

To: UltraViolet, GLAAD, Kairos, and Women's March
From: YouGov
Re: Hate Speech & Misinformation Survey Findings

UltraViolet, GLAAD, Kairos, and Women's March commissioned YouGov to conduct a study on attitudes about hate speech, harassment, and misinformation among Americans.¹ The survey was conducted online from July 7th to July 22th, 2022. This document summarizes the project's key findings.²

Key takeaways of this report

The broad online experience

- Respondents are broadly positive about a variety of aspects of the online experience – from the internet providing them a way to stay in touch with family and friends, space to pursue their hobbies, or have a voice.
- However, the plurality of respondents across all comparison groups are ambivalent about the risks of being online as well as on the internet's impact on mental health.

False information

- Respondents think that false information is a problem in online spaces and would like to hold platforms accountable for their role in promoting false information.
- A majority of Americans report that they frequently run into false information especially about politics, COVID-19, and public officials.
- Most respondents believe that false information leaves the public confused about the basic facts and impacts trust in the government and in one another, showing that the deleterious impacts of misinformation could have both immediate and downstream effects.
- Respondents see false information as an issue across multiple platforms, with Facebook and Twitter being seen as the most problematic.
- A very strong majority believe that platforms should be held accountable for various types of false information in their sites – about 9 out of 10 women and LGBTQ+ respondents agree with platform accountability when it comes to false information, while 8 out of 10 people of color do the same.

Discursive incivility and hate speech

- Discursive incivility and hate speech are both generally regarded as problematic in online spaces, even as large proportions of respondents report having witnessed specific disrespectful or hateful speech online.

¹ This survey consisted of a nationally-representative base sample (n=1,235), in addition to oversamples of people of color, women, and LGBTQ+ respondents (n=404). For weighting and analysis, the oversamples of people of color and women were combined with the responses from the base sample, yielding a sample size each of n=754 people of color and n=1134 women. For more details, see the methods statements at the end of this report.

² Making additional comparisons with the graphs and data in this report may not be statistically significant. Please contact info@weareultraviolet.org with questions about the data.

- Individuals who are LGBTQ+, POC, or women, are more likely to report having witnessed hate against people with their same identity.
- Respondents are critical of how platforms have addressed hate speech on their sites.

Harassment

- About one in five Americans report encountering harassment in online spaces either very frequently or frequently, while about one in three LGBTQ+ respondents report the same.
- Strong majorities of respondents say that society thinks it is socially unacceptable to harass respondents based on a host of individual characteristics. However, about a third of Americans say that it is socially acceptable to harass people online based on political opinions, views about current events, and COVID-19.
- A plurality of respondents think that online harassment is a major problem on Facebook and Twitter, a finding that mirrors the previous findings on false information and hate speech
- Nearly one in three Americans overall, women, and POC respondents said that social media platforms are doing a poor job at addressing online harassment on their sites, while almost two in five LGBTQ+ respondents said the same.
- Witnessing others be harassed online is a widespread experience among Americans overall, but LGBTQ+ respondents report seeing the same behaviors at increased rates. However, when asked to reflect about this, either pluralities or majorities of Americans across groups believe that harassment is more widespread than not.
- People aren't just harmed from direct harassment, they feel attacked when public figures are attacked based on their shared identity. This means, substantively speaking, that the problem of online harassment is not only one that affects the *victims* of harassment themselves, but the *witnesses* of this harassment.
- While a majority of Americans in the base sample have not been subject to harassment themselves, LGBTQ+ folk and women report higher rates of harassment than their respective comparison groups. This means that personal identity appears to play an important role in how one experiences harassment.

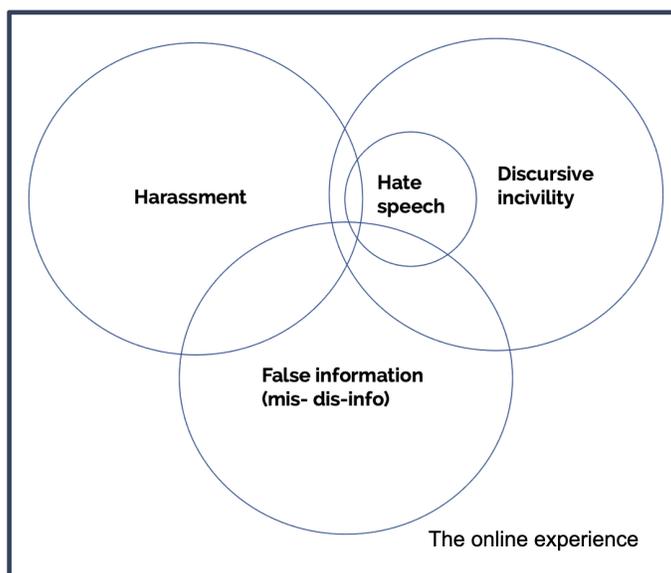
Solutions

- Repercussions and platform accountability performs the best across both the base sample and demographic oversamples, followed by platform-specific policies like improving content moderation or enforcing or creating new policies, and trailed by entirely banning hateful language, violence and disinformation from online platforms.
- Upon completing the survey, respondents grew more supportive of government regulating platforms than they were at the very start of the survey. This is particularly the case among POC, women, and LGBTQ+ respondents.

Online harassment is a growing issue in America today. At the same time, not all harassment is created equal—individuals with marginalized gender and racial identities tend to experience online harassment and hate speech at increased levels. The intersection of race

and gender is a common target of online harassers—scholarly research³ conducted using Twitter data showcases that derogatory messages are routinely directed toward BIPOC women. Deemed “intersectional aggression” by these scholars, this phenomenon is quite troubling given that research conducted in Norway⁴ also suggests that women are more likely than men to change their behavior online as a result of experiencing harassment. So not only do women, particularly BIPOC women, face a hostile experience online, but they also could respond to online harassment differently than men do.

If one aims to center the experience of individuals with marginalized identities, particularly gender and racial ones, in the work to help them experience the internet safely, then this is a phenomenon that ought to be thoroughly understood and quantified. This study uses a survey to help UltraViolet and partners understand the pervasiveness of online harassment, discursive incivility, and false information. As Figure 1 shows, these are interrelated phenomena that deserve their own in-depth exploration in order to get a thorough understanding of the broader online experience.



UltraViolet, GLAAD, Kairos, and Women’s March commissioned YouGov to study these interrelated phenomena using a national survey of social media users in the United States. Four distinct samples comprise this study:

- (1) a nationally-representative sample of social media users comprised of 1,235 respondents (referred to as the “Base Sample” throughout this report),
- (2) an oversample of self-identified people of color who are social media users
- (3) an oversample of self-identified LGBTQ+ respondents who are social media users (n=404), and
- (4) an oversample of self-identified women who are social media users.

³ Francisco, Sara C., and Diane H. Felmler. “What Did You Call Me? An Analysis of Online Harassment Towards Black and Latinx Women.” *Race and Social Problems* (2021): 1-13.

⁴ Nadim, Marjan, and Audun Fladmoe. “Silencing women? Gender and online harassment.” *Social Science Computer Review* 39.2 (2021): 245-258.

This study focuses on three broad comparisons in order to obtain clear and substantive insights about the online experience:

- LGBTQ+ oversample that gets compared to the entire base sample.
- All self-identified POC respondents⁵ (from OS + base sample) are compared to White respondents in the base sample.⁶
- All self-identified women respondents⁷ (from OS + base sample) to men respondents in the base sample.⁸

Respondents are ambivalent about the personal risks of being online, and view their own online experience positively

This survey asked respondents to evaluate a series of statements⁹ aimed at understanding their broad experience of being online.¹⁰ **Overall attitudes about the online experience tend to be positive.** A majority of respondents in the base sample said the internet helps them stay in touch with their social circle, report having a good balance in their online and offline lives, and that the internet provides them space for their hobbies. Specifically, three out of four respondents in the base sample said that the internet helps them stay in touch with family and friends, while two out of three in the same sample said they believe they have a good balance between their online and offline lives. **In none of these items or samples do respondents' disagreement outweigh their agreement.**

A worthwhile point is LGBTQ+ respondents feel that the internet provides them a space to be themselves and have a voice. While 36 percent of the base sample agrees with the

⁵ The survey weights combined the POC respondents from both the oversample *and* the base sample and reweighted this combination to a nationally-representative frame of POCs. More details are in the methods statement in the appendix of this report.

⁶ White respondents are subset from the overall base sample for the purposes of this comparison.

⁷ The survey weights combined the women respondents from both the oversample *and* the base sample and reweighted this combination to a nationally-representative frame of women. More details are in the methods statement in the appendix of this report.

⁸ Men respondents are subset from the overall base sample for the purposes of this comparison.

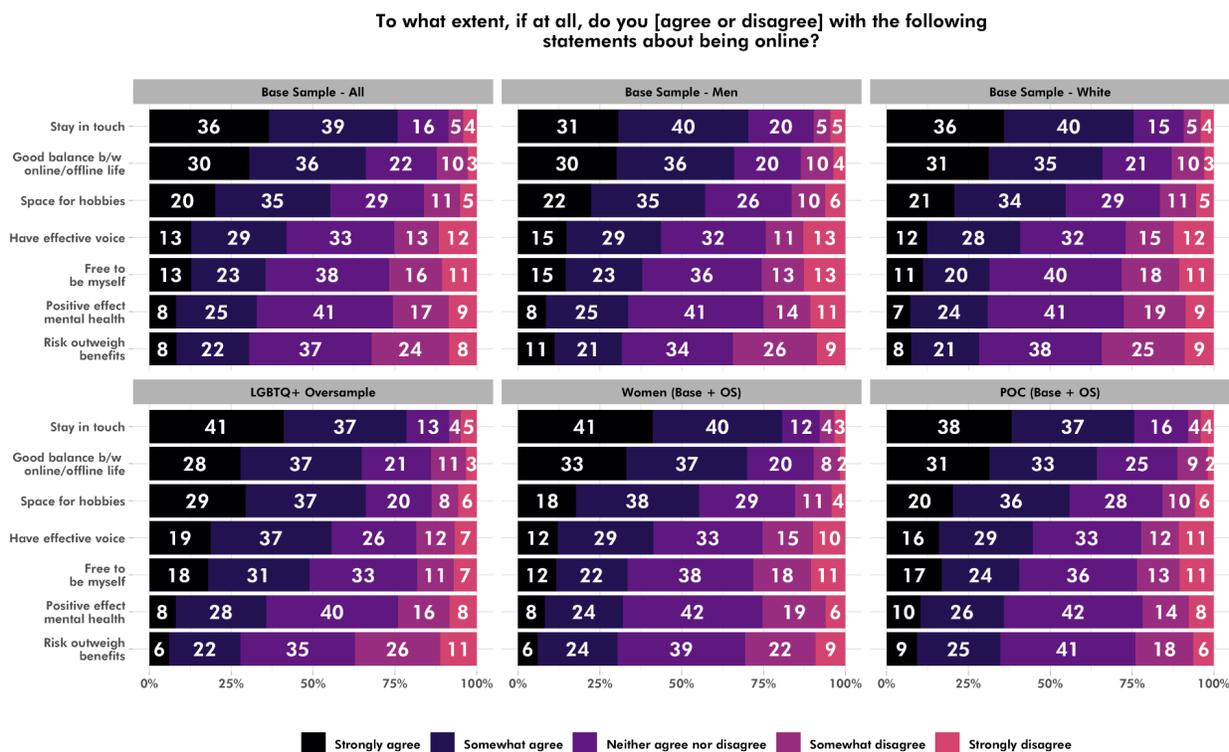
⁹ The full question read: To what extent, if at all, do you [agree or disagree] with the following statements about being online?

- I feel I have a good balance between my online and offline life
 - It gives me space to pursue my hobbies and interests in a way I couldn't do offline
 - I can share my opinions and have a voice online more easily or effectively than I can offline
 - Being online has an overall positive effect on my mental health
 - I feel more free to be myself online
 - It helps me to keep in touch with my friends and family
 - When it comes to being online, the risks outweigh the benefits
- <1> Strongly agree
<2> Somewhat agree
<3> Neither agree nor disagree
<4> Somewhat disagree
<5> Strongly disagree

¹⁰ This battery is a replication of a few items from a social media battery in the [2022 Online Nation report from Ofcom](#) (pg. 53, the online experience). That report is conducted among a national sample in the United Kingdom of respondents 13+ years old, so it is not directly comparable to our current study.

statement “I feel more free to be myself online”, 49 percent of LGBTQ+ people agreed with it. Similarly, while 43 percent of the base sample agreed that they can have a voice online more easily than they can offline, 56 percent of LGBTQ+ respondents said the same.

