ICLE International Center for Law & Economics

The Third Circuit's Oberdorf v. Amazon Opinion Offers a Good Approach to Reining in the Worst Abuses of Section 230
July 15, 2019
Gus Hurwitz

*[**Note:** A group of 50 academics and 27 organizations, including both myself and ICLE, recently released a statement of principles for lawmakers to consider in discussions of Section 230.]*

In a remarkable ruling issued earlier this month, the Third Circuit Court of Appeals held in *Oberdorf v. Amazon* that, under Pennsylvania products liability law, Amazon could be found liable for a third party vendor's sale of a defective product via Amazon Marketplace. This ruling comes in the context of Section 230 of the Communications Decency Act, which is broadly understood as immunizing platforms against liability for harmful conduct posted to their platforms by third parties (Section 230 purists may object to myu use of "platform" as approximation for the statute's term of "interactive computer services"; I address this concern by acknowledging it with this parenthetical). This immunity has long been a bedrock principle of Internet law; it has also long been controversial; and those controversies are very much at the fore of discussion today.

The response to the opinion has been mixed, to say the least. Eric Goldman, for instance, has asked "are we at the end of online marketplaces?," suggesting that they "might in the future look like a quaint artifact of the early 21st century." Kate Klonick, on the other hand, calls the opinion "a brilliant way of both holding tech responsible for harms they perpetuate & making sure we preserve free speech online."

My own inclination is that both Eric and Kate overstate their respective positions – though neither without reason. The facts of *Oberdorf* cabin the effects of the holding both to Pennsylvania law and to situations where the platform cannot identify the seller. This suggests that the effects will be relatively limited.

But, and what I explore in this post, the opinion does elucidate a particular and problematic feature of section 230: that it can be used as a liability shield for harmful conduct. The judges in *Oberdorf* seem ill-inclined to extend Section 230's protections to a platform that can easily be used by bad actors as a liability shield. Riffing on this concern, I argue below that Section 230 immunity be proportional to platforms' ability to reasonably identify speakers using their platforms to engage in harmful speech or conduct.

This idea is developed in more detail in the last section of this post – including responding to the obvious (and overwrought) objections to it. But first it offers some background on

Section 230, the *Oberdorf* and related cases, the Third Circuit's analysis in *Oberdorf*, and the recent debates about Section 230.

Section 230

"Section 230" refers to a portion of the Communications Decency Act that was added to the Communications Act by the 1996 Telecommunications Act, codified at 47 U.S.C. 230. (NB: that's a sentence that only a communications lawyer could love!) It is widely recognized as – and discussed even by those who disagree with this view as – having been critical to the growth of the modern Internet. As Jeff Kosseff labels it in his recent book, the key provision of section 230 comprises the "26 words that created the Internet." That section, 230(c)(1), states that "No 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." (For those not familiar with it, Kosseff's book is worth a read – or for the Cliff's Notes version see here, here, here, here, here, or here.)

Section 230 was enacted to do two things. First, section (c)(1) makes clear that platforms are not liable for user-generated content. In other words, if a user of Facebook, Amazon, the comments section of a Washington Post article, a restaurant review site, a blog that focuses on the knitting of cat-themed sweaters, or any other "interactive computer service," posts something for which that user may face legal liability, the platform hosting that user's speech does not face liability for that speech.

And second, section (c)(2) makes clear that platforms are free to moderate content uploaded by their users, and that they face no liability for doing so. This section was added precisely to repudiate a case that had held that once a platform (in that case, Prodigy) decided to moderate user-generated content, it undertook an obligation to do so. That case meant that platforms faced a Hobson's choice: either don't moderate content and don't risk liability, or moderate all content and face liability for failure to do so well. There was no middle ground: a platform couldn't say, for instance, "this one post is particularly problematic, so we are going to take it down – but this doesn't mean that we are going to pervasively moderate content."

Together, these two provisions stand generally for the proposition that online platforms are not liable for content created by their users, but they are free to moderate that content without facing liability for doing so. It recognized, on the one hand, that it was impractical (i.e., the Internet economy could not function) to require that platforms moderate all user-generated content, so section (c)(1) says that they don't need to; but, on the other hand, it recognizes that it is desirable for platforms to moderate problematic content to the best of their ability, so section (c)(2) says that they won't be punished (i.e., lose the immunity granted by section (c)(1) if they voluntarily elect to moderate content).

Section 230 is written in broad – and has been interpreted by the courts in even broader – terms. Section (c)(1) says that platforms cannot be held liable for the content generated by their users, full stop. The only exceptions are for copyrighted content and content that

violates federal criminal law. There is no "unless it is really bad" exception, or a "the platform may be liable if the user-generated content causes significant tangible harm" exception, or an "unless the platform knows about it" exception, or even an "unless the platform makes money off of and actively facilitates harmful content" exception. So long as the content is generated by the user (not by the platform itself), Section 230 shields the platform from liability.

