

**IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
AUSTIN DIVISION**

NETCHOICE, LLC d/b/a NetChoice,)
a 501(c)(6) District of Columbia organization,)
)
and)
)
COMPUTER & COMMUNICATIONS)
INDUSTRY ASSOCIATION d/b/a CCIA, a)
501(c)(6) non-stock Virginia Corporation,)
)
Plaintiffs,)
)
v.)
)
KEN PAXTON, in his official capacity as)
Attorney General of Texas)
)
Defendant.)
_____)

Civil Action No. 1:21-cv-00840-RP

**Exhibit C –
YouTube’s
Declaration**

**IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
AUSTIN DIVISION**

NETCHOICE, LLC, d/b/a NETCHOICE, a
501(c)(6) District of Columbia organization;
and COMPUTER & COMMUNICATIONS
INDUSTRY ASSOCIATION d/b/a CCIA, a
501(c)(6) non-stock Virginia corporation,

Plaintiffs,

v.

KEN PAXTON, in his official capacity as
Attorney General of Texas,

Defendant.

Civ. Action No. 21-cv-840

**DECLARATION OF YOUTUBE IN SUPPORT OF PLAINTIFFS'
MOTION FOR PRELIMINARY INJUNCTION**

I, Alexandra N. Veitch, declare as follows:

1. I am the Director of Public Policy for the Americas at YouTube. As part of my role, I lead a team that advises the company on public policy issues around online, user-generated content. My team advises on YouTube's content moderation policies and practices, identifies when changes to our policies or their application are required in response to new challenges, and assesses policy proposals and legislation, such as Texas's H.B. 20, that would affect YouTube's ability to moderate content.

2. The statements contained in this declaration are made upon my personal knowledge. I am over the age of 18 and am competent to make the statements herein. I make this Declaration in Support of Plaintiffs' Motion for Preliminary Injunction in the above-captioned matter. If called as a witness, I could and would testify under oath as follows.

3. YouTube is an online platform that allows users to create, upload, and share videos with others around the world. YouTube strives to be a community that fosters self-expression on an array of topics as diverse as its user base, and to nurture a thriving creative and informational ecosystem, as well as an engine of economic opportunity. Over two billion logged-in users worldwide visit each month, and over 500 hours of content are uploaded every minute by an extraordinarily diverse community of creators, who span over 100 countries and 80 languages. On a daily basis, users watch over a billion hours of video on YouTube.

4. YouTube is a part of Google LLC, a member of NetChoice and CCIA. YouTube does business in Texas and many of its users are located in Texas. Texas users have access generally to all content on YouTube that is available in the United States and worldwide.

Responsibility at YouTube

5. YouTube believes that the Internet is a force for creativity, learning, and access to information. Supporting the free flow of ideas is at the heart of YouTube's mission. We believe that the world is a better place when we listen, share, and build community through our stories. We strive to make YouTube as open as possible: to empower users to access, create, and share information. We believe that openness brings opportunity, community, and learning, and enables diverse and authentic voices to break through.

6. Yet an open platform means challenges, and it demands accountability to connect people with quality information. When you create a place designed to welcome many different voices, some will inevitably cross the line. Bad actors will try to exploit platforms for their own personal gain, even as we invest in the systems to stop and deter them. Harmful content on our platform makes YouTube less open, not more, by creating a space where creators and users may

not feel safe to share. We believe that, in order to have and protect openness, you must have responsibility. A commitment to openness is not easy. It sometimes means leaving up content that is outside the mainstream, controversial, or even offensive. But YouTube believes that hearing a broad range of perspectives ultimately makes us a stronger and more informed society, even if we disagree with some of those views. YouTube seeks to strike the right balance between fostering freedom of expression and decreasing the likelihood that users will encounter harmful content on our platform.

7. These beliefs and values drive the decisions we've made in building YouTube, and the editorial judgements we've made in crafting the content moderation tools and policies that protect our platform. We want YouTube to live up to the ideals of these values—despite challenges, complexities, and emerging threats. We work to maintain our community as a positive, open, and useful space on the Internet. Our balanced approach to content moderation, described below, represents these values. While important work remains to be done, this approach also represents years of ongoing conversations amongst YouTube and its users, creators, and advertisers, of the right balance for our products and businesses.

8. Responsibility is YouTube's number one priority. Indeed, our unique business model only works when our viewers, creators, and advertisers all have confidence that we are living up to our responsibility as a business. That responsibility has been critical to YouTube's success and essential to our continued growth, so we've invested heavily in hiring people and developing products, technology, and systems to apply our editorial discretion at scale.

