

December 11, 2023

The Honorable Charles E. Schumer
Majority Leader
United States Senate
322 Hart Senate Office Building
Washington, D.C. 20510

The Honorable Mitch McConnell
Minority Leader
United States Senate
317 Russell Senate Office Building
Washington, D.C. 2051

Re: S. 1993, “No Section 230 Immunity for AI Act”

Dear Majority Leader Schumer and Minority Leader McConnell:

We, the undersigned organizations and individuals, write to express serious concerns about the “No Section 230 Immunity for AI Act” (S. 1993). S. 1993 would threaten freedom of expression, content moderation, and innovation. Far from targeting any clear problem, the bill takes a sweeping, overly broad approach, preempting an important public policy debate without sufficient consideration of the complexities at hand.

Section 230 makes it possible for online services to host user-generated content, by ensuring that only users are liable for what they post—not the apps and websites that host the speech. S. 1993 would undo this critical protection, exposing online services to lawsuits for content whenever the service offers or uses any AI tool that is technically *capable* of generating any kind of new material. The now widespread deployment of AI for content composition, recommendation, and moderation would effectively render any website or app liable for virtually all content posted to them.

S. 1993 would preempt an important and necessary policy debate. As a threshold matter, even proponents of Section 230 disagree on whether and to what extent the law immunizes GenAI providers from treatment as the publisher of their tools’ outputs. While some argue that Section 230’s protections logically extend to the output of GenAI tools,¹ others—including Section 230’s authors—take the position that GenAI tools create new

¹ Jess Miers, *Yes, Section 230 Should Protect ChatGPT And Other Generative AI Tools*, TECHDIRT (Mar. 17, 2023, 11:59 AM), <https://www.techdirt.com/2023/03/17/yes-section-230-should-protect-chatgpt-and-others-generative-ai-tools/> (“ChatGPT (and similarly situated generative AI products) are functionally akin to ‘ordinary search engines’ and predictive technology like autocomplete.”).

content that the tools' purveyors are responsible for "developing," at least "in part."² Still others would argue that the question of whether Section 230 extends to GenAI output depends on the context in which it was used. The courts have not yet ruled on these questions. S.1993 would cut off this critical debate with overbroad language that could cause more problems than it fixes.

Carving out state law will lead to censorship. Section 230 was written to establish a consistent nationwide body of law for liability for content on the Internet. S. 1993 would effectively undo this benefit by carving out of Section 230 any civil claim or criminal charge brought under state law for conduct involving "the use or provision" of GenAI. This would enable politically motivated actors to censor online content they dislike.

Recent history illustrates the stakes: In March 2023, a bill was introduced in the Texas House of Representatives to criminalize providing "information on how to obtain an abortion-inducing drug," and create civil liability for any interactive computer service that "allows residents of [Texas] to access information or material that assists or facilitates efforts to obtain elective abortions or abortion-inducing drugs."³

Had this bill been enacted into law, Section 230 would have precluded its enforcement. But under S. 1993, it *would* be enforceable if the offending content was posted by anyone on any service that provides GenAI tools to its users—or even deploys GenAI for content moderation, as discussed below. S. 1993 would undoubtedly lead a wave of similar legislation targeting disfavored expression, from LGBTQ content to hate speech.⁴ At best, the result would be chaos and endless litigation. At worst, government officials will have been handed a ready-made tool to successfully fracture and censor the Internet.

² Cristiano Lima, *AI chatbots won't enjoy tech's legal shield, Section 230 authors say*, WASH. POST (Mar. 17, 2023, 9:03 AM), <https://www.washingtonpost.com/politics/2023/03/17/ai-chatbots-wont-enjoy-techs-legal-shield-section-230-authors-say/>. See also Matt Perault, *Section 230 Won't Protect ChatGPT*, LAWFARE (Feb. 22, 2023, 1:11 PM), <https://www.lawfaremedia.org/article/section-230-wont-protect-chatgpt>.

³ See Jennifer Pinsof, *This Texas Bill Would Systematically Silence Anyone Who Dares to Talk About Abortion Pills*, EFF (Mar. 13, 2023), <https://www.eff.org/deeplinks/2023/03/texas-bill-would-systematically-silence-anyone-who-dares-talk-about-abortion-pills>.