Conversely, **the plurality of respondents are ambivalent about the risks of being online as well as on the internet’s impact on mental health.** Overall, 37 percent of respondents in the base sample, 35 percent of LGBTQ+ respondents, 39 percent of women, and 40 percent of people of color said they neither agree nor disagree that the risks of being online outweigh the benefits. When it comes to the internet’s impact on mental health, 41 percent of respondents in the base sample, 40 percent of LGBTQ+ respondents, 42 percent of women, and 42 percent of people of color said they neither agree nor disagree with that statement.



Key Takeaways

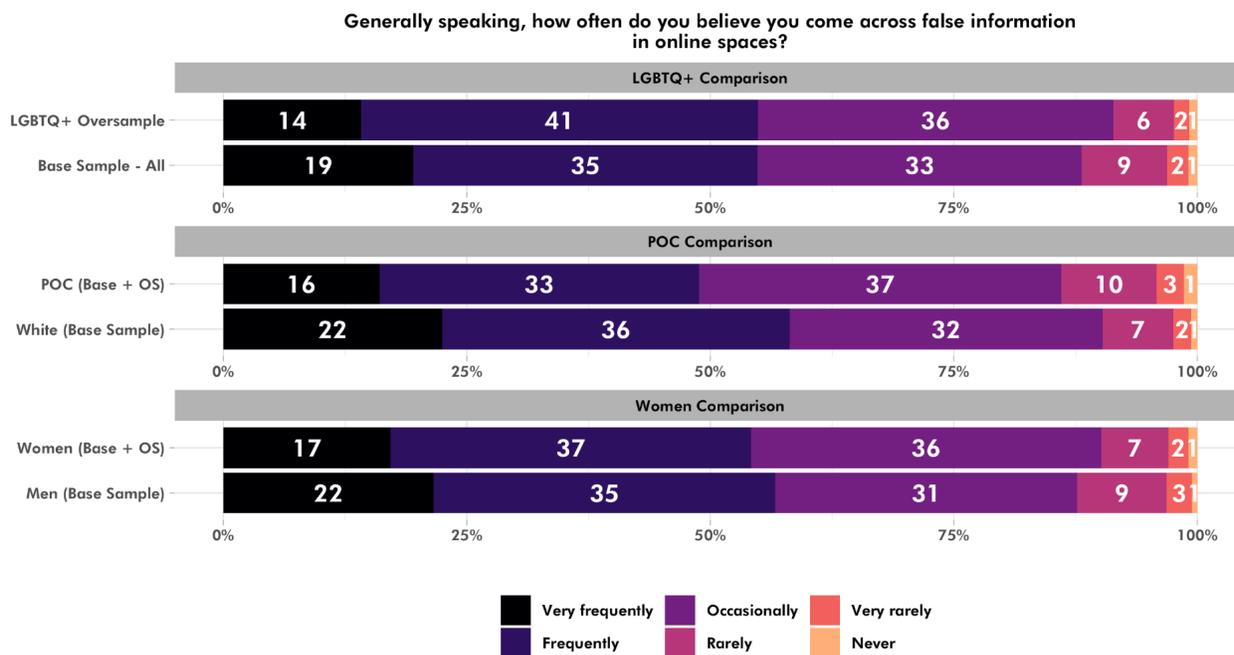
- Respondents are broadly positive about a variety of aspects of the online experience – from the internet providing them a way to stay in touch with family and friends, space to pursue their hobbies, or have a voice. LGBTQ+ respondents feel that the internet gives them a space to be themselves and allows them to have an effective voice.
- The plurality of respondents across all comparison groups are ambivalent about the risks of being online as well as on the internet’s impact on mental health.

Respondents' Experiences with False Information Online

This survey included several questions aimed at understanding respondents' experiences with witnessing false information in digital spaces.¹¹ While extensive scholarly work has already started to try to understand how widespread false information is online, this work is not necessarily evergreen—what narratives and problematic information are prevalent at a given point in time in digital spaces changes extremely quickly. Today's false information is, thus, different from tomorrow's. As a result, this survey included a host of questions aimed at measuring the public's *perceptions* of false information rather than its *prevalence* per se.

One such question read "Generally speaking, how often do you believe you come across false information in online spaces?". The response options ranged from "very frequently" to "never", in a 6 point scale.

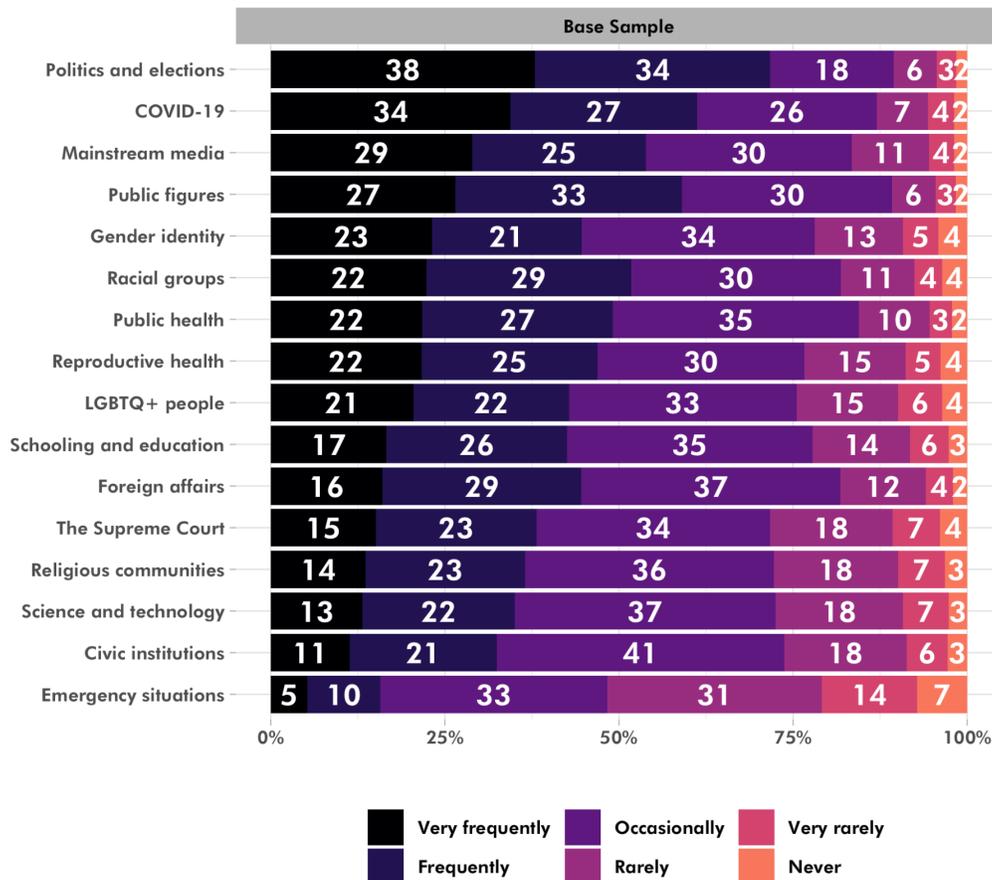
Americans think that false information is pervasive on the internet. A majority of Americans said that they come across false information in online spaces frequently, but about one in ten said they rarely or never come across it. Overall, 54 percent of respondents in the base sample report coming across this information frequently, while 55 percent of LGBTQ+ folk, 49 percent of POC, and 54 percent of women said the same. Conversely, about one out of ten respondents in each of the four distinct comparison groups of interest said they rarely or never come across false information online.



¹¹ This survey deliberately used the term "false information" in this survey instrument, as research has shown that Republicans and Democrats use the terms "misinformation", "disinformation", and "fake news" in different ways (Bello-Pardo, 2022). In order to avoid the potential for these terms to steer responses one way or the other – through priming partisanship or other potential mechanisms – this instrument used the term "false information" to encompass the phenomenon that UltraViolet is interested in. Additionally, the Pew Research Center's research (examples here, here, and here) has used the term "false information" in their polling for arguably similar reasons.

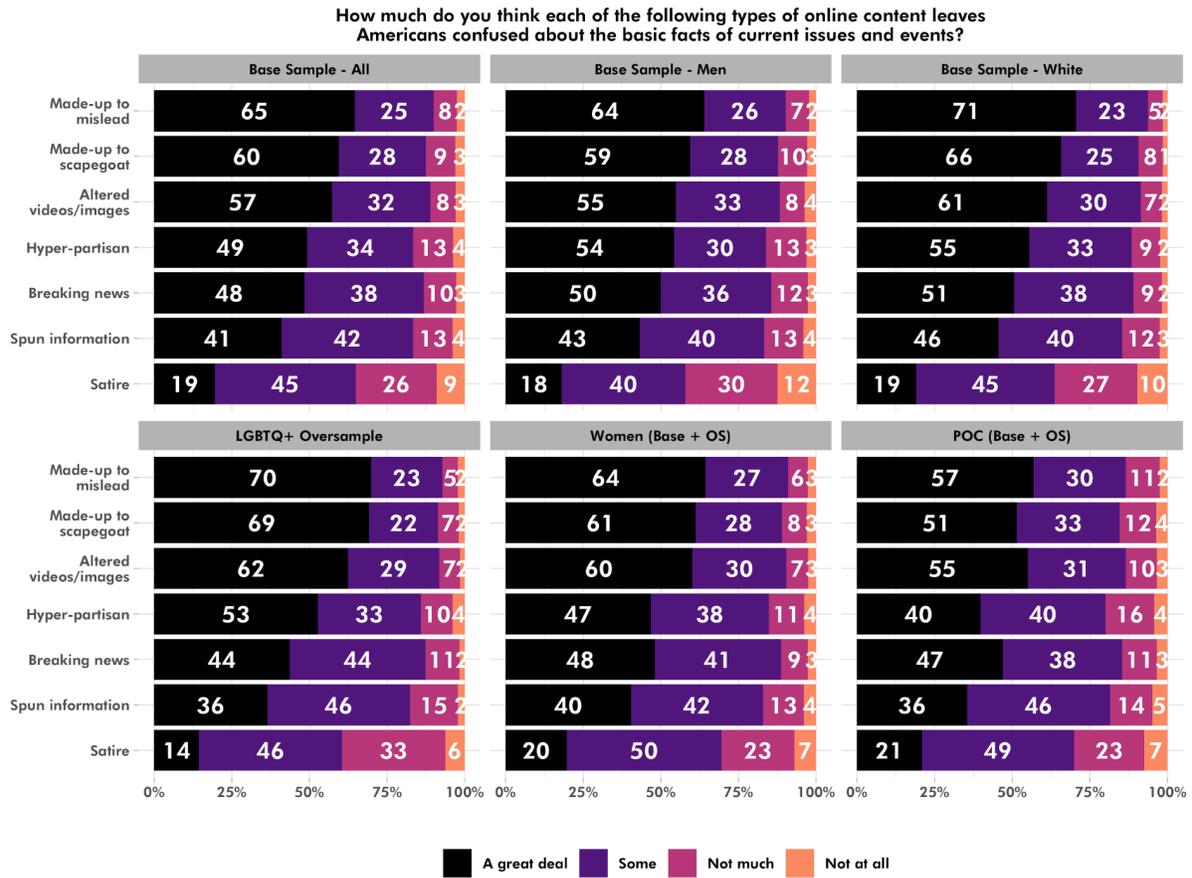
But not all false information is seen equally across online spaces. **Strong majorities of Americans report coming across false information about politics and elections** (72 percent in the base sample, along with 65 percent of LGBTQ+ folk, 62 percent of POC, and 70 percent of women), **COVID-19** (61 percent in the base sample, along with 60 percent of LGBTQ+ folk, 55 percent of POC, and 63 percent of women), **and public figures** (60 percent in the base sample, along with 58 percent of LGBTQ+ folk, 55 percent of POC, and 61 percent of women) **either very frequently or frequently.**

Now, more specifically. How often do you believe you come across false information in online spaces about each of the following items?



Either majorities or pluralities of the base sample report seeing false information frequently about prominent topics like racial groups (51 percent of the base sample), public health (49 percent of the base sample) reproductive health (47 percent of the base sample), gender identity (44 percent of the base sample), and LGBTQ+ people (43 percent of the base sample). When looking at the oversamples of individuals from these affected groups, there are similar proportions of witnessing misinformation about these topics. The one exception is the LGBTQ+ oversample. **About 56 percent of respondents from the LGBTQ+ oversample report frequently seeing misinformation about both LGBTQ+ people and gender identity.** Finally, respondents report seeing false information about emergency situations, civic institutions, and science/technology or religious communities the least.

