

Zeran v. AOL essay

Jonathan Zittrain

Twenty years after it was first litigated in earnest, the U.S. Communications Decency Act's Section 230 remains both obscure and vital. Section 230 nearly entirely eliminated the liability of Internet content platforms under state common law for bad acts, such as defamation, occasioned by their users. The platforms were free to structure their moderation and editing of comments as they pleased, without a traditional newspaper's framework in which to undertake editing was to bear responsibility for what was published. If the *New York Times* included a letter to the editor that defamed someone, the *Times* would be vulnerable to a lawsuit (to be sure, so would the letter's author, whose wallet size would likely make for a less tempting target). Not so for online content portals that welcome comments from anywhere – including the online version of the *New York Times*.

This strange medium-specific subsidy for online content platforms made good if not perfect sense in 1996. (My generally positive thinking about it from that time, including some reservations, can be found [here](#).) The Internet was newly mainstream, and many content portals comprised the proverbial two people in a garage. To impose upon them the burdens of traditional media would presumably require tough-to-maintain gatekeeping. Comments sections, if they remained at all, would have to be carefully screened to avoid creating liability for the company. What made sense for a newspaper publishing at most five or six letters a day amidst its more carefully vetted articles truly couldn't work for a small Internet startup processing thousands or even millions of comments or other contributions in the same interval. Over time, the reviews elicited by Yelp and TripAdvisor, the financial markets discussions on Motley Fool, the evolving articles on user-edited Wikipedia – all are arguably only possible thanks to that Section 230 immunity conferred in 1996.

The immunity conferred is so powerful that there's not only a subsidy of digital over analog, but one for third-party commentary over one's own -- or that of one's employees. Last year the notorious Gawker.com settled for \$31 million after being successfully sued for publishing a two-minute extract of a private sex video. If Gawker, instead of employing a staff whose words (and video excerpts) were attributable to the company, had simply let any anonymous user post the same excerpt -- and indeed worked to assure that that user's anonymity could not be pierced -- it would be immune from an identical invasion of privacy suit thanks to the CDA. From this perspective, Gawker's mistake wasn't to host the video, but to have its own employees be the ones to post it.

The Internet environment of 2017 is a lot different than that of 1997, and some of those two-people-in-a-garage ventures are now among the most powerful and valuable companies in the world. So does it make sense to maintain Section 230's immunities today?

1. [An infant industry has grown up.](#)

In 1997, it made sense on a number of fronts to treat the Internet differently from its analog counterparts. For example, there was debate from the earliest mainstreaming of Internet commerce about whether to make U.S. state sales tax collection apply to Internet-based faraway purchases. The fact that there was so little Internet commerce meant that there was not a lot of money foregone by failing to tax; that new companies (and, for that matter, existing ones) could try out e-commerce models without concerning themselves from the start with tax compliance in multiple jurisdictions; and that the whole Internet sector could gather momentum if purchasers were enticed to go online – which in turn would further entice more commerce, and other activity, online. I was among those who therefore [argued in favor of the de facto moratorium](#) on state sales tax. But that differential no longer makes sense. A single online company – Amazon – now accounts for about [5% of all U.S. retail sales](#), online or off. It's a good thing that Amazon's physical expansion has meant that it naturally has started collecting and remitting state sales tax around the country.

Perhaps the evolution of the merits of equal treatment for state sales tax provides a good model for a refined CDA: companies below a certain size or activity threshold could benefit from its immunities, while those who grow large enough to facilitate the infliction of that much more damage from defamatory and other actionable posts might also have the resources to employ a compliance department. That would militate towards at least some standard to meet in vetting or dealing with posts, perhaps akin to the light duties of booksellers or newsstands towards the wares they stock rather than the higher ones of newspapers towards the letters they publish. Apart from the first-order drawback of an incentive to game the system by staying just under whatever size or activity threshold triggers the new responsibilities, there's also the question of non-commercial communities that can become large without having traditional corporate hierarchies that lend themselves to direct legal accountability. Some of the most important computing services in the world rely on free and open source software, even as there remains a puzzle of how software liability would work when there's no organized firm singly producing it. This puzzle has remained unsolved even today, since liability for bugs or vulnerabilities in even corporate-authored software tends to be quite minimal. That might change as the line between hardware and software continues to blur with the Internet of Things.

