important nuance to theories of exposure to extremist content, ultimately focusing future work on online radicalization into QAnon and other movements. Our work supports calls for more mixed-methods research of digital trace data [38, 95, 97].

Lastly, this research has several implications for intervention and policy design. For example, our data make clear that in addition to social networks (e.g., Facebook, Twitter, YouTube [6, 69, 71]), everyday online spaces like email inboxes and search engines play a significant role in exposing users to extremist content, offering new intervention opportunities. Our data also suggest that the common strategy of deplatforming [31, 52, 81] is insufficient for reducing exposure as people shift to using alternative platforms. Interventions to reduce exposures to extremist content need to consider linkages to hubs central to people's online navigation that can serve as vectors to alternative platforms, such as search engines.

2 BACKGROUND AND RELATED WORK

2.1 Exposure to QAnon Content Online

Despite much speculation about the extent to which people encounter QAnon content while using the internet, there has been limited empirical research documenting the presence of QAnon content online and none to our knowledge about who encounters that content. Existing research could be characterized as examining the "public life" of QAnon, or where QAnon content appears across the internet. Scholars have used computational and/or content-analytic techniques to identify and analyze QAnon-related content on the platforms where Q posts Q drops (4chan/8kun [71]), as well as on other platforms, such as Twitter [9, 93, 99], Facebook [55], Reddit [31], Voat [72], Parler [5], Gab [86], and Telegram [103]. This work examines patterns in QAnon-related discourse, revealing that QAnon discourses relate to Q drops and current events [9, 72, 86] and that QAnon discourses differ in explicit support for QAnon [93, 99]. Other studies have examined how elections [103] and platform moderation efforts [31, 93] affect QAnon discourse online.

While this prior work examined how QAnon-related content appears online and has provided a foundation of scholarly knowledge, it speaks little to who consumed Qanon content. Kim & Kim [55] compare engagement with QAnon posts on Facebook with other posts but do not have information about the users who engaged with the posts or those who saw posts but did not like, share, or comment on them. Hanley et al. [49] report data on traffic to several QAnon websites as estimated by Amazon Alexa, but these estimates do not reveal any details about who visited the sites. In this way, while some existing research has documented the "public life" of QAnon,

we argue that existing work has not examined the “private life” of QAnon and how QAnon content fits into people’s media diets, many of whom see and may be influenced by the content but choose to leave no public traces of engagement or interaction (often referred to as “lurkers” in the context of social media platforms [65]).

Unlike other online content types for which there is literature documenting exposure, such as misinformation [45, 64] and hard news [33, 98], there has not yet been a systematic investigation of exposure to QAnon content. While work on misinformation and hard news has illustrated how often people are exposed to those types of content, who is more likely to be exposed, and how people come to be exposed, these all remain open questions in the context of QAnon content. As such, we pose the following research question:

RQ1: How were Americans exposed to QAnon websites online?

Recent but limited work examining support for QAnon may indicate who is more likely to be exposed to QAnon content online [30, 77, 90]. Specifically, Enders et al. [30] find that those with more extreme political orientations — both right and left — are most likely to support QAnon. Indeed, this pattern follows a relationship documented between political extremism and belief in conspiracy theories in general [76]. Although we treat questions about who is exposed to QAnon content as largely open, we hypothesize that political extremists should be more likely to be exposed to QAnon websites because they are more likely to support it. Specifically, we posit the following hypothesis:

H1: People with extreme political orientations are more likely to be exposed to QAnon websites than those without extreme political orientations.

2.2 Types of QAnon Content Consumers

Commonplace in media and scholarly depictions of the QAnon phenomenon is the idea that there are different ways that people come to be exposed to QAnon content online. While some talk of devout QAnon followers waiting for each new Q drop to be posted, others describe more incidental means by which non-followers could encounter QAnon-related content. With the prominence of algorithmic curation and QAnon disseminators’ strategic use of content features like hashtags to broaden the reach of their content, people unassumingly scrolling through their TikTok “For You Page” or perusing their YouTube recommendations may encounter QAnon content [34]. In other cases, people engaging in political spaces online (e.g., political subreddits, Facebook groups) may incidentally encounter the topic of QAnon, which could spur them to conduct additional research on Q.

The prior empirical work supports the idea that QAnon content exists on many platforms across the internet [49, 86]. Similar to other disinformation campaigns, QAnon supporters actively strive to post content across multiple platforms [55, 97]. The more QAnon content appears on different platforms, the more chances there are for different users to encounter it. While everyone’s online behaviors are undoubtedly idiosyncratic, there may be broader categories to describe groups of people who consume QAnon content in distinct ways. Indeed, prior work has identified distinct types of online supporters of QAnon that can be systematically identified via differences in their content posting activities [93, 99] (e.g., self-declared QAnon supporters, hyper-active conspiracy promoters [93]). Thus, it seems plausible that there are distinct systematic types of users who will encounter QAnon content (e.g., those who visit QAnon sites intentionally and regularly, those who visit QAnon sites incidentally and rarely). However,

what specifically characterizes different types of QAnon consumers and their consumption patterns remains an open question. Thus, we pose the following research question:

RQ2: Are there distinguishable types of QAnon website consumers?

3 DATA & MEASURES

3.1 Participants

The survey company YouGov passively collected web-browsing data (i.e., all URLs visited) from participants' (N = 1,238) computers, smartphones, and tablets from August 24, 2020, to December 7, 2020. YouGov collects browsing data via software (browser plugin on computers, app on mobile devices) called YouGov Pulse, which YouGov's survey panelists can install on their desktop/laptop computers, smartphones, and/or tablets to provide a record of websites they visit to YouGov in exchange for compensation. YouGov then provides these data to researchers for a fee (for examples of other research using YouGov Pulse web browsing data to explore the consumption of political content online, see [14, 15, 43–45, 64, 66], and for a discussion of the ethics surrounding these data, see section 5.5 of this paper). All participants provided informed consent to the terms of our specific research (this study was approved by the Stanford University Institutional Review Board [protocol no. IRB-53941]) and were compensated by YouGov for their participation. In total, we collected approximately 21 million URL visits from our participants, who also completed a survey via which we collected demographic information.

Of our participants, 29% (n = 358) were aged 65+, 48% (n = 599) were 45–64, 14% (n = 175) were 30–44, and 9% (n = 106) were under 30. Five-hundred fifty reported identifying as male and 683 as female. A majority of the participants, 80% (n = 989), reported their race as white, 9% reported being Black (n = 113), 5% Hispanic (n = 65), and 6% as other (n = 71). Furthermore, 89% (n = 1,103) said they follow politics most or some of the time, and 30% (n = 373) were considered highly knowledgeable about politics according to Pew Research Center's civic knowledge questionnaire [75]. In addition, 36% (n = 445) supported Donald Trump in the 2020 U.S. Presidential Election, and 59% (n = 726) supported Joe Biden. We classified individuals as having an extreme political ideology if they responded that their ideology was either "Very conservative" or "Very liberal" (n = 346, 27.9%) and as not having an extreme political ideology if they responded with any of "Conservative," "Moderate," "Not sure," or "Liberal" (n = 892, 72.1%). YouGov weighted participants to match a nationally representative sample of American adults, and we use these weights in our statistical analyses.¹

The time period for data collection is important. In the five months preceding the start of our data collection (August 2020), QAnon-related activity surged on social media platforms in response to political events such as COVID-19 lockdowns and Black Lives Matter protests [51, 103]. One such event captured during our data collection was the November 2020 U.S. presidential election. With increased activity also came increased scrutiny from social media companies, resulting in the widespread removal of QAnon users and groups [23, 85]. Whether this censorship was effective is contentious, as QAnon websites saw increased traffic and QAnon-related groups remained active — for instance, an election fraud group became one of the fastest-growing groups in Facebook's history [36, 49]. Altogether, this suggests that our data was collected during a contentious period in which visible QAnon activity likely decreased due to censorship. However,

¹ Data and materials supporting this manuscript are available at <https://osf.io/9mqew/>

QAnon-related activity still remained strong, especially with an election that could spur further participation.

3.2 QAnon Websites

To identify visits to QAnon websites in people's web browsing behavior, we use a database of 324 QAnon-related websites curated by Hanley et al. [49], the largest available list of QAnon sites at the time of this writing. Hanley et al. [49] identified these websites by crawling all content posted on 8kun and Voat, two "hotbeds" of QAnon content. After constructing a hyperlink graph based on those web crawls, the authors used automated and manual classification techniques to extract websites from the graph dedicated to QAnon, leading to 324 unique sites.

3.3 Misinformation and Hard News Websites

To contextualize exposure to QAnon, we also identified visits to hard news websites and misinformation websites in people's web browsing. To identify visits to hard news websites, we use websites rated by NewsGuard, an organization composed of former journalists and news editors that manually rates the quality of news websites. NewsGuard rates over 8,000 websites that they claim are "responsible for approximately 95% of all the news and information consumed and shared online in the U.S., U.K., Canada, Germany, France, and Italy" (see <https://www.newsguardtech.com/newsguard-faq/> and <https://www.newsguardtech.com/ratings/rating-process-criteria/> for more information about NewsGuard's selection and rating process). We consider all websites rated by NewsGuard to represent hard news websites. NewsGuard also assigns a quality score to each website ranging from 0 (the lowest quality) to 100 (the highest quality), which represents the extent to which a given news website adheres to 9 different journalistic principles and best practices (e.g., regularly corrects or clarifies errors, discloses ownership and financing). We use these quality scores to determine the "quality" of people's news diets. For examples of other research on digital trace data that relies on NewsGuard ratings, see [7, 29, 64, 83].

To identify visits to misinformation websites, we follow the approach of Moore et al. [64] and use sites that NewsGuard has labeled as repeatedly publishing false content in combination with sites previously identified by academic researchers as known purveyors of false news. The total size of this database of misinformation websites is 1,796 sites.

4 RESULTS

4.1 Overview of Analytical Approach

We analyze these web browsing data using both quantitative methods (answering RQ1, H1, and RQ2) and qualitative methods (answering RQ2). Our results are presented below, and a figure visualizing our data, measures, and analytical approaches is presented in Figure 1.





Fig. 1. Overview of methodological approach

4.2 RQ1: How Were Americans Exposed to QAnon Websites?

We look to prior research using URL-based data to examine exposure to misinformation [45, 64] as a guide for the metrics we calculate to answer RQ1. Overall, we estimate that during the 2020 U.S. Presidential election, 3.7% (95% CI = 2.1-5.4%) of American adults were exposed to at least one QAnon website. For reference, based on our data, 39.1% (95% CI = 35.0-43.1%) of American adults were exposed to a misinformation website during the same time period.