This survey also asked respondents about the impacts that they thought that different types of information, several of them that would fall under the umbrella of false information, has on the American public.¹² Overall, Americans are attuned to the pitfalls and dangers of made-up information. Here, **very strong majorities of Americans believe that information often used to mislead the public, whether deliberately made up or manipulated videos and images, leaves Americans confused about the basic facts of current events.**



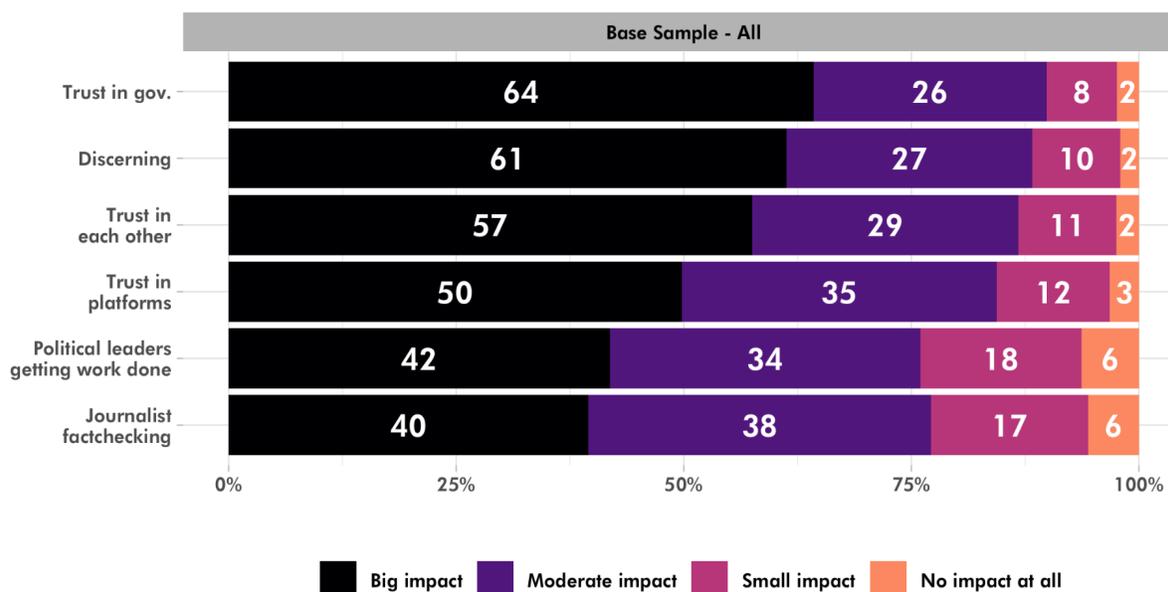
In the base sample, 92 percent of respondents said that *made-up information that misleads the public* confuses people either a great deal (64 percent) or some (26 percent). A similar pattern is found among LGBTQ+ respondents (93 percent), POC (87 percent), and women (91 percent). Similarly, 88 percent of respondents in the base sample believe that information that has been made up to scapegoat specific social groups leaves the public confused,

¹² The question asked... "How much do you think each of the following types of online content leaves Americans confused about the basic facts of current issues and events?"

- Made-up information that misleads the public
 - Made-up information that attacks or scapegoats a specific social group
 - Satire about an issue or event
 - Breaking news information that is published before everything is verified
 - Factual information presented to favor one side of an issue over another
 - Videos or images that are altered or made up to mislead the public
 - Hyper-partisan content
- <1> A great deal <2> Some <3> Not much <4> Not at all

while 89 percent believes that altered video or images do the same. On the other hand, satirical information, information that is spun to fit one narrative, hyperpartisan content, or breaking news information are less confusing, according to respondents across all groups.

False information, however, does not only impact the individual's *perception* of the facts, it can have deleterious downstream effects. A host of academic and public opinion research has theorized that misinformation could impact various attitudes like trust in institutions. In order to evaluate this possibility, this instrument included a question that asked respondents to evaluate the impact of false information on various outcomes¹³, including various measures of trust.



A majority of Americans think that false information affects trust in the government and in each other, as well as the public's ability to discern what is true information. Fully 64 percent of respondents in the base sample said that false information has a big impact on the public's trust in government (along with 68 percent of LGBTQ+ respondents, 66 percent of women, and 57 percent of POC). About 6 out of 10 in the base sample said both that false information has a big impact on the public's trust in each other and on the public's ability to

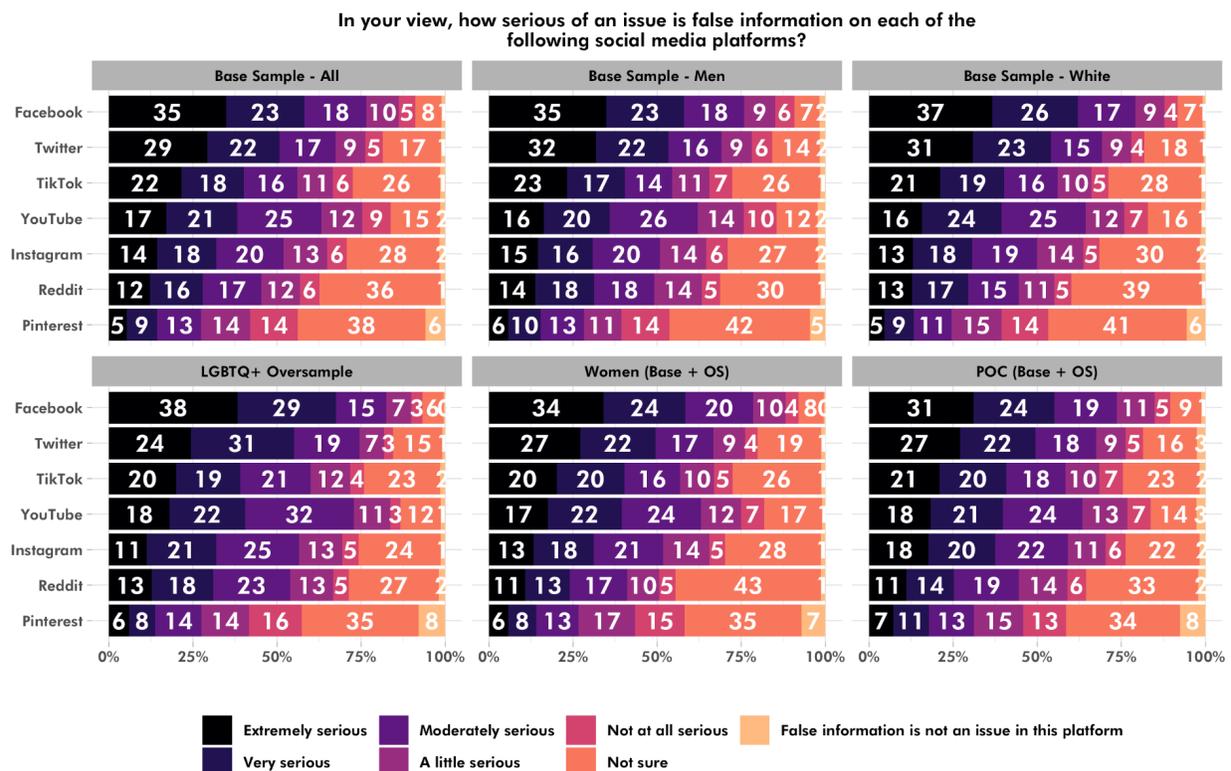
¹³ The question read:

How much do you personally think that false information impacts each of the following?

- The public's trust in government
 - The public's trust in each other
 - Political leaders' ability to get work done
 - Journalists' ability to get accurate information for their stories
 - The public's ability to tell what is true information
 - The public's trust in social media companies
- <1> Big impact <2> Moderate impact <3> Small impact <4> No impact at all

tell what is true information. Majorities across the three demographic groups of interest said the same.

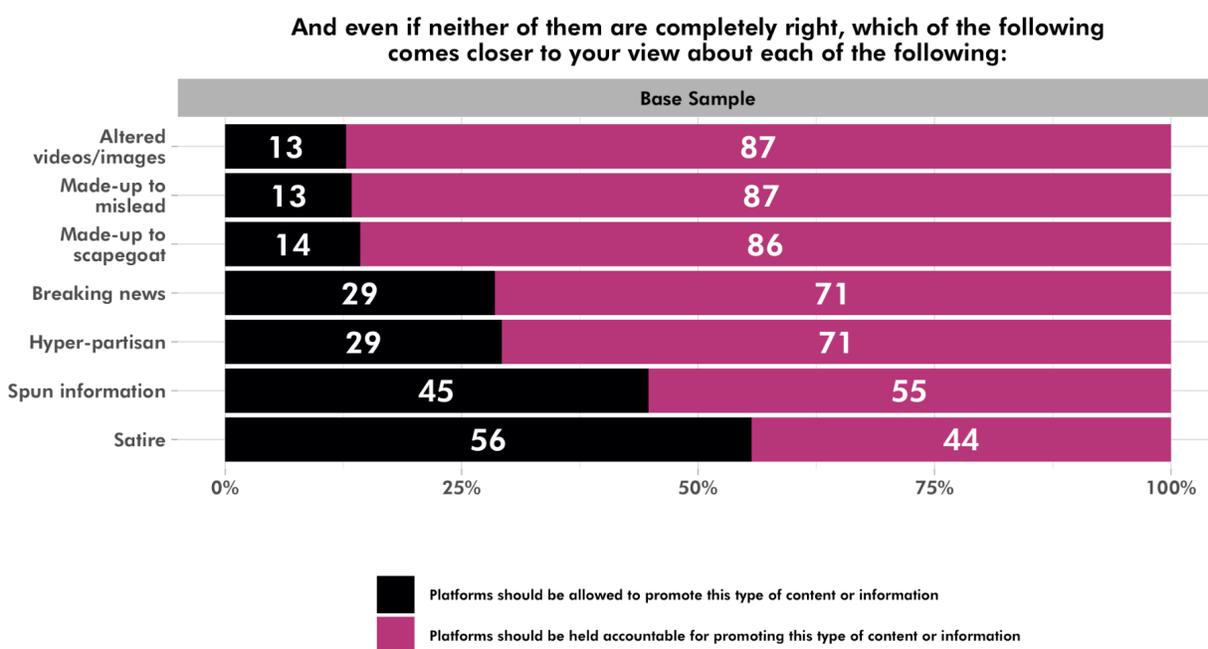
This survey also asked respondents about how serious of an issue they believe that false information is across various online platforms. **Either outright majorities or pluralities say that false information is a serious problem across social media platforms.** For instance, a majority of respondents in the base sample said that false information is either an extremely or very serious problem in Facebook (58 percent) and Twitter (51 percent). LGBTQ+ respondents (67 percent Facebook, 55 percent Twitter), women respondents (58 percent Facebook, 49 percent Twitter), and people of color (55 percent Facebook, 49 percent Twitter) exhibit similar patterns.



It is worthwhile to note that over a quarter of respondents across samples are unsure about how serious of a problem false information is on Pinterest, Reddit, and Instagram. However, this does not necessarily imply that false information is not an issue in these platforms – it simply means that respondents are less aware of how extensive the issue is in these platforms, perhaps due to lower rates of usage among these platforms. This is further evidenced by the fact that less than 3 percent of respondents across all platforms (except Pinterest) said that false information is not an issue in that site.

Respondents were also asked whether platforms should be held accountable for the spread of different types of false information on their sites or whether they should be allowed to promote it.¹⁴

The majority of respondents are on the side of platform accountability on six out of seven of these items – the only exception is satire. Almost 9 out of 10 respondents said that platforms should be held accountable for made up information that misleads the public (87 percent) or scapegoats a social group (86 percent), as well as video or images that are altered (87 percent). Additionally, 71 percent of respondents believe that platforms should be held accountable for breaking news information and hyperpartisan content, while 55 percent believe that they should be accountable for information spun to favor one side of an issue. Satire is the only item in this battery where a majority of respondents in the base sample agreed that platforms should be allowed to promote it.



¹⁴ The full item read:

And even if neither of them are completely right, which of the following comes closer to your view about each of the following:

- *Made-up information that misleads the public*
- *Made-up information that attacks or scapegoats a specific social group*
- *Satire about an issue or event*
- *Breaking news information that is published before everything is verified*
- *Factual information presented to favor one side of an issue over another*
- *Videos or images that are altered or made up to mislead the public*
- *Hyper-partisan content*
 - *Platforms should be allowed to promote this type of content or information*
 - *Platforms should be held accountable for promoting this type of content or information*

These findings are similar among respondents in the additional demographic samples, as shown in the following graph. Overall, **the overwhelming majority of respondents thinks platforms should not promote a variety of content**, including breaking news and hyper-partisan content, regardless of its truth content, suggesting a general preference for more control over what one is seeing online.



