



## SECTION 230 WON'T PROTECT CHATGPT

*Matt Perault\**

Introduction .....	363
I. Liability Protections in Current Law .....	364
II. Will ChatGPT Be Considered an Information Content Provider? .....	365
III. Should the Law Shield LLMs From Legal Liability? .....	367
IV. How Should LLMs Be Regulated? .....	369
Conclusion: What Will the Future Hold? .....	373

### INTRODUCTION

The emergence of products fueled by generative artificial intelligence (AI) such as ChatGPT will usher in a new era in the platform liability wars. Previous waves of new communication technologies—from websites and chat rooms to social media apps and video sharing services—have been shielded from legal liability for user-generated content posted on their platforms, enabling these digital services to rise to prominence. But with products like ChatGPT, critics of that legal framework are likely to get what they have long wished for: a regulatory model that makes tech platforms responsible for online content.

The question is whether the benefits of this new reality outweigh its costs. Will this regulatory framework minimize the volume and distribution of harmful and illegal content? Or will it stunt the growth of ChatGPT and other large language models (LLMs), litigating them out of mainstream use before their capacity to have a transformational impact on society can be understood? Will it tilt the playing field

---

\* Director of the Center on Technology Policy, UNC-Chapel Hill (matthew.perault@unc.edu). A version of this article previously appeared in *Lawfare*. Thanks to Eugene Volokh, Quinta Jurecic, Alan Rozenshtein, and Jess Miers for their helpful comments.

toward larger companies that can afford to hire massive teams of lawyers and bear steep legal fees, making it difficult for smaller companies to compete?

In this article, I explain why current speech liability protections do not apply to certain generative AI use cases, explore the implications of this legal exposure for the future deployment of generative AI products, and provide an overview of options for regulators moving forward.

### I. LIABILITY PROTECTIONS IN CURRENT LAW

The source of this conundrum is the text of Section 230, the 1996 law that scholars such as Jeff Kosseff cite as the basis for the rise of the internet.<sup>1</sup> The law enables platforms to host speech from users without facing legal liability for it. If a user posts something illegal on a website, you can hold the user liable, but not the website.

In recent years, Republicans and Democrats alike have criticized Section 230, dragging tech leaders to congressional hearings to lambast them about their content moderation practices and proposing dozens of bills with the goal of expanding accountability for tech platforms.<sup>2</sup> Policymakers outside the United States have pursued similar reforms—in Europe, for example, the Digital Services Act heralds a new era in expanded accountability for internet platforms that host user content.<sup>3</sup>

Under current law, an interactive computer service (a content host) is not liable for content posted by an information content provider (a content creator): Section 230 stipulates that “[n]o provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.”<sup>4</sup>

When a platform uses Section 230 to defend against liability in court, the judge uses a three-part test to evaluate whether the case should be dismissed:

- Is the defendant a provider or user of an interactive computer service?
- Is the plaintiff attempting to hold the defendant liable as a publisher or speaker?

---

<sup>1</sup> JEFF KOSSEFF, *THE TWENTY-SIX WORDS THAT CREATED THE INTERNET* (2019).

<sup>2</sup> Quinta Jurecic, *The Politics of Section 230 Reform: Learning from FOSTA’s Mistakes*, BROOKINGS, Mar. 1, 2022, <https://perma.cc/8FRR-NJRV>.

<sup>3</sup> European Commission, *The Digital Services Act Package*, <https://perma.cc/A6A3-G8LA>.

<sup>4</sup> 47 U.S.C. § 230(c)(1).

- Is the plaintiff's claim based on content posted by another information content provider?<sup>5</sup>

If the answer to any one of these questions is no, then the defendant cannot use Section 230 as a defense to liability.

Section 230 defines an “interactive computer service” as “any information service, system, or access software provider that provides or enables computer access by multiple users to a computer server, including specifically a service or system that provides access to the Internet and such systems operated or services offered by libraries or educational institutions.”<sup>6</sup>

The statute defines an “information content provider” as “any person or entity that is responsible, in whole or in part, for the creation or development of information provided through the Internet or any other interactive computer service.”<sup>7</sup> The law is not entirely clear on the meaning of “development” of content “in part,” but courts have held that a platform may not use Section 230 as a defense if it provides “pre-populated answers” so that it is “much more than a passive transmitter of information provided by others.”<sup>8</sup>

## II. WILL CHATGPT BE CONSIDERED AN INFORMATION CONTENT PROVIDER?

Given this definition, at least some judges will likely find that ChatGPT and other LLMs are information content providers. The result is that the companies that deploy these generative AI tools—like OpenAI, Microsoft, and Google—will be precluded in some circumstances from using Section 230 in cases arising from AI-generated content.