Among those who did visit at least one QAnon website, the median number of QAnon website visits was 2 (95% quantile CI = 1, 117.25). Visits were highly skewed (as seen in Figure 2A): while 43.1% of those exposed were only exposed to one QAnon website, one person had 1,483 visits. For the median exposed person, only 0.16% of their news diet consisted of QAnon websites (95% quantile CI = .005, 2.8%) (Figure 2B). Looking at engagement, the average visit to a QAnon website lasted for 45.5 seconds (95% CI = 41.1-50.0s), longer than the average visit to a misinformation website (38.6 seconds [64]; 95% CI = 34.5-42.7s) and longer ($t = 6.4$, $p < .001$) than the average duration of visits to non-QAnon websites in our data (31.0 seconds; 95% CI = 30.9-31.1s).

A Distribution of Total QAnon Website Visits






Fig. 3. QAnon exposure by age group and race/ethnicity
Note: Proportion of people in each subgroup who are exposed to at least one QAnon website. Bars represent 95% confidence intervals.

Table 1. Predictors of exposure to at least one QAnon website

	Exposure to at least one QAnon website (binary)	
	<i>OLS</i>	<i>logistic</i>
Trump supporter	0.051*** (0.012)	1.452*** (0.347)
Political knowledge	0.015** (0.005)	0.523** (0.172)
Political interest	−0.002 (0.006)	0.003 (0.184)
College	0.040** (0.012)	0.993** (0.326)
Female	−0.008 (0.011)	−0.195 (0.325)
Non-white	−0.001 (0.012)	−0.272 (0.416)
Age 30-44 years	0.015 (0.018)	0.186 (0.495)
Age 45-64 years	−0.010 (0.017)	−0.460 (0.495)
Age 65+	−0.016 (0.018)	−0.713 (0.553)
Constant	−0.019 (0.026)	−5.495*** (0.844)
Observations	1,238	1,238
R ²	0.040	
Adjusted R ²	0.033	
Log Likelihood		−146.830
Akaike Inf. Crit.		313.660
Residual Std. Error	0.189 (df = 1228)	
F Statistic	5.712*** (df = 9; 1228)	

Note: *p<0.05; **p<0.01; ***p<0.001

Note: OLS and logistic regression coefficients are shown with standard errors in parentheses. The dependent variable is a binary variable indicating whether an individual was exposed to at least one QAnon website during the data collection period (1) or not (0). “Trump supporter” variable: intending to vote for Trump in 2020 election = 1; not intending to vote for Trump in 2020 election = 0. “Political knowledge” variable: variable ranging from 0-4 representing the number of questions in Pew Research Center’s civic knowledge questionnaire answered correctly out of 4. “Political interest” variable: variable ranging from 1-4 where 4 = people who say they pay attention to what’s going on in government and politics “most of the time” and 1 = those who pay attention “hardly at all”. “College” variable: 1 = college graduate; 0 = not a college graduate. “Female” variable: 1 = indicated identifying as a female, 0 = did not indicate identifying as a female. “Non-white” variable: 1 = indicated identifying as a race other than white; 0 = indicated identifying as white. P-values are two-sided.

4.3 H1: Those with More Extreme Political Orientations are More Likely to be Exposed to QAnon Websites

Next, we tested H1: Political extremists are more likely to be exposed to QAnon websites. To do so, we first conducted a simple proportions test comparing the exposure rate among those with the most extreme ideologies to those without the most extreme ideologies. We find that 7.2% of participants with the most extreme ideologies were exposed to at least one QAnon site compared to 2.6% of those without extreme ideologies, representing exposure at a significantly higher rate ($\chi^2(1) = 12.4$, $p < .001$). Even after controlling for a set of sociodemographic covariates from prior research on misinformation exposure [45, 64], including which candidate a participant supported in the 2020 election, political knowledge, political interest, education, gender, race/ethnicity, and age, we find that those with extreme political ideologies are indeed significantly more likely to be exposed than those without an extreme political ideology, and this finding is robust to whether we fit an OLS regression ($p = .002$; see Table 2) or a logistic regression ($p = .004$). Overall, we take these results as evidence in support of H1.

To dig deeper into rates of QAnon website exposure among those on the ideological extremes, we conducted an exploratory proportions test to assess whether there were differences in exposure between those on the far left (i.e., identifying as “very liberal” when asked about their political ideology in our survey) and those on the far right (i.e., identifying as “very conservative”). We found that the proportion of participants identifying as “very conservative” who were exposed to QAnon sites (11.1%) was significantly higher than the proportion of those identifying as “very liberal” who were exposed to QAnon sites (3.0%) ($\chi^2(1) = 6.58$, $p = .01$). So, while we found evidence that those on the ideological extremes were more likely to be exposed to QAnon websites than those without extreme ideologies (supporting H1), the two ideological extremes themselves were not equally likely to be exposed to QAnon sites.

Table 2. Ideological extremity and exposure to QAnon websites

	Exposure to at least one QAnon website (binary)	
	<i>OLS</i>	<i>logistic</i>
Ideologically Extreme	0.039** (0.012)	0.905** (0.315)
Trump supporter	0.048*** (0.012)	1.438*** (0.350)
Political knowledge	0.016** (0.005)	0.547** (0.176)
Political interest	-0.003 (0.006)	-0.062 (0.186)
College	0.038** (0.012)	0.922** (0.332)
Female	-0.007 (0.011)	-0.153 (0.327)
Non-white	0.001 (0.012)	-0.205 (0.421)
Age 30-44 years	0.012 (0.018)	0.083 (0.502)
Age 45-64 years	-0.012 (0.017)	-0.553 (0.503)
Age 65+	-0.017 (0.018)	-0.764 (0.564)
Constant	-0.023 (0.026)	-5.613*** (0.853)
Observations	1,238	1,238
R ²	0.048	
Adjusted R ²	0.040	
Log Likelihood		-144.069
Akaike Inf. Crit.		310.137
Residual Std. Error	0.188 (df = 1227)	
F Statistic	6.137*** (df = 10; 1227)	

Note:

*p<0.05; **p<0.01; ***p<0.001

Note: OLS and logistic regression coefficients are shown with standard errors in parentheses. The dependent variable is a binary variable indicating whether an individual was exposed to at least one QAnon website during the data collection period (1) or not (0).

4.4 RQ2: Are There Distinguishable Types of QAnon Website Consumers?

We answer RQ2 using a mixed methods approach, which allows us to identify participants' encounters with QAnon websites in a way not constrained to a particular platform [97] or necessitating any engagement on their part other than a website visit. At a high level, our approach consists of first using quantitative techniques to identify clusters based on quantitative features of people's browsing data that the QAnon consumers in our dataset sorted into. Next, we qualitatively analyzed the browsing behaviors of the individuals sorted into the quantitatively-derived clusters, allowing us to deepen our understanding of how the people in these different

clusters consume QAnon websites and how QAnon websites fit into their online media diets. Our approach of combining quantitative and qualitative methods to analyze digital trace data follows calls from other scholars who argue that mixed methods analysis of large-scale online behavioral data yields a richer understanding of phenomena in those data than the use of qualitative or quantitative methods alone [38, 95, 97]. Given the lack of strong theoretical priors to derive specific hypotheses about the different types of QAnon users and their behaviors, we thought this approach was prudent in answering RQ2.

4.4.1 Quantitative Approach. To examine if there were quantitatively-identifiable behavioral differences among those who did visit QAnon websites, we calculated descriptive statistics of the browsing behaviors of those exposed to at least one QAnon website and created clusters based on those behaviors. Specifically, we calculated four main statistics that directly relate to how people interacted with QAnon content: (1) mean of the number of daily visits to QAnon websites, (2) standard deviation of the number of daily visits to QAnon websites, (3) the inertia of QAnon website visits, and (4) entropy of the distribution of visits to categories of websites. These metrics were inspired by prior research on social media and misinformation [18, 39, 56, 84].

An individual's mean number of daily visits to QAnon websites measures their total consumption of these websites. A person's standard deviation measures uncertainty in the calculation of their mean number of visits. A lower standard deviation could suggest that the individual is a consistent consumer of QAnon websites or low uncertainty in the mean calculation. In comparison, a higher standard deviation could suggest more fluctuation in their daily QAnon consumption or high uncertainty in the mean calculation. Inertia is the autocorrelation (acf) in a daily binary of QAnon website exposure at lag 1. In other words, does visiting a QAnon website one day predict visiting a QAnon website the next day? Higher inertia implies that visiting a QAnon website for that person is "sticky" and thus pulls them back to visit a QAnon website the next day. In contrast, lower inertia indicates that QAnon visits do not have as strong of a hold on an individual and that person is not as likely to visit a QAnon website the next day. While correlated, the mean number of visits and inertia are not perfectly correlated (see Supplementary Materials Figure S1) and measure distinct constructs. For example, a person could visit a large number of QAnon sites on a single day. As a result, they would have a high daily mean number of visits but low inertia. Still, individuals with high inertia, by definition, are likely to have a higher number of mean daily visits. Entropy measures the categorical distribution of visits across different website categories (misinformation, high-quality news, low-quality news, fact-checking, hate speech, Facebook, and Twitter [24]; in the case of a domain appearing in multiple lists, the domain is labeled by priority in the order listed). Higher entropy indicates that a person has a higher dispersion of visits across the categories. Conversely, lower entropy equates to more concentrated visits to a small number of categories of websites.

We use these metrics to cluster individuals exposed to at least one QAnon website. Following other research that clusters users based on their news browsing behaviors [60], we use a k-means clustering algorithm. In this case, a cluster represents individuals who are similar to one another based on the behavioral metrics we are analyzing. Before clustering, we removed one individual who had an extraordinary number of visits to QAnon websites (1,483; the person with the next highest number of visits had 125). Again, following past research, we selected the number of clusters using the elbow method [60], and our resulting clustering is in line with Dalmaijer et al.'s guidelines for statistical power for cluster analysis [25]. Furthermore, the number of participants we devise into clusters is larger than the number of participants in other research in human-

computer interaction which sorts participants into clusters based on behavioral features [26, 32, 41, 48, 58].

To create our clusters, we first calculated the mean number of daily visits to QAnon websites ($M = .11$, $SD = .26$), the standard deviation of the number of daily visits to QAnon websites ($M = .42$, $SD = .67$), the inertia of QAnon website visits ($M = .01$, $SD = .12$), and entropy ($M = .73$, $SD = .31$). The elbow method suggested that we should cluster individuals into four clusters. After the clustering, one individual was grouped into their own cluster and analyzed separately in the next section. The cluster means of the three remaining clusters are in Table 3.

Table 3. Quantitative cluster means

Cluster Name	entropy	inertia (acf)	daily mean	daily sd	percentage
1. Regular	0.77	0.19	0.55	1.62	10.2%
2. Diverse	0.93	-0.02	0.04	0.24	55.1%
3. Selective	0.39	-0.02	0.02	0.17	34.7%

Note: Means of the QAnon-related browsing statistics used to cluster individuals and percentage of individuals in the cluster. Entropy represents the diversity of the types of websites visited, with a higher entropy meaning higher diversity. Inertia is the autocorrelation of visiting a QAnon website one day with visiting a QAnon website the next day. Daily mean is the mean number of QAnon websites consumed per day. Daily sd is the standard deviation of the number of QAnon websites visited per day. Percentage is the percentage of individuals in the given cluster.