Key Takeaways

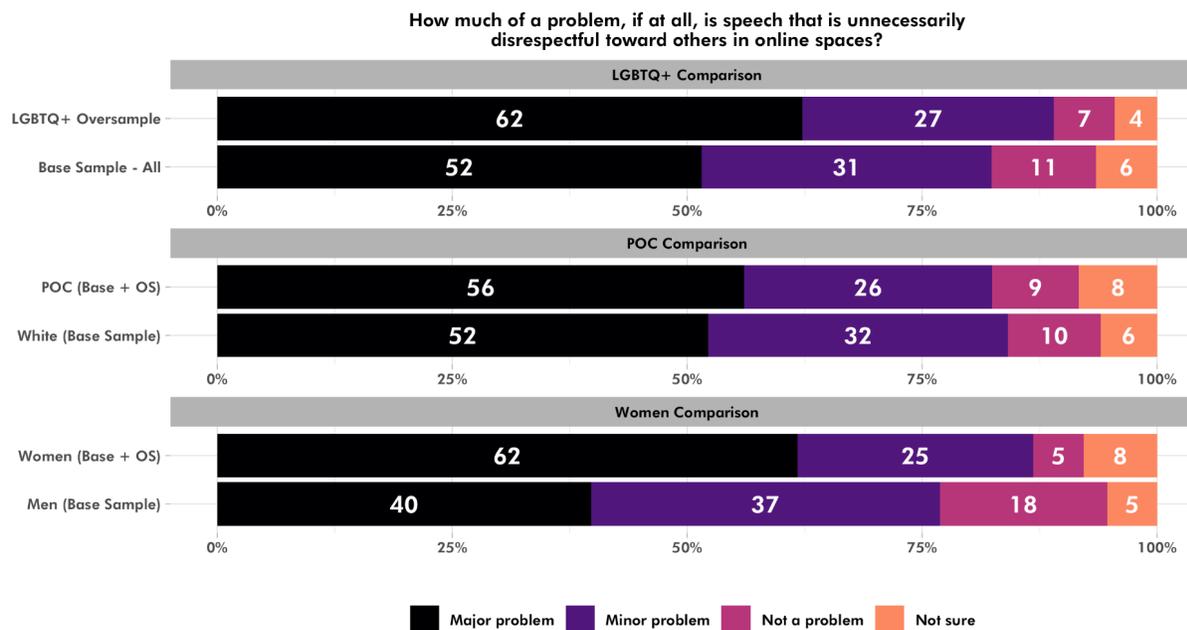
In short, respondents are broadly aligned with the idea that false information is a problem in online spaces and would like to hold platforms accountable for their role in promoting false information. The following points summarize the findings in this section:

- A majority of Americans report that they frequently run into false information especially about politics, COVID-19, and public officials.
- Most believe that made up information and altered videos/pictures leave the public confused about the basic facts of current events
- Most people believe that false information impacts trust, both in the government and in one another, showing that the deleterious effects of misinformation could have both immediate and downstream effects.
- Respondents see false information as an issue across multiple platforms, with Facebook and Twitter being seen as the most problematic.
- A very strong majority believe that platforms should be held accountable for various types of false information in their sites – about 9 out of 10 women and LGBTQ+ respondents agree with platform accountability when it comes to false information, while 8 out of 10 people of color do the same.

Discursive Incivility & Hate Speech: problematic in online spaces, and social media sites aren't doing enough to address them

Online harassment, hate, and incivility are growing issues in America today. This survey also included various questions about these phenomena in online spaces. This survey asked respondents about whether “speech that is unnecessarily disrespectful” and “hate speech” are a problem in online spaces.¹⁵

Overall, **Americans overwhelmingly agree that both disrespectful speech and hate speech are a problem in online spaces.** Among the base sample, 52 percent said that disrespectful speech is a major problem (and 11 percent not a problem, a net gap of 41 percentage points), while 61 percent said that hate speech itself is a major problem (and 10 percent said it's not, for a net gap of 51 percentage points).

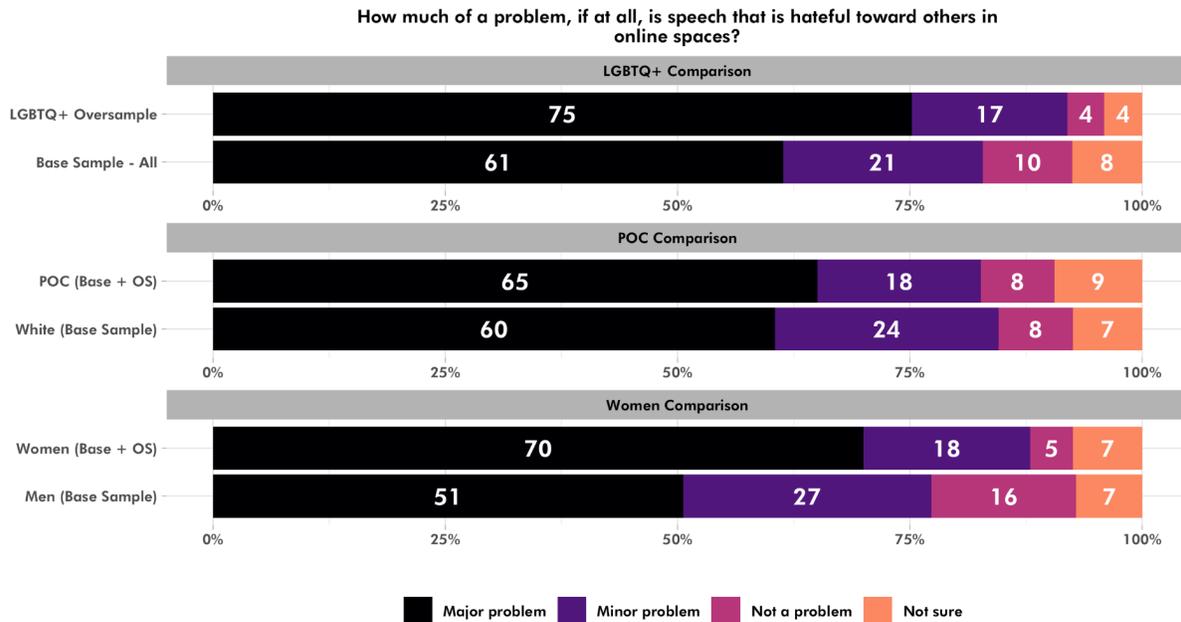


When asking the same questions of LGBTQ+, POC, and women respondents – individuals with marginalized identities – the gap grows even larger. Important highlights of these items:

- Among LGBTQ+ respondents, the gap is 55 percentage points for disrespectful speech and 71 percentage points for hate speech, while the corresponding gaps among POC (47 pp and 57pp, respectively) and women (57pp, and 65pp) respondents are similarly large.

¹⁵ The scholarly literature tends to differentiate between “discursive incivility” and “hate speech” – while the former tends to be speech that is generally deemed as offensive and disrespectful, it does not rise to the same level of vitriol and hatred as the latter. Therefore, this survey asked about both of these phenomena separately in this project.

- Additionally, LGBTQ+ people are 14 percentage points more likely than the general population to say that hateful speech is a problem in online spaces and 12 percentage points more likely to say that disrespectful speech is problematic.
- Strong gaps are also evident when comparing women to men (22 pp difference in disrespectful speech, and 19pp difference in thinking hate speech is problematic), while the gaps for POC when compared to white respondents are more equivocal (5pp and 4pp, respectively).

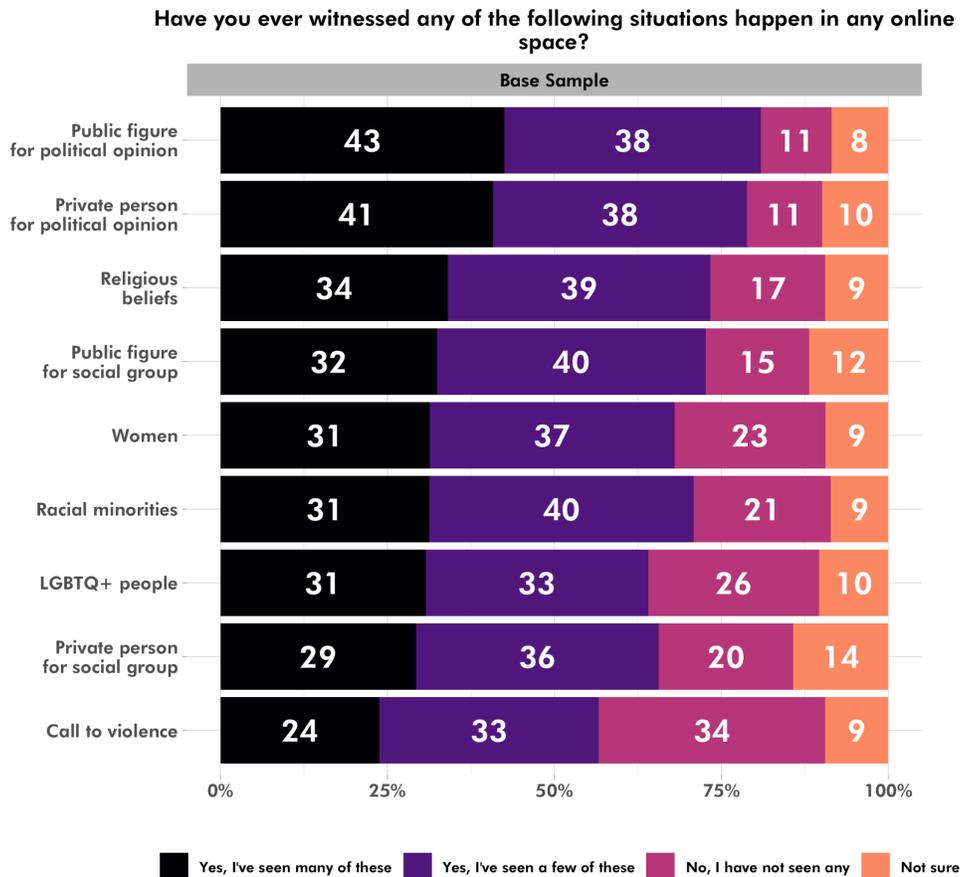


Large proportions of respondents report having witnessed specific instances of uncivil or hateful speech against another individual.¹⁶ **At least three-quarters of respondents in the base sample report having seen online posts that attack someone based on their political affiliation or opinion** – 83 percent have seen attacks against Republicans and 78 percent against Democrats, 81 percent have seen attacks against a public figure and 79 percent against a private person for their political opinion. Additionally, 71 percent have seen posts attacking a racial minority, 68 percent of respondents have seen posts that attack women, and 64 percent have seen a post that attacks LGBTQ+ people. Fully 57 percent of

¹⁶ Have you ever witnessed any of the following situations happen in any online space?

- A post that insults or attacks another specific private person for their political opinion
- A post that insults or attacks another specific private person for belonging to a social group
- A post that insults or attacks a public figure, like an actor, for their political opinion
- A post that insults or attacks a public figure, like an actor, for belonging to a social group
- A post that insults or attacks LGBTQ+ people
- A post that insults or attacks racial minorities
- A post that insults or attacks someone based on their religious beliefs
- A post that insults or attacks Republicans
- A post that insults or attacks Democrats
- A post that insults or attacks women
- A post that calls for violence against someone based on their race, gender, sexuality, or any other personal characteristic

respondents have seen a post that calls for actual violence against someone based on their race, gender, sexuality, or any other personal characteristic. In short, **at least half of Americans report having seen various types of hateful or harassing speech in digital spaces, with the incidence going upwards of 80 percent in some instances.**

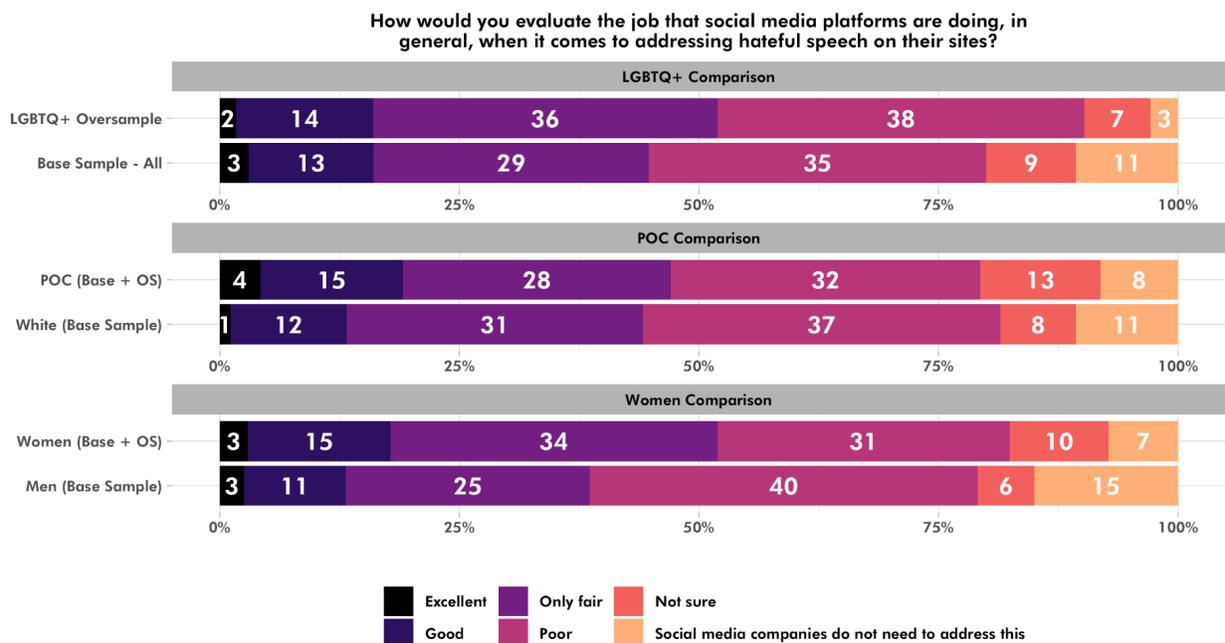


The demographic oversamples can be leveraged to further explore the experience of marginalized groups with uncivil or hateful speech in online spaces. Overall, **respondents who belong to a marginalized community tend to notice attacks against their community more than non-marginalized groups notice against that specific community.** For instance, while 88 percent of respondents in the LGBTQ+ sample report having seen a post that insults or attacks LGBTQ+ individuals, only 64 percent of respondents in the base sample do the same – a significant 24 percentage point gap. Likewise, 73 percent of LGBTQ+ respondents report having seen a call-to-violence against a marginalized group, while 57 percent of the base sample has (a 16 percentage point gap).

