

TESTIMONY OF BRAD R. CARSON

President, Americans for Responsible Innovation

United States Senate Committee on Commerce, Science, and Transportation

“Liability or Deniability? Platform Power as Section 230 Turns 30”

March 18, 2026

Chairman Cruz, Ranking Member Cantwell, and Members of the Committee:

Thirty years ago, Congress passed Section 230 to address a narrow but critical problem: whether platforms could moderate content without assuming publisher liability. Today, that law is widely - and rightly - criticized for enabling unchecked harms. As we confront artificial intelligence, we must ask, “Will we repeat the mistakes of Section 230, or will we learn from them?”

Some believe that Section 230 was ill-conceived from the very beginning and that ordinary tort and First Amendment case-by-case development through the common law would have produced a more nuanced and flexible body of law than did the broad statutory immunity of Section 230. On this view, Section 230 did, it should be acknowledged, accomplish one important task. Subsection (c)(2) of Section 230, which grants immunity to platforms for good faith content moderation decisions, directly addressed a real concern created by cases such as *Stratton Oakmont v. Prodigy*. In *Stratton Oakmont*, a trial court held that a platform’s attempt to curate content exposed it to liability as a publisher, thus penalizing the responsible content moderation. In correcting this erroneous court ruling, the protection provided by Section 230(c)(2) was very reasonable. However, subsection (c)(1) of Section 230, the broad immunity for third-party content, was not necessary. The common law and First Amendment were capable of producing the rules governing liability for third-party content, and the courts could have developed those rules over time in response to actual harm caused. Rather than allowing the development of these rules to occur, Congress in Section 230 froze the existing answers to questions that the law had not yet fully considered.

Another view is that Section 230 was defensible as enacted, however the courts interpreted it far beyond what Congress intended. On this view, the primary error occurred in the Fourth Circuit in *Zeran v. America Online* in 1997, where the court eliminated the long-standing distinction between publishers and distributors. Publishers are those entities that exercise editorial control over content and therefore bear full liability for the publication of that content. Distributors, on the other hand, are those entities that merely transmit content knowing its character and bear a limited liability based upon their role. Before the enactment of Section 230, a distributor that knew of particular harmful content and failed to take action against that content could be held liable for failing to act to prevent the distribution of the content. The Fourth Circuit interpreted Section 230 to eliminate that distinction altogether, holding that the statute granted platforms immunity from liability even when the platform had actual notice of the existence of specific harms and chose to take no action to prevent the continued distribution of the harmful content. Following the decision in *Zeran*, subsequent decisions extended that immunity further still, to

include algorithmic amplification, product design, and to business decisions that had nothing to do with the hosting of user-generated content. The statute was not the problem, according to this analysis, but rather the interpretation of the statute went awry.

Regardless of which is more accurate, both assessments of Section 230 recognize that the resulting legal regime has been unable to hold platforms accountable in proportion to the harms they have enabled. Families whose children are victimized by predators have been denied access to the courtroom. Parents whose teenage child experienced years of self-harm associated with foreseeable algorithmic targeting and engagement-optimization design choices discovered that the law too often offers them no realistic path to accountability. The social media platforms have borne few of the costs of the harms they have facilitated. Rather, those costs have fallen to the individuals that have sadly learned that the law affords them little protection when technology goes wrong.

I recount the controversial history of Section 230 because Congress now finds itself facing a structurally equivalent moment regarding artificial intelligence. Many have suggested that we freeze state law developments while enacting no comprehensive federal regulation of the technology. No matter whether the fault rests with the statute itself or with the interpretations that followed, the ultimate result of Section 230 was a meta-law - a law that determines who will govern an emerging industry instead of how that governance will occur - enacted for an industry that was in its infancy and whose consequences proved nearly uncorrectable once the industry matured. That is the error that Congress is being encouraged to repeat with respect to artificial intelligence. In practice, the first meta-law often becomes the last major law because it determines who has power to resist future corrections.

I. Section 230 Should Not Be Interpreted To Immunize Generative AI Outputs

I would like to make clear that Section 230 should not be interpreted to immunize the provider of a generative AI system for harms caused by the system's own outputs. Section 230(c)(1) grants immunity to platforms for the content "provided by *another* information content provider." Platforms receive this immunity because the content posted by users originates with the user; the platform is merely the host, not the author. The statute contemplates a world in which there exist active users and passive platforms. The latter are generally immunized from the actions of the former. (Of course, when a platform merely hosts content posted by users - including content those users generated with AI tools - Section 230 may still protect that hosting function. But the company that designs, trains, and deploys a model should not be treated as a passive host of the model's own speech-like outputs.)

