NOISY SPEECH EXTERNALITIES

Justin "Gus" Hurwitz[*]

INTRODUCTION

A central tenet of contemporary First Amendment law is the metaphor of the

[*] Justin (Gus) Hurwitz is a Senior Fellow and Academic Director of the Center for Technology, Innovation, and Competition at the University of Pennsylvania. He is also the Director of Law & Economics Programs at the International Center for Law & Economics, and was previously a professor of law and founding director of the Governance & Technology Center at the University of Nebraska. His work draws on his background in law, economics, and computer science to study the relationship between technology and society.

This Essay will be republished as a book chapter in MEDIA AND SOCIETY AFTER TECHNOLOGICAL DISRUPTION (Gus Hurwitz & Kyle Langvardt eds., forthcoming Cambridge Univ. Press 2023).

marketplace of ideas—that the solution to bad speech is more, better, speech.[1] This basic idea is well-established in both judicial and scholarly writing—but it is not without its critics. My contribution to this volume adds a new criticism of the marketplace-of-ideas metaphor. I argue that there are circumstances where ostensibly "good" speech may be indistinguishable by listeners from bad speech—indeed, that there are cases in which *any* incremental speech can actually make *other* good speech indistinguishable from bad speech. In such cases, seemingly "good" speech has the effect of "bad" speech. I call this process by which ostensibly good speech turns the effects of other speech bad "a noisy speech externality."

This thesis has important implications. First, it offers a poignant critique of the marketplace-of-ideas aphorism introduced by Justice Holmes in his *Abrams* dissent.[2] If the marketplace of ideas is subject to significant market failure, correctives *may* be justified. Market failures, after all, are a standard justification for regulatory intervention. But, second, my contribution goes a step farther, suggesting not only that there are circumstances in which good speech may fail as a corrective to bad speech but also that there are circumstances in which the addition of seemingly good speech may only yield more bad speech. In such cases, the only solution to bad speech may be *less* speech—encouraging more speech may actually be detrimental to our speech values. If that is the case, then correctives may be not only justified but needed to satisfy an important societal interest. And, third, this chapter presents solutions for content-neutral ways in which to implement such correctives.

The insight underlying this thesis builds on my prior work applying the insights of Claude Shannon's information theory to social media.[3] That piece applied Shannon's work to social media to argue that, at least at a metaphorical level and potentially at a cognitive level, our capacity to communicate is governed by Shannon's

---

[1] Whitney v. California, 274 U.S. 357, 377 (1927) (Brandeis, J., concurring) ("If there be time to expose through discussion the falsehood and fallacies, to avert the evil by the processes of education, the remedy to be applied is more speech, not enforced silence."), *overruled in part by* Brandenburg v. Ohio, 395 U.S. 444 (1969).

[2] Abrams v. United States, 250 U.S. 616, 630 (1919) ("[T]he best test of truth is the power of the thought to get itself accepted in the competition of the market. . . . That at any rate is the theory of our Constitution.").

[3] Justin "Gus" Hurwitz, *Madison and Shannon on Social Media*, 3 BUS. ENTREPRENEURSHIP & TAX L. REV. 249 (2019).

channel-capacity theorem.[4] This theorem tells us that the capacity of a communications channel is limited by that channel's *signal-to-noise* ratio. Critically, once that capacity is exceeded, any additional signal is indistinguishable from noise—and this has the effect of worsening the signal-to-noise ratio, further reducing the communications capacity. In other words, after a certain threshold, additional speech isn't merely ineffective: It creates a negative externality that interferes with other speech.

Other scholars have made similar arguments, which can casually be framed as exploring the effects of "too much information" or "information overload."[5] But the negative-externality element of this argument goes a step further. A "too much information" argument suggests that listeners are overwhelmed by the quantity of speech to which they may be subject. This argument suggests that speakers can—deliberately or otherwise—exercise a veto over other speakers by saturating listeners' information sources. For the listener, it is not merely a question of filtering out the good information from the bad (the signal from the noise): At the point of saturation, signal cannot be differentiated from noise and any filtering necessarily must occur upstream from the listener.

Filtering—reducing the overall amount of speech—has always been a key tool in fighting bad speech. All platforms must filter. Indeed, this is nothing new: Editorial processes have always been valuable to listeners. The question is how they do it, with a related question of the law surrounding that filtering. Under the current approach (facilitated through Section 230 of the Communications Decency Act and built upon First Amendment principles), platforms have substantial discretion over what speech they host. This chapter's normative contribution is to argue that liability shield should be contingent upon platforms using "reasonable best-available technology" to filter speech—a standard that most platforms, this chapter also argues, likely already meet.

The discussion in this chapter proceeds in four parts. It begins in Part I by introducing technical concepts from the field of information theory—most notably the ideas of channel capacity and the role of the signal-to-noise ratio in defining a

---

[4] Cognitive psychologists and neurobiologists have identified some of these limits. Some of this research, such as that showing that there is a roughly constant information density across spoken human languages despite their vastly different syntax, grammars, and word complexity, is considered in Hurwitz, *supra* note 3.

[5] *See infra* Part II.B.

channel's capacity. Part II then introduces the traditional "marketplace of ideas" understanding of the First Amendment and builds on lessons from information theory to argue that this "marketplace" may in some cases be subject to negative externalities—noisy speech externalities—and that such externalities may justify some forms of corrective regulation. It also considers other arguments that have a similar feeling ("too much information," "information overload," and "listeners' rights"), and explains how the negative-externalities consequence of exceeding a channel's carrying capacity presents an even greater concern than is advanced by those ideas. Parts III and IV then explore the First Amendment and regulatory responses to these concerns, arguing that the negative-externalities concern might justify limited regulatory response. In particular, Part IV argues that platforms can reasonably be expected to implement "reasonable best available technologies" to address noisy speech externalities.

