

**Exhibit 4 –**  
Declaration of  
Alexandra Veitch (YouTube)

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

COMPUTER & COMMUNICATIONS	)	
INDUSTRY ASSOCIATION, and	)	
	)	
NETCHOICE, LLC,	)	
	)	
<i>Plaintiffs,</i>	)	
	)	Civil Action No. _____
v.	)	
	)	
KEN PAXTON, in his official capacity as	)	
Attorney General of Texas,	)	
	)	
<i>Defendant.</i>	)	

**DECLARATION OF ALEXANDRA N. VEITCH 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 monitors and assesses the U.S. legislative and regulatory landscape, including policy proposals and legislation, such as Texas House Bill 18 (“HB18” or “Act”) that would affect YouTube’s ability to provide services and content to users in Texas—and perhaps worldwide. My team also leads external advocacy for YouTube’s policies and practices with government leaders and policy stakeholders, advises the company on public policy issues as it relates to online, user-generated content, and represents the company in public policy forums, across a broad region.

2. I submit this declaration in support of Plaintiffs’ Motion for Preliminary Injunction. I am over the age of 18 and am competent to make the statements herein. I have personal knowledge of the facts set forth in this declaration and, if called and sworn as a witness, could and would competently testify to them.

3. YouTube is an online service 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 a result, YouTube has a wealth of highly valuable and enriching user-generated content, providing users with a wide array of news, educational content, and entertainment.

4. People in Texas use YouTube’s services.

5. I understand that Google and YouTube may be subject to the Act, as one or both could meet the definition of “digital service provider” that own and operate a “digital service.” Tex. Bus. & Com. Code § 509.001(1)-(2).

#### **I. YouTube’s Principled Approach to Minors and Responsibility**

6. On YouTube, “[c]hildren and teenagers today can access a world of possibilities”: “Whether it’s exploring important topics around the world or looking up a video to help with algebra homework, they’ve never known a reality without this world at their fingertips.” Neal Mohan, *YouTube’s Principled Approach for Children and Teenagers*, YouTube Official Blog (Oct. 16, 2023), <https://perma.cc/D9ZQ-JRM9>. We believe that “[f]amilies everywhere deserve the same safe, high quality experience online, no matter where they live. And all children and teenagers ought to have the same access to the opportunities the internet provides.” *Id.*

7. YouTube has invested, and continues to invest, immense resources in developing and fostering age-appropriate services, features, and policies, making age-appropriate user-generated content accessible for its users. Five principles guide YouTube’s approach to minors:

- “The privacy, physical safety, mental health, and wellbeing of children and teenagers require special protections online”;
- “Parents and caregivers play an important role in setting the rules for their family’s online experiences, particularly for the youngest children”;
- “All children and teenagers deserve free access to high-quality and age appropriate content that meets their individual interests and needs”;

- “The developmental needs of children differ greatly from those of teenagers and should be reflected in their online experiences”; and
- “With appropriate safeguards, innovative technologies can benefit children and teenagers.”

*Id.*

8. Our Youth Principles are just one part of the work we do to foster a vibrant, safer, and open community for users. To successfully support this community, YouTube must constantly balance disseminating a wide range of user-generated content against addressing objectionable and harmful speech. On one hand, the Internet is a force for creativity, learning, and access to information. The world is a better place when we listen, share, and build community through our stories. We strive to empower users to access, create, and share information. We believe that a free flow of ideas fosters opportunity, community, and learning, and enables diverse and authentic voices to break through.

9. But on the other hand, an open Internet poses challenges. Bad actors exploit the Internet, including YouTube, for their own personal gain or for purposes of causing harm, even as we invest in the policies and systems to stop and deter them. Objectionable and harmful content on our website also makes YouTube less open, not more, by creating a space where creators and users may not feel safe to share, learn, and experience. We believe that, in order to have and protect openness, a service should attempt to strike the right balance between fostering freedom of expression and decreasing the likelihood that users will encounter harmful user-generated content on our service.

10. Consequently, working to keep harmful and illegal content off our services is core to the work of many different teams across Google and YouTube. When it comes to the information and user-generated content on our services, we discharge that responsibility based on clear, transparent policies and processes.