YouTube's Approach to Responsibility and Content Moderation

9. YouTube takes a multi-faceted and nuanced approach to exercising its discretion in setting its content-moderation policies, working to distinguish those posts that are truly problematic, those that are borderline, and those that contribute positively to the YouTube community. To that end, we have a diverse set of tools to help us enforce our content-moderation policies, including: age-gating, removing videos and comments, appending warnings, and suspending and/or terminating accounts. We also have other tools to help us provide authoritative information on our platform - such as the use of information panels. And we further limit when YouTube makes recommendations of borderline content to users. Because removing content is only part of the discussion, YouTube has chosen to develop and invest in this diverse set of tools that are essential in balancing free expression and responsibility on our platform. Simply put, these tools give us broader options than simply removing (or not removing) content from our platform.

10. Yet H.B. 20 would eliminate much of our ability to make these kinds of choices in setting our policies and would subject YouTube and its community to serious harm by frustrating our ongoing efforts to make YouTube a far more accessible and welcoming place.

11. YouTube has always had policies that govern how people may use the service, including restrictions on the types of content that they may post. These policies are designed and regularly updated to make YouTube a safer and more enjoyable place for users and creators, and reflect years of experience, investment, and an ongoing conversation between YouTube and its users. YouTube's approach has four pillars, set forth below.

12. **First, we remove content** that violates our **Community Guidelines**, a series of clear, publicly-facing policies governing what is allowed and not allowed on our platform.¹ We work closely with outside experts to help us craft these policies (and their enforcement), primarily focused on preventing real-world harms. The Community Guidelines prohibit a variety of harmful, offensive, and unlawful material, such as hate speech, pornography, terrorist incitement, false propaganda spread by hostile foreign governments, promotion of fraudulent schemes, spam, egregious violations of personal privacy like revenge pornography, violations of intellectual property rights, bullying and harassment, conspiracy theories, and dangerous computer viruses. A full list of YouTube’s Community Guidelines is available at: <https://bit.ly/3CbToFY>.

13. We employ an array of remedial actions when enforcing our policies, ranging from demonetization (i.e., removing a creator’s ability to earn advertising revenue) and warnings, to service-usage penalties such as temporary suspensions of uploading rights and permanent termination of accounts. When an account uploads content that violates the Community Guidelines, the content is removed and the account generally receives a warning. Subsequent violative content can result in a “strike,” which temporarily suspends the account’s ability to upload content. Generally, three strikes within 90 days leads to the account’s termination and deletion of all content uploaded from the account. In the case of severe abuse (such as predatory behavior, spam, or pornography), YouTube will immediately terminate accounts to protect the YouTube community.

¹ We communicate our practices to all users through YouTube’s Community Guidelines, which are incorporated into our Terms of Service. A user must agree to both the Terms and the Community Guidelines in order to create an account and upload materials to YouTube.

14. **Second, we reduce the spread of harmful misinformation and content that brushes up against our policy lines.** We refer to content that comes close to violating our Community Guidelines (but does not) as “borderline content”. Borderline content is just a fraction of 1% of what is watched on YouTube in the United States, and examples include videos promoting a phony miracle cure for a serious illness or conspiracy theory videos (e.g., “the moon landing was faked”).

15. Rather than remove such content outright, we’ve chosen to take steps to reduce the spread of such content using a variety of methods. Because such borderline content may be disturbing or otherwise inappropriate for some viewers, YouTube has chosen to take action (using algorithms) to reduce its availability, including updating YouTube’s recommendations system, and disabling features like sharing, commenting, and liking for the borderline content. We set a high bar for what videos we display prominently in our recommendations on the YouTube homepage or through the “Up next” panel.

16. **Third, we raise authoritative and trusted content.** For subjects such as news, science, and historical events, we believe that accuracy and authoritativeness are key and the quality of information and context matter most (as compared to other topics such as music or entertainment, where we look to relevance, newness, and popularity). Here, content moderation can include affirmatively providing users with information to help them make choices about whether or not to interact with certain kinds of content. It is sometimes helpful to provide viewers with additional context about the content they are watching.

- **Information Panels.** We display a variety of information panels that provide users with context on content relating to topics and news prone to misinformation, as well as context

about who submitted the content. One example is an information panel displayed on videos from a channel owned by a news publisher that is funded by a government.² Another example is National Suicide Prevention Hotline information that we display in response to search queries for terms related to suicide. Information panels, across all types, have been collectively shown billions of times. The COVID-19 information panels alone have been shown over 400 billion times.