These quantitatively-derived clusters represent unique media consumer types that vary in their media diets and consumption of QAnon websites. We refer to those in Cluster 1 as the “regular” consumers of QAnon websites. On average, they have a relatively high mean and standard deviation of daily visits. They also have middling entropy, meaning that the dispersion of categories of websites they visit is not as low or high as the other clusters. Cluster 1’s positive inertia means that when they visit a QAnon website on one day, that visit is “sticky,” and they are likely to visit a QAnon website the next day. Only 10.2% of our participants exposed to at least one QAnon website appear in this cluster, the fewest of the three clusters. Cluster 2 (55.1% of participants in the analysis), the “diverse” consumers, and Cluster 3 (34.7% of participants in the analysis), the “selective” consumers, represent infrequent consumers of QAnon websites. Both clusters contain people with low mean and standard deviations of exposure. They also have negative inertia – meaning that visiting a QAnon website on one day negatively correlates with visiting a QAnon website the next day. As seen in Figure 4, these two clusters differ in entropy levels. Cluster 2 has a higher level of entropy than Cluster 3, meaning that Cluster 2 consists of people who visit a greater diversity of different types of websites.




Fig. 4. Quantitative clusters of those exposed to QAnon websites

Note: Each point represents one person, colored by their assigned cluster. The x-axis represents the daily mean visits to QAnon sites for each person. The y-axis is people's average entropy.

To explore the underlying demographic and political characteristics of individuals in each cluster, we fit a multinomial logit regression in which the dependent variable was the cluster assignment and the sociodemographic variables from prior work on misinformation exposure [45, 64] were the independent variables. The breakdown of individuals in each cluster belonging to these sociodemographic groups is presented visually in Figure 5. The full multinomial regression output is in the Supplementary Materials (Table S1), but due to the relatively small sample size in some of the clusters we urge caution in interpreting the significance of the differences between sociodemographic groups in cluster membership, particularly in interpreting null results.




Fig. 5. Quantitative clusters of those exposed to QAnon websites

Note: Estimates of demographic and political variables for members in QAnon-exposed clusters. Dots represent point estimates and ranges represent 95% confidence intervals. Binary variable estimates calculated via one-sample proportions tests with participants being labeled as a “1” if the variable is applicable to them or “0” if not. Binary variables include: “Age”: participant’s age within the range = 1; age not in range = 0. “College” variable: 1 = college graduate; 0 = not a college graduate. “Female” variable: 1 = indicated identifying as a female, 0 = did not indicate identifying as a female. “Non-white” variable: 1 = indicated identifying as a race other than white; 0 = indicated identifying as white. “Trump supporter” variable: intending to vote for Trump in 2020 election = 1; not intending to vote for Trump in 2020 election = 0. Scale variable estimates calculated via one-sample t-tests. Scale variables include: “Political knowledge” variable: variable ranging from 0-4 representing the number of questions in Pew Research Center’s civic knowledge questionnaire [82] answered correctly out of 4. “Political interest” variable: variable ranging from 1-4 where 4 = people who say they pay attention to what’s going on in government and politics “most of the time” and 1 = those who pay attention “hardly at all”.

4.4.2 Qualitative Approach. To deepen our understanding of QAnon exposure among the individuals in the quantitatively-derived clusters, we performed a qualitative content analysis of a sample of 17 participants exposed to QAnon content. For many participants’ website visits, the full URL was available to researchers (e.g., twitter.com/POTUS/status/1351946842838347776)², making it possible to see the exact webpage a participant was visiting by loading the URL in a browser or web archive. Because the quantitative analysis was conducted only at the domain level (e.g., twitter.com), the qualitative content analysis provided rich, contextualizing information about each participant’s content consumption, browsing behavior, and demographics.

In order to pilot procedures for the qualitative review, three participants were selected in a random sample stratified by participants’ number of QAnon website visits (average, above

² For some website visits, YouGov does not provide the full URL to protect participant privacy as URLs can sometimes contain identifying information, see section 5.5 for a more detailed discussion.

average, and below average). Within each participant's web browsing history, researchers experimented with URL sampling strategies that would provide sufficient information about a participant's web activity without manually reviewing tens of thousands of URLs for each participant. Based on the results of this initial review, we finalized procedures for analyzing the browsing behavior of 14 additional participants who had been exposed to QAnon. We randomly selected these 14 participants from the quantitatively-generated clusters: "regular" consumers ($n = 5$), "diverse" consumers ($n = 4$), and "selective" consumers ($n = 5$). Researchers individually reviewed (1) one full day of each participant's URL visits, (2) each participant's individual QAnon exposures, and (3) 10% of participants' URL visits divided into randomly sampled batches of 100 sequential URLs. In addition, researchers reviewed each participant's 100 most visited, longest visited, most conservative, most liberal, highest journalistic quality (according to NewsGuard quality scores), and lowest journalistic quality URLs.


These sampling strategies reflect both the nature of the research objectives and the data itself. First, in selecting our pilot participants based on QAnon exposure and our review participants based on QAnon cluster assignment, we are, in a sense, selecting participants on a dependent variable. This selection is because we are not interested in qualitatively distinguishing QAnon-exposed participants from unexposed participants — such a comparative design would be exceedingly difficult to execute qualitatively, given the magnitude of the data. Instead, we seek qualitative breadth and depth among the participants exposed to QAnon by stratifying based on quantitative metrics. Thus, these sampling strategies are intended primarily to complement the quantitative analysis rather than supplant or stand apart from it. We are specifically interested in the relationship between the quantitative exposure metrics and participants' actual browsing and exposure experiences. Given this motivation, we believe these strata are best equipped to give in-depth information on QAnon-exposed participants, especially as this strategy captures a high proportion of the most frequently QAnon-exposed participants — a comprehensive review of all 48 exposed participants is impractical and unnecessary, given that the vast majority of participants were exposed to QAnon only once.

For analysis, researchers followed iterative, open coding procedures based on a narrow set of high-level codes directly related to participants' (1) QAnon exposure, (2) browsing behavior, and (3) cluster assignments. These three high-level codes were chosen in order to ensure that the thematic analysis complemented the quantitative analysis and addressed the research questions, and they were generated by the research team as a whole prior to thematic analysis. Given the high-level codes, sub-codes (e.g., "intentional exposure" vs "incidental exposure" for the high-level code "QAnon exposure") were developed iteratively by comparing content that was coded under the high-level code. This coding was performed by one researcher who is trained in qualitative coding and most familiar with QAnon and alt-right online content. Due to the nature of the data, this researcher produced summative "field notes" about participants that highlighted patterns in web activity, such as the political messages of visited URLs, details revealed about participants' lives, and which sites participants used to search or navigate the web. The researcher applied qualitative codes to these descriptive notes rather than to individual URLs. For example: Descriptions of how a participant was exposed to QAnon were captured under the high-level code "QAnon exposure"; then these descriptions were compared against one another to find similarities and differences; this led to the development of the sub-codes "intentional exposure" and "incidental exposure;" finally, the appropriate sub-code was applied to all descriptions coded "QAnon exposure." This coding strategy focused the qualitative content analysis to reach

saturation, as the vast number of URLs lends itself to endless insights about the minutiae of individual participants. The results of this review are discussed below.

QAnon Exposure. QAnon exposures in our data were identified using the QAnon domains provided by Hanley et al. [49]. While Hanley et al. developed a more detailed taxonomy of QAnon websites, we reduce it to a binary distinction: *QAnon-adjacent websites* and *QAnon-explicit websites*. QAnon-adjacent websites dominated exposures: These websites often repeated QAnon tropes and reported on news through a conspiratorial lens, but they were usually multi-purpose conservative news aggregators (e.g., thedcpatriot.com, nationalvanguard.com, newspunch.com).³ By design, these sites produce low-quality, incendiary conservative news articles. QAnon-adjacent websites sometimes explicitly support or discuss QAnon, but visiting these sites without encountering QAnon content is possible. Instead, their content works in service of multiple far-right ideologies — including QAnon followers, white supremacists, Trump supporters, or anti-vaxxers. These websites are often no less radical or extreme in their beliefs than QAnon-explicit websites. QAnon-explicit websites, on the other hand, acknowledge their relationship to QAnon by hosting discussions of QAnon, writing approvingly of QAnon, and compiling the latest information on the conspiracy theory (e.g., qntmpkts.keybase.pub, dustinnemos.com, theqanons.com). These terms (*QAnon-adjacent* and *QAnon-explicit*) will be used to differentiate QAnon exposure among participants. Examples of such sites are provided in Figure 6.

QAnon-explicit Website



QAnon-adjacent Website



Fig. 6. Examples of QAnon-explicit and QAnon-adjacent websites

Note: The screenshot on the left (*QAnon-explicit Website*) is taken from qanon.pub, one of many generic sites that aggregated posts from Q for wider distribution. The screenshot on the right (*QAnon-adjacent Website*) is taken from thepeoplesvoice.tv, formerly newspunch.com, a far-right news site that occasionally expressed support for QAnon beliefs among many other extremist, anti-establishment convictions.

Of the 14 participants subject to qualitative review, 13 participants' movement from non-QAnon websites to QAnon websites can be understood as a *compatible* ideological shift; that is to say, the far-right content they regularly consumed prior to QAnon exposure was generally no less radical or conspiratorial. These 13 participants were consuming information that predisposed them to QAnon exposure such that their eventual exposure was not an aberration from typical news consumption. In practice, this often presented itself as typical browsing: They would be on Facebook reading conservative articles, and then either the Facebook page or the news website

³ The content found on websites may be distressing to researchers. For example, some content found on the websites we examined included support for white supremacist ideologies, anti-semitic imagery and humor, calls for violence against minority groups, and unanticipated explicit sexual materials. Researchers should consider the potential harms and benefits of qualitatively analyzing such content and carefully consider whether it is worth the potential psychological distress to have members of the research team consume such content.

would lead to a QAnon-adjacent website that described a similar news story. For the remaining participant, QAnon exposure was *incompatible* with their regular left-wing news consumption and stated political affiliation, indicating that QAnon exposure was *intentional* – driven perhaps by curiosity. They consumed a wide range of content, including popular socialist podcasts, inflammatory far-right blogs, and tweets of all political allegiances; they frequented libertarian and communist gaming forum discussions; and it is possible that this browsing was related to their job, which, based on visited URLs, was in content moderation. Their QAnon exposure was mixed indiscriminately with other far-right content. For these reasons, this participant has been excluded from the analysis. Of the 13 remaining participants, their QAnon exposure can be categorized as compatible and *intentional* ($n = 3$) or compatible and *incidental* ($n = 10$), differentiated by whether they actively sought QAnon content or seemed to happen upon it.