**A. YouTube’s Age-Appropriate Experiences and Policies for Minors**

11. YouTube has long invested in research, policies, and practices to offer age-appropriate ways for minors to explore, learn, and participate in the online world. Such efforts are designed with consideration of our youngest users’ wellbeing. We also provide tools and guidance that offer families flexibility to choose the right experience for them, enabling them to create healthy and positive digital habits at home and in school. We believe that appropriate safeguards can empower young people and help them learn, grow, and prepare for the future. Accordingly, we (1) invest in age-appropriate services and features that align with children’s and teens’ developmental stages and needs; (2) offer tools that give families flexibility to manage their relationships with technology; (3) implement policies, protections, and programs that aim to increase online safety for children and teens; and (4) provide resources that teach the fundamentals of digital citizenship and media literacy to allow for confident, safer online exploration.

12. We also prioritize making robust resources and guides available for children, teens, and their parents so they can better understand YouTube’s services. *E.g.*, YouTube Help, Understand Your Choices as a Family, <https://perma.cc/JV49-8YVC>; YouTube Help, What is a Supervised Experience on YouTube?, <https://perma.cc/D8KB-MX5D> (“YouTube Help, What is a Supervised Experience”).

13. Users known to be under the age of 13 are required to have their parent’s consent to their having an account. Further, the parent’s account remains linked to their child’s account so that they can help manage their child’s YouTube experience. When a parent sets up a child’s Google Account, the parent is given the option to either allow the child to access only YouTube Kids (a standalone service described in detail below) or set up a “Supervised Experience” on YouTube. YouTube offers varying content levels and features that are meant to align with kids’ and teens’ developmental stages and give families tools to supervise children’s use of the website.

Developed in partnership with child-development experts, these services and experiences reach more than 100 million active viewers worldwide every month:

a. **YouTube Kids:** YouTube’s standalone service “YouTube Kids” is a filtered version of YouTube built specifically to let children under the age of 13 explore age-appropriate content, with additional tools for parents and caregivers to moderate the content children see. YouTube Kids, A Safer Online Experience for Kids, <https://perma.cc/3XW9-82UL>. YouTube Kids features a carefully curated selection of “content that is age-appropriate, adheres to our quality principles, and is diverse enough to meet the varied interests of kids globally.” *Id.* YouTube works “with a variety of external specialists in child development, children’s media, and digital learning and citizenship” to determine what content should be available on YouTube Kids. *Id.* As a result, kids have access to a broad range of high-quality content. *See* James Beser, *Enabling a High Quality Kids Experience & Helping Kids Creators Thrive on YouTube*, YouTube Official Blog (Dec. 16, 2022), <https://perma.cc/XVM4-ZNHF> (“Beser, *Enabling a High Quality Kids Experience & Helping Kids Creators Thrive on YouTube*”). In addition to other tools discussed below, our built-in timer in YouTube Kids lets parents or caregivers, at their discretion, “limit screen time by telling kids when it’s time to stop watching.” *See* YouTube for Families Help, Limit Screen Time on YouTube Kids, <https://perma.cc/WH7E-RZTP>. “The timer will display a friendly alert and stop the app when the session is over” so that parents or caregivers do not have to. *Id.*

b. **Supervised Accounts:** YouTube also provides “supervised accounts” that are linked to parents’ accounts for children (up to the relevant age of consent) whose parents decide their children are ready to explore some of YouTube’s vast universe of user-

generated content. In these parent-managed accounts, and at the parent or caregiver's discretion, "parents [can] select a content setting that limits the videos and music [their] children under 13 can find and play." YouTube Help, What is a Supervised Experience. For example, parents can, among other things, (1) "block specific channels"; (2) change their minor child's permissible "content level settings," including "Explore" (generally for viewers 9+), "Explore More" (generally for viewers 13+), and "Most of YouTube" (almost all videos on YouTube except for videos marked only for adults or otherwise not appropriate); and (3) review, pause, and clear their "child's watch history." YouTube Help, Parental Controls & Settings for Supervised Experiences on YouTube, <https://perma.cc/8HUG-HFVM>. "Supervised accounts also change the features they can use, the default account settings, and the ads they see." YouTube Help, What is a Supervised Experience. For instance, functionalities like creating content and posting content are disabled on the Supervised Experience. *Id.* Further, when parents create a Google Account for their child with Family Link and use YouTube's Supervised Experience, parents and caregivers can set screen time limits on their Android device or Chromebook. Google Family Link, FAQ, <https://perma.cc/GR6Q-XBDE>. Parents and caregivers can set their child's Android device or Chromebook to lock after the child has used it for a certain amount of time or when the parents or caregivers think their child needs downtime. *Id.* Parents set up their own accounts to serve as managing accounts. YouTube for Families Help, Get Started with Supervised Accounts, <https://perma.cc/8F6M-P6QY>.