- **Breaking News.** Similarly, after a breaking news event, it takes time to verify, produce, and publish high-quality videos. Journalists often write articles first to break the news rather than produce videos. So YouTube has chosen to prioritize these articles and provides a short preview of news articles in search results on YouTube that link to the full article during the initial hours of a major news event.

17. **Fourth, we reward trusted, eligible creators by setting a higher bar for ads/monetization.** Users must meet additional eligibility requirements³ for the privilege of earning advertising revenue (“monetization”) on videos they upload. They must be eligible for, and join, the YouTube Partner Program (“YTPP”) and follow YTPP guidelines.⁴ Just over 2 million users worldwide, out of the 2 billion monthly users generally, are part of the YTPP and monetize their videos. Such users and their monetized videos also must meet more restrictive criteria, including the Ad-friendly Content Guidelines, because advertisers typically do not want to be associated with controversial or sensitive content on YouTube.⁵ Violations of the guidelines may result in a range of actions, such as (1) ads being disabled on a particular video, (2)

² <https://bit.ly/3fpnHzu>.

³ <https://www.youtube.com/howyoutubeworks/policies/monetization-policies/>

⁴ https://support.google.com/youtube/answer/72851?hl=en&ref_topic=9153826

⁵ <https://bit.ly/3ojt7B9>.

suspending or permanently disabling a user’s eligibility to monetize ads, (3) or, in exceptional circumstances, suspending or disabling a user’s account altogether to protect the integrity of the platform or protect our users from harm.

18. **Scale.** In Q2 2021 alone, YouTube removed over 4 million channels (or accounts), over 6 million videos, and over 1 billion comments, for violations of YouTube’s Community Guidelines alone.⁶ In Q2 2021, 29.9% of the videos removed were due to child safety issues. 55% of removed comments were due to spam.⁷ Further statistics (including others discussed in this declaration) may be found in the YouTube Community Guidelines enforcement report, updated quarterly.⁸

19. H.B. 20 significantly limits these ongoing efforts to prevent harm to our users and to make YouTube an accessible and welcoming place.

The Evolution of YouTube’s Content Moderation

20. YouTube has always had rules of what speech we permit on the platform, and we have never claimed that YouTube would host all user-generated content. YouTube has never allowed pornography, incitement to violence, or content that would harm children, for example.

21. The harms of user-generated content are ever-evolving and often unpredictable, and YouTube’s content moderation policies have necessarily had to evolve to address them. Each of our policies is carefully thought through (so they are consistent, well-informed, and can be applied to content from around the world), and often developed in partnership with a wide range

⁶ Of those videos, more than 30,000 contained misinformation about the COVID-19 vaccine. This was part of YouTube’s larger effort to remove medical misinformation about the virus, which resulted in the removal of over 1,000,000 videos related to dangerous or misleading COVID-19 information since February 2020.

⁷ YouTube uses automated systems to identify comments that are likely spam.

⁸ <https://bit.ly/2VhAsVG>.

of external industry and policy experts. We revise them regularly to account for new and different content or behavior that YouTube deems unacceptable, unsafe, or unwelcome on its service. YouTube has also invested significantly in being able to detect and respond quickly to emerging harms. YouTube's Intelligence Desk, an internal team, monitors news, social media, and user reports to detect these new trends—such as the unpredictable viral 'dares' that risk significant physical harm by, for instance, encouraging viewers to ingest Tide Pods—so as to address them before they become a larger issue. YouTube has over 100 people working to develop new content-moderation policies and improve existing ones.

22. This approach and investment has given YouTube flexibility to build and maintain responsible practices to handle legal but potentially harmful speech. In 2020, for instance, YouTube updated its policies related to medical misinformation alone more than ten times, which is in line with historical trends. In 2019, YouTube made over 30 updates to its content moderation policies generally—on average, once every 12 days. We saw a similar pace in 2018. And when necessary, YouTube is able to react quickly to promote the safety of its users in changing and emerging contexts. For example, when mobile phone towers in the U.K. were set on fire after a conspiracy theory video blamed COVID-19 on 5G wireless networks, we updated our Community Guidelines in a single day to ban and remove that harmful content.

23. YouTube's judgments evolve over time as social and cultural conditions change or unforeseen threats and challenges arise. For instance, after a recent violent military coup in Myanmar, YouTube took action against five existing YouTube channels run by the Myanmar military, terminating the channels to prevent the military from promoting political propaganda.⁹

⁹ <https://nyti.ms/3xoq0IW>.