Oberdorf v. Amazon

This background leads us to the Third Circuit's opinion in *Oberdorf v. Amazon*. The opinion is remarkable because it is one of only a few cases in which a court has, despite Section 230, found a platform liable for the conduct of a third party facilitated through the use of that platform.

Prior to the Third Circuit's recent opinion, the best known previous case is the 9th Circuit's Model Mayhem opinion. In that case, the court found that Model Mayhem, a website that helps match models with modeling jobs, had a duty to warn models about individuals who were known to be using the website to find women to sexually assault.

It is worth spending another moment on the *Model Mayhem* opinion before returning to the Third Circuit's *Oberdorf* opinion. The crux of the 9th Circuit's opinion in the Model Mayhem case was that the state of Florida (where the assaults occurred) has a duty-to-warn law, which creates a duty between the platform and the user. This duty to warn was triggered by the case-specific fact that the platform had actual knowledge that two of its users were predatorily using the site to find women to assault. Once triggered, this duty to warn exists between the platform and the user. Because the platform faces liability directly for its failure to warn, it is not shielded by section 230 (which only shields the platform from liability for the conduct of the third parties using the platform to engage in harmful conduct).

In its opinion, the Third Circuit offered a similar analysis – but in a much broader context.

The *Oberdorf* case involves a defective dog leash sold to Ms. Oberdorf by a seller doing business as The Furry Gang on Amazon Marketplace. The leash malfunctioned, hitting Ms. Oberdorf in the face and causing permanent blindness in one eye. When she attempted to sue The Furry Gang, she discovered that they were no longer doing business on Amazon Marketplace – and that Amazon did not have sufficient information about their identity for Ms. Oberdorf to bring suit against them.

Undeterred, Ms. Oberdorf sued Amazon under Pennsylvania product liability law, arguing that Amazon was the seller of the defective leash, so was liable for her injuries. Part of Amazon's defense was that the actual seller, The Furry Gang, was a user of their Marketplace platform – the sale resulted from the storefront generated by The Furry Gang and merely hosted by Amazon Marketplace. Under this theory, Section 230 would bar Amazon from liability for the sale that resulted from the seller's user-generated storefront.

The Third Circuit judges had none of that argument. All three judges agreed that under Pennsylvania law, the products liability relationship existed between Ms. Oberdorf and Amazon, so Section 230 did not apply. The two-judge majority found Amazon liable to Ms. Oberford under this law – the dissenting judge would have found Amazon's conduct insufficient as a basis for liability.

This opinion, in other words, follows in the footsteps of the Ninth Circuit's *Model Mayhem*opinion in holding that state law creates a duty directly between the harmed user and the platform, and that that duty isn't affected by Section 230. But *Oberdorf* is potentially much broader in impact than *Model Mayhem*. States are more likely to have broader product liability laws than duty to warn laws. Even more impactful, product liability laws are generally strict liability laws, whereas duty to warn laws are generally triggered by an actual knowledge requirement.

The Third Circuit's Focus on Agency and Liability Shields

The understanding of *Oberdorf* described above is that it is the latest in a developing line of cases holding that claims based on state law duties that require platforms to protect users from third party harms can survive Section 230 defenses.

But there is another, critical, issue in the background of the case that appears to have affected the court's thinking – and that, I argue, should be a path forward for Section 230. The judges writing for the Third Circuit majority draw attention to

> the extensive record evidence that Amazon fails to vet third-party vendors for amenability to legal process. The first factor [of analysis for application of the state's products liability law] weighs in favor of strict liability not because The Furry Gang cannot be located and/or may be insolvent, but rather because Amazon enables third-party vendors such as The Furry Gang to structure and/or conceal themselves from liability altogether.

This is important for analysis under the Pennsylvania product liability law, which has a marketing chain provision that allows injured consumers to seek redress up the marketing chain if the direct seller of a defective product is insolvent or otherwise unavailable for suit. But the court's language focuses on Amazon's design of Marketplace and the ease with which Marketplace can be used by merchants as a liability shield.

This focus is unsurprising: the law generally does not allow one party to shield another from liability without assuming liability for the shielded party's conduct. Indeed, this is pretty basic vicarious liability, agency, first-year law school kind of stuff. It is unsurprising that judges would balk at an argument that Amazon could design its platform in a way that makes it impossible for harmed parties to sue a tortfeasor without Amazon in turn assuming liability for any potentially tortious conduct.