14. YouTube also offers a suite of digital wellbeing tools that encourage children and teens to be mindful of their screen time. These include:

a. **Autoplay:** YouTube’s autoplay feature is turned off by default for all users we know to be under 18 across all of YouTube’s services, and parents or caregivers can choose to disable autoplay on YouTube Kids and YouTube’s Supervised Experiences. YouTube Help, Autoplay Videos, <https://perma.cc/4Q44-LRWN>.

b. **Recommendations on YouTube for Teens:** We have worked to refine our recommendation systems so that teens on YouTube are not overly exposed to user-generated content that, while perhaps acceptable as a single video, could be considered differently if viewed in high quantities. Through consultation with our Youth & Families Advisory Committee and with input from academic research, we “define[d] categories of videos that are innocuous in a single view, but that could be problematic if viewed in repetition for some young viewers” which “include content that”: (1) “Compares physical features and idealizes some types over others”; (2) “Idealizes specific fitness levels or body weights”; and (3) “Features social aggression (non-contact fights and intimidation).” YouTube, Building Content Recommendations to Meet the Unique Needs of Teens and Tweens at 3, <https://perma.cc/K283-W863>. We have implemented guardrails for teens to limit repeated recommendations of videos related to these topics.

c. **“Take a Break” Reminders:** These reminders pause a video until a user dismisses or resumes playing the video. This feature is turned on by default for all users we know to be under 18 across all of YouTube’s services, including on Supervised Experiences. YouTube Help, Take a Break Reminder, <https://perma.cc/Q39X-65AW>.



d. **Bedtime Reminders:** Users can set a specific time to receive a reminder to stop watching YouTube and go to bed. This reminder is turned on by default for all users we know to be under 18 across all of YouTube’s services, including on Supervised Experiences. YouTube Help, Set a Bedtime Reminder, <https://perma.cc/WGM9-WDM2>.

## **B. Partnerships With Experts**

15. YouTube has ongoing partnerships with independent experts in fields such as child development, emerging media, and digital wellbeing. For example, we have created a set of best practices for kids and family content. *See, e.g.*, YouTube Help, Best Practices for Kids & Family Content, <https://perma.cc/Y6CJ-XVPD>. These best practices help guide creators to create quality content that supports child development and wellbeing by promoting things like respect, healthy habits, learning and curiosity, and play and imagination. Beser, *Enabling a High Quality Kids Experience & Helping Kids Creators Thrive on YouTube*. They also help determine which content is recommended in the main YouTube experience and is included in YouTube Kids. Last year, YouTube updated our Community Guidelines to update our approach to eating disorder-related content, including by limiting the visibility of certain content to adults. Dr. Garth Graham, *An Updated Approach to Eating Disorder-Related Content*, YouTube Official Blog (Apr. 18, 2023), <https://perma.cc/3MHA-7PT6> (“Dr. Graham, *An Updated Approach*”). This change was informed by independent experts such as the National Eating Disorder Association. *Id.* YouTube’s Youth and Family Advisory Committee also advises YouTube on the developmental stages of teens and how content consumed online can affect their wellbeing. James Beser, *Continued Support for Teen Wellbeing and Mental Health on YouTube*, YouTube Official Blog (Nov. 2, 2023), <https://perma.cc/DGK2-4ES3>.

### C. YouTube's Minor-Specific Content and Safety Policies

16. YouTube strives to provide positive online experiences for our users, and that is especially true for children and teens. Over the years, we have developed a robust set of policies and enforcement mechanisms to ensure that YouTube remains a welcoming community for our users. For minors specifically, that might require limiting their access to certain user-generated content—in addition to developing different services that better strive to reflect minors' needs.

17. YouTube has always maintained a set of Community Guidelines that outline categories of content that are prohibited on YouTube. These Guidelines apply to all types of content on our services, including videos, comments, links, thumbnails, video descriptions, stories, posts, live streams, playlists, external links contained in uploaded content (including clickable URLs or verbal directions to users to visit other sites), and any other user-generated content. Our Community Guidelines are a key part of our broader suite of policies, and we regularly update them to keep pace with emerging challenges. *See* YouTube, Community Guidelines, <https://perma.cc/NV9C-9VPG>. Our policies for prohibited or restricted content address a variety of issues, including:

- a. **Age-Restricted Content:** We age-restrict user-generated content that might not violate our Community Guidelines but that would otherwise be inappropriate for minors. These include categories of content such as (1) “nudity and sexually suggestive content”; (2) content that depicts “adults participating in dangerous activities that minors could easily imitate”; (3) “violent or graphic content”; and (4) “vulgar language.” YouTube Help, Age-Restricted Content, <https://perma.cc/5JC9-JZA3>. Age-restricted videos are not viewable to users who are known to be under 18 or any signed-out user. *Id.*
- b. **Child Safety:** YouTube has a standalone “Child Safety Policy” that prohibits “[s]exually explicit content featuring minors and content that sexually exploits minors.”