Algorithms and Machine Learning

24. YouTube’s engineers have designed and built sophisticated software systems using machine learning—a type of algorithm—to moderate content in two key ways: 1) to proactively identify and flag potentially harmful content uploaded to the site, and 2) to automatically remove content that is identical or substantially similar to violative content that was previously removed. Machine learning is the product of human decision-making and is used to implement the standards set in our Community Guidelines, thereby reflecting YouTube’s editorial judgments. Our engineers design these systems to identify certain types of content. We then use data inputs (reflecting the judgment of human reviewers) to train these machine learning systems to identify patterns in content—both the rich media content in videos, as well as textual content like metadata and comments—so that our systems can make predictions and find new examples to match the identified types of content. Machine learning is well-suited to detecting patterns, which helps us to identify new content similar to that we have already removed, even before it is ever viewed. We also use hashes (or “digital fingerprints”) to automatically identify copies of known violative content before they are ever made available for viewing.¹⁰ These systems automatically remove content only where there is high confidence of a policy violation—*e.g.*, spam—and flag the rest for human review. Algorithmic detection identifies the vast majority of content deemed to violate the Community Guidelines.

25. Machine learning is critical to implementing all aspects of YouTube’s approach to content moderation and keeping our users safe. YouTube relies heavily on technology and algorithms to moderate content and cannot feasibly do otherwise, since over 500 hours of video

¹⁰ In Q1 2021, 27.8% of removed videos were taken down before a single view. A further 39% of removed videos had between 1 and 10 views. <https://bit.ly/3fpoLmY>

are uploaded to YouTube every single minute of every day. At this massive scale, it would be virtually impossible to remove content that violates our Community Guidelines without the use of algorithmic tools, even with tens of thousands of reviewers watching newly uploaded videos 24 hours a day, 7 days a week. Due to large multi-year investments in machine learning algorithms, since 2017 we have seen a 70% drop in the quarterly estimate of the number of views for video deemed violate to our policies (known as the violative view rate, “VVR”).¹¹

26. The vast majority of Community Guidelines violations were flagged by algorithms. In Q2 2021, YouTube removed 6,278,771 videos that violated the Community Guidelines. The vast majority—5,927,201, or 94% of the total removals—were automatically flagged for moderation by YouTube’s algorithms. About 5%—351,570 videos—were removed based on initial flags by a user or other human. This removal system is highly efficient: the majority of removed videos were removed before accumulating more than 10 views. Similarly in Q2 2021, YouTube also removed over 1 billion comments, 99.5% of which were flagged for moderation by YouTube’s automated systems.

27. Our machine learning and human reviewers work hand in hand: machine learning is effective for scale and volume, whereas human reviewers can evaluate context for more nuanced enforcement of our policies. Once our machine learning systems flag a potentially violative video without high confidence of a policy violation, human reviewers assess whether the content does indeed violate our policies, and remove those that do. In making those judgment calls, the reviewers seek to protect content that has an educational, documentary, scientific, or artistic purpose, keeping such videos on the platform. These human decisions and judgments are

¹¹ See <https://bit.ly/38noixm>.

in turn used as data inputs to improve the accuracy of our automated detection systems so that we are constantly updating and improving the system's ability to identify potentially violative content. Using that human review, our machine learning systems can automatically remove re-uploads of content that has already been reviewed and determined to violate our policies. In addition, when we introduce a new policy or alter an existing one, it takes our systems time to improve detection rates and begin accurately detecting violative content at scale. Our enforcement of new policies improves over time.

Further Examples of Our Values Embodied in YouTube's Content Moderation Processes.

28. During summer 2020, YouTube faced a dilemma when confronting the tension that arises between 1) accuracy when enforcing content policies and 2) the need to limit potentially harmful content accessible on the site. In response to COVID-19 lockdowns worldwide, YouTube took steps to protect the health and safety of our extended workforce and reduced in-office staffing. As a result of reduced human review capacity, YouTube had to choose between limiting enforcement while maintaining a high degree of accuracy, or relying on automated systems and algorithms to cast a wider net to remove potentially harmful content quickly but with less accuracy. Because of YouTube's belief that responsibility is critical, YouTube chose the latter, despite the risks that automation would lead to over-enforcement—in other words, removing more content that may not violate our policies for the sake of removing more violative content overall.

29. For certain sensitive high-risk policy areas, such as violent extremism and child safety, YouTube chooses to accept a lower level of accuracy to remove as many pieces of violative content as possible (again, to protect the health and safety of our extended workforce

and reduced in-office staffing). This also means that, in these areas specifically, a higher amount of non-violative content was removed. YouTube’s decision to over-enforce in these policy areas—out of an abundance of caution—has led to a more than 3x increase in removals of content that our systems suspected was tied to violent extremism or potentially harmful to children. These include dares, challenges, or other posted content that may endanger minors. Moreover, YouTube will immediately suspend users for egregious violations (rather than allowing a user multiple ‘strikes’).