The relevant question will be whether LLMs “develop” content, at least “in part.” It is difficult to imagine that a court would find otherwise if an LLM drafts text on a topic in response to a user request or develops text to summarize the results of a search inquiry (as ChatGPT can do). In contrast, Twitter does not draft tweets for its users, and most Google Search results simply identify existing websites in response to user queries.

---

<sup>5</sup> Valerie C. Brannon & Eric N. Holmes, Congressional Research Service, Section 230: An Overview (R46751), 12 (Apr. 7, 2021), <https://perma.cc/4XHM-PKT8>.

<sup>6</sup> 47 U.S.C. § 230(f)(2).

<sup>7</sup> 47 U.S.C. § 230(f)(3).

<sup>8</sup> Fair Housing Council of San Fernando Valley v. Roommates.com, LLC, 521 F.3d 1157, 1166 (9th Cir. 2008) (en banc).

Of course, determining whether a defendant is entitled to Section 230 protection is a fact-intensive analysis that depends on how a particular technology is used in a particular case.<sup>9</sup> ChatGPT could be found to be an information content provider in some product integrations and not in others. If, for instance, the technology is used to simply determine which search results to prioritize or which text to highlight from underlying search results, then a court might find that it is entitled to Section 230 protection. That said, many of the LLM use cases that have been discussed since ChatGPT's release seem likely to sit outside the current boundaries of Section 230 protection.

In addition, the scope of product possibilities for generative AI tools will narrow significantly even if only a minority of judges find that they are not entitled to Section 230 protections. Once some judges have reached that determination, companies—and particularly smaller companies who are more sensitive to the potential for increased legal fees—will be wary of releasing products that could land them in lengthy, costly litigation. Lawyers will be empowered to push back on new product features proposed by engineers, and will play a more assertive oversight role in the product review process. The result will be less innovative products, more tailored to avoiding legal risk and less oriented toward maximizing the technology's potential.

Increasing legal liability for platforms is the outcome many policymakers are hoping for. Members of Congress have proposed dozens of reforms to Section 230 in recent years, and the underlying theme is consistent: Lawmakers want to increase platform liability for content moderation decisions.

It may also align with an evolution in courts' interpretation of existing law. In February, the Supreme Court heard arguments in two cases that address whether platforms should be held accountable for the content they recommend to users.<sup>10</sup> Just as with content generated by LLMs, the question lies in how to apportion liability when platforms use algorithmic decision-making to determine what information people see. If the Court finds in favor of the plaintiffs, it will establish a

---

<sup>9</sup> Other scholars have offered far more extensive analyses of the various situations where an LLM might be found to be protected by Section 230. See Derek Bambauer & Mihai Surdeanu, *Authorbots*, 3 J. FREE SPEECH L. 375 (2023); Jess Miers, *Yes, Section 230 Should Protect ChatGPT and Other Generative AI Tools*, TECHDIRT, Mar. 17, 2023, 11:59 am, <https://perma.cc/ZH73-N3XA>.

<sup>10</sup> Transcript of Oral Argument, *Gonzalez v. Google*, No. 21-1333 (U.S. argued Feb. 21, 2023); Transcript of Oral Argument, *Twitter v. Taamneh*, No. 21-1496 (U.S. argued Feb. 22, 2023).

beachhead for liability when a platform uses an algorithm to surface content to users.

In the course of the argument, the justices seemed to make the connection between the case before them and the emergence of new technologies like LLMs. Justice Neil Gorsuch stated that LLMs would likely not receive Section 230 protections. “Artificial intelligence generates poetry. It generates polemics today that would be content that goes beyond picking, choosing, analyzing, or digesting content. And that is not protected. Let’s assume that’s right. Then the question becomes, what do we do about recommendations?”<sup>11</sup>

### III. SHOULD THE LAW SHIELD LLMs FROM LEGAL LIABILITY?

Courts will likely find that ChatGPT and other LLMs are excluded from Section 230 protections because they are information content providers, rather than interactive computer services. But is that a good thing?

In some cases, liability may seem desirable. An LLM could make it easier to break the law, such as if it provides instructions on how to build a bomb, or manufacture and distribute illegal drugs.<sup>12</sup> It could make it easier for a student to cheat.<sup>13</sup> It could use racist or misogynistic language, or help fuel disinformation campaigns.<sup>14</sup> It could cause harm by providing inaccurate medical information, which might result in people making adverse health decisions or doctors recommending problematic courses of treatment.<sup>15</sup> Making LLMs liable for problematic content

---

<sup>11</sup> Transcript of Oral Argument at 49, *Gonzalez v. Google*, No. 21-1333 (U.S. argued Feb. 21, 2023).