A similar pattern is evident among POC and women, even when the gaps are slightly less statistically clear than with the LGBTQ+ comparison. While 73 percent of POC report having seen a post insulting or attacking a racial minority, only 68 percent of White respondents did, a 5 percentage point gap. And while 62 percent of POC respondents report having seen a call to violence, only 54 percent of White folk did. Fully 70 percent of women report having

seen a post insulting or attacking a woman, while only 66 percent of men did, and 53 percent of women report seeing a call-to-violence online (while 58 percent of men did). While these results are less statistically clear-cut than the LGBTQ+ comparison, they are all directionally consistent with the general idea that individuals tend to notice more attacks against a community that they are a part of.

Finally, **Americans give bad marks to the job that social media platforms are doing to address hateful speech on their sites.** A plurality of respondents in the base sample (35 percent), as well as in the LGBTQ+ (38 percent) and POC (32 percent) groups said that social media platforms are doing a poor job at addressing hateful speech on their sites, while approximately three out of ten women said the same. Across all samples, less than one fifth of respondents said that platforms are doing either an excellent or good job at addressing hateful speech on their sites.



Key Takeaways

People think hate speech and incivility are ubiquitous in online spaces. The following points briefly summarize the findings of this section:

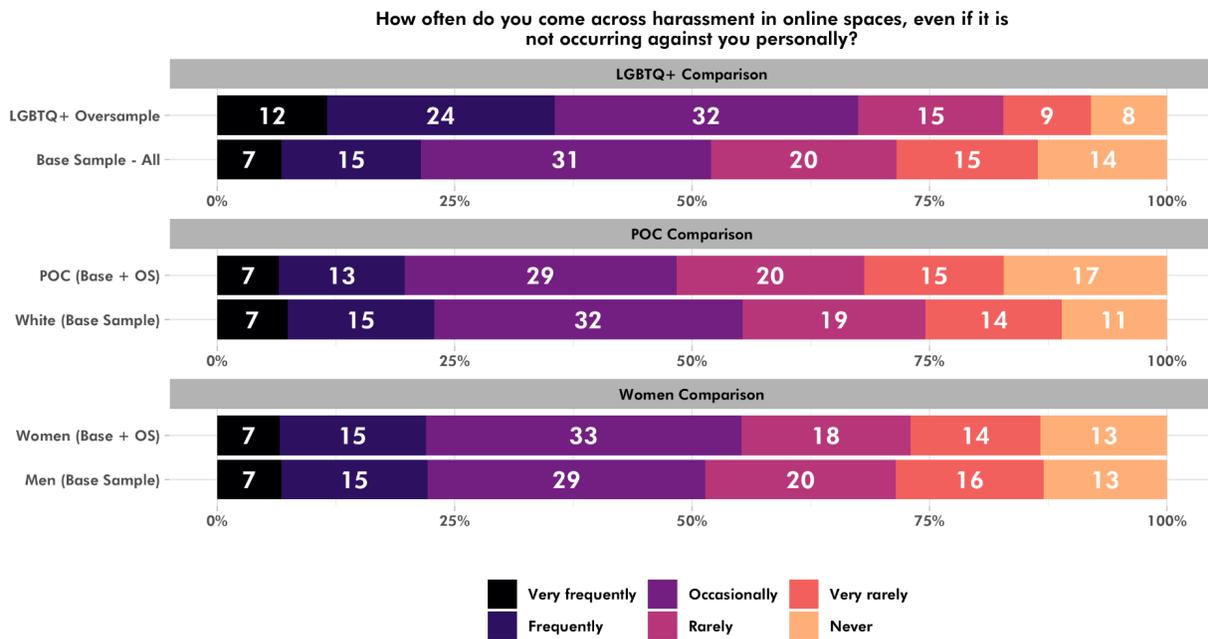
- Large proportions of respondents report having witnessed specific instances of disrespectful or hateful speech online
- Individuals who hold a marginalized identity, like LGBTQ+ folk, POC, or women, are more likely to report having witnessed hate against people with their same identity. This could be perhaps because they are more attuned to recognizing attacks on their own community.
- Americans give bad marks to the job that social media platforms are doing to address hateful speech on their site.

Quantifying Experiences Of Online Harassment

This survey had three sections aimed to further quantify respondents' experiences with online harassment. There are a multiplicity of ways in which a person can experience harassment in digital spaces – personal experiences with witnessing *others* be harassed to suffering harassment themselves. The following subsections address each of these phenomena, as well as general views on harassment.

General views on harassment

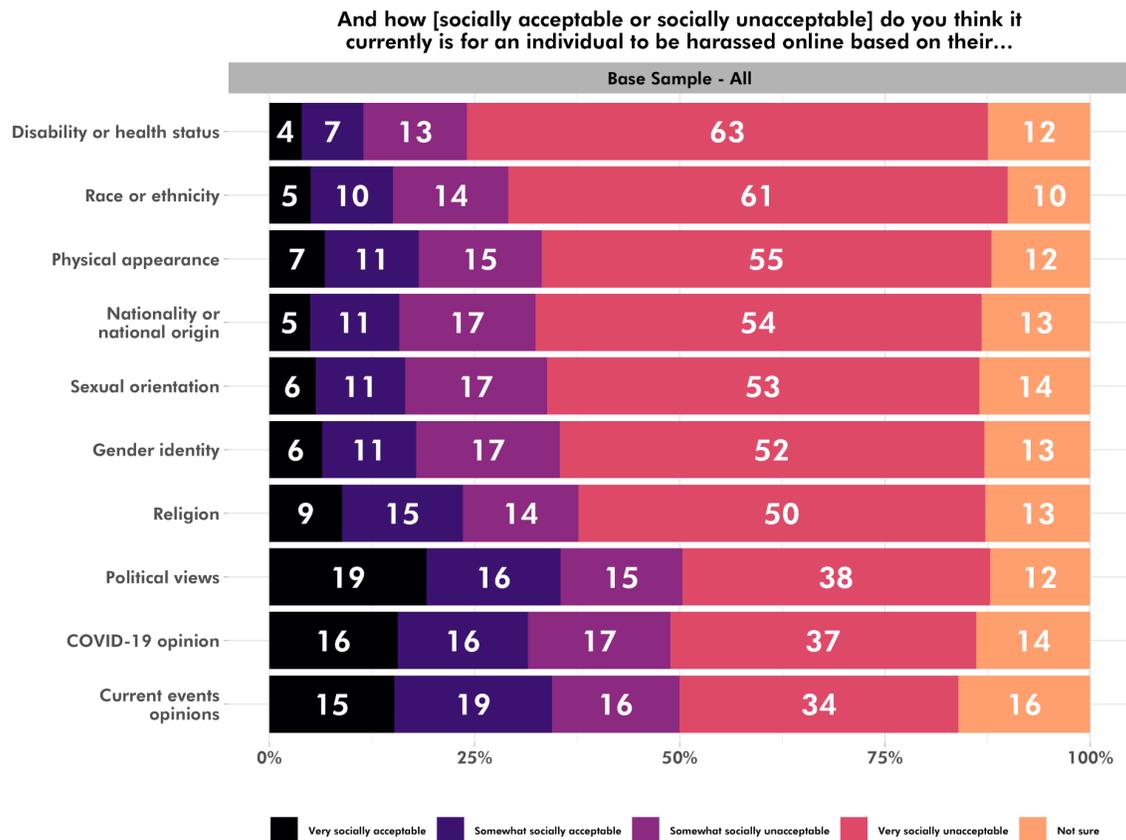
To explore the incidence of harassment, this survey asked respondents how often they come across harassment in online spaces “even if it is not occurring against you personally”. The following graph summarizes the findings across the different comparisons of interest.



About one in five Americans, POC, and women, report encountering harassment in online spaces either very frequently or frequently, while one third of LGBTQ+ respondents report the same.

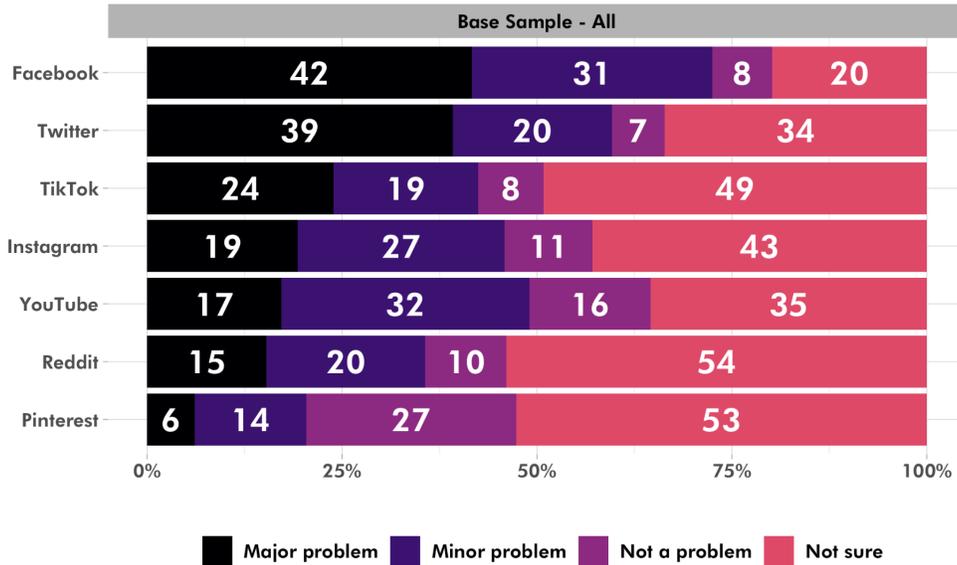
To understand how respondents further evaluate their general experience with harassment in online spaces, this survey asked them to share whether they think that society thinks it is socially acceptable to harass an individual based on various personal characteristics. **Strong majorities of respondents say that society thinks it is socially unacceptable to harass respondents based on a host of individual characteristics.** For example, 76 percent of respondents in the base sample said it is socially unacceptable to harass someone based on their disability or health status, 75 percent said the same about race or ethnicity, 71 about nationality, 70 percent sexual orientation or physical appearance, and 69 percent said the same about gender identity. **However, about a third of Americans say that it is viewed as**

generally socially acceptable to harass people online based on political opinions, views about current events, and COVID-19. For example, 35 percent of Americans say that it is generally socially acceptable to harass someone based on their political opinions, 34 percent said the same about current events, and 32 percent said so about COVID-19 opinions. The ordering and percentages are substantively the same to this across demographic comparisons of interest.