30. **EDSA.** Because YouTube values creativity and learning, our content policies have an exception for videos that would otherwise be in violation if there is a compelling educational, documentary, scientific, or artistic reason that is apparent in the content or context of the video. YouTube refers to this exception as “EDSA,” which is a critical way to make sure that important speech remains on YouTube, while simultaneously protecting the wider YouTube ecosystem from harmful content.¹² These decisions depend on a variety of factors that depend on context and require nuanced judgments, and the bar varies by video and policy category. For example, hate speech and encouragement of violence violate our policies but a documentary about WWII that features speeches from Nazi leaders may be allowed if the documentary provides historical context and does not aim to support the despicable views promoted by the Nazis. There are also certain types of content where we don't allow an EDSA exception under any circumstances because of the sensitivity and egregiously harmful nature of the content, or when it violates the law. For example, content that endangers children or any content with footage of deadly violence filmed by the perpetrator is not allowed on YouTube regardless of the context.

¹² <https://bit.ly/2VhM7DW>

Transparency

31. Given YouTube’s scale, we sometimes make mistakes, which is why creators can appeal video removal decisions. YouTube generally notifies creators when their video is removed, and we provide a link with instructions on how to appeal the removal decision. If a creator chooses to submit an appeal, the video goes to human review, and the decision is upheld, reversed, or modified (modification leads to reinstatement of the video but with restricted access). We provide transparency about our appeals process. As reported in our most recent Transparency Report, in Q2 2021, creators appealed approximately 217,446 videos, or 3.5% of all videos removed. Of those, more than 52,696 were reinstated.

The Burdens Posed by H.B. 20

32. I understand that on September 9, 2021, the State of Texas enacted H.B. 20, which will go into effect on December 2, 2021.

33. The restrictions of H.B. 20 would fundamentally burden and undermine YouTube’s ability to operate responsibly and enforce the content-moderation policies described above. The statute has a broad definition of “censorship” (“to block, ban, remove, deplatform, demonetize, de-boost, restrict, deny equal access or visibility to, or otherwise discriminate against expression.”) that covers YouTube’s broad portfolio of content-moderation tools (reflecting our judgment and discretion) across a broad variety of topics.

34. “Expression” is defined broadly by H.B. 20, and would include any and all user-generated content on YouTube.

35. For instance, YouTube simply “blocks” or “removes” certain speech like hate speech that violates our Community Guidelines’ policy on hate speech. But because hate speech

expresses “viewpoints”—as abhorrent as those viewpoints are—H.B. 20 would bar YouTube from taking any content moderation action against such content, such as removing it, age-restricting it, or demonetizing it.

36. H.B. 20’s “censorship” prohibition will directly prevent YouTube from enforcing critical standards designed to prevent the degradation of our users’ experiences on the platform and to ensure their safety, including for children. YouTube needs discretion and flexibility when designing, building, and maintaining our content-moderation policies because it encounters such a broad range of content, and at such high volumes. As described above, YouTube’s Terms of Service, Community Guidelines, and other content-moderation rules include flexible terms that allow YouTube to exercise its judgment about specific uses or pieces of content in order to provide a better and safer user experience.

37. While H.B. 20 contains certain content exceptions chosen by the Texas Legislature under Section 143A.006, the state’s narrow choices mean that the broad restrictions on content moderation would still eliminate wholesale many of the categories of content (in both our Community Guidelines and Advertising policies) that YouTube has chosen to moderate.

38. YouTube currently has numerous viewpoint-based policies against many kinds of harmful content, for which H.B. 20 has no applicable exception. For example, YouTube’s Community Guidelines has a Violent Criminal Organizations policy¹³ under which YouTube currently removes content produced by violent criminal or terrorist organizations (“VCTOs”), content praising or justifying violent acts carried out by VCTOs, content aimed at recruiting members for VCTOs, or hostage videos. In order to comply with H.B. 20, YouTube would have

¹³ bit.ly/3m0tMVo.

to stop removing such violent extremist content. Similarly, paragraphs 41-45 below discuss examples of additional categories ranging from dangerous pranks risking imminent harm, drug use, suicide/self harm, animal abuse, and medical misinformation.