Table 4. Qualitative typology of participants exposed to QAnon content

	Compatibility with QAnon Content	
	Compatible	Incompatible
<i>Intentional</i> exposure	3	1
<i>Incidental</i> exposure	10	0

Note: Participants were deemed “compatible” with QAnon content if their prior content consumption history and stated political affiliation signaled openness to QAnon ideas, whereas “incompatible” participants lacked consumption patterns and/or political affiliations that aligned with QAnon beliefs. Participant exposures were deemed “intentional” if participants actively sought out QAnon content, while “incidental” exposures appeared to occur more-or-less at random.

All compatible, *incidental* participants supported former president Donald Trump; however, their political consumption is diverse. This group includes a flat Earth, Evangelical conspiracy theorist; a white supremacist; and several Trump supporters who stay abreast of crime, culture wars, Joe Biden, pedophilia, and COVID-19 vaccines. In each case, exposure to QAnon was brief, emerging from a link found on social media, email, or another conservative news aggregator. However, these participants were exposed almost entirely to QAnon-adjacent websites. In many cases, these QAnon websites were no more conservative or conspiratorial than other content – and in the case of the white supremacist and flat-Earther, these websites were *less* conspiratorial. Overall, these QAnon exposures occurred within predominantly far-right browsing ecosystems, QAnon or otherwise.

Compatible, *intentional* participants are far-right Trump supporters who sought out QAnon content on purpose. Like other participants, determining intent required seeing the patterns in their browsing behavior, identifying moments of clear intent (e.g., Googling a specific term), and then seeing how browsing behavior incorporated QAnon websites. This pattern often meant QAnon website visits had no antecedent (e.g., they visited it regularly like one might visit Facebook or Google) or a clear antecedent (e.g., they visited it after a specific Google search or news story). These participants regularly visited QAnon websites, watched videos produced by QAnon influencers, and searched for news on alternative search engines and fringe websites. However, these participants had distinct relationships with QAnon:

Participant 909 (see Figure 7) was a fervent consumer of far-right, conservative news. They followed various pathways between Facebook, Twitter, and alternative search engine swisscows.com to arrive at increasingly fringe sites. Before their intentional QAnon exposures, they visited several QAnon-adjacent websites within their far-right media ecosystem. After searching for the term “qanon,” they became a regular visitor to a QAnon-explicit site, theqanons.com, which dramatically dropped the average quality of their news consumption in combination with more far-right news.




Fig. 7. Participant 909’s QAnon website consumption

Note: Panel A shows the average news quality per day for the sites visited by Participant 909. The y-axis is the average NewsGuard score of websites they visited with 0 indicating the lowest quality news sources and 100 indicating the highest quality sources (<https://www.newsguardtech.com/ratings/rating-process-criteria/>). Panel B shows the percentage of websites visited by the participant per day that were QAnon websites. These sites are divided into QAnon-adjacent domains (dotted line) and QAnon-explicit domains (solid line), with accompanying Loess regressions. The vertical red line in each panel marks the day that the participant searched for “QAnon.”

Participant 405 (see Figure 8) was already a QAnon supporter at the start of data collection, evidenced partly by the time they spent shopping for QAnon merchandise. They gravitated toward far-right personalities, including early QAnon influencers, popular libertarian journalists, and public Evangelicals. Their QAnon consumption notably increased following the removal of one of their favorite influencers from YouTube. While the average quality of this participant's news diet dropped, this is likely because they moved from watching QAnon content on YouTube to BitChute, a video site that is rated as worse quality than YouTube.




Fig. 8. Participant 405's QAnon website consumption

Note: Panel A shows the average news quality per day for the sites visited by Participant 405. The y-axis is the average NewsGuard score. Panel B shows the percentage of websites visited by the participant per day that were QAnon websites (solid line) and the percentage of visits that were to the alternative video platform BitChute (dotted line). There are accompanying Loess regressions for each line. The vertical red line in each panel marks the day that a QAnon YouTube channel was suspended, leading the participant to the same channel on BitChute.

Participant 371 (see Figure 9) began as a run-of-the-mill Trump supporter with incidental exposures to a QAnon-explicit site. However, in the leadup to and aftermath of the 2020 presidential election, they began actively seeking out QAnon content on the most explicit QAnon sites while experiencing a sharp decline in the overall quality of their news diet.




Fig. 9. Participant 371’s QAnon website consumption

Note: Panel A shows the average news quality per day for the sites visited by Participant 371. The y-axis is the average NewsGuard score. Panel B shows the percentage of websites visited by the participant per day that were QAnon websites, with an accompanying Loess regression. The vertical red line in each panel marks the 2020 U.S. general election day.

We previously noted the exclusion of two participants with high exposure to QAnon who were removed from the cluster analysis. These participants were QAnon followers before the start of data collection, and they incorporated QAnon-explicit websites into their general browsing behavior — one used QAnon to keep up with the news, and the other visited QAnon websites in tandem with sites dedicated to “Doomsday prepping.” In all five of these cases, it is impossible to pinpoint the exact moment in which a participant became a QAnon content consumer — if their consumption did not predate the start of data analysis, it is still impossible to definitively state that whatever trends we see in their consumption are related to offline exposure or exposure/consumption hidden within inaccessible browsing (e.g., Facebook and Instagram). Instead, we have only indicators of their engagement with QAnon. Finally, we recognized that no Biden supporters were included in our 14-participant stratified random sample of QAnon website consumers. We conducted a brief, purposive review of four Biden supporters who primarily consumed liberal news. We found that their exposure to QAnon was *incompatible* with their content consumption and represented an *incidental* encounter with QAnon. For example, one stumbled on QAnon while attempting to fact-check a QAnon-related news story. Another Googled QAnon after reading a mainstream news article that mentioned it and landed on a QAnon-explicit website. It is important to note that, regardless of whether exposure was *compatible* or *incompatible* with participants’ political content consumption, for most participants, it was entirely *incidental* — often representing only one brief exposure.

Browsing Behavior. Each participant navigated the vastness of the internet differently. In every instance, participants had certain “hubs” from which the rest of their browsing originated. Specific hubs, such as search engines, email inboxes, and social media platforms, usually moved participants out of the hub and into other sites but were inevitably returned to. These hubs facilitated initial exposure to QAnon websites. Other hubs, such as news websites and video-sharing platforms, facilitated binging behavior, including the binging of QAnon content. Additionally, some participants engaged in adversarial/confirmatory searching, searching for terms to produce results that conformed to their presuppositions. These searches often contained negatively connotated keywords that parroted right-wing arguments, prompting similar results (e.g., “George Floyd overdose” or “Kamala Harris innocent men jail”), especially for content that had yet to be fact-checked. This pattern in our data is corroborated by Aslett et al. [12], who found that conducting online searches about misinformation tends to increase belief in misinformation, especially among those for whom search engines return lower-quality results.

Cluster Assignments. The qualitative review revealed where the clusters’ quantifiable distinctions succeeded and failed to accurately group participants together. Cluster 1, the “regular” QAnon supporters, accurately captured the 3 participants who intentionally sought out QAnon content because they supported it. However, this cluster also incorrectly captured the aforementioned excluded participant who, while a frequent consumer of QAnon content, appeared to find the conspiracy theory entertaining or interesting rather than believable. Finally, this cluster included a consumer of QAnon-adjacent content who, while certainly aligned with the extreme far-right, including known white supremacist groups, was not a rank-and-file QAnon supporter.

As for Clusters 2, the “diverse” consumers, and 3, the “selective” consumers, qualitative evidence supports participants’ categorization based on negative inertia following QAnon exposure. These participants were primarily exposed to QAnon-adjacent websites that appeared incidentally but aligned with their general news consumption. It is accurate to say that these incidental exposures do not predict, and may even negatively predict, continued exposure to

QAnon. However, it would be misleading to suggest that participants in Clusters 2 and 3 are avoiding QAnon exposure. The “selective” cluster’s low entropy is partly explained by the presence of a central hub in participants’ browsing. This hub (Google for three participants, the social media platform Tagged for one participant) dominates browsing history and thus appears to reduce browsing diversity overall. Meanwhile, the “diverse” cluster members had several central hubs that were roughly equal in size, increasing browsing diversity and entropy.

5 DISCUSSION

5.1 Empirical Contributions

We found that a small percentage (3.7%) of U.S. adults were exposed to QAnon websites during the 2020 election. Consistent with existing research on belief in QAnon and conspiracy theories in general, we found that those on the political extremes were more likely to be exposed to QAnon websites. Most people who visited QAnon websites visited very few during the three-month period we collected data, although some were prolific QAnon consumers.

As the first nationally representative study of QAnon exposure, our findings raise key questions about the prevalence of QAnon support within the United States. While one poll has put support for the *ideas behind* QAnon as high as 41%, only 2% of participants in the same poll self-identified as QAnon supporters [101]. In triangulating support for QAnon, it becomes increasingly important to recognize the gaps between media coverage, self-identification with the conspiracy theory, belief in the principles of the conspiracy theory, and exposure to the conspiracy theory itself.

Our typographical analysis and qualitative review highlight this gap. For most participants reviewed, exposure to QAnon was a *compatible, incidental* shift occurring within a right-wing media ecosystem filled with *QAnon-adjacent* websites that already promoted the principles of QAnon — even if QAnon was not explicitly named. For those who experienced *compatible, intentional* exposure to QAnon, their immersion in this right-wing ecosystem likely primed their exposure; even without access to their social media browsing data or offline media consumption, the content they consumed was indicative of someone who shares beliefs that are compatible with tenets of QAnon. This finding supports the notion that QAnon support “may be deeper than it is wide” [30]. While a lack of broad support for QAnon may appear positive — as much as deep, fervent belief sounds negative — the qualitative investigation ultimately points to the sophistication of a right-wing media ecosystem that allows distinct forms of far-right extremism to co-exist in parallel, as the line between the sites visited by flat-Earthers, white supremacists, and QAnon supporters becomes thin. When we consider that these news sites often appear to work in tandem and that even people with high exposure to “mainstream” press seek out these fringe sites, it is evident that a more thorough investigation of how people arrive at these sites is warranted.

The three *compatible, intentional* cases also indicate trends that warrant greater attention. The average quality of Participant 405’s news consumption, as seen in Figure 7, was inflated by the amount of time spent on YouTube rather than BitChute — the fringe platform they were pushed to following YouTube’s suspension of their preferred QAnon influencer. Again, this suggests that engagement with QAnon was deep, and attempts to censor QAnon on mainstream platforms likely do not prevent hardline believers from seeking such content. Participant 909, who found QAnon through a simple web search on a non-mainstream platform speaks to how the right-wing media ecosystem may enable incidental exposure to become intentional exposure. Participant 909’s and Participant 371’s experiences correspond most closely with the “rabbit hole” theory of

radicalization — wherein people exposed to misinformation fervently consume more misinformation content and plunge deeper into fringe beliefs [19]. While our data is not robust enough to substantiate such theories, these participants' movement towards the more extreme content of QAnon aligns with current empirical evidence on rabbit holes [82]. Crucially, though, these participants' movements to intentional QAnon consumption seem driven more by personal interest than algorithmic manipulation and do not appear to be happening within the context of a single platform, which again highlights the need for multi-platform studies that incorporate qualitative information about participants' online and offline lives.