I. Information Theory, Channel Capacity, and the Signal-to-Noise Ratio

Initially developed by Claude Shannon at AT&T Bell Labs in the 1940s to study how, and how much, information could be transmitted over the communications channels making up the telephone network, information theory studies how we encode and transmit usable information over communications channels.[6] While mathematical and abstract in its characterization of information and communication, it is quite literally at the foundation of all modern communications networks.[7]

To understand the questions that information theory answers, we can start with a counterfactual. Imagine a perfect, noiseless, communications medium being used by two people to share meaningful information between them—say a professor wants to transmit a 90,000-word "article" to a journal editor. Assuming no limits on the part of the two communicating individuals, how quickly can one transmit that information to the other? We can ask the same question in a slightly different way: Because we assume the speakers do not impose any constraints (that is, we assume that each can speak or listen at any speed), we want to know how much information the communications medium can carry per unit of time—its "channel capacity."[8]

---

[6] *See* Hurwitz, *Madison and Shannon*, *supra* note 3, at 259.

[7] C. E. Shannon, *A Mathematical Theory of Communications*, 27 Bell Sys. Tech. J. 379 (1948).

[8] A communications medium is referred to as a "channel"—as in a communications channel.

To answer this to a first approximation, one could imagine that a professor could read the article aloud in 3 hours. So, the channel capacity is at least 30,000 words per hour. But we have assumed that the communicating individuals aren't the constraint. So, in principle, the professor could read faster, and the editor could transcribe faster—say, 90,000 words per hour, or even 180,000 words per hour, or even 180,000,000 words per hour. In the limit case, because we have assumed that the endpoints (the speaker and listener) don't impose any constraints *and* that the communications channel is a perfect, noiseless medium, the professor and the editor could communicate instantaneously.

This of course is not the case. But it illustrates two distinct limits we need to be aware of: the ability of the endpoints (speaker and listener) to encode and decode information at a given speed, and the ability of the communications channel to transmit, or carry, that information at a given speed.

Shannon studied both the encoding and carrying questions. We are focused on the carrying question, which for Shannon boiled down to two factors: the strength of the information-carrying signal and the amount of background noise. Taken together, these define the *signal-to-noise ratio* of the communications channel (mathematically, signal divided by noise, or signal/noise). Increasing this ratio increases the channel capacity. This means that you can increase the channel capacity either by increasing the signal strength *or* by decreasing the noisiness of the channel.

This should make intuitive sense to anyone who has ever had conversation in a noisy room. It is hard to have a conversation at a loud party—you generally need to speak more slowly and loudly to be heard clearly. You speak more slowly because the carrying capacity of the room (*qua* communications channel) is reduced by the noise; you speak more loudly to increase the strength of your communications signal.

The example of the noisy room also demonstrates three key takeaways from information theory. What is the source of the "noise" in the room? Mostly other people having their own conversations, or perhaps music is playing in the background. This is not meaningless "noise." Noise is not merely static or unintelligible sound (though static would be noise). Rather, noise is any signal that is not carrying

---

As defined by Shannon, a "channel is merely the medium used to transmit the signal from transmitter to receiver." *Id.* at 381.

meaningful information *for the recipient*. The recipient needs to expend mental energy trying to differentiate signal from noise (if there is enough signal available to reconstruct the intended communication), which slows down her ability to receive information.

Second, all noise is, therefore, reciprocal. Your conversation is noise to everyone else in the room! This means that when you speak more loudly so that your interlocutor can make out what you are saying, you are increasing the amount of noise that everyone else in the room must deal with—you are worsening their signal-to-noise ratio.

Something similar happens when you exceed the carrying capacity of a given communications channel, even if there is no one using that channel other than the speaker and listener. If you think about having a conversation on a phone line with a lot of noise in the background, there is a maximum speed at which you can talk and be understood. What happens when you speak faster than this? The sounds you make become unintelligible—they become noise. This is another lesson quantified by Shannon: When a channel's carrying capacity is exceeded, any additional information put onto that channel is interpreted not as signal but as noise.[9]

This last observation illustrates a third key takeaway: Signal and noise are interpreted by, and at, the receiver. There could be a thousand conversations going on in the room. Only those that reach the given individual's ears contribute to the signal and noise she must decipher. Similarly, in the online context, the signal-to-noise ratio is a function only of the message that an individual receives, not of the universe of messages that a platform carries. If a platform that carries a billion messages each day only delivers the relevant, meaningful ones to its users, it will have a very high signal-to-noise ratio; if another platform carries tens or hundreds of messages each day but delivers them all to a user regardless of their relevance, requiring her to sift through all of the messages in order to find those of relevance to her, that platform will have a relatively low signal-to-noise ratio.

---

[9] Shannon, *supra* note 7, at 410 ("If an attempt is made to transmit at a higher rate than C, say $C + R_1$, then there will necessarily be an equivocation equal to or greater than the excess R_1. Nature takes payment by requiring just that much uncertainty, so that we are not actually getting any more than C through correctly.").

II. Externalities and Speech Regulation

The discussion above tells us something additional about shared communications channels: At a certain point, information added to a communications channel creates a negative externality, reducing the capacity of that channel for everyone using it. But additional foundation is needed before we can look at why this carries an important lesson for how we think about speech regulation. The discussion below revisits the metaphor of the marketplace of ideas, looks at other scholars who have considered the challenges of limited communications capacity, and then introduces the idea that the negative externality created when a channel's carrying capacity is exceeded can justify regulation. It concludes by discussing some examples to illustrate these concerns.