39. H.B. 20 seems to allow moderation of content that “directly incites criminal activity or consists of specific threats of violence targeted against a person or group.” But this limited exception actually excludes many categories found in YouTube’s Community Guidelines hate speech policy.¹⁴ 143A.006(3). For example, YouTube removes content “promoting violence or hatred against individuals or groups based on,” among other things, veterans status or sexual orientation. H.B. 20 would stop YouTube from taking action against, for example, content promoting violence or hatred against veterans. Even in the categories that H.B. 20 enumerates, H.B. 20 would still bar YouTube from taking action against content “promoting violence or hatred” without a specific threat of violence.

40. Reflecting our view of the nuance involved in balancing freedom of expression and responsibility, YouTube has chosen to build systems and processes that apply different standards for different content-moderation actions. For example, we apply the Community Guidelines for removals, and the Ad-friendly Content Guidelines for demonetization. We also age-restrict, reduce availability or functionality, or restrict other borderline content (which otherwise remains available on our platform). By treating all these actions as prohibited “censorship,” H.B. 20 will eliminate YouTube’s discretion to find the right balance between free expression on YouTube and responsibility for fostering a safe community for its users. The

¹⁴ https://support.google.com/youtube/answer/2801939?hl=en&ref_topic=9282436

following are examples showing the nuance and complexity of YouTube's content moderation policies applied in contexts where H.B. 20 would prohibit YouTube from taking action.

41. Dangerous Pranks. Under our Community Guidelines we remove videos depicting extremely dangerous challenges that pose an imminent risk of physical injury, such as the well-known "Tide Pod" challenge. Another example is the "No Lackin'" challenge, where people post videos of themselves pointing guns at others.¹⁵ Because YouTube is concerned that minors could easily imitate such challenges, we may allow, but age-restrict, content that explains these challenges in an educational or documentary way. However, YouTube may allow, without restriction, a video warning minors against performing such challenges. H.B. 20 would require YouTube to treat each of these examples of dangerous prank-related content equally and leave all of them up on our platform.

42. Drug Use. Under our Community Guidelines, we remove videos with depictions of the use of hard drugs (like intravenous heroin injection), and depictions of minors using *any* alcohol or drugs (using vaporizers, e-cigarettes, tobacco, or marijuana). Still, we may allow videos that discuss the scientific effects of drug use, content that does not promote or glorify drug usage (e.g., a personal story about the opioids crisis), or news reports about drug busts (with no visible consumption or distribution). Such content, especially if it shows the injection of drugs, may still be age-restricted. H.B. 20 would require YouTube to treat each of these different examples of drug use-related content equally and leave all of them up on our platform.

43. Suicide. Our Community Guidelines prohibit (1) videos promoting or glorifying suicide, (2) providing instructions on how to self-harm or die by suicide, and (3) graphic images

¹⁵ News articles report that this challenge was involved in one 2019 death in the Houston area. <https://abc13.com/no-lackin-challenge-teen-shooting-killed-playing-with-guns/5009272/>

of self-harm posted to shock or disgust viewers. Still, we may permit, without advertising, videos with first-person accounts (e.g., a biography or detailed interview on survivors and their pasts) and detailed descriptions of suicide. Further, for searches for terms related to suicide, YouTube shows authoritative content helping users connect with the National Suicide Prevention Hotline. H.B. 20 would require YouTube to treat each of these different examples of suicide-related content equally and leave all of them up on our platform.

44. Animals. Under our current Community Guidelines, we remove depiction of content that includes a human maliciously causing an animal to experience suffering, or where animals are encouraged or coerced to fight by humans. Under our Ad-Friendly Content Guidelines, we demonetize, but allow, videos with graphic depictions of skinning or slaughtering animals. We permit advertising on videos portraying animal preparation for eating by professionals focusing on the trade and act of cutting animals, or the preparation of meat or fish (such as BBQ cooking techniques). H.B. 20 would require YouTube to treat each of these different examples of animal-related content equally and leave all of them up on our platform.

45. Medical Misinformation. YouTube does not allow certain types of misleading or deceptive content with serious risk of egregious harm, like medical misinformation (such as content claiming that harmful substances or treatments can have health benefits). This includes content about COVID-19 that poses a serious risk of egregious harm, such as treatment misinformation. One example is content that promotes drinking “mineral miracle solution (MMS)” as a treatment for COVID-19. The FDA has warned that “MMS Consumers Are Drinking Bleach” since “when mixed according to package directions, [MMS products] become

a strong chemical that is used as bleach.”¹⁶ H.B. 20 would bar YouTube from taking any content moderation action against content expressing these viewpoints.

46. More generally, much of what YouTube does is to vary “access or visibility” to certain pieces of content—or certain classes of content—according to subjective judgments about the viewpoint expressed in the speech in accordance with its policies and what YouTube believes will be most relevant to individual users.