5.2 Methodological Contributions

Our work makes methodological contributions that we argue are highly relevant to scholars studying people's journeys through the online (dis)information ecosystem. In particular, we highlight the utility of using web browsing data to study people's exposure to online extremist content such as QAnon content. These passive measures of information consumption do not rely on people self-reporting their exposure to certain content online, which people may be unwilling or unable to report accurately [46, 47, 73, 80]. Even more importantly, compared with past work on QAnon, which relies on data about people's publicly-viewable behaviors on specific platforms (e.g., posts on Reddit, retweets on Twitter), web browsing data cover people's exposure regardless of if they engage with content in a way that leaves a publicly visible signature and also is not constrained to a particular platform; instead, it tracks journeys as individuals move from platform to platform. That is, while prior work on people's engagement with QAnon content is largely limited to the "public life" of QAnon, web browsing data sheds light on the "private life" of QAnon. Furthermore, the ability to observe users' encounters with content as they move across platforms on the web is especially valuable in the context of disinformation and related content where deplatforming of users and organization frequently occurs [31, 52, 63] and cross-platform content strategies are an important part of information operations [63, 96, 97].

Relatedly, we demonstrate the value of analyzing web browsing data using a mixed-methods approach. While quantitative techniques are typically used to analyze web browsing data [67], our mixed-methods analysis shows how qualitative techniques can add context and richness to quantitative findings. Our quantitative findings answer questions about overall levels of exposure to QAnon websites and whether certain people were more likely to be exposed than others. We further leverage quantitative methods to identify different clusters of users exposed to QAnon websites based on metrics about their media diets (e.g., diversity of website categories visited, rarity of QAnon websites in their diet). Critically, though, our qualitative analysis sheds light on the role of QAnon content in users' media diets. On their own, these qualitative methods can paint compelling narratives that explain the context, behaviors, and personal histories that may have contributed to QAnon exposure. When used in tandem with quantitative analysis, the qualitative results deepen our understanding of the quantitative findings: They both validate and complicate the cluster analysis, they contextualize exposure stratified by political affiliation and make apparent how "exposure" can take on distinct meanings for distinct users. Finally, qualitative methods reveal patterns in web browsing behavior (such as indicators of "rabbit hole" radicalization) that can guide future quantitative analysis. Future work can use these qualitative insights to inform quantitative metrics that are used to identify patterns in browsing behavior and extremist exposure, and such an iterative process can yield a better and more nuanced understanding of how people come to encounter and engage with QAnon and extremist content on the internet. More generally, our approach demonstrates the value of mixed-methods

approaches to analyzing online behaviors [38, 95, 97], whether through web browsing data or other digital traces. While work in CSCW emphasizes the challenges of using conventional quantitative and qualitative methods to analyze large-scale digital trace data [54, 87], our study provides a blueprint for how both approaches can be integrated in the analysis of a specific type of digital trace data, web browsing logs.

5.3 Design Implications

The results of this research demonstrate the value of taking an ecological approach to the monitoring and combating of online extremism [74] and have implications for the design of interventions and policies to limit exposure to QAnon and other extremist content online. Specifically, by grounding our approach in the cross-platform browsing behavior of individuals exposed to QAnon — rather than aggregated public activity from users of isolated platforms — our research sheds light on how users' navigational patterns within a broad online ecosystem should be considered when attempting to track the spread of extremist content and develop interventions for stopping the spread of extremist content. First, this research highlights the function of network hubs, which aggregate information and provide entry points into exposure. While social networking sites have previously banned hyperlinks that redirect users to extremist content [6], other platforms such as search engines, mail inboxes, news aggregators, and forums are important nodes within the digital ecosystems which may provide a space for exposure to extremist content. Determining which nodes facilitate the greatest flow of extremist content within online ecosystems is important for monitoring the growth of extremist communities, online infrastructures, and movements.

Moreover, this research highlights the potential effectiveness of efforts designed to combat extremism coordinated across multiple online platforms that play roles within extremist sites, including mainstream sites such as Facebook, Google, and Outlook Mail. While deplatforming efforts localized to individual platforms have reduced the influence of extremist communities and spokespeople on mainstream social media networks [31, 52, 81], certain extremist communities may be resilient to deplatforming [1, 20, 63]. The tendency for extremists to be exposed to content and congregate on platforms outside of mainstream social networking and news sites reiterates the willingness of extremists to both maintain and occupy alternative platforms which are more difficult to regulate [74]. Rather, efforts to disconnect alternative platforms from popular hubs (such as search engines, social networks, and advertisements) could prevent incidental exposure to extremist sites. If the promises of a federated, decentralized social internet are fulfilled, such efforts become all the more important, as mutual cooperation between platforms to prevent the amplification of extremist content becomes critical.

5.4 Limitations

Of course, our study has limitations and our results should only be viewed in light of them. First, our URL-based data constrains how much of people's online behavior we capture. While we collected the websites visited by our participants in their web browsers, we do not have information about participants' use of apps, which are increasingly important for news and political information consumption, especially on smartphones [42]. While our study, which captures web browsing across computers, smartphones, and tablets, is an improvement over past research using web browsing data limited only to desktop and laptop computers (e.g., [45]), we still are unable to observe participants' use of apps, in which QAnon content exposure could be happening. Furthermore, many social media platforms (e.g., Facebook, Twitter) expose people to

information in feeds reflected by static URLs (e.g., www.facebook.com). We observe instances where people click on posts in feeds and are taken to external websites but do not know the information people were exposed to on their feeds that they did not click on. Future work on QAnon content consumption can use device-based (e.g., Screenomics [79]) and platform-specific (e.g., YouTube recommended videos tracker [19]) data collection tools to fill in these important gaps in browser-based URL collection.

Second, our data is limited in its representativeness. These data were collected by people living in the U.S. who installed the YouGov Pulse plugin in exchange for financial compensation. This self-selection bias can impact the generalizability of the results [27, 40]. However, this affects other web-browsing data collection methods which require participants to install software or send researchers their browsing history [14, 15, 21, 45, 64, 89, 94, 98]. Prior research indicates that people without a high school diploma are underrepresented in the YouGov Pulse panel and that the panel tilts in favor of the Democratic Party. Still, along many dimensions, the Pulse panel closely resembles the American adult population [45]. Moreover, there may be concerns that our participants modified their browsing behavior because they knew they were being observed. However, past research using digital trace data has not found evidence of strong Hawthorne Effects [13, 100], and the prevalence of website visits in our data with sensitive material, such as pornography sites, suggests that this self-censorship likely did not occur at a significant scale.

Third, the generalizability of our findings is limited by the time frame during which we collected data. In the time leading up to the 2020 election, the presence of QAnon content online surged in response to the emergence of the COVID-19 pandemic in the spring and the Black Lives Matter protests in the summer [51, 103]. While this increased presence led the QAnon movement to face some moderation on social media platforms, traffic to QAnon websites and QAnon groups on major platforms like Facebook increased or remained essentially unchanged [49, 55]. While all this means that the time period during which we collected our data was well-poised to demonstrate QAnon content consumption, QAnon has faced severe setbacks since 2020 that have likely changed its online presence and, thus, people's consumption of QAnon content. Notably, after the January 6, 2021 insurrection, QAnon faced severe moderation on mainstream social media platforms. The movement retreated to more obscure platforms, like Telegram and Parler [5]. While these platforms may have fewer users, researchers have found that the QAnon-related discourse appearing on these platforms is more toxic than discourse that appeared on more mainstream platforms [86]. Thus, as the media ecosystem and political environment continuously evolve, researchers need to collect new data to update knowledge on the prevalence of extremist content. Such updates can illuminate changes in who is consuming QAnon content, how, and to what effect, and more generally contribute to our understanding of the mechanisms that underlie the processes of online radicalization and conspiratorial ideation (for examples of such "updates" in the domain of misinformation, see [44, 64]).

5.5 Ethical Considerations

While the web browsing data we leveraged in this paper provide an incredibly rich and valuable portrait of people's online information consumption behaviors, these data contain sensitive information, and care must be taken to protect participant privacy during all stages of data collection and analysis [27]. In our case, we collected browsing data from participants using an enterprise-grade application (YouGov's Pulse software) administered by a professional survey company (YouGov), who preprocessed the URL-level data to remove sensitive identifying information from the data that were ultimately delivered to us. Participants who provided these

data provided informed consent, were compensated for their participation, and were free to cease participating in the study at any time. The research team received the browsing data from YouGov via secure transfer and all participant data was stored on an encrypted server and was only accessible to members of the research team. That said, our qualitative analysis indicated that — despite measures taken by the data provider to anonymize URLs — digital trace data contain a significant amount of both deeply private and personally-identifying information (e.g., Google Maps searches for specific addresses). While this data helped to situate each participant in rich personal contexts, we ensured that this personally-identifying information was not reproduced in notes prepared for content analysis. As interest in digital trace data as a way to measure human behaviors on the internet has grown, several publicly-available (e.g., Web Historian [62]) and custom-made (e.g., Google search results page tracker [83]; YouTube recommended videos tracker [19]) browser extensions have been developed which allow researchers to collect web browsing data. While these extensions improve the accessibility of digital trace data, care needs to be taken by researchers who ask participants to install any software to ensure that participants are informed about the data they are providing to researchers, that that data is collected and stored securely, and that participants' privacy is maintained during all parts of the data collection, processing, and analysis stages (for an in-depth discussion of these issues, see Munzert et al. [67] and Duncan & Camp [28]).

6 CONCLUSION

To our knowledge, our work is the first to document exposure to QAnon websites online. Our findings help bring context to ongoing policy debates, contribute to our understanding of QAnon content consumption compared to consumption of other types of nefarious online content (e.g., misinformation), and illuminate fruitful future directions for researchers of online behavior in general and radicalization in particular. While we find that relatively few Americans were exposed to QAnon websites, we do not want to dismiss this number. By extrapolation, our 3.7% exposure estimate means that over 9.5 million American adults were exposed to a QAnon website during the 2020 election. While some of these exposures may have been incidental or ineffectual, some exposures may have translated into belief, which could have led to participation in radical protest or even violence. Regardless of the number of people identifying as QAnon supporters, the conspiracy theory established a significant online presence that has translated into offline actions, such as the January 6th insurrection, poll-watching, and running for local offices. With the ease of political organization online, it will be essential for researchers to study how fringe actors can establish sizable infrastructure, reach new audiences, and mobilize offline. Understanding the social, technological, and political landscape that enables fringe movements to reach consumers is essential for protecting democracy.