YouTube Help, Child Safety Policy, <https://perma.cc/FK6U-7X7S> (“YouTube’s Child Safety Policy”). We report content on YouTube that contains child sexual abuse material (“CSAM”) to the National Center for Missing and Exploited Children (“NCMEC”). *Id.* This policy also prohibits content: (1) showing a minor participating in “harmful or dangerous acts”; (2) that could encourage minors to participate in dangerous activities; (3) showing any infliction of physical, emotional, or sexual abuse on a child; or (4) concerning cyberbullying and harassment involving minors. *Id.* The Child Safety Policy also prohibits content that targets minors and families, but contains problematic content including, but not limited to, sexual themes, violence, and other inappropriate themes intended to shock young audiences. *Id.*

#### **D. YouTube’s General Content Moderation Policies**

18. I understand that the Act requires YouTube to “develop and implement a strategy to prevent the known minor’s exposure to harmful material,” defined separately by Texas law, and “content that promotes, glorifies, or facilitates: (1) suicide, self-harm, or eating disorders; (2) substance abuse; (3) stalking, bullying, or harassment; or (4) grooming, trafficking, child pornography, or other sexual exploitation or abuse.” Tex. Bus. & Com. Code § 509.053(a).

19. YouTube already has in place robust policies to moderate objectionable and harmful content, including content identified in the Act. Specifically:

- a. “Harmful Material”: YouTube prohibits “[e]xplicit content meant to be sexually gratifying.” YouTube Help, Nudity & Sexual Content Policy, <https://perma.cc/N3WS-LPRK>. Under this policy, users may not post, for instance, (1) the “depiction of clothed or unclothed genitals, breasts, or buttocks that are meant for sexual gratification”; and (2) “[p]ornography, the depiction of sexual acts, or fetishes that are meant for sexual gratification.” *Id.*

- b. Suicide, Self-harm, or Eating Disorders: YouTube prohibits “content on YouTube that promotes suicide, self-harm, or eating disorders, that is intended to shock or disgust, or that poses a considerable risk to viewers.” YouTube Help, Suicide, Self-Harm, and Eating Disorders Policy, <https://perma.cc/7SBK-57YA>. Any content that is related to suicide, self-harm, or eating disorders and that is targeted to minors or encourages minors to participate in suicide or self-harm is strictly prohibited on YouTube. *Id.*
- c. Substance Abuse: YouTube prohibits (1) “non-educational” “displays of hard drug uses”; (2) “making hard drugs”; (3) “minors using alcohol or drugs”; and (4) “selling or facilitating the sale of hard or soft drugs.” YouTube Help, Illegal or Regulated Goods or Services Policies, <https://perma.cc/2HR3-VJ5B> (capitalization altered). Additionally, YouTube prohibits expression “aim[ed] to directly sell, link to, or facilitate access to” “alcohol,” “controlled narcotics and other drugs,” “nicotine, including vaping products,” and “pharmaceuticals without a prescription.” *Id.* (capitalization altered). YouTube prohibits expression “[s]howing minors involved in dangerous activities,” including “using a controlled substance like alcohol or nicotine.” YouTube’s Child Safety Policy.
- d. Stalking, Bullying, and Harassment: YouTube’s “Harassment & Cyberbullying Policies” prohibits expression “that threatens an identifiable individual or their property” and targets an individual with “prolonged insults or slurs based on [the individual’s] intrinsic attributes.” YouTube Help, Harassment & Cyberbullying Policies, <https://perma.cc/6T9H-7M4G>. YouTube prohibits any content that intends to “shame, deceive or insult a minor,” “reveals personal information” about a minor, “contains

sexualization” of a minor, or “encourages others to bully or harass” a minor. YouTube’s Child Safety Policy (capitalization altered).