47. Because H.B. 20’s definition of “censor” includes “restrict” and “deboost,” H.B. 20 would prohibit YouTube’s approach to borderline content—content that, in our judgement, comes close to violating our Community Guidelines. Rather than remove this content entirely, YouTube currently takes steps to reduce the spread and restrict its availability (rather than remove the content outright). In 2019, we changed our recommendation system to reduce suggesting such borderline content to users.

48. YouTube has designed our search ranking systems and algorithms to prioritize different factors depending on the search term requested. In areas such as music or entertainment, we often use relevance, freshness, or popularity to rank search results. In other areas where veracity and credibility are key, including news, politics, and medical or scientific information, our search systems prioritize surfacing authoritative content from trusted sources. For example, when you proactively search for news-related topics, a Top News section will appear near the top of search results, which raises relevant results from authoritative voices including news sources like CNN and Fox News.

¹⁶ <https://bit.ly/3kNf8BF>.

49. So H.B. 20 will forbid YouTube from making both *individualized decisions* that perhaps one user will prefer certain content relative to other content because of the “viewpoints” expressed in that content; and *broad decisions* that certain content should be emphasized or deemphasized across all users.

50. H.B. 20’s definition of censorship includes action to “demonetize” based on viewpoint. Currently, YouTube requires that users wishing to monetize their content comply with Community Guidelines, but also an additional set of viewpoint-based guidelines, the Advertiser-friendly Content Guidelines. H.B. 20 would bar YouTube from enforcing these guidelines, and prevent YouTube from demonetizing harmful/offensive content. YouTube would be forced to continue to let a harmful content creator earn advertising revenue off YouTube’s platform and thus encourage that creator to upload as much harmful and offensive content as quickly as possible.

51. Finally, H.B. 20 prohibits YouTube from engaging in its *own speech* because it prohibits YouTube from “otherwise discriminat[ing]” against user-submitted expression. This provision—as vague and broad as it is—encompasses situations in which YouTube appends its own expression to user-submitted content, whether to express disagreement with or disapproval of that expression, or to add context YouTube believes is necessary for certain topics prone to misinformation. For certain content (e.g., potential hate speech) that is both close to the Community Guidelines line for removal and is offensive to viewers, YouTube adds a warning message before viewers can watch the video. Because YouTube will only append its own expression based on the “viewpoint” expressed in the content, that would constitute censorship under H.B. 20. Similarly, YouTube displays a variety of information panels that provide users

with context on content relating to topics and news prone to misinformation, as well as context about the publishers of the content.

52. Therefore, YouTube will face an impossible choice between (1) risking liability by moderating content identified to violate its standards or (2) subjecting YouTube's community to harm by allowing violative content to remain on the site.

Other Impact

53. **Age Gating, Restricted Mode, and YouTube Kids.** YouTube provides features, tools, and age-gated offerings to sensitive users and organizations (such as libraries and families with young children). These features are a way for YouTube to balance free expression with responsibility. For example, YouTube uses age-gating, a process whereby certain content—such as material featuring sexual situations, heavy profanity, or graphic depictions of violence—is made inaccessible to users under age 18. In order to view this content, users coming to YouTube must be signed-in and the age associated with their account must be 18 or older in order to view the video. YouTube also has a feature called Restricted Mode, an optional setting that sensitive users can choose to use to limit the content they see on YouTube. It is also used by libraries, schools, and public institutions. Videos containing potentially adult content like drugs or alcohol use, sexual situations, or violence are not shown to users in Restricted Mode.¹⁷ YouTube also produces an app called YouTube Kids, which includes only videos that are determined to be suitable for children through a combination of human and algorithmic review, and which blocks access to comments more suitable for adults. For example, YouTube Kids does not show videos

¹⁷ <https://bit.ly/3jiTW11>.

with paid product placements or endorsements, nor overly commercial or promotional videos. Over 35 million weekly viewers in more than 100 countries use YouTube Kids.

54. H.B. 20's prohibition on "censorship" includes "restricting" content. Complying with that requirement would force Restricted Mode and YouTube Kids to display all content, even if that content would otherwise be violative of YouTube's policies, or is content that YouTube (and a reasonable user would) believe in its judgment to be inappropriate for those audiences. Similarly, YouTube would have to stop age-gating such content. These changes would contradict the purpose of these features and products to give parents options for increased safety, forcing YouTube to make age-inappropriate content available to minors generally, and to other users choosing to use Restricted Mode.