Acknowledgements

The authors thank participants at the 2021 QAnonference and the 2023 International Communication Association conference, as well as the Stanford Social Media Lab, for feedback on this work. We also would like to thank the participants who provided the data to make this study possible. This project was funded by a MURI award no. W911NF-20-1-0252 and support from Stanford's Freeman Spogli Institute for International Studies. R.C.M. is supported by a Stanford Interdisciplinary Graduate Fellowship. R.D. is supported by graduate fellowship awards from Knight-Hennessy Scholars and Stanford Data Science at Stanford University. The

funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

References

- [1] Saharsh Agarwal, Uttara M. Ananthakrishnan, and Catherine E. Tucker. 2022. Content Moderation at the Infrastructure Layer: Evidence from Parler. <https://doi.org/10.2139/ssrn.4232871>
- [2] Zhila Aghajari. 2023. Adopting an Ecological Approach to Misinformation: Understanding the Broader Impacts on Online Communities. In *Companion Publication of the 2023 Conference on Computer Supported Cooperative Work and Social Computing (CSCW '23 Companion)*, October 14, 2023. Association for Computing Machinery, New York, NY, USA, 417–420. <https://doi.org/10.1145/3584931.3608915>
- [3] Zhila Aghajari, Eric P. S. Baumer, and Dominic DiFranzo. 2023. Reviewing Interventions to Address Misinformation: The Need to Expand Our Vision Beyond an Individualistic Focus. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1 (April 2023), 87:1–87:34. <https://doi.org/10.1145/3579520>
- [4] Syeda Zainab Akbar, Anmol Panda, Divyanshu Kukreti, Azhagu Meena, and Joyojeet Pal. 2021. Misinformation as a Window into Prejudice: COVID-19 and the Information Environment in India. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW3 (January 2021), 249:1–249:28. <https://doi.org/10.1145/3432948>
- [5] Max Aliapoulos, Emmi Bevensee, Jeremy Blackburn, Barry Bradlyn, Emiliano De Cristofaro, Gianluca Stringhini, and Savvas Zannettou. 2021. An Early Look at the Parler Online Social Network. <https://doi.org/10.48550/arXiv.2101.03820>
- [6] Bobby Allyn. 2020. Twitter Removes Thousands Of QAnon Accounts, Promises Sweeping Ban On The Conspiracy. NPR. Retrieved July 16, 2023 from <https://www.npr.org/2020/07/21/894014810/twitter-removes-thousands-of-qanon-accounts-promises-sweeping-ban-on-the-conspir>
- [7] Sacha Altay, Rasmus Kleis Nielsen, and Richard Fletcher. 2022. Quantifying the “infodemic”: People turned to trustworthy news outlets during the 2020 coronavirus pandemic. *J. Quant. Descr. Digit. Media* 2, (August 2022). <https://doi.org/10.51685/jqd.2022.020>
- [8] Scott L. Althaus, Anne M. Cizmar, and James G. Gimpel. 2009. Media Supply, Audience Demand, and the Geography of News Consumption in the United States. *Polit. Commun.* 26, 3 (July 2009), 249–277. <https://doi.org/10.1080/10584600903053361>
- [9] Ahmed Anwar, Haider Ilyas, Ussama Yaqub, and Salma Zaman. 2021. Analyzing QAnon on Twitter in Context of US Elections 2020: Analysis of User Messages and Profiles Using VADER and BERT Topic modeling. In *DG.O2021: The 22nd Annual International Conference on Digital Government Research (DG.O'21)*, June 09, 2021. Association for Computing Machinery, New York, NY, USA, 82–88. <https://doi.org/10.1145/3463677.3463718>
- [10] Ahmer Arif. 2018. Designing to Support Reflection on Values & Practices to Address Online Disinformation. In *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '18 Companion)*, October 30, 2018. Association for Computing Machinery, New York, NY, USA, 61–64. <https://doi.org/10.1145/3272973.3272974>
- [11] Ahmer Arif, John J. Robinson, Stephanie A. Stanek, Elodie S. Fichet, Paul Townsend, Zena Worku, and Kate Starbird. 2017. A Closer Look at the Self-Correcting Crowd: Examining Corrections in Online Rumors. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*, February 25, 2017. Association for Computing Machinery, New York, NY, USA, 155–168. <https://doi.org/10.1145/2998181.2998294>
- [12] Kevin Aslett, Zeve Sanderson, William Godel, Nathaniel Persily, Jonathan Nagler, and Joshua A. Tucker. 2023. Online searches to evaluate misinformation can increase its perceived veracity. *Nature* (December 2023), 1–9. <https://doi.org/10.1038/s41586-023-06883-y>
- [13] Christopher A. Bail, Lisa P. Argyle, Taylor W. Brown, John P. Bumpus, Haohan Chen, M. B. Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. 2018. Exposure to opposing views on social media can increase political polarization. *Proc. Natl. Acad. Sci.* 115, 37 (September 2018), 9216–9221. <https://doi.org/10.1073/pnas.1804840115>
- [14] Saumya Bhadani. 2021. Biases in Recommendation System. In *Proceedings of the 15th ACM Conference on Recommender Systems (RecSys '21)*, September 13, 2021. Association for Computing Machinery, New York, NY, USA, 855–859. <https://doi.org/10.1145/3460231.3473897>
- [15] Saumya Bhadani, Shun Yamaya, Alessandro Flammini, Filippo Menczer, Giovanni Luca Ciampaglia, and Brendan Nyhan. 2022. Political audience diversity and news reliability in algorithmic ranking. *Nat. Hum. Behav.* 6, 4 (April 2022), 495–505. <https://doi.org/10.1038/s41562-021-01276-5>
- [16] Shannon Bond. 2021. Unwelcome On Facebook And Twitter, QAnon Followers Flock To Fringe Sites. NPR. Retrieved October 20, 2022 from <https://www.npr.org/2021/01/31/962104747/unwelcome-on-facebook-twitter-qanon-followers-flock-to-fringe-sites>

- [17] Shelley Boulianne and Sangwon Lee. 2022. Conspiracy Beliefs, Misinformation, Social Media Platforms, and Protest Participation. *Media Commun.* 10, 4 (October 2022), 30–41. <https://doi.org/10.17645/mac.v10i4.5667>
- [18] Çiğdem Bozdağ and Suncem Koçer. 2022. Skeptical Inertia in the Face of Polarization: News Consumption and Misinformation in Turkey. *Media Commun.* 10, 2 (May 2022), 169–179. <https://doi.org/10.17645/mac.v10i2.5057>
- [19] Megan A. Brown, James Bisbee, Angela Lai, Richard Bonneau, Jonathan Nagler, and Joshua A. Tucker. 2022. Echo Chambers, Rabbit Holes, and Algorithmic Bias: How YouTube Recommends Content to Real Users. <https://doi.org/10.2139/ssrn.4114905>
- [20] Cody Buntain, Martin Innes, Tamar Mitts, and Jacob Shapiro. 2023. Cross-Platform Reactions to the Post-January 6 Deplatforming. *J. Quant. Descr. Digit. Media* 3, (March 2023). <https://doi.org/10.51685/jqd.2023.004>
- [21] Andreu Casas, Ericka Menchen-Trevino, and Magdalena Wojcieszak. 2022. Exposure to Extremely Partisan News from the Other Political Side Shows Scarce Boomerang Effects. *Polit. Behav.* (February 2022). <https://doi.org/10.1007/s11109-021-09769-9>
- [22] Xunru Che, Danaë Metaxa-Kakavouli, and Jeffrey T. Hancock. 2018. Fake News in the News: An Analysis of Partisan Coverage of the Fake News Phenomenon. In *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '18 Companion)*, October 30, 2018. Association for Computing Machinery, New York, NY, USA, 289–292. <https://doi.org/10.1145/3272973.3274079>
- [23] Kate Conger. 2020. Twitter Takedown Targets QAnon Accounts. *The New York Times*. Retrieved October 31, 2022 from <https://www.nytimes.com/2020/07/21/technology/twitter-bans-qanon-accounts.html>
- [24] Ross Dahlke, Ryan C. Moore, Danya Adib-Azpeitia, Johan Ugander, Matthew Gentzkow, and Jeffery T. Hancock. 2022. Understanding Social Media's Direct and Indirect Role in Guiding People to Good, Bad and Ugly Information Sources. March 2022. Norman, OK.
- [25] Edwin S. Dalmajer, Camilla L. Nord, and Duncan E. Astle. 2022. Statistical power for cluster analysis. *BMC Bioinformatics* 23, 1 (May 2022), 205. <https://doi.org/10.1186/s12859-022-04675-1>
- [26] Ignacio X. Domínguez, Rogelio E. Cardona-Rivera, James K. Vance, and David L. Roberts. 2016. The Mimesis Effect: The Effect of Roles on Player Choice in Interactive Narrative Role-Playing Games. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, May 07, 2016. ACM, San Jose California USA, 3438–3449. <https://doi.org/10.1145/2858036.2858141>
- [27] Susan Dumais, Robin Jeffries, Daniel M. Russell, Diane Tang, and Jaime Teevan. 2014. Understanding User Behavior Through Log Data and Analysis. In *Ways of Knowing in HCI*, Judith S. Olson and Wendy A. Kellogg (eds.). Springer, New York, NY, 349–372. https://doi.org/10.1007/978-1-4939-0378-8_14
- [28] John Duncan and L. Jean Camp. 2012. Conducting an Ethical Study of Web Traffic. 2012. . Retrieved January 16, 2024 from <https://www.usenix.org/conference/cset12/workshop-program/presentation/duncan>
- [29] Laura Edelson, Minh-Kha Nguyen, Ian Goldstein, Oana Goga, Damon McCoy, and Tobias Lauinger. 2021. Understanding engagement with U.S. (mis)information news sources on Facebook. In *Proceedings of the 21st ACM Internet Measurement Conference (IMC '21)*, November 02, 2021. Association for Computing Machinery, New York, NY, USA, 444–463. <https://doi.org/10.1145/3487552.3487859>
- [30] Adam M. Enders, Joseph E. Uscinski, Casey A. Klofstad, Stefan Wuchty, Michelle I. Seelig, John R. Funchion, Manohar N. Murthi, Kamal Premaratne, and Justin Stoler. 2022. Who Supports QAnon? A Case Study in Political Extremism. *J. Polit.* 84, 3 (July 2022), 1844–1849. <https://doi.org/10.1086/717850>
- [31] Kristen Engel, Yiqing Hua, Taixiang Zeng, and Mor Naaman. 2022. Characterizing Reddit Participation of Users Who Engage in the QAnon Conspiracy Theories. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW1 (April 2022), 53:1–53:22. <https://doi.org/10.1145/3512900>
- [32] Cristina Fiani, Robin Bretin, Mark McGill, and Mohamed Khamis. 2023. Big Buddy: A Simulated Embodied Moderating System to Mitigate Children's Reaction to Provocative Situations within Social Virtual Reality. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, April 19, 2023. ACM, Hamburg Germany, 1–7. <https://doi.org/10.1145/3544549.3585840>
- [33] Seth Flaxman, Sharad Goel, and Justin M. Rao. 2016. Filter Bubbles, Echo Chambers, and Online News Consumption. *Public Opin. Q.* 80, S1 (January 2016), 298–320. <https://doi.org/10.1093/poq/nfw006>
- [34] Peter L. Forberg. 2021. From the Fringe to the Fore: An Algorithmic Ethnography of the Far-Right Conspiracy Theory Group QAnon. *J. Contemp. Ethnogr.* (September 2021), 08912416211040560. <https://doi.org/10.1177/08912416211040560>
- [35] Peter L. Forberg. 2023. “No Cult Tells You to Think for Yourself”: Discursive Ideology and the Limits of Rationality in Conspiracy Theory QAnon. *Am. Behav. Sci.* 67, 5 (May 2023), 649–664. <https://doi.org/10.1177/00027642221091199>
- [36] Sheera Frenkel. 2020. The Rise and Fall of the ‘Stop the Steal’ Facebook Group. *The New York Times*. Retrieved October 31, 2022 from <https://www.nytimes.com/2020/11/05/technology/stop-the-steal-facebook-group.html>
- [37] R. Kelly Garrett and Brian E. Weeks. 2013. The promise and peril of real-time corrections to political misperceptions. In *Proceedings of the 2013 conference on Computer supported cooperative work (CSCW '13)*, February 23, 2013. Association for Computing Machinery, New York, NY, USA, 1047–1058. <https://doi.org/10.1145/2441776.2441895>