- e. Grooming, Trafficking, Child Pornography, or Other Sexual Exploitation or Abuse: YouTube has a “Child Safety Policy” that prohibits “[s]exually explicit content featuring minors and content that sexually exploits minors,” including “[s]exualization of minors.” *Id.* Any content on YouTube that contains CSAM is reported to NCMEC. *Id.* We also report instances of grooming when we are aware there is an exigent threat to a minor.

20. YouTube’s policies go further than the Act’s requirements and protect users by prohibiting hate speech and violent extremist or criminal organization content, among other things. YouTube forbids any content that promotes violence or hatred against individuals based on protected attributes, including: “age, caste, disability, ethnicity, gender identity and expression, nationality, race, immigration status, religion, sex/gender, sexual orientation, victims of a major violent event and their kin, and veteran status[.]” YouTube Help, Hate Speech Policy, <https://perma.cc/AYE8-SBRB> (capitalization altered). Similarly, any content that “praise[s], promote[s], or aid[s] violent extremist or criminal organizations is” prohibited. YouTube Help, Violent Extremist or Criminal Organizations Policy, <https://perma.cc/7EUV-VDQL>. And “[v]iolent or gory content intended to shock or disgust viewers or content that encourages others to commit violent acts are not allowed on YouTube.” *See* YouTube Help, Violent or Graphic Content Policies, <https://perma.cc/62Q4-3D5U>.

21. YouTube has also developed policies around “educational, documentary, scientific, and artistic content” (EDSA), dealing with topics that might otherwise violate our Community Guidelines. Michael Grosack, *A Look at How We Treat Educational, Documentary, Scientific, and*

*Artistic Content on YouTube*, YouTube Official Blog (Sept. 17, 2020), <https://perma.cc/X53C-M8LU>. Decisions about whether a piece of content qualifies for an EDSA exception “are nuanced and context is important.” *Id.* As an example, YouTube “do[es] not allow content targeting minors with insults or bullying, but [YouTube] may allow content that shows this as part of an **educational** anti-bullying campaign provided the minors are actors or their identity [is] hidden.” *Id.* “EDSA exceptions are a critical way we make sure that important speech stays on YouTube, while protecting the wider YouTube ecosystem from harmful content.” *Id.*; see YouTube Help, How YouTube Evaluates Educational, Documentary, Scientific and Artistic (EDSA) Content, <https://perma.cc/S7WN-CD3F>.

22. YouTube has billions of monthly logged-in users, and over 500 hours of content are uploaded every minute by an extraordinarily diverse community of creators, who span more than 100 countries and 80 languages. On a daily basis, users watch more than a billion hours of video on YouTube and generate billions of views. See YouTube for Press, YouTube Official Blog, <https://perma.cc/VE2L-PW6Z>.

23. YouTube has developed robust means of addressing user-generated content that violates our policies. YouTube relies on a mix of human and automated effort to detect violative content and remove it. YouTube’s first line of defense is automated systems that are trained with machine learning and/or incorporate hash-matching technology to quickly identify and take action against violative content. YouTube also relies on global teams to review flagged content and remove it, restrict who can view it, or leave the content on the site if it does not violate any policies. YouTube employs expert teams around the world to investigate sophisticated bad actors who are adept at circumventing automated defenses. YouTube, Information Quality & Content Moderation at 12, <https://perma.cc/3UW2-J4C7>. Content moderation can be human-resource intensive, even

for CSAM and known violative content. YouTube verifies with human review all CSAM hashes to confirm that they are CSAM before they are reported to NCMEC. And human review becomes even more time-intensive and important for other illegal conduct like “grooming,” trafficking, or harassment that requires more context to understand if illegal activity is taking place.

24. While no system of reviewing such a high volume of content is foolproof, YouTube expends great effort and resources in removing objectionable and harmful user-generated content before most users see it. In fact, the vast majority of violative videos are seen by fewer than 10 people before they are removed. In the first quarter of 2024, approximately 57.53% of removed videos had no views before they were removed, and approximately 25.18% had between 1-10 views. *See* Google, YouTube Community Guidelines Enforcement Report – Removals, <https://perma.cc/Q257-XV7T> (“Google, Community Guidelines Enforcement Report – Removals”). Over that same period, YouTube removed 8,295,304 videos. *Id.* YouTube has also created a metric for determining what percentage of views on YouTube comes from content that violates our policies. This metric, known as the Violative View Rate (VVR), helps us measure our responsibility work. *See* Google, YouTube Community Guidelines Enforcement Report – Views, <https://perma.cc/W5HH-PPCY>. YouTube’s VVR is calculated by taking a sample of viewed videos to a team for review. The team then determines which videos in the sample violate our policies, and YouTube uses these decisions to estimate a VVR. *Id.* In the first quarter of 2024, YouTube estimated a 0.12-0.13% VVR. *Id.* This means that out of every 10,000 views on YouTube, 12-13 come from violative content. This is down by more than 81% when compared to the final quarter of 2017 (when our teams started tracking this metric), in large part thanks to our investments in robust content moderation, including automated technologies and human review.