Even for companies suited for new, light responsibilities under a modified CDA, there might be a distinction made between damages for past acts and duties for future ones. The toughest part of the *Zeran* case even for those sympathetic to the CDA is that apparently AOL was repeatedly told that the scandalous advertisement purporting to be from Ken Zeran was in fact not at all related to him – and the company was in a comparatively good position to confirm that. Even then the company did nothing. It's one thing to have permitted some defamatory content to come through amidst millions of messages; it's another to be fully aware of it once it's posted, and to still not be charged with any responsibility to deal with it. A more refined CDA might underscore such a distinction, favoring the kind of knowledge of falsehood that's at the heart of the heightened *New York Times v. Sullivan* barrier that public figures must meet in establishing defamation by a newspaper, and also cover knowledge that might come about after publication rather than before – leading only to responsibility once the knowledge is gained and not timely acted upon.

2. The AI thicket.

Even massive online speech-mediating companies can only hire so many people. With [thousands of staffers](#) around the world apparently committed to reviewing complaints arising over Facebook posts, the company still relies on algorithms to sift helpful from unhelpful content. And here the distinction between pre- and post-publication becomes blurred, because services like Facebook and Twitter not only host content – as a newspaper website does by permitting comments to appear in sequence after an article – but they also help people navigate it. A post might reach ten people or a billion, depending on whether it's placed in no news feeds or many.

The CDA as it stands allows maximum flexibility for salting feeds, since no liability will attach for spreading even otherwise-actionable content far and wide. A refined CDA could take into account the fact that Facebook and others know exactly whom they've reached: perhaps a more reasonable and fitting remedy for defamation would less be to assess damages against the company for having abetted it, but rather to require a correction or other followup to go out to those who saw – and perhaps came to believe – the defamatory content. (To be sure, this solution doesn't work for other wrongs such as invasion of privacy; no correction can “uninvade” it among those who saw the content in question.)

Such corrective, rather than compensatory, remedies may be more fitting both for the wronged party and for the publisher, but it could in turn make content elision much more common. For example, in the context of traditional book publishing, including for non-interactive digital books like those within a Kindle, the CDA does not protect the publisher against the author's defamation. With a threat of liability remaining, I've worried that [in addition to damages, a litigant might demand a digital retraction](#): a forced release of a new version of an e-book to all e-readers that omits the defamatory content.

Of course, if the challenged words are really defamatory that might be thought of as an improvement for both injured party and for the reader. But if done without notice to the reader, [it smacks of propaganda](#), and to the extent lawsuits or threats of same can induce defendant publishers to cave – when caving doesn't entail paying out damages but rather altering the content they've stewarded – it could come to happen all too frequently, and with the wrong incentives. Similarly, an AI trained to avoid controversial subjects – perhaps defined as subjects that could give rise to threats of litigation – might be very much against the public interest. This would mirror some of the [damaging incentives](#) of Europe's “right to be forgotten” as developed against search engines. Any refinement of the CDA that could inspire AI-driven content shaping runs this risk, with the perverse solace that even with today's CDA the major content platforms are already shaping content in ways that are not understandable or reviewable outside the companies.

Related to the power of AI is the refined power to personalize content in 2017, including by jurisdiction. If a Texas court finds something defamatory under Texas law, such as [maligning](#)

[certain food products](#), it might not be defamatory under, say, Massachusetts law. Any diminution of CDA 230's immunities might in the first order impel online platforms like Facebook to have to police away any food disparagement – even if it's posted and read by Facebook users in food-indifferent Massachusetts. If there were to be exposure under Texas law, perhaps it should only arise if the content were shown (or continued to be shown) in Texas. This could also provide a helpful set of pressures on the substantive doctrine: Texas citizens, including legislators, might rue being excluded from certain content online that's available in other states.

The Internet's development over the past twenty years has benefited immeasurably from the immunities conferred by Section 230. We've been lucky to have it. But any honest account must acknowledge the collateral damage it has permitted to be visited upon real people whose reputations, privacy, and dignity have been hurt in ways that defy redress. Especially as that damage becomes more systematized – now part of organized campaigns to shame people into silence online for expressing opinions that don't fit an aggressor's propaganda aims – platforms' failures to moderate become more costly, both to targets of harassment and to everyone else denied exposure to honestly-held ideas.

As our technologies for sifting and disseminating content evolve, and our content intermediaries trend towards increasing power and centralization, there are narrow circumstances where a path to accountability for those intermediaries for the behavior of their users might be explored. Incrementalism gets a bad rap, but it's right to proceed slowly if at all here, with any tweaks subject to rigorous review of how they impact the environment. The vice from the indiscriminate nature of Section 230's broad immunity is somewhat balanced by a virtue of everyone knowing exactly where matters stand -- line-drawing carries its own costs and distortions.