- [38] Rui Gaspar, Cláudia Pedro, Panos Panagiotopoulos, and Beate Seibt. 2016. Beyond positive or negative: Qualitative sentiment analysis of social media reactions to unexpected stressful events. *Comput. Hum. Behav.* 56, (March 2016), 179–191. <https://doi.org/10.1016/j.chb.2015.11.040>
- [39] Rumi Ghosh, Tawan Surachawala, and Kristina Lerman. 2011. Entropy-based Classification of “Retweeting” Activity on Twitter. Retrieved January 15, 2024 from <https://arxiv.org/abs/1106.0346v1>
- [40] Teresa Gil-López, Clara Christner, Ernesto de León, Mykola Makhortykh, Aleksandra Urman, Michaela Maier, and Silke Adam. 2023. Do (Not!) Track Me: Relationship Between Willingness to Participate and Sample Composition in Online Information Behavior Tracking Research. *Soc. Sci. Comput. Rev.* (February 2023), 08944393231156634. <https://doi.org/10.1177/08944393231156634>
- [41] José González Cabañas, Ángel Cuevas, and Rubén Cuevas. 2017. FDVT: Data Valuation Tool for Facebook Users. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, May 02, 2017. ACM, Denver Colorado USA, 3799–3809. <https://doi.org/10.1145/3025453.3025903>
- [42] Sandra González-Bailón and Michael A. Xenos. 2022. The blind spots of measuring online news exposure: a comparison of self-reported and observational data in nine countries. *Inf. Commun. Soc.* 0, 0 (May 2022), 1–19. <https://doi.org/10.1080/1369118X.2022.2072750>
- [43] Andrew Guess. 2021. (Almost) Everything in Moderation: New Evidence on Americans’ Online Media Diets. *Am. J. Polit. Sci.* 65, 4 (2021), 1007–1022. <https://doi.org/10.1111/ajps.12589>
- [44] Andrew Guess, Benjamin Lyons, Jacob M. Montgomery, Brendan Nyhan, and Jason Reifler. 2019. Fake news, Facebook ads, and misperceptions: Assessing information quality in the 2018 U.S. midterm election campaign.
- [45] Andrew M. Guess, Brendan Nyhan, and Jason Reifler. 2020. Exposure to untrustworthy websites in the 2016 US election. *Nat. Hum. Behav.* 4, 5 (May 2020), 472–480. <https://doi.org/10.1038/s41562-020-0833-x>
- [46] Andrew Guess, Kevin Munger, Jonathan Nagler, and Joshua Tucker. 2019. How Accurate Are Survey Responses on Social Media and Politics? *Polit. Commun.* 36, 2 (April 2019), 241–258. <https://doi.org/10.1080/10584609.2018.1504840>
- [47] Katherine Haenschen. 2020. Self-Reported Versus Digitally Recorded: Measuring Political Activity on Facebook. *Soc. Sci. Comput. Rev.* 38, 5 (October 2020), 567–583. <https://doi.org/10.1177/0894439318813586>
- [48] Kyungsik Han, Patrick C. Shih, Mary Beth Rosson, and John M. Carroll. 2014. Enhancing community awareness of and participation in local heritage with a mobile application. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing (CSCW ’14)*, February 15, 2014. Association for Computing Machinery, New York, NY, USA, 1144–1155. <https://doi.org/10.1145/2531602.2531640>
- [49] Hans W. A. Hanley, Deepak Kumar, and Zakir Durumeric. 2022. No Calm in the Storm: Investigating QAnon Website Relationships. *Proc. Int. AAAI Conf. Web Soc. Media* 16, (May 2022), 299–310.
- [50] Y. Linlin Huang, Kate Starbird, Mania Orand, Stephanie A. Stanek, and Heather T. Pedersen. 2015. Connected Through Crisis: Emotional Proximity and the Spread of Misinformation Online. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW ’15)*, February 28, 2015. Association for Computing Machinery, New York, NY, USA, 969–980. <https://doi.org/10.1145/2675133.2675202>
- [51] Sam Jackson, Brandon Gorman, and Mayuko Nakatsuka. 2021. QAnon on Twitter: An Overview. The George Washington University. Retrieved from <https://iddp.gwu.edu/sites/g/files/zaxdzs3576/f/downloads/QAnon%20on%20Twitter%3B%20Jackson.pdf>
- [52] Shagun Jhaver, Christian Boylston, Diyi Yang, and Amy Bruckman. 2021. Evaluating the Effectiveness of Deplatforming as a Moderation Strategy on Twitter. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (October 2021), 381:1–381:30. <https://doi.org/10.1145/3479525>
- [53] Perna Juneja. 2021. Identifying, Measuring and Contesting Algorithmically Curated Misinformation. In *Companion Publication of the 2021 Conference on Computer Supported Cooperative Work and Social Computing (CSCW ’21 Companion)*, October 23, 2021. Association for Computing Machinery, New York, NY, USA, 256–258. <https://doi.org/10.1145/3462204.3481788>
- [54] Brian C. Keegan, Shakked Lev, and Ofer Arazy. 2016. Analyzing Organizational Routines in Online Knowledge Collaborations: A Case for Sequence Analysis in CSCW. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW ’16)*, February 27, 2016. Association for Computing Machinery, New York, NY, USA, 1065–1079. <https://doi.org/10.1145/2818048.2819962>
- [55] Sojong Kim and Jisu Kim. 2023. The Information Ecosystem of Conspiracy Theory: Examining the QAnon Narrative on Facebook. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1 (April 2023), 150:1–150:24. <https://doi.org/10.1145/3579626>
- [56] Naimisha Kolli, N. Balakrishnan, and K. R. Ramakrishnan. 2017. On Quantifying Predictability in Online Social Media Cascades Using Entropy. In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017 (ASONAM ’17)*, July 31, 2017. Association for Computing Machinery, New York, NY, USA, 109–114. <https://doi.org/10.1145/3110025.3110071>
- [57] Loukas Konstantinou and Evangelos Karapanos. 2023. Nudging for Online Misinformation: a Design Inquiry. In *Companion Publication of the 2023 Conference on Computer Supported Cooperative Work and Social Computing*

- (CSCW '23 Companion), October 14, 2023. Association for Computing Machinery, New York, NY, USA, 69–75. <https://doi.org/10.1145/3584931.3607015>
- [58] Heera Lee and Andrea Kleinsmith. 2019. Public Speaking Anxiety in a Real Classroom: Towards Developing a Reflection System. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, May 02, 2019. ACM, Glasgow Scotland Uk, 1–6. <https://doi.org/10.1145/3290607.3312875>
 - [59] Gionnieve Lim and Simon Tangi Perrault. 2021. Local Perceptions and Practices of News Sharing and Fake News. In *Companion Publication of the 2021 Conference on Computer Supported Cooperative Work and Social Computing (CSCW '21 Companion)*, October 23, 2021. Association for Computing Machinery, New York, NY, USA, 117–120. <https://doi.org/10.1145/3462204.3481767>
 - [60] Mykola Makhortykh, Claes de Vreese, Natali Helberger, Jaron Harambam, and Dimitrios Bountouridis. 2021. We are what we click: Understanding time and content-based habits of online news readers. *New Media Soc.* 23, 9 (September 2021), 2773–2800. <https://doi.org/10.1177/1461444820933221>
 - [61] Alice E Marwick and William Clyde Partin. 2022. Constructing alternative facts: Populist expertise and the QAnon conspiracy. *New Media Soc.* (May 2022), 14614448221090201. <https://doi.org/10.1177/14614448221090201>
 - [62] Ericka Menchen-Trevino. 2016. Web historian: enabling multi-method and independent research with real-world web browsing history data. *ICConference 2016 Proc.* (March 2016). <https://doi.org/10.9776/16611>
 - [63] Corrado Monti, Matteo Cinelli, Carlo Valensise, Walter Quattrociochi, and Michele Starnini. 2023. Online conspiracy communities are more resilient to deplatforming. <https://doi.org/10.48550/arXiv.2303.12115>
 - [64] Ryan C. Moore, Ross Dahlke, and Jeffrey T. Hancock. 2023. Exposure to untrustworthy websites in the 2020 US election. *Nat. Hum. Behav.* (April 2023), 1–10. <https://doi.org/10.1038/s41562-023-01564-2>
 - [65] Michael Muller. 2012. Lurking as personal trait or situational disposition: lurking and contributing in enterprise social media. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work (CSCW '12)*, February 11, 2012. Association for Computing Machinery, New York, NY, USA, 253–256. <https://doi.org/10.1145/2145204.2145245>
 - [66] Simon Munzert, Sebastian Ramirez-Ruiz, Pablo Barberá, Andrew M. Guess, and JungHwan Yang. 2022. Who's cheating on your survey? A detection approach with digital trace data. *Polit. Sci. Res. Methods* (November 2022), 1–9. <https://doi.org/10.1017/psrm.2022.42>
 - [67] Simon Munzert, Sebastian Ramirez-Ruiz, Oliver Watteler, Johannes Breuer, Veronika Batzdorfer, Christina Eder, Deborah A. Wiltshire, Pablo Barberá, Andrew Guess, and JungHwan Yang. 2023. Publishing Combined Web Tracking and Survey Data. <https://doi.org/10.31219/osf.io/y4v8z>
 - [68] Shuo Niu and Kathleen Palm Reed. 2023. Towards Understanding Substance Abuse Misinformation in YouTube Videos. In *Companion Publication of the 2023 Conference on Computer Supported Cooperative Work and Social Computing (CSCW '23 Companion)*, October 14, 2023. Association for Computing Machinery, New York, NY, USA, 64–68. <https://doi.org/10.1145/3584931.3606974>
 - [69] Shuo Niu, Keegan Veazey, Phoenix Pagan, and Abhisar Ghimire. 2022. Understanding Hate Group Videos on YouTube. In *Companion Publication of the 2022 Conference on Computer Supported Cooperative Work and Social Computing (CSCW'22 Companion)*, November 08, 2022. Association for Computing Machinery, New York, NY, USA, 32–36. <https://doi.org/10.1145/3500868.3559465>
 - [70] Brendan Nyhan and Jason Reifler. 2010. When Corrections Fail: The Persistence of Political Misperceptions. *Polit. Behav.* 32, 2 (June 2010), 303–330. <https://doi.org/10.1007/s11109-010-9112-2>
 - [71] Antonis Papasavva, Max Aliapoulos, Cameron Ballard, Emiliano De Cristofaro, Gianluca Stringhini, Savvas Zannettou, and Jeremy Blackburn. 2022. The Gospel according to Q: Understanding the QAnon Conspiracy from the Perspective of Canonical Information. *Proc. Int. AAAI Conf. Web Soc. Media* 16, (May 2022), 735–746.
 - [72] Antonis Papasavva, Jeremy Blackburn, Gianluca Stringhini, Savvas Zannettou, and Emiliano De Cristofaro. 2021. “Is it a Qoincidence?": An Exploratory Study of QAnon on Voat. In *Proceedings of the Web Conference 2021 (WWW '21)*, April 19, 2021. Association for Computing Machinery, New York, NY, USA, 460–471. <https://doi.org/10.1145/3442381.3450036>
 - [73] Douglas A. Parry, Brittany I. Davidson, Craig J. R. Sewall, Jacob T. Fisher, Hannah Mieczkowski, and Daniel S. Quintana. 2021. A systematic review and meta-analysis of discrepancies between logged and self-reported digital media use. *Nat. Hum. Behav.* 5, 11 (November 2021), 1535–1547. <https://doi.org/10.1038/s41562-021-01117-5>
 - [74] Irene V. Pasquetto, Alberto F. Olivieri, Luca Tacchetti, Gianni Riotta, and Alessandra Spada. 2022. Disinformation as Infrastructure: Making and Maintaining the QAnon Conspiracy on Italian Digital Media. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW1 (April 2022), 84:1–84:31. <https://doi.org/10.1145/3512931>
 - [75] pew Pew Research Center. 2018. 10. Political engagement, knowledge and the midterms. Pew Research Center - U.S. Politics & Policy. Retrieved November 6, 2020 from <https://www.pewresearch.org/politics/2018/04/26/10-political-engagement-knowledge-and-the-midterms/>
 - [76] Jan-Willem van Prooijen, André P. M. Krouwel, and Thomas V. Pollet. 2015. Political Extremism Predicts Belief in Conspiracy Theories. *Soc. Psychol. Personal. Sci.* 6, 5 (July 2015), 570–578. <https://doi.org/10.1177/1948550614567356>