A. Recapitulating the Marketplace of Ideas

Justice Holmes's dissent in *Abrams v. United States* introduced one of the most enduring metaphors of American law: the marketplace of ideas.[10] The concept of the marketplace of ideas is more intuitive than it is appealing: Just as better products (in terms of either price or quality) brought to market will sell better than inferior ones, so too will better ideas curry more favor with the public than lesser ones. And, in dynamic terms, just as overpriced or low-quality products will encourage new entrants into the market, lesser ideas will create an opportunity for better ideas to prevail.

This metaphor does important work toward vindicating the First Amendment's protection of individuals' speech against government interference—indeed, this is its true appeal, rather than the idea that speech will work as a marketplace. It promises that there is a mechanism to arbitrate between competing speech in the place of the government. Even where there may be some social need for speech to be moderated, state actors can take a step back and rely on this alternative mechanism to moderate in their stead—a need that might otherwise create demand for government intervention.

The marketplace-of-ideas metaphor has monopolized understandings of the First Amendment's protection of speech for the past century. While one can, and many do, debate its propriety and fidelity to the Amendment, I will posit that it was

---

[10] 250 U.S. 616, 630 (1919) (Holmes, J., dissenting) ("[T]he best test of truth is the power of the thought to get itself accepted in the competition of the market. . . . That at any rate is the theory of our Constitution.").

fit to task for most of this era. And the reason for this is that the listener-to-speaker ratio was relatively high. This was an era of rapidly changing technologies during which innovation ensured that new entrants and media were regularly entering the marketplace, but the high capital costs of those technologies inhibited entry, largely limiting it to those with the resources and ability to sincerely engage as a participant in this marketplace. A relative few broadcast platforms competed for market share based on the quality of their reporting, and number of local media outlets pruned these broadcasters' speech even further as a means to reach local communities. And throughout much of the 20th century, where media failed a community's needs, entry was both possible and often occurred.

B. Other Characterizations of Speech Regulation

As we entered the modern era of communications—with the widespread adoption of cable television and explosive growth of talk radio in the 1980s and the rapid digitalization of consumer-focused communications in the early 1990s[11]—increasing attention was paid to the idea of "too much information."[12] Indeed, television had famously been described as a "vast wasteland" as early as 1961.[13] By the early 1990s, the number of channels that cable systems could carry exceeded the number of channels of content being produced, satellite systems that could carry several times that many channels were being developed, and the Internet had come to the attention of sophisticated commentators.

Around the turn of the century, for instance, Cass Sunstein and Richard Posner both considered how the changing media landscape might affect our understanding of media regulation.[14] Sunstein, for instance, juxtaposed the marketplace-of-

---

[11] *Compare* Note, *The Awareness Doctrine*, 135 HARV. L. REV. 1907, 1911 (2022) (citing Steve Rendall, *The Fairness Doctrine: How We Lost It, and Why We Need It Back*, FAIR (Jan. 1, 2005), https://perma.cc/P557-C8FA), *with* Marvin Ammori, *The Fairness Doctrine: A Flawed Means to Attain a Noble Goal*, 60 ADMIN. L. REV. 881 (2008) (discussing issues with the Fairness Doctrine in practice).

[12] *See, e.g.*, CASS SUNSTEIN, TOO MUCH INFORMATION: UNDERSTANDING WHAT YOU DON'T WANT TO KNOW (2020).

[13] Newton N. Minow, *Television and the Public Interest*, 55 FED. COMMC'NS L.J. 395, 397 (2003).

[14] Cass R. Sunstein, *The First Amendment in Cyberspace*, 104 YALE L.J. 1757 (1995); Richard Posner, *Bad News*, N.Y. TIMES BOOK REV. (July 31, 2005) (reviewing and discussing eight recent books on the changing media landscape).

ideas approach to free speech with a Madisonian perspective, under which the purpose of the First Amendment is not merely to protect private speakers from government intrusion into their speech, but also affirmatively to promote and facilitate deliberative democracy.[15] Under the marketplace model, regulatory intervention had generally only been understood as appropriate in the face of scarcity—a lack of sufficient communications channels that prevented competition within the marketplace. As newer technologies increased the capacity of communications channels, and decreased the cost of deploying new ones, this rationale for regulating speech diminished. But, Sunstein argued, the Madisonian perspective suggested that regulation might nonetheless be appropriate if the new, emerging marketplace of ideas was not conducive to a functioning deliberative democracy.[16]

A decade later, Richard Posner considered many of the same issues that result from the decreasing costs of entering the market and sharing information in the information ecosystem.[17] He presented a different perspective than Sunstein, however, arguing that most consumers of information had always primarily wanted entertainment—not droll information—and that increased competition in the marketplace was catering to this interest.[18] In typical contrarian Posnerian fashion, he argued this was possibly not a bad thing: Just as increased competition in the marketplace catered to those citizens who were more interested in entertainment than in information, it would also better cater to those citizens who were more interested in information.[19]

---

[15] Sunstein, *supra* note 14, at 1759.

[16] *Id*. at 1804.

[17] Posner, *supra* note 14.

[18] *Id.* ("But increased competition has not produced a public more oriented toward public issues, more motivated and competent to engage in genuine self-government, because these are not the goods that most people are seeking from the news media. They are seeking entertainment, confirmation, reinforcement, emotional satisfaction; and what consumers want, a competitive market supplies, no more, no less.").

[19] *Id*. ("Yet what of the sliver of the public that does have a serious interest in policy issues? Are these people less well served than in the old days? Another recent survey by the Pew Research Center finds that serious magazines have held their own and that serious broadcast outlets, including that bane of the right, National Public Radio, are attracting ever larger audiences. And for that sliver of a sliver that invites challenges to its biases by reading The New York Times and The Wall Street Journal, that watches CNN and Fox, that reads Brent Bozell and Eric Alterman and everything in between, the increased polarization of the media provides a richer fare than ever before.").