A large language model does not provide third-party content in the way contemplated by Section 230(c)(1). A user provides a prompt to the model. The model, having been trained on a dataset selected by the company, having been fine-tuned by the company, and having been deployed with parameters determined by the company, produces the output. Therefore, the company is responsible, in significant measure, for the creation of the content. Pursuant to the statute's literal text, the immunity of Section 230 does not apply to artificial intelligence; the output of Claude or ChatGPT is not the "information content" of "another." I believe, though, that Congress should clarify this before lower courts, following the *Zeran* precedent, extend immunity further than the text of Section 230 permits and establish an erroneous precedent that will require yet another generation to correct. Legislation introduced by Senators Hawley and Durbin would make clear

that Section 230 does not apply to generative AI, and I would urge Congress to move swiftly to enact it.

I ask Congress to act even though the textual answer seems plain. The lesson of Section 230 is precisely that textually plain answers do not always survive contact with courts under pressure to resolve novel cases quickly and in the absence of congressional guidance. *Zeran* was not an unreasonable decision given the legal vacuum in which the Fourth Circuit found itself. It was, however, a consequential one. Congress should not again create the conditions in which courts are forced to improvise and then spend a generation arguing about whether the improvisation was correct.

II. Congress Should Not Create An Industry-Wide Immunity For An Emerging Industry That It Cannot Yet Predict

While the question of whether Section 230 applies to AI can be clearly answered in the negative, the more relevant lesson of Section 230 is that Congress should not establish broad, essentially unchangeable immunity for an industry that it does not yet understand and whose future development can scarcely be charted.

When Congress passed Section 230 in 1996, it could not have imagined that a few websites would soon control virtually all public communication, influence election results, and have market capitalization larger than many countries' entire GDPs. Indeed, in 1996, just 7% of Americans had access to the internet, Netscape was the dominant browser, Google did not exist, and Facebook's launch was eight years away.

Yet Section 230 shielded an entire infant industry from legal responsibility. The immunity has been difficult to repeal as the infant industry has grown beyond anyone's imagination. In retrospect, creating broad immunity for an emerging industry should be seen as a category error disguised as prudence. We do not permanently exempt a child from the law because the child appears harmless today; we calibrate accountability to capability and impose it consistently as capability grows. That is what a responsible legal system does. Section 230 did not do that.

III. Preemption Without A Federal Framework Would Repeat Section 230's Mistakes

The discussion regarding federal preemption of state laws related to artificial intelligence is worthy of separate treatment here, because it is not merely analogous to the historical development of Section 230; it is a repeat of Section 230's history, replete with the errors of the former, and without any of the latter's corrective measures.

Section 230 created two types of legal stagnation. The first was interpretive. Because Section 230 created a broad immunity, courts never developed common law governing platform liability for actual harms. The second type of stagnation was legislative. Once Section 230 became entrenched, amending or replacing it required a political consensus that proved nearly impossible to assemble. State legislatures found their efforts preempted before they could be tested. For its part, Congress could not create new legal structures without first eliminating an existing statute that had become a "third rail" politically. As such, Section 230 remains frozen in place as technology continues to revolutionize the world.

Similarly, a federal immunity toward AI – which is what the preemption proposals offered in Congress to date really are - will lead to the same stagnation, but more rapidly, and for a technology whose potential implications are much larger than those of social media. Consider what preemption without a federal framework actually accomplishes. It removes a state’s authority to apply its tort law, consumer protection statutes, and civil liability regime to AI systems. It prevents plaintiffs from having access to courtrooms prior to the development of a federal alternative to receive them. It informs courts that they are under no obligation to develop a body of common law regarding AI liability, while providing no alternative for them to apply. In short, it is a legal vacuum created by Congress to which neither the courts nor the legislative branch can respond. This is the Section 230 model, deliberately reproduced.

If Congress wishes to preempt state laws governing AI, then the cost of that preemption must include a genuine federal framework that replaces the space that the preemption vacates. This federal framework could include, first, a flexible federal structure that assigns liability standards to agencies based on categories of risk, allowing those standards to evolve as the technology does rather than locking them into statute. Second, the framework could help develop markets for independent AI auditing and safety certification, established through procurement preferences and safe harbors conditional on third-party review, to create accountability incentives in the marketplace. Third, the framework could insist on mandatory liability insurance calibrated to a system’s capability and risk profile, compensating victims while creating pricing disciplines that accumulate with experience. None of these approaches require Congress to forecast the nature of AI in ten years. To the contrary, that is one of the advantages they possess over broad immunity. Such a federal framework establishes accountability proportional to capability, rather than immunity inversely proportional to foresight.

IV. Conclusion

Section 230 was an understandable, if lamentable, case of poor foresight exacerbated by a failure to permit the law to evolve as the technology evolved. The people who have suffered the most as a result of that failure have been the families, the teenagers, and the individual victims who discovered that the law offers them nothing when technology harms them.

We should not repeat that story with artificial intelligence. Meta-laws and broad immunities for emerging industries rarely age well, and the lesson of Section 230 is that they age worst of all for the people least able to protect themselves. Congress has better tools available, including frameworks that scale with capability, standards that evolve with the technology, and accountability that grows as power grows. The choice is between mechanisms that establish accountability proportional to capability and immunities that are inversely proportional to foresight.

Thank you. I am happy to respond to your questions.