⁴ States are already targeting politically disfavored content. See, e.g., *New York can't target protected online speech by calling it 'hateful conduct'*, FIRE (Dec. 1, 2022), <https://www.thefire.org/news/lawsuit-new-york-cant-target-protected-online-speech-calling-it-hateful-conduct/> ("The law forces internet platforms of all stripes to publish a policy explaining how they will respond to online expression that could 'vilify, humiliate, or incite violence' based on a protected class, like religion, gender, or race."). And some states plan to use proposed federal laws to shut down LGBTQ speech. See, e.g., Mike Masnick, *Heritage Foundation Says That Of Course GOP Will Use KOSA To Censor LGBTQ Content*, TECHDIRT (May 24, 2023, 12:26 PM), <https://www.techdirt.com/2023/05/24/heritage-foundation-says-that-of-course-gop-will-use-kosa-to-censor-lgbtq-content/>; Jared Eckert & Mary McCloskey, *How Big Tech Turns Kids Trans*, THE HERITAGE FOUNDATION (Sept. 15, 2022), <https://www.heritage.org/gender/commentary/how-big-tech-turns-kids-trans> ("[W]e must guard against the harms of sexual and transgender content.").

S. 1993 will benefit vexatious litigants. The bill’s definition of GenAI (“an artificial intelligence system that is capable of generating novel [content] based on prompts or other forms of data provided by a person”) is broad enough to encompass tools as basic and commonplace as predictive text (autocomplete), autocorrect, and potentially even search autocomplete suggestions or grammar and spellchecking features, as well as any other AI-generated content.

A core function of Section 230 is to provide for the early dismissal of claims and avoid the “death by ten thousand duck-bites” of costly, endless litigation.⁵ This bill provides an easy end-run around that function: simply by plausibly alleging that GenAI was somehow involved with the content at issue, plaintiffs could force services into protracted litigation in hopes of extracting a settlement for even meritless claims.

The bill misallocates liability and rewards malicious actors. S. 1993 would, inexplicably, reverse Section 230’s sensible allocation of legal liability to the party ultimately responsible for the wrongfulness of content. Again, under S. 1993, the provision or use of any AI tool technically capable of generating some form of content, from predictive text to content moderation tools, would effectively expose platforms and online services to liability for any content it hosts or enables the creation of.

Consider a musician who utilizes a platform offering a GenAI production tool to compose a song including synthesized vocals with lyrics expressing legally harmful lies (libel) about a person. Even if the lyrics were provided wholly by the musician, the conduct underlying the ensuing libel lawsuit would undoubtedly “involve the use or provision” of GenAI—exposing the tool’s provider to litigation. In fact, the tool’s provider could lose immunity even if it did *not* synthesize the vocals, simply because the tool is *capable* of doing so.⁶

Like any tool, GenAI can be misused by malicious actors, and there is no sure way to prevent such uses—every safeguard is ultimately circumventable. Stripping immunity from services that offer those tools irrespective of their relation to the content does not just ignore this reality, it incentivizes it. The ill-intentioned, knowing that the typically deep pockets of GenAI providers are a more attractive target to the plaintiffs’ bar, will only be further encouraged to find ways to misuse GenAI.

Still more perversely, malicious actors may find themselves immunized by the same protection that S. 1993 strips from GenAI providers. Section 230(c)(1) protects both providers of interactive computer services *and* users from being treated as the publisher of third-party content. But S. 1993 only excludes the former from Section 230 protection. If Section 230 does indeed protect GenAI output to at least some degree as the proponents of

⁵ Fair Hous. Council of San Fernando Valley v. Roommates.com, LLC, 521 F.3d 1157, 1174 (9th Cir. 2008).

⁶ S. 1993, 118 Cong. (2023), <https://www.congress.gov/bill/118th-congress/senate-bill/1993/text>.

this bill fear, the malicious user who manipulates ChatGPT into providing a defamatory response⁷ would be immunized for re-posting that content, while OpenAI would face liability.

S. 1993 will make content moderation harder, worse—and more biased. AI-powered content moderation tools are now ubiquitous; they have the potential to increase consistency, decrease bias, and provide a measure of relief to beleaguered human moderators who sift through the worst content imaginable at great cost to their well-being.⁸ But because most lawsuits regarding content moderation decisions are dismissed under Section 230(c)(1), this bill threatens the viability of the development and use of such tools.

OpenAI, as just one example, has deployed GPT-4 to assist in revising its content policies by prompting it with a policy and feeding it sample content, examining the labels and reasoning assigned by GPT-4 to that content, and then clarifying the policy until the AI achieves satisfactory results.⁹ This process itself may preclude OpenAI from invoking Section 230(c)(1) in a lawsuit over its content moderation decisions: the creation of the policy applied to moderated content would appear to be part and parcel of the “conduct underlying the claim.”

AI tools are also increasingly used by a variety of platforms and services to perform day-to-day content moderation functions, which would similarly strip moderation decisions of their Section 230(c)(1) immunity. An AI-generated content flag, accompanied by a generated explanation of the relevant policy’s application, is surely a “use” of GenAI. But even if the moderation tool did not provide a generated explanation (“novel text”), it would still lose immunity; the GenAI need not actually *generate anything*—the mere fact that it is *capable* of doing so (which GPT-4 plainly is) would bring it under S. 1993’s exclusion.