More recently, there have been arguments about "too much information" and "information overload."[20] The general theme of these arguments is apparent on their face: Consumers of information face a glut of information that overwhelms their ability to process it all. A generation or two ago, there were relatively few sources of information. Consumers could reasonably assume that these sources had gone through some kind of vetting process and were therefore basically trustworthy. Indeed, should they so desire, an interested consumer could at least somewhat meaningfully undertake to investigate the quality of those competing information sources. The "too much information" argument says that neither of these is as possible today, if it is possible at all, as in it was prior generations—that the sheer quantity of information we encounter on a day-to-day basis undermines media sources' authority and interferes with listeners' purposes. And this resonates with the "marketplace of ideas" frame as well, for markets are driven, in part, by the consumer's ability to make informed choices—if that is not possible, the marketplace may not work.

There is another, more recent, argument that, again, challenges the marketplace orthodoxy: listeners' rights.[21] The listeners'-rights idea echoes the Madisonian (i.e., democracy-oriented) perspective on the First Amendment, though it may or may not be aligned with the marketplace concept. Under this view, the purpose of the First Amendment is not merely to ensure individuals' unfettered ability to speak without government interference, but also to ensure that individuals have access to (viz., the opportunity to listen to) information without undue government interference. Thus, if listeners want certain types of information but speakers interfere with their ability to obtain that information, the government may have some role in mediating that conflict and, when it does so, it should preference the listeners' choices about what information they want to receive over the speakers' efforts to influence the speakers.

C. An Externalities Argument for Speech Regulation

We can now return to the ideas introduced with information theory. The discussion in Part I concluded with the idea that any additional speech added to a saturated communications channel is interpreted as noise, not signal, by all parties to that communications channel. This has the effect of worsening the signal-to-noise

---

[20] SUNSTEIN, *supra* note 12.

[21] James Grimmelmann, *Listeners' Choices*, 90 COLO. L. REV. 65 (2019).

ratio, which reduces the overall channel capacity for everyone using that communications channel. In effect, this combines both the "too much information" construct and the listeners'-rights understanding of the First Amendment.

More important, it introduces a fundamentally different justification for (and, as will be discussed in Parts III and IV, a different approach to) speech regulation: externalities. Both the "too much information" and listeners'-rights perspectives present an information-asymmetry rationale for regulating the marketplace of ideas. Information asymmetries are a traditional justification for intervening to regulate a market: When one side of the market systematically has better information than the other, we might regulate to prevent harmful exploitation of that information.[22] For instance, we may require nutrition or energy-usage labels on products where consumers are not in a position to ascertain that information on their own. So too one could imagine requiring disclosures about the sources of information that speakers communicate to listeners, either as a way of helping consumers to meaningfully make use of the glut of information communicated to them or as a way of vindicating their rights to receive meaningful information as balanced against speakers' rights to share information.[23]

Externalities are another traditional justification for regulating markets.[24] Externalities occur where one party's private conduct has impacts on one or more third parties. Those impacts are "external" to the primary private conduct—as such, parties engaging in that conduct have little incentive to take them into account. Perhaps the most standard example of an externality is pollution: If I burn coal to generate electricity and no one has told me that I cannot put smoke into the air, I will not factor the environmental, health, or other costs of that pollution into my prices. The same can be said for many other types of activity (sometimes even

---

[22] *See* THOMAS LAMBERT, HOW TO REGULATE: A GUIDE FOR POLICYMAKERS 185–218 (2017). *See also* George A. Akerlof, *The Market for "Lemons": Quality Uncertainty and the Market Mechanism*, 84 Q.J. ECON. 488 (1970); Howard Beales, Richard Craswell & Steven C. Salop, *The Efficient Regulation of Consumer Information*, 24 J.L. & ECON. 491 (1981).

[23] This is not meant to advocate for either of those outcomes. One could also imagine a Posnerian market-based approach in which information providers compete to be trustworthy by making the information they provide easy for consumers to verify. It is entirely possible that the market could address concerns about either listeners being overwhelmed by information or vindicating their rights against speakers better than regulatory approaches.

[24] *See* LAMBERT, *supra* note 22, at 22–59. *See also* Garrett Hardin, *The Tragedy of the Commons*, 162 SCI. 1243 (1968); Ronald Coase, *The Problem of Social Cost*, 3 J.L. & ECON. 1 (1960).

individual conduct) and can be positive or negative. A neighborhood in which many people have dogs that they need to take on walks regularly (a private activity) may not be as welcoming to individuals who are scared of dogs (a negative externality) and also may have less crime (a positive externality).

Importantly, because the impacts of negative externalities are usually dispersed among many people and are difficult to measure except in aggregate, it may not be possible for injured parties to bring a lawsuit to recover for the injuries, either practically or as a matter of law. Lawmakers therefore might step in to address externalities, such as by prohibiting the underlying private conduct, requiring the parties to it to take care to prevent the externalities, or imposing taxes or fees on those parties that can be used to compensate any injured third parties to the case.

Additional speech—even ostensibly productive speech—added to a saturated communications channel has the characteristics of a negative externality. Because the carrying capacity of the channel is already saturated, the additional speech is interpreted by all who hear it as noise. This worsens the signal-to-noise ratio, further decreasing the carrying capacity of the channel. In a very real sense, noise is like air pollution: Just as pollution reduces the usability of air for all who breathe it, noise reduces the usability of a communications channel for all who communicate over it. Thinking back to the example of the loud room at a party, it is intuitive that if someone walks into a room in which several people are having conversations and turns up the volume on a stereo, this act will negatively impact the ability of all those in the room to continue their conversations.