- [77] PRRI. 2021. Understanding QAnon's Connection to American Politics, Religion, and Media Consumption. PRRI. Retrieved October 20, 2022 from <https://www.prii.org/research/qanon-conspiracy-american-politics-report/>
- [78] Will Rahn and Dan Patterson. 2021. What is QAnon? What does WWG1WGA mean? The conspiracy theory that explains everything and nothing. CBS News. Retrieved July 17, 2023 from <https://www.cbsnews.com/news/what-is-the-qanon-conspiracy-theory/>
- [79] Byron Reeves, Thomas Robinson, and Nilam Ram. 2020. Time for the Human Screenome Project. *Nature* 577, 7790 (January 2020), 314–317. <https://doi.org/10.1038/d41586-020-00032-5>
- [80] Melanie Revilla, Carlos Ochoa, and Germán Loewe. 2017. Using Passive Data From a Meter to Complement Survey Data in Order to Study Online Behavior. *Soc. Sci. Comput. Rev.* 35, 4 (August 2017), 521–536. <https://doi.org/10.1177/0894439316638457>
- [81] Manoel Horta Ribeiro, Shagun Jhaver, Jordi Cluet i Martinell, Marie Reigner-Tayar, and Robert West. 2024. Deplatforming Norm-Violating Influencers on Social Media Reduces Overall Online Attention Toward Them. <https://doi.org/10.48550/arXiv.2401.01253>
- [82] Manoel Horta Ribeiro, Raphael Ottoni, Robert West, Virgílio A. F. Almeida, and Wagner Meira. 2020. Auditing radicalization pathways on YouTube. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*, January 27, 2020. Association for Computing Machinery, New York, NY, USA, 131–141. <https://doi.org/10.1145/3351095.3372879>
- [83] Ronald E. Robertson, Jon Green, Damian J. Ruck, Katherine Ognyanova, Christo Wilson, and David Lazer. 2023. Users choose to engage with more partisan news than they are exposed to on Google Search. *Nature* 618, 7964 (June 2023), 342–348. <https://doi.org/10.1038/s41586-023-06078-5>
- [84] Marten Scheffer, Denny Borsboom, Sander Nieuwenhuis, and Frances Westley. 2022. Belief traps: Tackling the inertia of harmful beliefs. *Proc. Natl. Acad. Sci.* 119, 32 (August 2022), e2203149119. <https://doi.org/10.1073/pnas.2203149119>
- [85] Craig Silverman, Craig Timberg, Jeff Kao, and Jeremy B. Merrill. 2022. Facebook groups topped 10,000 daily attacks on election before Jan. 6, analysis shows. *Washington Post*. Retrieved October 31, 2022 from <https://www.washingtonpost.com/technology/2022/01/04/facebook-election-misinformation-capitol-riot/>
- [86] Andrea Sipka, Aniko Hannak, and Aleksandra Urman. 2022. Comparing the Language of QAnon-Related Content on Parler, Gab, and Twitter. In *14th ACM Web Science Conference 2022 (WebSci '22)*, June 26, 2022. Association for Computing Machinery, New York, NY, USA, 411–421. <https://doi.org/10.1145/3501247.3531550>
- [87] Robert Soden, David Ribes, Seyram Avle, and Will Sutherland. 2021. Time for Historicism in CSCW: An Invitation. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (October 2021), 459:1–459:18. <https://doi.org/10.1145/3479603>
- [88] Kate Starbird, Ahmer Arif, and Tom Wilson. 2019. Disinformation as Collaborative Work: Surfacing the Participatory Nature of Strategic Information Operations. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW (November 2019), 127:1–127:26. <https://doi.org/10.1145/3359229>
- [89] Sebastian Stier, Frank Mangold, Michael Scharrow, and Johannes Breuer. 2022. Post Post-Broadcast Democracy? News Exposure in the Age of Online Intermediaries. *Am. Polit. Sci. Rev.* 116, 2 (May 2022), 768–774. <https://doi.org/10.1017/S0003055421001222>
- [90] Jack Thompson and Sierra Davis Thomander. 2021. What drives support for QAnon? Evidence from a survey experiment. <https://doi.org/10.31219/osf.io/23qaj>
- [91] Yu-Chia Tseng, Nanyi Bi, Yung-Ju Chang, and Chien Wen (Tina) Yuan. 2022. Investigating Correction Effects of Different Modalities for Misinformation about COVID-19. In *Companion Publication of the 2022 Conference on Computer Supported Cooperative Work and Social Computing (CSCW'22 Companion)*, November 08, 2022. Association for Computing Machinery, New York, NY, USA, 54–58. <https://doi.org/10.1145/3500868.3559455>
- [92] Joseph E. Uscinski. 2022. Getting QAnon Wrong and Right. *Soc. Res. Int. Q.* 89, 3 (2022), 551–578. <https://doi.org/10.1353/sor.2022.0038>
- [93] Emily L. Wang, Luca Luceri, Francesco Pierri, and Emilio Ferrara. 2022. Identifying and Characterizing Behavioral Classes of Radicalization within the QAnon Conspiracy on Twitter. <https://doi.org/10.48550/arXiv.2209.09339>
- [94] Brian E Weeks, Ericka Menchen-Trevino, Christopher Calabrese, Andreu Casas, and Magdalena Wojcieszak. 2021. Partisan media, untrustworthy news sites, and political misperceptions. *New Media Soc.* (July 2021), 14614448211033300. <https://doi.org/10.1177/14614448211033300>
- [95] Eoin Whelan, Robin Teigland, Emmanuelle Vaast, and Brian Butler. 2016. Expanding the horizons of digital social networks: Mixing big trace datasets with qualitative approaches. *Inf. Organ.* 26, 1 (June 2016), 1–12. <https://doi.org/10.1016/j.infoandorg.2016.03.001>
- [96] Tom Wilson and Kate Starbird. 2020. Cross-platform disinformation campaigns: Lessons learned and next steps. *Harv. Kennedy Sch. Misinformation Rev.* 1, 1 (January 2020). <https://doi.org/10.37016/mr-2020-002>
- [97] Tom Wilson and Kate Starbird. 2021. Cross-platform Information Operations: Mobilizing Narratives & Building Resilience through both “Big” & “Alt” Tech. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (October 2021), 345:1–345:32. <https://doi.org/10.1145/3476086>

- [98] Magdalena Wojcieszak, Sjifra de Leeuw, Ericka Menchen-Trevino, Seungsu Lee, Ke M. Huang-Isherwood, and Brian Weeks. 2021. No Polarization From Partisan News: Over-Time Evidence From Trace Data. *Int. J. Press.* (November 2021), 19401612211047194. <https://doi.org/10.1177/19401612211047194>
- [99] Wentao Xu and Kazutoshi Sasahara. 2022. A Network-Based Approach to QAnon User Dynamics and Topic Diversity During the COVID-19 Infodemic. *APSIPA Trans. Signal Inf. Process.* 11, 1 (June 2022). <https://doi.org/10.1561/116.000000055>
- [100] Leo Yeykelis, James J. Cummings, and Byron Reeves. 2018. The Fragmentation of Work, Entertainment, E-Mail, and News on a Personal Computer: Motivational Predictors of Switching Between Media Content. *Media Psychol.* 21, 3 (July 2018), 377–402. <https://doi.org/10.1080/15213269.2017.1406805>
- [101] YouGov. 2022. Conspiracy Theories | Economist/YouGov Poll: August 7 - 9, 2022 | YouGov. Retrieved October 25, 2022 from <https://today.yougov.com/topics/politics/articles-reports/2022/08/10/conspiracy-theory-economist-yougov-poll-august-7>
- [102] J. Zaller. 1992. *The nature and origins of mass opinion*. Cambridge University Press, Cambridge, England.
- [103] Saifelddeen Zihiri, Gabriel Lima, Jiyoung Han, Meeyoung Cha, and Wonjae Lee. 2022. QAnon shifts into the mainstream, remains a far-right ally. *Heliyon* 8, 2 (February 2022), e08764. <https://doi.org/10.1016/j.heliyon.2022.e08764>
- [104] Ethan Zuckerman. 2019. QAnon and the Emergence of the Unreal. *J. Des. Sci.* 6 (July 2019). <https://doi.org/10.21428/7808da6b.6b8a82b9>

Received January 2024; revised April 2024; accepted May 2024.