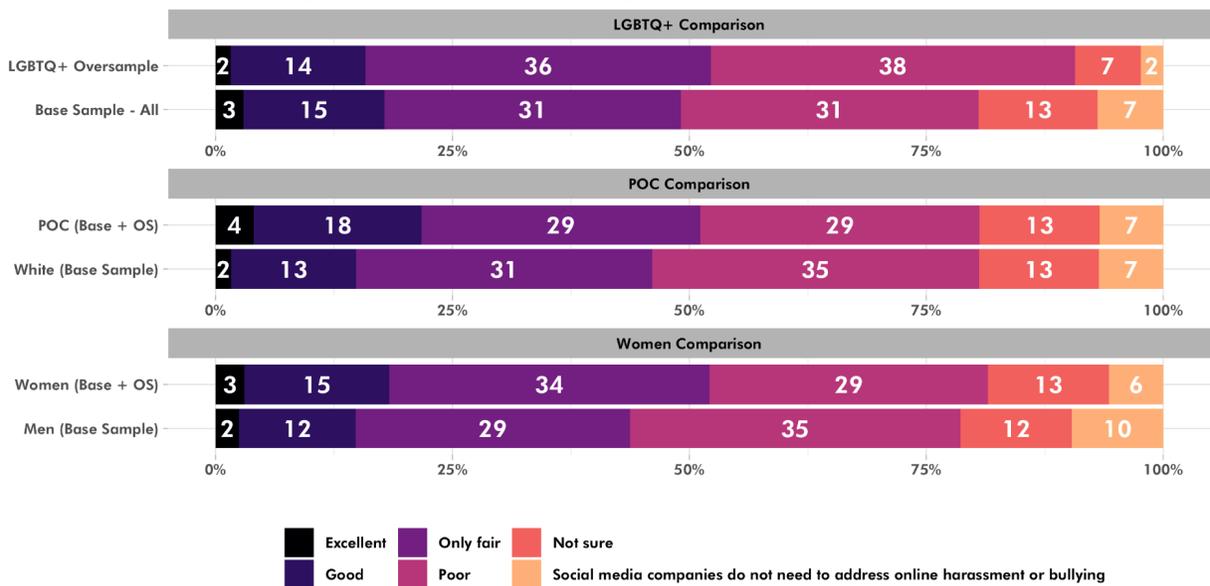
This survey also asked respondents about how serious of a problem they think that online harassment or bullying is in various online platforms. **A plurality of respondents in the base sample said that online harassment is a major problem on Facebook and Twitter, a finding that mirrors the previous findings on false information and hate speech.** Less than 10 percent of the base sample said that harassment is *not* a problem in these platforms and TikTok, while 11 percent said it is not a problem in Instagram, 16 percent said the same about YouTube, 10 percent about reddit, and 27 percent about pinterest. **However, this does not necessarily mean that these other platforms are in the clear – outright majorities or large pluralities of respondents said they are unsure about whether harassment is a problem in reddit (54 percent), pinterest (53 percent), TikTok (49 percent), and Instagram (43 percent).** This suggests either that further information about harassment in these platforms is not widespread or that most public attention about the issues of online harassment has been focused on Facebook and Twitter, while not addressing other platforms where problematic behavior could also be endemic.

Thinking of some experiences that might happen to people when they use the internet, how much of a problem, if at all, is regular people being harassed or bullied on each of the following social media platforms:



Furthermore, this survey asked respondents to evaluate the job that social media companies have been doing at addressing online harassment in their platforms. **Nearly a third of Americans overall, women, and POC respondents said that social media platforms are doing a poor job at addressing online harassment on their sites, while almost two out of five LGBTQ+ respondents said the same.** Less than five percent of respondents in all samples said that social media sites are doing an excellent job at addressing the situation, while only about one of out six respondents across all groups said the sites are doing a good job with it. Less than ten percent of all respondents across groups said that social media platforms do not need to address harassment on their platforms.

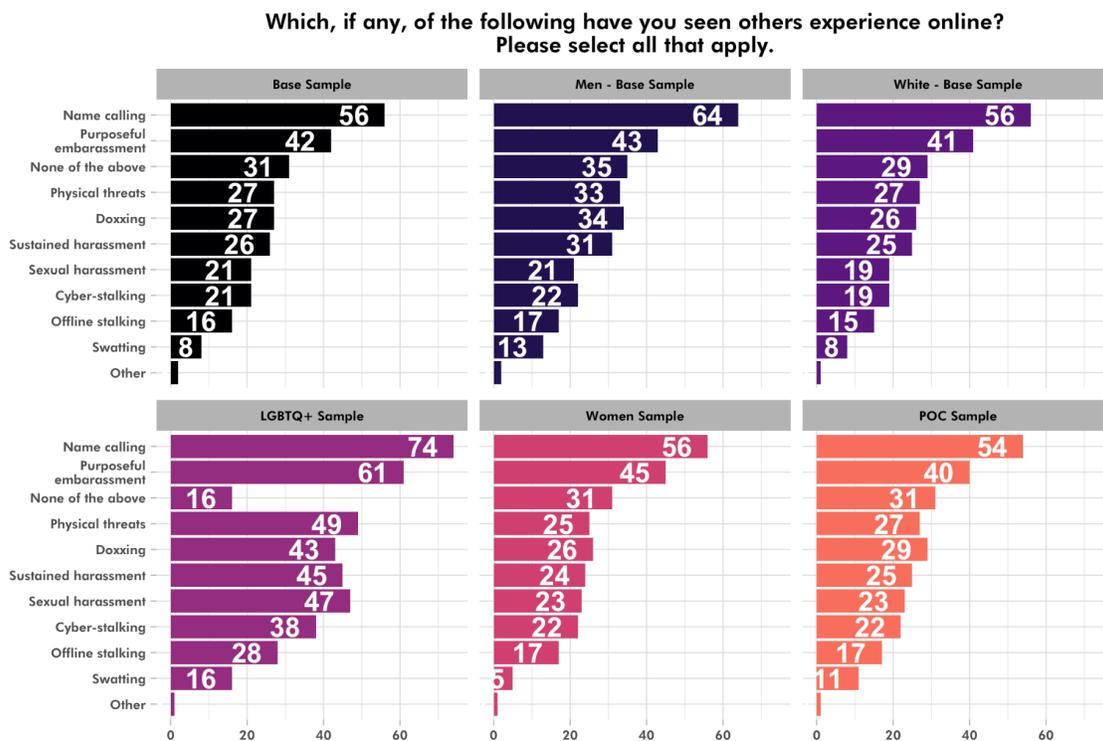
How would you evaluate the job that social media companies are doing, in general, when it comes to addressing online harassment or bullying on their platforms?



Vicarious experiences with harassment

Individuals can experience harassment in a variety of ways in digital spaces, including seeing others be harassed (what this section deems a “vicarious” experience with harassment).¹⁷

Overall, a third of respondents in the base sample report not having seen others experience any of the specific instances of harassment, but only about 16 percent of LGBTQ+ people report the same. However, **witnessing harassment online is widespread among Americans overall**: 56 percent of respondents in the base sample report seeing others being called offensive names, 42 percent reported seeing others be embarrassed on purpose, and 27 percent saw others be physically threatened.



¹⁷ To gauge what harassment respondents have seen *others* endure online, this survey asked respondents:

Which, if any, of the following have you seen others experience online? Please select all that apply.

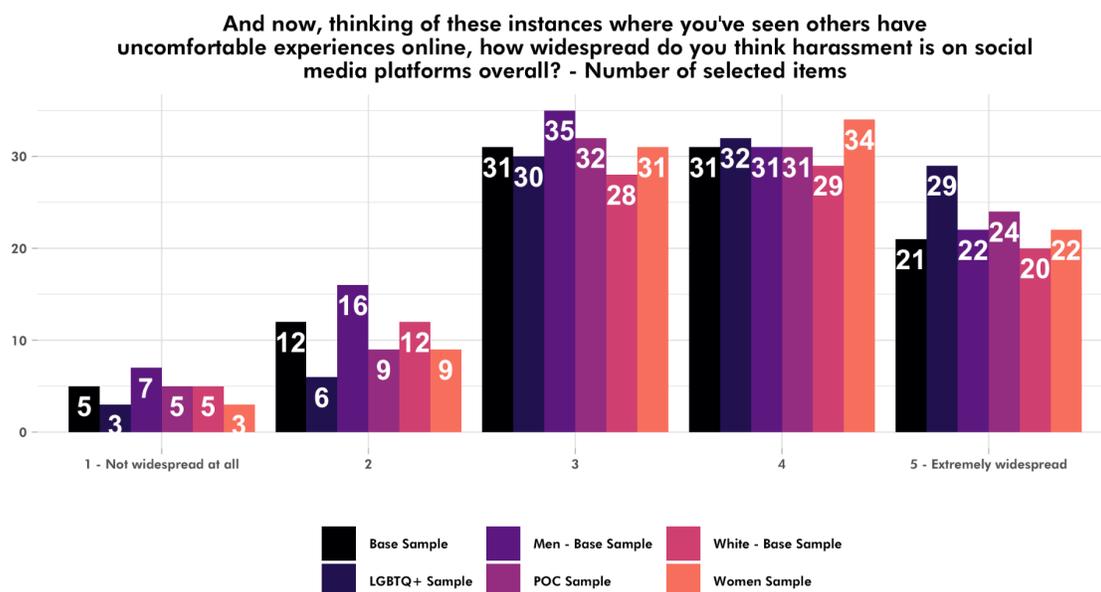
- <1> Been called offensive names
- <2> Been physically threatened
- <3> Been harassed for a sustained period of time
- <4> Been cyber-stalked (online only)
- <5> Been stalked offline, as an extension of online harassment
- <6> Had someone try to embarrass them on purpose
- <7> Been sexually harassed
- <8> Had someone release private information about them (doxxing)
- <9> Had an emergency service dispatched to your home as a hoax (swatting)
- <10> Other (please specify)
- <11> I have not seen any of the above

At the same time, **LGBTQ+ respondents report seeing the same behaviors at increased rates** when compared to the broader American public – 74 percent of LGBTQ+ folk report having seen others be called offensive names (a 18 percentage point difference), 61 percent saw others be embarrassed on purpose (a 19 percentage point difference), and 49 percent saw others be physically threatened (a 22 percentage point difference). Women and POC do not necessarily report similar patterns in this item, suggesting that LGBTQ+ respondents might perhaps be more attuned to recognizing these instances of harassing behavior.

This survey also asked respondents to reflect, after asking the previous question, on how widespread they think harassment is in social media overall. The item read:

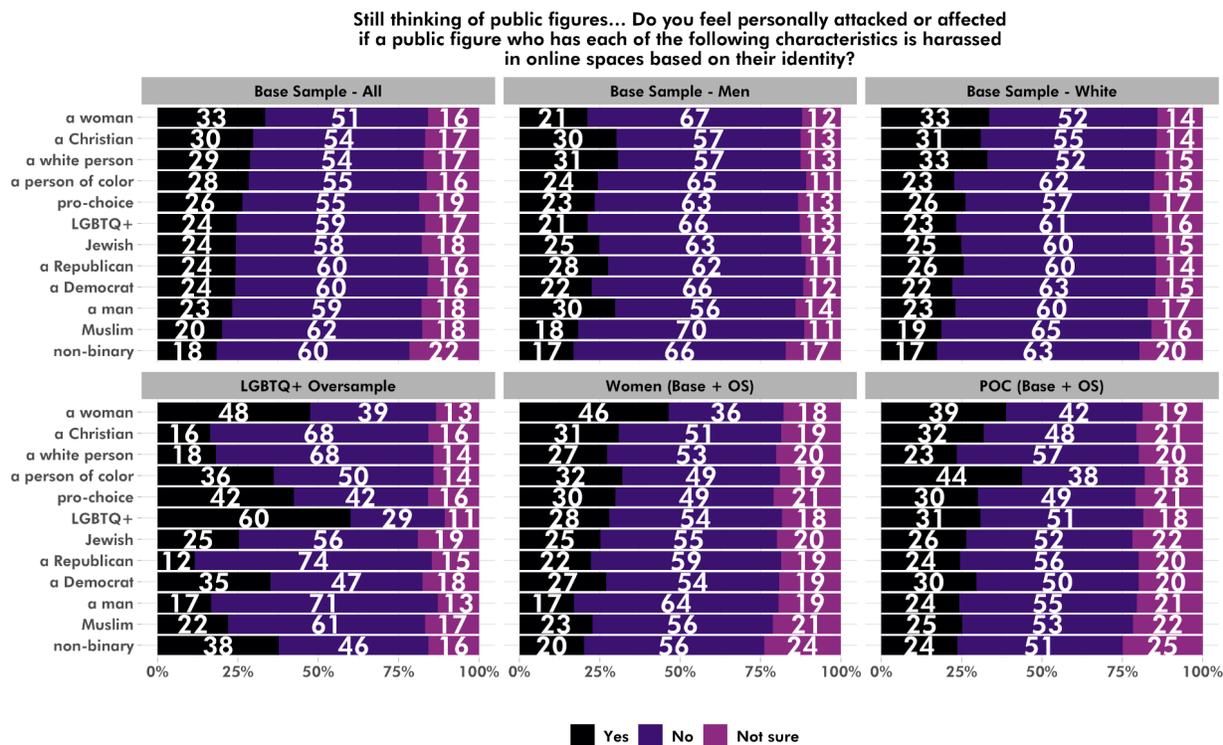
And now, thinking of these instances where you've seen others have uncomfortable experiences online, how widespread do you think harassment is on social media platforms overall?

Respondents could choose a response in a 5 point scale ranging from “not widespread at all” to “extremely widespread”. Overall, **either pluralities or majorities of Americans across groups believe that harassment is more widespread than not** – 52 percent of Americans in the base sample said either a 4 or 5 in this response scale, while 61 percent of LGBTQ+ respondents, 55 percent of POC, and 56 percent of women said the same. The overwhelming majority of respondents believe harassment on social media is at least somewhat widespread.