D. Some Examples of Noisy Speech Externalities

Pointing to concrete examples of noisy speech externalities is challenging because the concept itself is somewhat abstract, and because the impacts may not be readily identifiable as discrete events.

The example of the noisy room presents a case study: At some point, a quiet room becomes too noisy to comfortably have a conversation; at a further point, it becomes impractical to have a conversation; at a further point, it becomes impossible to have a conversation. This transition charts the increasing harms that stem from noisy speech externalities. These harms most clearly need remedy when conversation becomes impossible. But in practice, the forum is likely to be abandoned by most participants before that point.

Useful forums have to solve the signal-to-noise problem somehow, then, and

they differentiate themselves by addressing the problem in different ways. Newspapers are as much a filter of information as a source of information; so are television and radio stations. Bookstores sort their books into sections and by topics; publishers select books for publication, perhaps filtering by subject, genre, or audience, and ensuring a quality threshold. Noisy forums are not usually sought after as platforms for information-sharing. One would not ordinarily negotiate a contract at a rock concert or debate politics with strangers in a busy subway station.

Social media presents the clearest setting for examples of noisy speech externalities—and mis- and disinformation are likely the clearest examples. The very promise of social-media platforms is that they allow users to communicate directly with one another—their defining feature is that they do *not* filter or select the information shared between users. And they also lack the traditional indicia of being too noisy to serve as a forum for information to be exchanged. A rock concert is a poor forum because all of the sound is heard at once—attendees attempting to have a conversation necessarily hear the loud music at the same time they are trying to hear their interlocutor—making it difficult to differentiate signal from noise. But in the social-media setting, content is presented in individual pieces, creating a perception for users that they have the ability to meaningfully engage with it.

We see the effects of noisy speech externalities, both intentionally and unintentionally created, with dis- and misinformation. For instance, there is some agreement that Russia weaponized disinformation around the 2016 U.S. election.[25] Commentators such as Bruce Schneier have argued that the purpose of Russian disinformation campaigns has been less to influence specific outcomes than to attack the American information ecosystem.[26] An adversary can win just as much by attacking our ability to separate fact from fiction as by convincing us to accept falsity as truth. But we do not need to turn to deliberate efforts to cause harm for examples. The inventor of Twitter's "retweet" button, for instance, has described the retweet feature as "hand[ing] a 4-year-old a loaded weapon."[27] Speaking of an example, he said: "Ask any of the people who were targets at that time, retweeting helped them

---

[25] *See* Henry Farrell & Bruce Schneier, *Democracy's Dilemma*, BOS. REV. (May 15, 2019), https://perma.cc/MG27-J4K4; YOCHAI BENKLER, ROBERT FARIS & HAL ROBERTS, NETWORK PROPAGANDA: MANIPULATION, DISINFORMATION, AND RADICALIZATION IN AMERICAN POLITICS (2018).

[26] *Id.*

[27] Alex Kantrowitz, *The Man Who Built the Retweet: "We Handed A Loaded Weapon to 4-Year-Olds"*, BUZZFEED (July 23, 2019), https://perma.cc/5LQ8-QFWC.

get a false picture of a person out there faster than they could respond. We didn't build a defense for that. We only built an offensive conduit."[28]

It is important to note here—and to draw attention to a theme to which I will return in Part IV—that the point of these examples is not to say that they demonstrate that Twitter (or any other particular platform) ought to be regulated. To the contrary, platforms like Twitter have always tried to develop new features to address concerns like these. Even the retweet feature was initially envisaged as a way to improve the signal-to-noise ratio (the theory being that making it easier to quickly share good information would help users get more information that they wanted). Platforms like Reddit have comprehensive user-based moderation technologies and norms.[29] Facebook and Twitter have invested substantially in addressing mis- and disinformation. The argument I make in Part IV is that it may be reasonable and permissible for Congress to require firms to engage in such efforts, and that efforts such as these should satisfy any regulatory obligations.

III. Preliminaries of Addressing Noisy Speech Externalities

Whether there is need and legally justifiable reason to support regulating speech in the digital era has prompted substantial debate in recent years. Arguments against such regulations most often sound, on both fronts, in the metaphor of the marketplace of ideas: The ostensible problem we face today is that the cost of speech is too low and the barriers to expression are too few, both of which would typically support the functioning of a robust marketplace.[30] The discussion below argues that concern about negative externalities derived from an information-theory-based understanding of channel capacity is sufficient to justify and overcome fundamental legal obstacles (that is, First Amendment concerns) to regulation. It then considers what such regulation could look like, drawing from other settings where the law addresses externalities.

A. Speech Regulation and the First Amendment

The first question is simply, as a matter of law, whether noisy speech externalities provide a cognizable legal basis for speech regulation—that is, whether such regulation could survive First Amendment scrutiny. As above, debates about such

---

[28] *Id.*

[29] *See Reddit Mod Education*, Reddit, https://perma.cc/R8QX-V5WV.

[30] *See* Posner, *supra* note 14; Eugene Volokh, *Cheap Speech and What It Will Do*, 104 Yale L.J. 1805 (1995).

speech regulation are often framed in terms of Madisonian vs Holmesian principles, whether the purpose of the First Amendment is to support robust democratic engagement or to foster a robust marketplace of ideas. But it may be useful to anchor the discussion in more doctrinal terms: Would regulation intended to address noisy speech externalities survive First Amendment scrutiny under existing doctrine?