55. **Disclosure and Notice Requirements.** The "disclosure" and operational restrictions will likewise burden YouTube's discretion in designing its content-moderation systems and processes. While YouTube endeavors to be transparent with its users and creators, this law would impose ambiguous and wide-ranging transparency requirements on all of YouTube's decisions to remove content of any kind. For example, these transparency requirements would apply to all types of content—not just videos—on YouTube. When removing videos under the Community Guidelines, YouTube generally provides users with notice, a complaint system, and an ability to appeal—but it does not currently provide any of this when removing comments.

56. To comply with H.B. 20, YouTube would have to expand these systems' capacity by over 100X—from a volume handling millions of removals to that of over a billion removals: during the last quarter (Q2 2021), YouTube removed 9.5 million videos and well over 1.16

billion comments. YouTube would have to provide notice of each of these 1.16 billion decisions to remove a comment. When any users receiving notice complain about, or appeal, those 1.16 billion removal decisions, YouTube will have to handle those requests within an accelerated response period.

57. Though YouTube endeavors to be transparent about its Terms of Service, Community Guidelines, and other content moderation practices generally, H.B. 20 does not explain the level of “specific information” required by the public disclosures section. For example, it seeks public disclosure of “search, ranking, or other algorithms or procedures.” Public disclosure of that aspect (and others) of YouTube’s content moderation would risk revealing its trade secrets and other confidential intellectual property to our competitors, since YouTube relies on sophisticated proprietary software systems, including machine learning algorithms, in which YouTube has invested significant resources to build and develop. Moreover, detailed disclosure of technical details of our enforcement methods would risk empowering the unscrupulous users seeking gaps and weaknesses in our systems for exploitation and to evolve their tactics to evade our efforts. For these reasons, YouTube does not publicly disclose these kinds of technical details.

58. H.B. 20 requires a biannual transparency report calling for expansive though ambiguous disclosure including, for example, whenever YouTube took action including “any other action taken in accordance with the platform’s acceptable use policy,” including detailed breakdowns by rule violated and source of alert. At the immense scale that YouTube operates, this level of granular reporting of every content-moderation decision would be extremely burdensome.

59. The specter of liability from countless private lawsuits (only for the anti-editorial-discretion provisions) and Attorney General enforcement (for all of the provisions) will substantially chill YouTube's use of editorial discretion to moderate content.

60. **User Scope.** H.B. 20 prohibits "censoring" a Texas "user's ability to receive the expression of another person," and that "person" need not be in Texas. YouTube has no way to comply without altering its editorial policies platform-wide, because YouTube's Community Guidelines are enforced consistently across the globe, regardless of where the content is uploaded. When content is removed for violating YouTube's Community Guidelines, it is removed globally.

61. **Harm to YouTube.** To comply with this law, YouTube would have to eliminate many, if not most, of our content-moderation standards that currently apply to any video and comment posted platform-wide. Users will leave YouTube for platforms that are able to responsibly moderate their platforms. Controversial content generally does not perform well with users on YouTube (compared to other categories like music or comedy). Advertisers do not want their brands associated with problematic content and actors. We've seen first-hand that when advertisers lack trust in our systems, they scale back their spend on YouTube. In response to several prior incidents involving extremist, child exploitation, and other harmful content, advertisers (who do not want their advertisements next to objectionable content) have stopped advertising on YouTube. Loss of advertiser trust negatively impacts creator earnings (since that revenue is dependent upon the willingness of advertisers to associate their brands with YouTube content), causing creators, too, to seek alternative platforms. The cost of not taking sufficient action over the long term results in lack of trust from our users, advertisers, and creators. Past

egregious actions of just a handful of creators have harmed the reputation of YouTube and the creator community among advertisers, the media industry and most importantly, the general public. When just one creator does something particularly blatant—like conducts a heinous prank where people are traumatized, promotes violence or hate toward a group, demonstrates cruelty, or sensationalizes the pain of others in an attempt to gain views or subscribers—we have seen how it can cause lasting damage to the community, including viewers, creators and the outside world.

62. This harm is why responsibility is critical to YouTube’s success, and is our number one priority. YouTube has responded to these past incidents by updating the way we moderate content with stricter policies, better controls, and greater transparency. We’ve made much progress to earn trust, recognizing more can and should be done. Yet H.B. 20 would unilaterally replace much of this entire framework to content moderation and runs contrary to user safety and enjoyment of the user experience.

I declare under penalty of perjury under the laws of the United States of America, pursuant to 28 U.S.C. § 1746, that the foregoing to be true and correct to the best of my knowledge. Executed on this September 30, 2021 in Washington, DC.



Alexandra N. Veitch