Vicarious harassment has a potentially deleterious effect: seeing someone with your shared identity be harassed could spill over into being personally affected by that

harassment.¹⁸ For instance, if a LGBTQ+ person sees another LGBTQ+ person be harassed in an online space, both the person facing the harassment themselves and the witness could get affected by the harassment.



Among groups with marginalized identities, there is some evidence of group consciousness and solidarity when they see others being attacked online who share their own personal identity. For example, while only 24 percent of the base sample feel personally affected when an LGBTQ+ person does, 60 percent of respondents in the LGBTQ+ sample feel affected when an LGBTQ+ public person gets attacked online. Similarly, while

¹⁸ To explore this, this survey asked respondents:

Still thinking of public figures... Do you feel personally attacked or affected if a public figure who has each of the following characteristics is harassed in online spaces based on their identity?

- A public figure who is LGBTQ+
- A public figure who is a person of color
- A public figure who is a woman
- A public figure who is a man
- A public figure who is a white person
- A public figure who is non-binary
- A public figure who is a Christian
- A public figure who is Jewish
- A public figure who is Muslim
- A public figure who is pro-choice
- A public figure who is a Democrat
- A public figure who is a Republican

only 18 percent of the base sample feels personally affected when they see a non-binary person get attacked, 38 percent of the LGBTQ+ sample feels affected by this.

There is a similar pattern among women and people of color: while 21 percent of men feel affected when a woman gets attacked, 46 percent of women feel the same, and while 23 percent of white people feel affected when a person of color gets attacked online, 44 percent of POC respondents said they feel personally affected by the same.

In short, vicarious harassment can translate to feelings of being affected by seeing others with one's own identity be attacked in online spaces. This means, substantively speaking, that the problem of online harassment is not only one that affects the *victims* of harassment themselves, but the *witnesses* of this harassment.

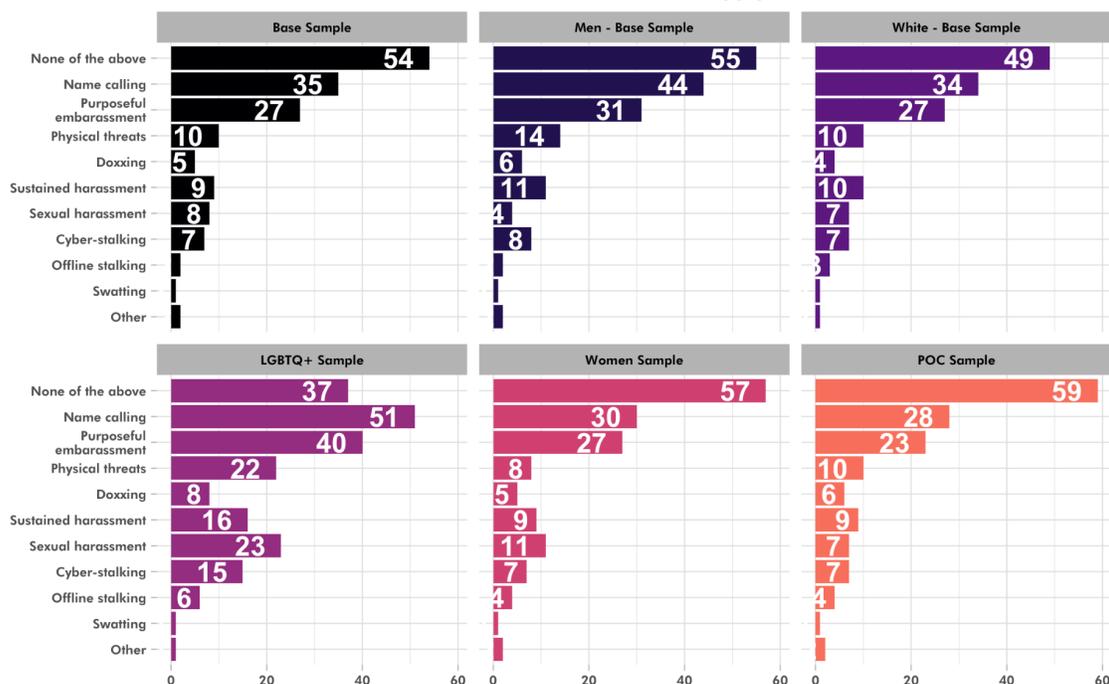
Own experiences with harassment

The final instance of harassment this survey explored is individuals' own experience with having been harassed themselves.¹⁹ Fully 35 percent of respondents in the base sample said they have been called offensive names online, 27 percent had someone embarrass them, and 10 percent have been physically threatened. **LGBTQ+ folk and women report higher rates of harassment than their respective comparison groups.** For example, 51 percent of LGBTQ+ people have been called offensive names (a 16 percentage point gap when compared to the base sample), 40 percent had someone try to embarrass them (a 13 percentage point gap), 23 percent have been sexually harassed, and 22 percent physically threatened. Women, on the other hand, report higher rates of sexual harassment than men – 11 percent were sexually harassed (compared to 4 percent of men).

¹⁹ Which, if any, of the following have you personally experienced online? Please select all that apply.

- <1> Been called offensive names
- <2> Been physically threatened
- <3> Been harassed for a sustained period of time
- <4> Been cyber-stalked (online only)
- <5> Been stalked offline, as an extension of online harassment
- <6> Had someone try to embarrass them on purpose
- <7> Been sexually harassed
- <8> Had someone release private information about them (doxing)
- <9> Had an emergency service dispatched to your home as a hoax (swatting)
- <10> Other (please specify)
- <11> I have not experienced any of the above

**Which, if any, of the following have you personally experienced online?
Please select all that apply.**



The questionnaire then asked respondents a follow-up question²⁰ asking them to reflect on what could have been the cause for any of the harassment they personally experienced online.²¹ Overall, 56 percent of respondents in the base sample who had reported being subjected to harassment said that this happened due to their political views, while about a third report the harassment having been due to their COVID-19 opinions, or their opinions about other current events.

However, personal identity appears to play an important role in how one experiences harassment. For example, a majority of LGBTQ+ respondents (52 percent) said that the

²⁰ Were any of the experiences you mentioned in the previous question a result of any of your following characteristics? Please check all that apply.

- <1> Sexual orientation
- <2> Gender identity
- <3> Nationality or national origin
- <4> Political views
- <5> Physical appearance
- <6> Disability or health status
- <7> Religion
- <8> Race or ethnicity
- <9> Opinions about COVID-19
- <10> Opinion about other current events
- <11> Other
- <12> No, these experiences were not related to any of these reasons

²¹ This question was only asked of respondents who reported having been subject to at least one type of harassment in the previous question.

harassment had been due to their sexual orientation, while only about 14 percent of the base sample said the same. Moreover, more LGBTQ+ respondents report having been subject to harassment due to their physical appearance (40 percent compared to 20 percent of the base sample) or gender identity (31 percent compared to 12 percent of the base sample). POC face more race-based harassment (38 percent) than white respondents do (15 percent), as well as more harassment based on physical appearance (28 percent vs. 19 percent of whites) and on national origin (19 percent vs. 7 percent of whites). Similarly, women face more appearance-based harassment than men do (25 percent vs. 17 percent of men), while men report facing more harassment than women based on their political opinion (61 percent of men vs. 52 percent of women report harassment due to this).

Key Takeaways

- About one in five Americans, POC, and women, report encountering harassment in online spaces either very frequently or frequently, while one third of LGBTQ+ respondents report the same.
- A plurality of respondents think that online harassment is a major problem on Facebook and Twitter, a finding that mirrors the previous findings on false information and hate speech
- Nearly a third of Americans overall, women, and POC respondents said that social media platforms are doing a poor job at addressing online harassment on their sites, while almost two out of five LGBTQ+ respondents said the same.
- When it comes to solutions, Americans are deeply divided on the question of whether individuals should be allowed to sue platforms where they have experienced online harassment.
- Witnessing others be harassed online is a widespread experience among Americans overall, but LGBTQ+ respondents report seeing the same behaviors at increased rates. However, when asked to reflect about this question, either pluralities or majorities of Americans across groups believe that harassment is more widespread than not.
- People aren't just harmed from direct harassment, they feel attacked when public figures are attacked based on their shared identity. This means, substantively speaking, that the problem of online harassment is not only one that affects the *victims* of harassment themselves, but the *witnesses* of this harassment.
- While a majority of Americans in the base sample have not been subject to harassment themselves, LGBTQ+ folk and women report higher rates of harassment than their respective comparison groups. This means that personal identity appears to play an important role in how one experiences harassment.

Potential Solutions MaxDiff

This module sought to understand how the American public and different demographic groups of interest would rank a suite of potential solutions to the issues of online harassment and disinformation. To explore this, this survey used an analytic strategy called MaxDiff, which is commonly used to rate a set of items on some dimension: quality, importance, persuasiveness, favorability, and so on. Researchers often use this method as an alternative to standard response (e.g., 5-point) scales, because it forces respondents to make trade-offs between options, which may be necessary to generalize to real-world applications. It also is easier for respondents than ranking questions, which can be difficult when the set of items becomes more than a few options.

A MaxDiff consists of a series of questions, each presenting a subset of the total set of items. Instead of asking respondents to rank the *whole* set, this survey only asked respondents to indicate which potential solution to online harassment and disinformation they were *most* and *least* likely to support. These series of questions will present items more than once and in different combinations to best get at a respondent's underlying attitudes about the items.

Specifically, this survey asked respondents:

So far, in this survey you've been asked to consider the issues of false information, hate speech, and online harassment. Now, please consider actions that online platforms could take to curb these issues. In the next 5 screens, you will see various ways that have been proposed as potential solutions to this issue. In each screen, please indicate which solution you **support the most** and which one you **support the least**.

They then saw five different subsets of the following potential ways to address false information, hate speech, and online harassment in online platforms:

Better enforcement: Better enforcement of existing policies against hate, violence, or disinformation

Content moderation: Train content moderators to better detect content that promotes hate, violence, or disinformation. Provide resources for moderators and workers as they engage in emotionally disturbing content moderation.

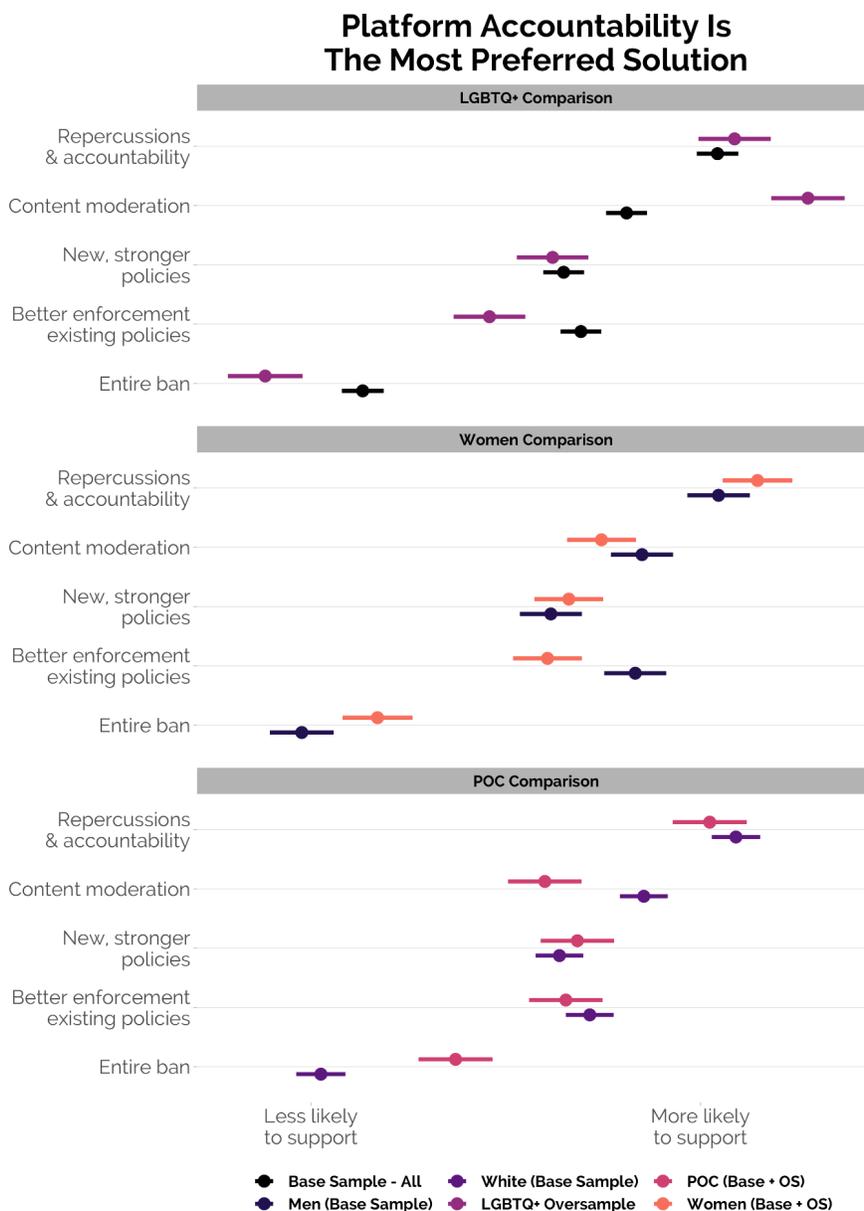
Entire ban: Entirely ban hateful language, violence, and disinformation from the platform

Repercussions & accountability: Create repercussions and accountability for social media companies who promote hateful language, violence, or disinformation

New, stronger policies: Create new, better, and stronger policies to deter hate, violence, or disinformation"

The MaxDiff scores for each item—across the base sample and demographic oversamples—are plotted below. The graph portrays the overall rank-ordering of items as well as the magnitude of the differences between them.