There are two distinct lines of cases most relevant to this question: broadcast-regulation cases (e.g., *Red Lion*, *Pacifica*, *Turner*)[31] and the noise-regulation cases (such as *Ward v. Rock Against Racism*).[32] The broadcast cases are the foundation for media regulation in the United States—and, as foundations go, they are notably weak. They derive from midcentury understandings of spectrum and technological ability to use spectrum as a broadcast medium. The central concept of these cases is *scarcity*.[33] Because spectrum was a scarce resource that only supported relatively few television or radio broadcast stations in any geographic area, there was sufficient justification for the government to regulate who had access to that spectrum in order to ensure "the right of the public to receive suitable access to social, political, esthetic, moral and other ideas and experiences."[34] These cases expressly discuss this right in terms of the marketplace of ideas—importantly, scarcity is one of the most traditional justifications for regulation to intervene in the operation of markets—though it is arguably the case that this "right" is compatible with the Madisonian view of the First Amendment.

Ward v. Rock Against Racism is best known for its treatment of content-neutral "time, place, and manner" restrictions on speech and its clarification that such regulations need only be narrowly tailored to address a legitimate government purpose, not the least restrictive means of doing so.[35] Curiously, *Ward v. Rock* itself dealt with noise regulations that limited the volume of music at a concert in New

---

[31] Red Lion Broad. Co. v. Fed. Commc'ns Comm'n, 395 U.S. 367 (1969); Fed. Commc'ns Comm'n v. Pacifica Found., 438 U.S. 726 (1978); Turner Broad. Sys., Inc. v. Fed. Commc'ns Comm'n, 520 U.S. 180 (1997).

[32] Ward v. Rock Against Racism, 491 U.S. 781 (1989).

[33] Nat'l Broad. Co. v. United States, 319 U.S. 190 (1943).

[34] *Red Lion*, 395 U.S. at 390.

[35] *Ward*, 491 U.S. at 791.

York's Central Park—though the Court's concern about "noise" in that case is understood in the colloquial sense of "disturbing or distracting sound" rather than in information theory's more technical sense. But even without resorting to information theory and the concern that noise reduces a channel's information carrying capacity, this type of noise is a classic example of a negative externality.[36]

Where a communications channel is at its carrying capacity, we may see both scarcity and negative externalities coming into play. Scarcity does not necessarily implicate externalities, because the lack of options may not affect third parties. And externalities alone do not necessarily meaningfully implicate scarcity because they can adversely affect every option that a consumer may reasonably have. But if a communications channel reaches the point of saturation, that suggests that its users do not have a robust set of alternative channels available to them (as otherwise they would switch to a less congested channel) and so face scarcity, and that any additional information added to that channel worsens the signal-to-noise ratio for all users, creating a negative externality.

It is therefore likely, or at least plausible, that the government would have a substantial interest in narrowly tailored regulations intended to lessen these impacts on a content-neutral basis.

B. Technical Responses to a Poor Signal-to-Noise Ratio

Mathematically, the way to address a poor signal-to-noise ratio is to increase the strength of the signal or decrease the amount of noise in the signal. This guides technical approaches to addressing a poor signal-to-noise ratio. In settings where the communications channel's capacity is not exceeded, the signal strength can be directly increased (akin to speaking more loudly). Alternatively, filters can be added to reduce noise, either at the transmitter or receiver. Importantly, filters can both reject true "noise" (e.g., background static), or they can reject "unwanted signal." For instance, a radio receiver might use a filter to reject signal from adjacent radio stations (e.g., a radio tuned to FM station 101.3 might filter out signal from stations 101.1 and 101.5).

---

[36] It bears emphasis that the negative externality of noise alone is a less substantial than that which occurs when incremental speech is added to a saturated communications channel. In the latter case, the incremental speech is a negative impact on its own, but it has the more substantial effect of making otherwise-meaningful speech also on the in the communications channel indistinguishable from noise.

In more complex settings, such as cellular telephone networks, solutions are more sophisticated. In cellular networks, capacity can be increased by adding more "cells" (antennas spread across the network) and decreasing the power at which the antennas within them transmit. Adding more cells brings antennas closer to cell phones—the reduced distance decreases the minimum signal strength that is needed for communications. This, in turn, reduces the extent to which one's phone interferes with another's. And there is always the least-sophisticated solution to a poor signal-to-noise ratio: decreasing the speed of communications.[37]

C. Legal Responses to Externalities

On the legal side of the ledger, there are both public and private legal institutions that address externalities. On the public law side of the ledger, environmental regulation to reduce pollution presents the clearest analogy. Here, there are a few standard tools in the regulatory toolbox. Environmental regulations, for instance, might implement direct command-and-control style regulations, prohibiting certain types of conduct (such as the emission of certain pollutants beyond a threshold, and an absolute prohibition on the use or emission of certain pollutants or chemicals).[38] In other cases, the EPA uses "best-available control technology" (BACT) or similar requirements, under which the agency will undertake a regulatory process to ascertain the state-of-the-art in pollution-control technologies and require sources of pollutants to use those technologies.[39]

There are also private law institutions that address externalities—though they are relatively rare to see implemented at scale. In common law, these are most often seen in the context of new or changing technologies that lead to new conflicts to

---

[37] While less sophisticated, this approach is nonetheless important. It is probably one of the two intuitive responses to overcoming a poor signal-to-noise ratio (the other being to increase signal strength—that is, to speak more loudly or clearly). It also finds analogy in discussions about online speech regulation, for instance with suggestions that platforms insert "friction" into speech.

[38] *See, e.g.*, Bruce A. Ackerman & Richard B. Stewart, *Reforming Environmental Law*, 37 STAN. L. REV. 1333 (1985); Rena Steinzor, *Reinventing Environmental Regulation: The Dangerous Journey from Command to Self-Control*, 22 HARV. ENV'T L. REV. 103 (1998) ("Command and control rules impose detailed, legally enforceable limits, conditions, and affirmative requirements on industrial operations, generally controlling sources that generate pollution on an individual basis.").