Section 230 is having a bad day

As most who have read this far are almost certainly aware, Section 230 is a big, controversial, political mess right now. Politicians from Josh Hawley to Nancy Pelosi have suggested curtailing Section 230. President Trump just held his "Social Media Summit." And countries around the world are imposing near-impossible obligations on platforms to remove or otherwise moderate potentially problematic content – obligations that are anathema to Section 230 as they increasingly reflect and influence discussions in the United States.

To be clear, almost all of the ideas floating around about how to change Section 230 are bad. That is an understatement: they are potentially devastating to the Internet – both to the economic ecosystem and the social ecosystem that have developed and thrived largely because of Section 230.

To be clear, there is also a lot of really, disgustingly, problematic content online – and social media platforms, in particular, have facilitated a great deal of legitimately problematic conduct. But deputizing them to police that conduct and to make real-time decisions about speech that is impossible to evaluate in real time is not a solution to these problems. And to the extent that some platforms may be able to do these things, the novel capabilities of a few platforms to obligations for all would only serve to create entry barriers for smaller platforms and to stifle innovation.

This is why a group of 50 academics and 27 organizations released a statement of principles last week to inform lawmakers about key considerations to take into account when discussing how Section 230 may be changed. The purpose of these principles is to acknowledge that some change to Section 230 may be appropriate – may even be needed at this juncture – but that such changes should be careful and modest, carefully considered so as to not disrupt the vast benefits for society that Section 230 has made possible and is needed to keep vital.

The Third Circuit offers a Third Way on 230

The Third Circuit's opinion offers a modest way that Section 230 could be changed – and, I would say, improved – to address some of the real harms that it enables without undermining the important purposes that it serves. To wit, Section 230's immunity could be attenuated by an obligation to facilitate the identification of users on that platform, subject to legal process, in proportion to the size and resources available to the platform, the technological feasibility of such identification, the foreseeability of the platform being used to facilitate harmful speech or conduct, and the expected importance (as defined from a First Amendment perspective) of speech on that platform.

In other words, if there are readily available ways to establish some form of identify for users – for instance, by email addresses on widely-used platforms, social media accounts, logs of IP addresses – and there is reason to expect that users of the platform could be

subject to suit – for instance, because they're engaged in commercial activities or the purpose of the platform is to provide a forum for speech that is likely to legally actionable – then the platform needs to be reasonably able to provide reasonable information about speakers subject to legal action in order to avail itself of any Section 230 defense. Stated otherwise, platforms need to be able to reasonably comply with so-called unmasking subpoenas issued in the civil context to the extent such compliance is feasible for the platform's size, sophistication, resources, &c.

An obligation such as this would have been at best meaningless and at worst devastating at the time Section 230 was adopted. But 25 years later, the Internet is a very different place. Most users have online accounts – email addresses, social media profiles, &c – that can serve as some form of online identification.

More important, we now have evidence of a growing range of harmful conduct and speech that can occur online, and of platforms that use Section 230 as a shield to protect those engaging in such speech or conduct from litigation. Such speakers are clear bad actors who are clearly abusing Section 230 facilitate bad conduct. They should not be able to do so.

Many of the traditional proponents of Section 230 will argue that this idea is a non-starter. Two of the obvious objections are that it would place a disastrous burden on platforms especially start-ups and smaller platforms, and that it would stifle socially valuable anonymous speech. Both are valid concerns, but also accommodated by this proposal.

The concern that modest user-identification requirements would be disastrous to platforms made a great deal of sense in the early years of the Internet, both the law and technology around user identification were less developed. Today, there is a wide-range of low-cost, off-the-shelf, techniques to establish a user's identity to some level of precision – from logging of IP addresses, to requiring a valid email address to an established provider, registration with an established social media identity, or even SMS-authentication. None of these is perfect; they present a range of cost and sophistication to implement and a range of offer a range of ease of identification.

The proposal offered here is not that platforms be able to identify their speaker – it's better described as that they not deliberately act as a liability shield. It's requirement is that platforms implement reasonable identity technology in proportion to their size, sophistication, and the likelihood of harmful speech on their platforms. A small platform for exchanging bread recipes would be fine to maintain a log of usernames and IP addresses. A large, well-resourced, platform hosting commercial activity (such as Amazon Marketplace) may be expected to establish a verified identity for the merchants it hosts. A forum known for hosting hate speech would be expected to have better identification records – it is entirely foreseeable that its users would be subject to legal action. A forum of support groups for marginalized and disadvantaged communities would face a lower obligation than a forum of similar size and sophistication known for hosting legally-actionable speech.

This proportionality approach also addresses the anonymous speech concern. Anonymous

speech is often of great social and political value. But anonymity can also be used for, and as made amply clear in contemporary online discussion can bring out the worst of, speech that is socially and politically destructive. Tying Section 230's immunity to the nature of speech on a platform gives platforms an incentive to moderate speech – to make sure that anonymous speech is used for its good purposes while working to prevent its use for its lesser purposes. This is in line with one of the defining goals of Section 230.