Overall, respondents felt that **repercussions and platform accountability** were the best solutions to the challenge of online hate performs the best across both the base sample and demographic oversamples, followed by platform-specific policies like improving content moderation or enforcing or creating new policies, and trailed by entirely banning hateful language, violence and disinformation from online platforms.



Respondents grew more supportive of regulation after they completed the survey

This survey started by asking respondents about their views on government regulation of social media, and this same question was repeated at the end of the survey. This design allows within-participant comparisons that help understand whether respondents moved either in favor or in opposition of government regulation of social media after taking the survey.²²

After the survey, respondents grew more supportive of platform regulation. This happens both in the base sample and in the substantive comparison groups. The graph below shows the percentage of respondents who agreed with each of the response options in the first ask (asked at the very start of the survey) and in the second ask (asked at the end of the survey).

At the start of the survey, 48 percent of respondents in the base sample said that social media companies should be regulated more than they currently are. After answering the survey, 56 percent of respondents in the base sample said that social media companies should be more regulated, a net increase of 8 percentage points.²³ A similar 8 percentage point increase in support for regulation happened in the LGBTQ+ oversamples, while women had a 7 percentage point increase and POC respondents had a 10 percentage point increase in their support for government regulation of social media companies.

In short, **the survey itself moved respondents to agree with increasing regulation, going from a plurality agreeing with this position to a majority agreeing with it after the survey** (and this movement is consistently observed across the four samples).

²² The full wording of the questions read as follows:

First ask (start of survey): Thinking about the role of the government in regulating social media companies, do you think they should be regulated...

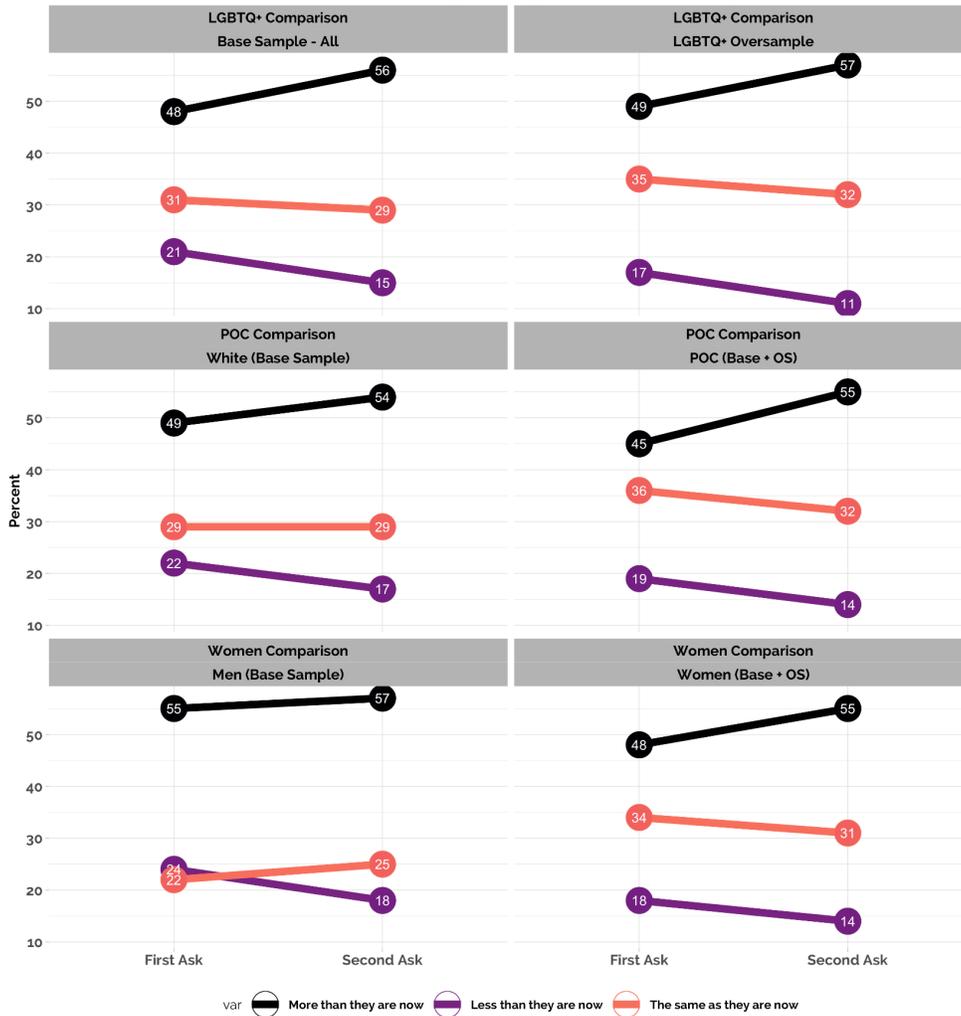
- More than they are now
- Less than they are now
- The same as they are now

Second ask (end of survey): After thinking about a topic during a survey, some people change their minds, while others do not. Thinking again about the role of the government in regulating social media companies, do you think they should be regulated...

- More than they are now
- Less than they are now
- The same as they are now

²³ When looking at the base sample alone, moreover, Hispanic and Black respondents, respondents who did not vote in the 2020 election, and those who use more social media sites were more likely to change their minds throughout the survey (regardless of direction). On the other hand, college-educated respondents were less likely to have changed their minds by the end of the survey. Partisanship, gender, and age had no effect here.

Survey Movement In Views About The Role Of The Government In Regulating Social Media Companies



Thinking about the role of the government in regulating social media companies, do you think they should be regulated.

Key takeaway

- Asking individuals to deeply think about their experiences with online harassment, incivility, and false information in this survey made them more open to supporting government regulation of social media companies. This movement is apparent both at the general population level, but also in the demographic groups of interest here.

Conclusion

The interrelated experiences of online harassment, hate speech, and false information are widespread across both the general American public and specific groups like LGBTQ+ folk, women, and people of color. Respondents in this survey generally recognize that these three phenomena are systemic problems within online environments, and see the potential ramifications of these across a number of domains, even when they report they recognize positive aspects of the online experience (like the ability to stay more connected with friends/family). Specifically, respondents in this survey believe that false information, which they see as widespread, leaves the public confused about basic facts and impacts overall trust in government and social trust in each other.

At the same time, respondents see others go through harassment often, and report having been subject to it themselves. When considering personal identities, moreover, the deleterious effects of the online environment become even more apparent – for example, LGBTQ+ respondents, POC, and women feel personally affected when they see public figures who share their identity get attacked online. This suggests that the harms of online harassment have impacts beyond the individual experiencing the actual harassment – in fact, such harassment can have a ripple effect on that individual's community.

Overall, this project also finds that Americans give bad marks to social media platforms when it comes to addressing hateful speech, harassment, and false information online. When exploring a host of potential solutions to the situation at hand, platform accountability outperforms other options. Respondents overall are in favor of platform accountability and government regulation of these sites. Indeed, after being asked to reflect deeply throughout this survey about the topics of hateful speech, harassment, and false information, respondents grew more supportive of platform regulation than they were at the very start of the survey.

Methods

Base sample methods statement

This survey is based on 1,235 interviews conducted by YouGov on the internet of respondents who are representative of the US population who are social media users. The sample was weighted according to gender, age, race/ethnicity, education, and US Census region based on voter registration lists, the U.S. Census American Community Survey, and the U.S. Census Current Population Survey, as well as 2020 Presidential vote. Respondents were selected from YouGov to be representative of the US population who are social media users. The weights range from 0.23 to 5.89 with a mean of 1 and a standard deviation of 0.52.

The margin of error (a 95% confidence interval) for a sample percentage p based upon the subsetting sample is approximately 3.1%. It is calculated using the formula:

$$\hat{p} \pm 100 \times \sqrt{\frac{1 + CV^2}{n}}$$

where CV is the coefficient of variation of the sample weights and n is the sample size used to compute the proportion. This is a measure of sampling error (the average of all estimates obtained using the same sample selection and weighting procedures repeatedly). The sample estimate should differ from its expected value by less than margin of error in 95 percent of all samples. It does not reflect non-sampling errors, including potential selection bias in panel participation or in response to a particular survey.

LGBTQ+ methods statement

This survey is based on 404 interviews conducted by YouGov on the internet of respondents who are LGBTQ+ and are social media users. The sampling process for this survey involved sampling to a universe of nationally-representative respondents, capturing responses based on key demographic variables to ensure we had a representative sample of the US population. Then, respondents who indicated they are not LGBTQ+ or do not use social media sites were screened out. For respondents who screened out, their demographic data was kept. A total of 6,618 respondents provided demographic information, 404 of whom qualified to take the full survey. All 6,618 respondents were weighted according to gender, age, race/ethnicity, education, and US Census region based on voter registration lists, the U.S. Census American Community Survey, and the U.S. Census Current Population Survey, as well as 2020 Presidential vote. Then, the respondents who were not LGBTQ+ social media users were dropped and weights were recentered. The oversample weights range from 0.39 to 2.29 with a mean of 1 and a standard deviation of 0.29.

The margin of error (a 95% confidence interval) for a sample percentage p based upon the subsetted sample is approximately 5.1%. It is calculated using the formula:

$$\hat{p} \pm 100 \times \sqrt{\frac{1 + CV^2}{n}}$$

where CV is the coefficient of variation of the sample weights and n is the sample size used to compute the proportion. This is a measure of sampling error (the average of all estimates obtained using the same sample selection and weighting procedures repeatedly). The sample estimate should differ from its expected value by less than margin of error in 95 percent of all samples. It does not reflect non-sampling errors, including potential selection bias in panel participation or in response to a particular survey.

POC methods statement

This survey is based on 403 interviews conducted by YouGov on the internet of respondents who are people of color who are social media users. For weighting and analysis, this project combined people of color from this oversample with the people of color who completed the survey in the base sample, to yield a total of 754 people of color. This sample was weighted according to gender, age, race/ethnicity, education, and US Census region based on voter registration lists, the U.S. Census American Community Survey, and the U.S. Census Current Population Survey, as well as 2020 Presidential vote. Respondents were selected from YouGov to be representative of people of color who are social media users. The weights range from 0.28 to 3.98 with a mean of 1 and a standard deviation of 0.50.

The margin of error (a 95% confidence interval) for a sample percentage p based upon the subsetted sample is approximately 4.0%. It is calculated using the formula:

$$\hat{p} \pm 100 \times \sqrt{\frac{1 + CV^2}{n}}$$

where CV is the coefficient of variation of the sample weights and n is the sample size used to compute the proportion. This is a measure of sampling error (the average of all estimates obtained using the same sample selection and weighting procedures repeatedly). The sample estimate should differ from its expected value by less than margin of error in 95 percent of all samples. It does not reflect non-sampling errors, including potential selection bias in panel participation or in response to a particular survey.

Women methods statement

This survey is based on 438 interviews conducted by YouGov on the internet of respondents who are women who are social media users. For weighting and analysis, this project combined the women from this oversample with the women who completed the survey from the base sample, to yield a total of 1,154 women. The sample was weighted according to gender, age, race/ethnicity, education, and US Census region based on voter registration lists, the U.S. Census American Community Survey, and the U.S. Census Current Population Survey, as well as 2020 Presidential vote. Respondents were selected from YouGov to be representative of women who are social media users. The weights range from 0.27 to 5.66 with a mean of 1 and a standard deviation of 0.51.

The margin of error (a 95% confidence interval) for a sample percentage p based upon the subsetted sample is approximately 3.3%. It is calculated using the formula:

$$\hat{p} \pm 100 \times \sqrt{\frac{1 + CV^2}{n}}$$

where CV is the coefficient of variation of the sample weights and n is the sample size used to compute the proportion. This is a measure of sampling error (the average of all estimates obtained using the same sample selection and weighting procedures repeatedly). The sample estimate should differ from its expected value by less than margin of error in 95 percent of all samples. It does not reflect non-sampling errors, including potential selection bias in panel participation or in response to a particular survey.