[39] Steinzor, *supra* note 38, at 114; Ackerman & Stewart, *supra* note 38, at 1335 (discussing best-available technology regulations); 42 U.S.C § 7475(a)(4) ("[T]he proposed facility is subject to the best available control technology for each pollutant subject to regulation under this chapter").

arise between individuals. For instance, when once-distant residential communities and industrial farming operations expand to the point that they are near-neighbors,[40] noise-generating machinery is installed in areas where it was not previously used,[41] cement plants expand to serve the needs of growing communities,[42] new construction obstructs longstanding enjoyment of the sun,[43] or new technologies such as mills alter the character and landscape of a community,[44] judges may be called in to adjudicate the uncertain rights that exist between individuals engaging in these activities and the third parties affected by the externalities resulting from them. Most often, legal claims arising from these changing uses are styled as regarding public or private nuisance—although where they implicate rights that have been clearly established under the common law or statute, they may be treated as trespass or statutory violations.

IV. How We Should Regulate Noisy Speech Externalities

This brings us to this chapter's ultimate question: How should we regulate online speech in response to noisy speech externalities? My answer is that we should adopt a model similar to that used by the EPA for pollution control—a best-available control technology—but rely on customary industry practices to determine whether such a standard is being met. Unlike the EPA model, the baseline requirement in this setting ought to be a "reasonable best-available technology" requirement, recognizing the vast variation between the capabilities of various platforms and needs of their users. One option for implementing this requirement is to make Section 230's liability shield contingent upon the use of such technologies.[45]

Part III.C introduced standard legal approaches to addressing externalities, including public law approaches such as environmental regulation and private law approaches such as nuisance and trespass claims. So far, neither of these approaches

---

[40] Spur Industries v. Del E. Webb Development Co., 108 Ariz. 178 (1972).

[41] Sturges v. Bridgman, 11 Ch. D. 852 (1879).

[42] Boomer v. Atlantic Cement Co., 257 N.E.2d 870 (N.Y. 1970).

[43] Fontainebleau H. Corp. v. Forty-Five Twenty-Five, Inc., 114 So. 2d 357 (Fla. Ct. App. 1960); Prah v. Maretti, 108 Wis. 2d 223 (1982).

[44] Rylands v. Fletcher, L.R. 3 H.L. 330 (1868).

[45] The most trenchant response to proposals such as this is that weakening Section 230's liability shield creates a legitimately worrisome possibility of harming platforms—especially smaller platforms that cannot easily absorb the cost of litigation—by exposing them to potential claims sufficient to survive a motion to dismiss. This concern is addressed near the end of this section.

have been put into practice in the online environment. Rather, Section 230 of the Communications Decency Act creates a permissive self-regulatory environment.[46] Section 230 frees up platforms to moderate users' speech while making clear that they are under no obligation to do so.[47] Under this approach, platforms are shielded from liability for any harms caused by speech generated by their users, including to the speech interests of their other users.

Part III.A, however, suggested that the government may have a sufficient interest in regulating this speech with narrowly tailored regulations intended to lessen the impacts of noisy speech externalities. We might look to aspects of both public and private law approaches to regulating externalities for a model of how the government could respond to noisy speech externalities. In the final analysis, this approach would combine the self-regulatory approach embraced by Section 230 with an affirmative requirement that platforms use best-available content-moderation technologies as suitable for their scale.[48]

Pollution and pollution control are the most traditional legal analogies for thinking about noisy speech externalities and their regulation—and the control-technologies analogy maps onto the concept of content-moderation technologies. A requirement that a platform uses "best-available content-moderation technology" to ensure as best as possible that its users have meaningful access to content on the platform is a content-neutral policy. More-prescriptive policies would run the risk of making content-based distinctions, especially if they were based in concerns that some types of speech on platforms were more or less in need of protection.[49] And while policies such as a common-carrier obligation would likely qualify

---

[46] 47 U.S.C. § 230.

[47] § 230(c)(2) provides that platforms shall not be held liable on account of their moderation activities and § 230(c)(1) provides that they shall not be held liable for information shared by their users.

[48] That last proviso, "as suitable for their scale," is likely redundant. A small platform likely has little need for significant content moderation practices because users are likely able to replicate the experience of that platform on any number of other platforms. As the scale of the platform grows, the value of content moderation grows, and the costs to the platform of poor practices potentially decreases as users face fewer competitive alternatives. In addition to potential competition considerations, it is potentially the case that as the signal-to-noise ratio decreases, so too will users' ability to evaluate the quality of the platform. This creates an endogenous challenge that might on its own justify some amount of regulatory intervention.

[49] *See, e.g.*, Turner Broad. Sys., Inc. v. Fed. Commc'ns Comm'n, 520 U.S. 180, 234 (O'Connor,

as facially neutral, their effect on the signal-to-noise ratio of platforms would be to render the platforms useless for the vast majority of communications.[50]

Content-moderation techniques and technologies are akin to technological filters or amplifiers that reduce noise or increase the strength of desirable signal to improve the signal-to-noise ratio. These techniques or technologies may come in many forms (indeed, they need not be technological or algorithmic, but could result, for instance, from cultivating community norms or market mechanisms). Their defining characteristic is that they improve a platform's signal-to-noise ratio, making it easier for users to engage with desired information (signal) or less likely that they will encounter undesired information (noise). As with filters and amplifiers, these technologies can be misconfigured—content moderation can have the effect of amplifying harmful speech or filtering desirable speech. But this is a specific question of how a technology is implemented (including whether it is a "best-available" technology), not a question of the viability or desirability of the underlying technology.