<sup>12</sup> See, e.g., Max Daly, *ChatGPT Gave Me Advice on How To Join a Cartel and Smuggle Cocaine Into Europe*, VICE, Feb. 3, 2023, <https://perma.cc/NM23-46T2>.

<sup>13</sup> See, e.g., Kala Rama, *ChatGPT Banned for Students in New York City Public Schools*, PIX11, Feb. 13, 2023; Valerie Flugge, *Stopping Artificial Intelligence From Impairing the Development of Student Intelligence: Do AI Companies Have Liability if Students Use Them to Cheat?* (working paper).

<sup>14</sup> See, e.g., Ax Sharma, *OpenAI’s New ChatGPT Bot: 10 Dangerous Things It’s Capable Of*, BLEEPING COMPUTER, Dec. 6, 2022, <https://perma.cc/ZC9F-2FAB>.

<sup>15</sup> See Jane Bambauer, *Negligent AI Speech: Some Thoughts About Duty*, 3 J. FREE SPEECH L. 343 (2023).

they generate in response to a user prompt will push engineers to design the products to minimize the likelihood they produce harmful content.<sup>16</sup>

But this liability regime would also come with costs. If a company that deploys an LLM can be dragged into lengthy, costly litigation any time a user prompts the tool to generate text that creates legal risk, companies will narrow the scope and scale of deployment dramatically. Without Section 230 protection, the risk is vast: Platforms using LLMs would be subject to a wide array of suits under federal and state law.<sup>17</sup> Section 230 was designed to allow internet companies to offer uniform products throughout the country, rather than needing to offer a different search engine in Texas and New York or a different social media app in California and Florida. In the absence of liability protections, platforms seeking to deploy LLMs would face a compliance minefield, potentially requiring them to alter their products on a state-by-state basis or even pull them out of certain states entirely.

With such legal risk, platforms would deploy LLMs only in situations where they could bear the potential costs. This could cause platforms to limit the integration of LLMs into certain products or impose significant restrictions on LLM functionality in order to ensure that it cannot be used to deploy text that might land the platform in court. Platforms might restrict usage to pre-vetted users in an effort to exercise additional control over content. The result would be to limit expression—platforms seeking to limit legal risk will inevitably censor legitimate speech as well. Historically, limits on expression have frustrated both liberals and conservatives,

---

<sup>16</sup> Eugene Volokh posits that this regime might be similar to the notice and takedown system currently used in intellectual property. Eugene Volokh, *Large Libel Models? Liability for AI Output*, 3 J. FREE SPEECH L. 489, 514–21 (2023). Note that even for interactive computer services, platforms cannot currently use Section 230 as a defense in cases brought under IP law. 47 U.S.C. § 230(e)(2).

<sup>17</sup> See Jon M. Garon, *An AI's Picture Paints a Thousand Lies: Designating Responsibility for Visual Libel*, 3 J. FREE SPEECH L. 425 (2023) (describing generative AI platforms' potential liability for "visual libel"); see Nina Brown, *Bots Behaving Badly: A Products Liability Approach to Chatbot-Generating Defamation*, 3 J. FREE SPEECH L. 389 (2023) (describing holding generative AI platforms liable by using product liability pleading as a means of circumventing the mens rea element in defamation claims); see Volokh, *supra* note 16, at 548–54 (describing generative AI platforms' potential liability for a range of torts).

with those on the left concerned that censorship disproportionately harms marginalized communities,<sup>18</sup> and those on the right concerned that censorship disproportionately restricts conservative viewpoints.<sup>19</sup>

The risk of liability could also impact competition in the LLM market. Because smaller companies lack the resources to bear legal costs like Google and Microsoft may, it is reasonable to assume that this risk would reduce startup activity.<sup>20</sup> Large tech companies already have more resources to pour into engineering complex AI technologies, and giving them a compliance advantage in emerging technologies increases the likelihood that the tech sector will become even more concentrated.

Critics of current Section 230 jurisprudence will consider many of these limitations to be welcome. Yet regardless of one's stance on current Section 230 jurisprudence, it is important to recognize that a liability regime limiting the deployment of LLMs will come with costs of its own. If LLMs have the potential to expand expression, broaden access to information, improve productivity, and enhance the quality of a wide range of existing tech products, then there are costs to a liability regime that threatens to constrain their development and deployment.