25. Our practices also involve taking action against users and channels that continue to upload violative content. A YouTube channel is terminated if it accrues three Community Guidelines strikes in 90 days, has a single case of severe abuse (such as predatory behavior), or is determined to be wholly dedicated to violating our guidelines (as is often the case with spam accounts). When a channel is terminated, all of its videos are removed. In the first quarter of 2024, YouTube terminated 15,799,880 total channels. *See* Google, Community Guidelines Enforcement Report – Removals. The majority of channel terminations are a result of accounts being dedicated to spam (96%) or adult sexual content (1%) in violation of our guidelines. *Id.*

26. Our efforts are not limited to removing objectionable and harmful user-generated content in the videos uploaded on our site. They also extend to user-generated comments. In the first quarter of 2024, YouTube removed approximately 1,443,821,162 comments. *Id.* Approximately 99.7% percent of comments removed are first detected by YouTube’s automatic flagging system, and the majority of actions taken against comments (83.9%) are for spam. *Id.*

27. Enforcing our policies takes a tremendous amount of effort. YouTube’s global teams work continually to remove user-generated content from YouTube that violates our policies. We endeavor to give our teams visibility into emerging issues before they reach, or become widespread on, our website. That valuable visibility is driven by our Intelligence Desk, a team within YouTube’s Trust & Safety organization. These specialized analysts identify trends and the risks they pose. They also regularly monitor ongoing threats, both tracking their prevalence across media and evaluating how they morph over time. Our teams also work to provide appropriate age restrictions and other appropriate warnings on content.

28. Given the open nature and scale of YouTube and because bad actors are constantly refining how they attempt to evade detection, detecting violative user-generated content cannot be



done with perfect accuracy nor upon the moment of upload. YouTube, Information Quality & Content Moderation at 12. Nevertheless, YouTube continuously works to combat new challenges and bad actors through a multi-faceted and nuanced approach to exercising its discretion in setting its content moderation policies, working to distinguish among content that is truly objectionable or harmful, borderline content (which is content that comes close to violating, but does not quite violate our Community Guidelines), and content that contributes positively to the YouTube community. To that end, we have a diverse set of means to help us keep our community safe, including: (1) appending warning interstitials for sensitive content; (2) surfacing helpful resources for crisis management; and (3) suspending and/or terminating channels or accounts. YouTube also encourages creators to age-restrict content when appropriate. The YouTube Team, *Using Technology to More Consistently Apply Age Restrictions*, YouTube Official Blog (Sept. 22, 2020), <https://perma.cc/KRF6-NBGC>. We apply age restrictions on content ourselves, when needed. We have other tools to help us provide authoritative information on our service—such as the use of information panels to show basic background information, sourced from independent, third-party partners, that give more context on a topic. For example, in partnership with third-party expert organizations and global experts in the space, we expanded our crisis resource panels to introduce a “pause” page for queries related to suicide, self-harm, and eating-disorder topics that provides users with options with how to proceed before they are able to see search results related to these topics. These include selecting to contact third-party crisis-resource-hotline partners for support or searching for various self-help tools. We further prioritize the recommendations of authoritative content and limit the spread of borderline content to users through recommendations. Because removing content is only one way of managing content, YouTube has developed and invested in this diverse set of tools that are essential in balancing free expression and responsibility on our

website. Simply put, these tools provide additional information and options compared to simply removing (or not removing) user-generated content from our website.

## **II. YouTube’s Compliance Obligations Under the Act**

29. Google and YouTube support regulations designed to protect the wellbeing of minors online while ensuring that users’ rights to access speech are not impeded. *See generally* Google, Legislative Framework to Protect Children and Teens Online, <https://perma.cc/TBA3-73NA> (“Google, Legislative Framework”). But the Act here is not well-suited to protect minors. It fails to take into account differences among online services, fails to take into account the differences among minors at different developmental levels, and ultimately will only impede minors’ access to valuable speech online. Furthermore, it will discourage and stymie innovation that makes YouTube’s services better for all.