The harder question is who decides what content-moderation technologies are reasonably considered "best available" and how regulation based upon those technologies may be implemented. In the environmental-regulation context, this is done through a regulatory process in which the regulator gathers information about industry practices and dictates what technologies to use. This is not a desirable approach in the speech-moderation setting. As an initial matter, different content-moderation technologies may have different effects on different speech or speakers. Unlike in the pollution context, this potentially creates substantial issues, including embedding content-based distinctions into regulations. That could bring us into the domain of strict scrutiny and concerns about government interference in private speech—a central concern against which the First Amendment is meant to protect.

J., dissenting) ("But appellees' characterization of must-carry as a means of protecting [local broadcast] stations, like the Court's explicit concern for promoting "'community self-expression'" and the "'local origination of broadcast programming,'" reveals a content-based preference for broadcast programming.").

[50] Indeed, one could argue that a common-carrier obligation would amplify certain types of speech over others in a social-media environment. If we need carriage guaranteed to ensure that some types of speech are viable on the platform, that suggests that the regulations are not, in fact, content-neutral. *Id.*

An additional challenge is the range of speech and the range of platforms hosting that speech. This is a more dynamic environment than the environmental-pollution setting. The EPA regulates a small number of pollutants produced by a small number of chemical processes, which can only be addressed by a small number of control technologies. This makes assessing the best available among those control technologies a tractable task. Courts and regulators are unlikely to be able to keep up to speed with changing needs and capabilities of content moderation—indeed, they are likely to lack the sophistication needed to understand how the technologies even work. And different technologies may be better- or worse-suited to different types of speech or different types of communication platforms.

It is not unusual for courts to look to industry custom in the face of changing technologies or scientifically complex settings.[51] This is a setting where deference to customary industry practices, as opposed to prescriptive command-and-control regulation, makes good sense. And it is an important margin along which online speech platforms compete today. Indeed, platforms invest substantially in their content-moderation operations and are continually innovating new techniques to filter undesired speech, amplify desired speech, and generally to give users greater control of the information that they receive. To be sure, not all of these technologies succeed, and platforms often need to balance the effectiveness of these technologies with the business needs of the platform. To take one example of the former effect, the initial theory behind adding the ability to "like" and "retweet" content on Twitter was to amplify desirable content—but its greater effect was to substantially worsen the platform's signal-to-noise ratio by increasing the velocity of lower-value content and superficial engagement with that content.[52] On the other hand, tools like verified accounts and the ability for high-reputation users to help moderate posts through systems like Twitter's Birdwatch or Reddit's moderator system provide useful amplification and filters that help manage the platforms' signal-to-noise ratios.

Of course, any regulation of online platforms' speech practices needs to confront Section 230's liability shield. Today, Section 230 permits but does not require

---

[51] The classic case demonstrating the role of industry custom in judicial decision-making is *The T.J. Hooper*, 60 F.2d 737 (2d Cir. 1932). Medical-malpractice cases are an instance where courts rely on customary practices as a tool for understanding complex scientific settings. *See, e.g.*, Richard N. Pearson, *The Role of Custom in Medical Malpractice Cases*, 51 IND. L.J. 528 (1976).

[52] Kantrowitz, *supra* note 27.

platforms to adopt content-moderation policies. As suggested above, if a platform is exceeding its channel-carrying capacity without losing users, this suggests there is a market failure keeping those users beholden to the platform. At that point, justification for Section 230's permissive no-moderation provisions are at their nadir and it becomes reasonable to expect the platform to adopt improved moderation practices.

This is no small recommendation, and I do not make it lightly. Weakening the protections of Section 230's liability shield significantly increases the cost of litigation, especially for smaller platforms. In its current form, it is difficult for a plaintiff in a Section 230 case to survive a motion to dismiss. Making that shield contingent upon a nebulous "reasonable best-available technology" requirement invites suits that would survive a motion to dismiss. Critically, this raises significant concerns that litigation could deprive platforms' users of their chosen venue for the exercise of their constitutionally protected speech rights. Any alteration to Section 230's liability shield should thus be accompanied by specific requirements to counterbalance these concerns, such as sanctions for sham or strategic litigation, fee-shifting requirements, specific pleading or discovery requirements, and safe harbors for smaller platforms. These should be accompanied by a presumption that any industry-standard content-moderation techniques satisfy a "reasonable best-available technology" benchmark. Policymakers might also consider a two-pronged requirement that a) concerns about a platform's content-moderation practices must be reported to a state attorney general prior to the commencement of a suit and b) that a private suit can only be brought after that attorney general non-prejudicially declines to conduct their own investigation.

Conclusion

A central tenet of contemporary First Amendment law is the metaphor of the marketplace of ideas—that the solution to bad speech is more, better speech. But this is built upon an assumption that more, better speech is possible. Information theory tells us that there are circumstances where any additional speech is necessarily bad speech. This is analytically equivalent to an externality, a common form of market failure and traditional regulatory intervention; and it is analogically equivalent to a market failure in the marketplace of ideas. Indeed, examples of regulation in the face of such failures is common in cases such as pollution and nuisance law—as well as in the First Amendment setting.

This chapter has argued that regulation may be justified, and may survive First Amendment challenges, in cases where noisy speech externalities are likely to occur. It also argues that such regulation should draw from the examples of pollution control's use of best-available control technologies to mitigate these externalities—but that unlike the pollution setting, courts should look to customary industry practices to evaluate whether a particular platform is using such technologies. This chapter suggests that Section 230's liability shield, which currently allows but does not require platforms to implement any content-moderation technologies, should require platforms to adopt such technologies. However, this is likely a less radical suggestion than it may seem, as most platforms already actively use and develop content-moderation technologies in the standard course of business; such efforts should be sufficient to satisfy a "best-available content-moderation technology" requirement. Rather, only platforms that actively eschew content-moderation practices, or that otherwise neglect these technologies, would risk the loss of Section 230's liability shield.