The challenge, of course, has been how to do this without exposing platforms to potentially crippling liability if they fail to effectively moderate speech. This is why Section 230 took the approach that it did, allowing but not requiring moderation. This proposal's user-identification requirement shifts that balance from "allowing but not requiring" to "encouraging but not requiring." Platforms are under no legal obligation to moderate speech, but if they elect not to, they need to make reasonable efforts to ensure that their users engaging in problematic speech can be identified by parties harmed by their speech or conduct. In an era in which sites like 8chan expressly don't maintain user logs in order to shield themselves from known harmful speech, and Amazon Marketplace allows sellers into the market who cannot be sued by injured consumers, this is a common-sense change to the law.

It would also likely have substantially the same effect as other proposals for Section 230 reform, but without the significant challenges those suggestions face. For instance, Danielle Citron & Ben Wittes have proposed that courts should give substantive meaning to Section 230's "Good Samaritan" language in section (c)(2)'s subheading, or, in the alternative, that section (c)(1)'s immunity require that platforms "take[] reasonable steps to prevent unlawful uses of its services." This approach is problematic on both First Amendment and process grounds, because it requires courts to evaluate the substantive content and speech decisions that platforms engage in. It effectively tasks platforms with undertaking the task of the courts in developing a (potentially platform-specific) law of content moderations – and threatens them with a loss of Section 230 immunity is they fail effectively to do so.

By contrast, this proposal would allow, and even encourage, platforms to engage in such moderation, but offers them a gentler, more binary, and procedurally-focused safety valve to maintain their Section 230 immunity. If a user engages in harmful speech or conduct and the platform can assist plaintiffs and courts in bringing legal action against the user in the courts, then the "moderation" process occurs in the courts through ordinary civil litigation.

To be sure, there are still some uncomfortable and difficult substantive questions – has a platform implemented reasonable identification technologies, is the speech on the platform of the sort that would be viewed as requiring (or otherwise justifying protection of the speaker's) anonymity, and the like. But these are questions of a type that courts are accustomed to, if somewhat uncomfortable with, addressing. They are, for instance, the sort of issues that courts address in the context of civil unmasking subpoenas.

This distinction is demonstrated in the comparison between Sections 230 and 512. Section 512 is an exception to 230 for copyrighted materials that was put into place by the 1998

Digital Millennium Copyright Act. It takes copyrighted materials outside of the scope of Section 230 and requires platforms to put in place a "notice and takedown" regime in order to be immunized for hosting copyrighted content uploaded by users. This regime has proved controversial, among other reasons, because it effectively requires platforms to act as courts in deciding whether a given piece of content is subject to a valid copyright claim. The Citron/Wittes proposal effectively subjects platforms to a similar requirement in order to maintain Section 230 immunity; the identity-technology proposal, on the other hand, offers an intermediate requirement.

Indeed, the principal effect of this intermediate requirement is to maintain the pre-platform status quo. IRL, if one person says or does something harmful to another person, their recourse is in court. This is true in public and in private; it's true if the harmful speech occurs on the street, in a store, in a public building, or a private home. If Donny defames Peggy in Hank's house, Peggy sues Donny in court; she doesn't sue Hank, and she doesn't sue Donny in the court of Hank. To the extent that we think of platforms as the fora where people interact online – as the "place" of the Internet – this proposal is intended to ensure that those engaging in harmful speech or conduct online can be hauled into court by the aggrieved parties, and to facilitate the continued development of platforms without disrupting the functioning of this system of adjudication.

Conclusion

Section 230 is, and has long been, the most important and one of the most controversial laws of the Internet. It is increasingly under attack today from a disparate range of voices across the political and geographic spectrum — voices that would overwhelming reject Section 230's pro-innovation treatment of platforms and in its place attempt to co-opt those platforms as government-compelled (and, therefore, controlled) content moderators.

In light of these demands, academics and organizations that understand the importance of Section 230, but also recognize the increasing pressures to amend it, have recently released a statement of principles for legislators to consider as they think about changes to Section 230.

Into this fray, the Third Circuit's opinion in *Oberdorf* offers a potential change: making Section 230's immunity for platforms proportional to their ability to reasonably identify speakers that use the platform to engage in harmful speech or conduct. This would restore the status quo ante, under which intermediaries and agents cannot be used as litigation shields without themselves assuming responsibility for any harmful conduct. This shielding effect was not an intended goal of Section 230, and it has been the cause of Section 230's worst abuses. It was allowed at the time Section 230 was adopted because the used-identity requirements such as proposed here would not have been technologically reasonable at the time Section 230 was adopted. But technology has changed and, today, these requirements would impose only a moderate burden on platforms today.

View Article