### **A. Monitoring and Censorship Requirement (Tex. Bus. & Com. Code § 509.053)**

30. The Act would also require YouTube to redefine how its services work, and it fundamentally burdens and undermines YouTube’s ability to operate responsibly, enforce important policies, and innovate and improve its services. Moreover, compliance with the Act would be enormously difficult and burdensome, requiring large upfront investments and continuing maintenance in terms of time, engineering, and operations resources. Independently and in combination, therefore, compliance with the Act’s provisions would make it much more difficult for YouTube to offer its services to those in Texas. Below are examples of the burdens the Act would impose.

31. I understand that the Act provides: “In relation to a known minor’s use of a digital service, a digital service provider shall develop and implement a strategy to prevent the known minor’s exposure to harmful material,” defined separately by Texas law, and “content that promotes, glorifies, or facilitates: (1) suicide, self-harm, or eating disorders; (2) substance abuse;

(3) stalking, bullying, or harassment; or (4) grooming, trafficking, child pornography, or other sexual exploitation or abuse.” Tex. Bus. & Com. Code § 509.053(a).

32. As I understand this provision, it will require alterations to YouTube’s content moderation—not because YouTube’s content moderation is currently deficient, but because of uncertainty of how the Texas Attorney General will interpret the Act. As explained above, YouTube has—and vigorously enforces—policies preventing minors’ exposure to objectionable and harmful content, including the broad topics identified by the Act. In fact, we have publicly stated that “content policies should address and respond to risks related to content that promotes eating disorders, self-harm, or bullying.” Google, Legislative Framework at 3. Yet this provision creates multiple compliance questions. For example, while YouTube has developed its own internal understandings of what, for example, “promotes, glorifies, or facilitates . . . suicide, self-harm, or eating disorders”—and can effectively identify and address such content—YouTube has no assurance that the Texas Attorney General will agree. Furthermore, it is unclear whether YouTube must simply have a “strategy” in place, or whether it must engage in what the Attorney General considers perfect content moderation preventing all exposure of all prohibited content (as defined by the Attorney General) to all minors.

33. To ensure compliance with this provision, YouTube may need to alter its policies and enforcement to take an overinclusive approach to content moderation, resulting in unrecoverable compliance costs. For example, YouTube runs the risk of liability every time a minor sees any content that even references self-harm. Thus, the only way to avoid the Act’s large penalties is to decline to disseminate all content mentioning self-harm to known minors, despite the fact that certain types of related content can aid in recovery. *See* Dr. Graham, *An Updated Approach*. Similar concerns about uncertainty apply to all of the categories of prohibited speech under the Act,

with the end result that compliance with the Act may demand that YouTube cease disseminating large categories of speech.

34. That will have clear harms for users as well, who may be denied access to protected and valuable speech as a result of YouTube’s overinclusive moderation.

35. In addition, the required *means* of moderation will also require YouTube to alter its content moderation enforcement techniques. Under the Act, our strategy “must include”: “(A) creating and maintaining a comprehensive list of harmful material or other content described by Subsection (a) to block from display to a known minor; (B) using filtering technology and other protocols to enforce the blocking of material or content on the list under Paragraph (A); (C) using hash-sharing technology and other protocols to identify recurring harmful material or other content described by Subsection (a); (D) creating and maintaining a database of keywords used for filter evasion, such as identifiable misspellings, hash-tags, or identifiable homoglyphs; (E) performing standard human-performed monitoring reviews to ensure efficacy of filtering technology; (F) making available to users a comprehensive description of the categories of harmful material or other content described by Subsection (a) that will be filtered; and (G) except as provided by Section 509.058, making available the digital service provider’s algorithm code to independent security researchers[.]” Tex. Bus. & Com. Code § 509.053(b)(1) (emphasis added). The permissive content moderation tools “may include”: “(A) engaging a third party to rigorously review the digital service provider’s content filtering technology; (B) participating in industry-specific partnerships to share best practices in preventing access to harmful material or other content described by Subsection (a); or (C) conducting periodic independent audits to ensure: (i) continued compliance with the digital service provider’s strategy; and (ii) efficacy of filtering technology and protocols used by the digital service provider.” *Id.* § 509.053(b)(2).