#### IV. HOW SHOULD LLMs BE REGULATED?

If Section 230 is left intact as LLMs emerge, then critics of the liability protection will get what they have long desired: cracks in the shield. But to explore the real potential benefits of LLMs and to understand their capacity both for harm and for good, Section 230 reform is necessary.

The rapid emergence of ChatGPT will likely put pressure on lawmakers to reform Section 230 in a way that is largely antithetical to their recent efforts: That is, it may push Congress to limit liability for generative AI products that generate content in response to user inputs. In order to flourish, LLMs will need a liability regime that allows them to be deployed without crippling them with lawsuits and court fees.

---

<sup>18</sup> Senator Mark Warner, *The SAFE TECH Act (Safeguarding Against Fraud, Exploitation, Threats, Extremism and Consumer Harms Act)*, <https://perma.cc/TLV4-MNNU>.

<sup>19</sup> Press Release, Senator Ted Cruz, *Sen. Cruz: Big Tech Believes There Is No Power That Can Constrain Them* (Dec. 11, 2020), <https://perma.cc/JF3M-ZW79>.

<sup>20</sup> Press Release, Engine, *Statement on Engine Filing Section 230 Briefs with Supreme Court* (Jan. 19, 2023), <https://perma.cc/LA5L-DL4D>.

Ideal reform would not immunize LLM companies entirely but would instead align the costs they bear with social harm created by their products. That goal is easy to articulate but incredibly difficult to achieve.

One option would be to reform Section 230 to remove the “in part” language from the definition of an information content provider: “any person or entity that is responsible, in whole, for the creation or development of information provided through the Internet or any other interactive computer service.” That change would immunize LLMs in cases where they generate illegal content in response to a user prompt. The obvious downside is that it would provide even broader liability protection for tech platforms, contrary to the aim of most policymakers.

An alternative approach would be to reform Section 230 to explicitly carve out LLMs from the definition of an information content provider: “any person or entity that is responsible, in whole or in part, for the creation or development of information provided through the Internet or any other interactive computer service. A large language model shall not be considered an information content provider when it generates text in response to a query from another information content provider.” While this revision would maintain the existing scope of liability for the tech sector, it would provide too broad an immunity for LLMs: The law should provide sufficient incentive for LLMs to avoid generating illegal content.

These two alternatives might be more palatable to policymakers if they were combined with other Section 230 reforms that have been pushed by both parties.<sup>21</sup> For instance, Democrats might be more willing to consider an LLM carveout if it were accompanied by provisions that prohibit tech platforms from using Section 230 in cases involving state civil rights law. Republicans might be more willing to consider LLM-related reforms if they were combined with revisions that expand liability when platforms remove conservative speech. While such reforms might

---

<sup>21</sup> Beatriz Botero Arcila draws on European law to propose an alternative framing, where a broad safe harbor is provided, but only if generative AI companies implement risk mitigation obligations. That framing draws on European models of intermediary liability regulation, including the Digital Services Act. Beatriz Botero Arcila, *Is It a Platform? Is It a Search Engine? It's ChatGPT! A Commentary on the Liability Regime of Online Content Generated by Large Language Models in Europe*, 3 J. FREE SPEECH L. 455 (2023).



make reform more politically feasible, they come with their own downsides. Experts contend they are likely to increase censorship in some cases and increase the proliferation of harmful content in others.<sup>22</sup>

But any recommendation to expand liability protection for LLMs seems infeasible in the current political environment. In fact, the recommendation that seems to be gaining the most traction is to require platforms to “slow down.” A large group of more than 13,000 engineers, experts, and researchers signed an open letter calling for AI labs to pause the training of AI systems more powerful than ChatGPT-4 for at least six months.<sup>23</sup> The letter recommends that if the private sector does not commit to this pause, then the government should “institute a moratorium.”<sup>24</sup>

While the idea of a pause may be appealing in the face of rampant speculation about the potential harms of the technology, it is not a course that companies or the government should pursue. First, it seems to assume that a pause is cost-free, simply because it involves inaction rather than action. But this omission bias glosses over the fact that inaction can have costs as well.<sup>25</sup> If the technology confers benefits that outweigh harms, then ceasing development will produce a worse outcome. A pause might also cause companies to slow development of new tools that might help to make the technology safer, less prone to abuse, and more accurate.

Second, it ignores that good actors are likely to abide by a pause, but bad actors are not. So even if the largest, leading generative AI providers that have been in the spotlight on this issue—companies like Google, Microsoft, and OpenAI—obey a pause order, users will likely be able to access tools from a range of other providers who might be less inclined to develop it responsibly.