⁷ See Adam Thierer & Shoshana Weissmann, *Without Section 230 Protections, Generative AI Innovation Will Be Decimated*, R STREET INSTITUTE (Dec. 6, 2023), <https://www.rstreet.org/commentary/without-section-230-protections-generative-ai-innovation-will-be-decimated/> (“The person typing in the request was the one intending to create libel, but the AI company would be liable too.”).

⁸ See, e.g., Andrew Arsht & Daniel Etcovitch, *The Human Cost of Online Content Moderation*, JOLT DIGEST (Mar. 2, 2018), <https://jolt.law.harvard.edu/digest/the-human-cost-of-online-content-moderation/> (“Some journalists, scholars, and analysts have noted PTSD-like symptoms and other mental health issues arising among moderators.”); Jaspreet Singh, *OpenAI says AI tools can be effective in content moderation*, REUTERS (Aug. 15, 2023), <https://www.reuters.com/technology/openai-says-ai-tools-can-be-effective-content-moderation-2023-08-15/>; Aditya Jain, *Impact of Generative AI on Content Moderation*, AVASANT (July 2023), <https://avasant.com/report/impact-of-generative-ai-on-content-moderation/#:~:text=AI%20Content%20vs.-,Traditional%20Content,cultural%20references%2C%20and%20subtle%20nuances.>

⁹ See Lilian Weng et al., *Using GPT-4 for content moderation*, OPENAI (Aug. 15, 2023), <https://openai.com/blog/using-gpt-4-for-content-moderation>. See also Kyle Wiggers, *OpenAI proposes a new way to use GPT-4 for content moderation*, TECHCRUNCH (Aug. 15, 2023, 2:15 PM), <https://techcrunch.com/2023/08/15/openai-proposes-a-new-way-to-use-gpt-4-for-content-moderation/>.

This unfortunate result may indeed be the intended one. S. 1993 leaves untouched Section 230(c)(2)(A), which Section 230 critics argue should be the sole protection for content moderation decisions. Under Section 230(c)(2)(A), defendants must show that they “voluntarily” removed objectionable content “in good faith.” This standard is highly fact-dependent; as such, defendants would no longer be able to resolve lawsuits on motion to dismiss. This, in turn, would allow plaintiffs to exact heavy discovery costs on any platform attempting to defend its moderation decisions. Indeed, it is unclear how the “good faith” of AI could ever be established. What *is* clear is that the development and use of valuable AI-based content moderation tools will be disincentivized by the high costs imposed by S. 1993.

S. 1993’s breadth disincentivizes all GenAI tools. As noted above, the bill’s definition of GenAI (“an artificial intelligence system that is capable of generating novel [content] based on prompts or other forms of data provided by a person”) would encompass commonplace tools like predictive text (autocomplete), autocorrect, and potentially even grammar and spellchecking features.

This extensive definition is particularly troubling because S. 1993 is not limited to instances where GenAI contributed to the tortious or illegal nature of content. Rather, S. 1993 excludes from Section 230(c)(1)’s protection any claim based on conduct that “involves the use or provision of [GenAI].” Thus, a social media platform could find itself facing liability for all its users’ posts simply because it provided predictive text or grammar suggestions (both forms of GenAI) to aid users in expressing their own ideas—or even because it utilizes GenAI for content recommendation and moderation.

Moreover, while S. 1993 would only exclude GenAI use or provision “by the interactive computer service,” in practice, a social media platform has no reliable way to discern whether a piece of content posted to it was created using one of their own GenAI tools; users might have saved or copied GenAI output for later use. In this way, too, platforms would have to choose between not offering any GenAI tools or risking liability for every piece of content posted on their service. The latter result would effectively be tantamount to a full repeal of Section 230.

GenAI has become increasingly important in the creation of online content, and it promises to make our communications more effective, inexpensive, and accessible. Congress should not inhibit these exciting advancements by forcing online services to choose between foregoing use of GenAI technology or exposing themselves to crushing liability.

—

Generative Artificial Intelligence is a complex issue that deserves careful thought and nuanced, precise legislation—not a rigid, heavy-handed overreaction that threatens to

undermine free speech, user safety, and American competitiveness in the AI marketplace. We urge Congress to consider a more thoughtful approach.

If you have any questions or would like to discuss the issues in this letter further, please contact Ari Cohn at acohn@techfreedom.org.

Sincerely,

Organizations

American Civil Liberties Union

Americans for Prosperity

Association of Research Libraries

Center for Democracy and Technology

Chamber of Progress

Competitive Enterprise Institute

Computer & Communications Industry Association

Copia Institute

Electronic Frontier Foundation

Foundation for Individual Rights and Expression

R Street Institute

Software & Information Industry Association

Taxpayers Protection Alliance

TechFreedom

Individuals

Joshua Levine, American Action Forum*

*Affiliation listed for identification purposes only