36. The Act’s checklist-style enumeration of mandatory and permissive content moderation tools threatens to limit YouTube’s broad portfolio of content moderation tools, and may render YouTube less effective at content moderation. Among other things, we have no way to know (1) if those are the exclusive means of moderation tools the statute permits, to the exclusion of other tools like machine learning; or (2) whether YouTube must use each enumerated tool in Section 509.053(b)(1) for each category of enumerated content. For instance, if the Act requires us to use all tools for all categories of prohibited content, then we will have to divert resources into using ineffectual means of moderation for certain categories of content—which may render our content moderation efforts over-inclusive. For instance, except for CSAM and terrorism, we cannot “hash” most violative content, such as substance-abuse promotion. Likewise, *behavioral* offenses such as grooming and harassment cannot be hashed. And we try to limit the use of filtering to content where it is effective, and not applying it in situations where it may be overinclusive. Neither hashing nor filtering are effective at detecting nuance, such as content that *discusses* eating disorders in an educational context and content that “glorifies” eating disorders.

37. Even if the statute’s checklist is not exclusive, it still threatens to hinder YouTube’s ability to successfully moderate content. That is because the checklist replaces YouTube’s highly customized content moderation practices with a one-size-fits-all mandate. But YouTube’s community and its operations are unique. To be successful, YouTube’s content moderation practices must also be unique in certain respects. But HB18 makes developing and enforcing unique content moderation practices much more difficult. These difficulties will only grow larger with time. New methods may render, and in some cases have already rendered, HB18’s checklist outmoded and incomplete. Furthermore, the statute does not take into account the role of contextual determinations in content moderation: “Once models are trained to identify potentially violative content, the

role of content moderators remains essential throughout the enforcement process. . . . For example, a speech by Hitler at the Nuremberg rallies with no additional context may violate our hate speech policy. But if the same speech was included in a documentary that decried the actions of the Nazis, it would likely be allowed under our EDSA guidelines. . . . This distinction may be more difficult for a model to recognize, while a content moderator can more easily spot the added context.” Matt Halprin & Jennifer Flannery O’Connor, *On Policy Development at YouTube*, YouTube Official Blog (Dec. 1, 2022), <https://perma.cc/DE2X-82RD>.

38. Of particular note, requiring YouTube to “mak[e] available the digital service provider’s algorithm code to independent security researchers” would require YouTube to disclose highly commercially sensitive and trade secret information. *Id.* § 509.053(b)(1)(G). That would both harm YouTube and could create the potential that malicious actors could access information that would allow them to evade YouTube’s moderation efforts—harming users.

39. Although difficult to estimate with certainty, YouTube would need to dedicate budget, engineering resources, and other staff investment to address this provision. And to come into compliance with this provision, YouTube would need to start developing its means of compliance now. Further, YouTube would likely need to alter its content moderation service wide to comply with this requirement.

**B. Age-Registration and Parental-Dispute Requirements  
(Tex. Bus. & Com. Code §§ 509.051, 509.101)**

40. I understand the Act also mandates that YouTube adopt procedures to register the ages of users creating accounts, allow adults to dispute the ages of people who register as adults (entailing investigation into the real age of those users), and verify the identity of parents relative to their minor children on the websites. *Id.* §§ 509.051, 509.101. These requirements facilitate both the content moderation requirements discussed above and a set of requirements to provide

parental-access and parental-control tools for parents of minor children. *Id.* §§ 509.052-.054, 509.101-03.


41. YouTube currently does not have systems to verify the custodial arrangements of its teenage users, nor the means to even identify the individuals to contact as necessary to comply with the Act's requirements for obtaining consent prior to allowing a minor user to share their data. Developing those systems will require a large upfront investment of resources and would overly restrict the ability of teenagers to participate in online speech. Furthermore, operating and maintaining those systems will require substantial resources, including a large ongoing investment in the human resources necessary to evaluate the identities and familial relationships of our users.

42. In sum, if the Act takes effect on September 1, 2024, Google and YouTube will suffer irreparable harm in the form of unrecoverable compliance costs, diminished ability to provide valuable user-generated content, and ongoing regulatory uncertainty about whether its services comply with the Act's broad prohibitions. Moreover, YouTube's users may experience a degraded service that frustrates their ability to engage with the wealth of speech available on YouTube.

\* \* \*

I declare under penalty of perjury, pursuant to 28 U.S.C. § 1746, that the foregoing is true and correct to the best of my knowledge.

Executed on July 24, 2024, in WASHINGTON, D.C.

DocuSigned by:  
  
346A76BAA93G431...  
ALEXANDRA N. VEITCH