Third, other countries would likely use a pause to try to gain a competitive advantage relative to American firms. As has been well-documented, China is invest-

---

<sup>22</sup> Brief for TechFreedom as Amici Curiae Supporting Respondents, *Gonzalez v. Google*, No. 21-1333 (U.S. argued Feb. 21, 2023).

<sup>23</sup> Future of Life Institute, *Pause Giant AI Experiments: An Open Letter*, Mar. 22, 2021, <https://perma.cc/SKH2-JB3L>.

<sup>24</sup> *Id.*

<sup>25</sup> The Decisional Lab, *Why Don't We Pull the Trolley Lever: The Omission Bias, Explained*, <https://perma.cc/97BU-YPW4>.

ing aggressively in AI, and in the wake of the launch of ChatGPT, Chinese companies have been rushing to build competing products.<sup>26</sup> US companies now seem to hold an advantage over their rivals abroad, but if they pause development, their competitors would likely use that time to try to close the gap.

Instead of slowing down, the optimal approach at this stage would be to move in the other direction: accelerate experimentation with the technology, but require platforms to share data to enable learning and improvement. To this end, policymakers could pass an LLM-specific liability law that provides narrow, time-limited immunity for LLM platforms, while enabling judges, policymakers, and academics to gather data about the kinds of content generated by LLMs and its impact on society. The objective would be to learn about the costs and benefits of LLMs, to better understand the types of situations where legal liability might arise, and to use this information to develop a smarter, stronger liability regime in the future.

A new law could be crafted to provide some temporary immunity for LLMs, while also ensuring that LLMs could be held liable in certain cases, such as when they are actively involved in violations of federal criminal law. To facilitate sufficient transparency, the law should enable platforms to provide data to researchers and should protect both platforms and researchers when data is shared consistent with privacy and security best practices. Such a law would not open up the text of Section 230 and therefore would not alter current law for the rest of the internet.

The risks of this approach could be minimized by making it time bound and transparent. The law could sunset in two years so as to ensure that the limited liability for LLMs does not continue indefinitely, and it could create government-funded audit committees composed of technologists, academics, and community organizations to review the impact of LLMs as they are deployed more broadly. These committees might be required to report on specific aspects of LLM performance, such as whether they disproportionately harm marginalized communities or disproportionately censor specific political viewpoints. They might also be required to produce policy recommendations at the end of the two-year period so as to provide lawmakers with a road map for reform.

Of course, none of these options is likely to be politically feasible. In the current political environment, lawmakers are unlikely to support legislation that would

---

<sup>26</sup> Zeyi Yang, *Chinese Tech Giant Baidu Just Released Its Answer to ChatGPT*, MIT TECH. REV., Mar. 16, 2023, <https://perma.cc/4WUM-RVS9>.

provide liability protections for new technologies. It is far more likely that Congress will hold hearings about use cases that receive negative media attention and propose messaging bills that are aimed more toward generating headlines than toward establishing tangible new rules to guide the emergence of generative AI products.

If Congress does not take action and LLMs are held liable as information content providers, then it will fall to the tech research ecosystem to ensure that the impact of this regulatory regime is quantified in some form.

In many situations, we measure the impact of action we take, but not the impact of inaction.<sup>27</sup> In the case of LLMs, a failure to act will have costs. The latest technology will be deployed more narrowly than it would be otherwise. Valuable features will never be built. Fearing legal fees and burdens, startups will be less likely to enter the market. Academics and tech experts should measure those costs, with the aim of developing a better understanding of the costs and benefits of the status quo that can then inform future policy.

#### CONCLUSION: WHAT WILL THE FUTURE HOLD ?

Policymakers have a range of policy tools they might use to manage the emergence of LLMs, but it's hard to be optimistic that they will do it well. Tech reform has stalled in Congress for years, despite increasing political consensus that reform is necessary. Given partisan disagreements on the core objectives of Section 230 reform, political compromise is likely to remain elusive.

The most likely outcome is that Congress will stand still while courts take the lead in dictating the liability regime for LLMs. Since they must make decisions based on existing law, rather than crafting what a more desirable liability regime might be, judicial bodies are not well positioned to design an approach that adequately balances the potential benefits of LLMs against their potential costs. As a result, Section 230's critics will get what they have long wanted and its proponents what they have long feared: expanded liability for new technologies, alongside diminished opportunity for expression and innovation.

---

<sup>27</sup> See The Decisional Lab, *supra* note 25.

