

---

## Constructing AI Speech

*Margot E. Kaminski & Meg Leta Jones*

**ABSTRACT.** Artificial Intelligence (AI) systems such as ChatGPT can now produce convincingly human speech, at scale. It is tempting to ask whether such AI-generated content “disrupts” the law. That, we claim, is the wrong question. It characterizes the law as inherently reactive, rather than proactive, and fails to reveal how what may look like “disruption” in one area of the law is business as usual in another. We challenge the prevailing notion that technology inherently disrupts law, proposing instead that law and technology co-construct each other in a dynamic interplay reflective of societal priorities and political power. This Essay instead deploys and expounds upon the method of “legal construction of technology.” By removing the blinders of technological determinism and instead performing legal construction of technology, legal scholars and policymakers can more effectively ensure that the integration of AI systems into society aligns with key values and legal principles.

Legal construction of technology, as we perform it, consists of examining the ways in which the law’s objects, values, and institutions constitute legal sensemaking of new uses of technology. For example, the First Amendment governs “speech” and “speakers” toward a number of theoretical goals, largely through the court system. This leads to a particular set of puzzles, such as the fact that AI systems are not human speakers with human intent. But other areas of the law construct AI systems very differently. Content-moderation law regulates communications platforms and networks toward the goals of balancing harms against free speech and innovation; risk regulation, increasingly being deployed to regulate AI systems, regulates risky complex systems toward the ends of mitigating both physical and dignitary harms; and consumer-protection law regulates businesses and consumers toward the goals of maintaining fair and efficient markets. In none of these other legal constructions of AI is AI’s lack of human intent a problem.

By going through each example in turn, this Essay aims to demonstrate the benefits of looking at AI-generated content through the lens of legal construction of technology, instead of asking whether the technology disrupts the law. We aim, too, to convince policymakers and scholars of the benefits of the method: it is descriptively accurate, yields concrete policy revelations, and can in practice be deeply empowering for policymakers and scholars alike. AI systems do not in some abstract sense disrupt the law. Under a values-driven rather than technology-driven approach to technology policy, the law can do far more than just react.

## INTRODUCTION

Artificial Intelligence (AI) systems such as ChatGPT can now produce convincingly human speech, at scale.<sup>1</sup> That speech can take many forms, typically in response to user prompts and queries. AI systems can produce political manifestos, poetry, and fascinating facts. If it were produced by a human, such “high-value speech” would sit at the heart of the First Amendment’s protections.<sup>2</sup> AI systems can also produce, unprompted or with very little nudging, wildly untrue hallucinations, recipes for explosives, and harmful lies about individual people. These are examples of “low-value speech,” some of which are unprotected by the Constitution and thus readily regulated, at least in theory.<sup>3</sup>

It is tempting to ask whether AI speech “disrupts” the law. That is, is there something special about AI-generated speech? Will the law be able to keep pace with it? Will the harms caused by AI speech at scale drive the creation of new and *sui generis* areas of law and regulatory systems? Or will regulation of AI speech be an example of Frank H. Easterbrook’s infamous “law of the horse”: old law applied, with minimal alteration, to newish things and behavior?<sup>4</sup>

These are the questions frequently asked in law-and-technology scholarship. And these are the wrong questions to ask. They lead to an approach that examines how the law responds to technology, which then leads to several major mistakes, including underexamination of the role existing law plays and underscrutinization of the often surprising conditions under which technology has developed. Instead of asking whether the law can keep up with technology, we should be asking how the law shapes and intervenes in technology and how it can do so – or refrain from doing so – in ways that further important values. A

- 
1. For a comprehensive description of how such systems work, see, for example, Nina Brown, *Bots Behaving Badly: A Products Liability Approach to Chatbot-Generated Defamation*, 3 J. FREE SPEECH L. 389, 392-97 (2023).
  2. See Genevieve Lakier, *The Invention of Low-Value Speech*, 128 HARV. L. REV. 2166 (explaining how courts have constructed categories of speech that carry different degrees and forms of constitutional protection). Low-value speech is speech that threatens serious violence, disorder, and “dominant norms of civility, decency, and piety[,]” like obscenity and libel, and has, since the New Deal Court, been given only limited First Amendment protection. *Id.* at 2168. High-value speech is speech that is related to matters of public interest; further, speech outside of historical low-value categories, including vulgarity, hate speech, and lying, is also protected as high-value speech. *See id.* at 2229.
  3. *Id.* at 2168. *See also* United States v. Alvarez, 567 U.S. 709, 717-18 (2012) (listing categories of speech not historically subject to First Amendment protection). For a full discussion of this issue, *see infra* Section II.B.
  4. Frank H. Easterbrook, *Cyberspace and the Law of the Horse*, 1996 U. CHI. LEGAL F. 207, 207, 215; Lawrence Lessig, *The Law of the Horse: What Cyberlaw Might Teach*, 113 HARV. L. REV. 501 (1999).

values-first approach to technology law would start not by asking whether the law is capable of keeping up with the development of technological or scientific expertise (it almost always is), but by asking what values technology, and the social practices of technology, further. It would ask how the law and its systems of sensemaking can and should be shaped to further values.

This Essay therefore examines what we have elsewhere termed the legal construction of technology, and more specifically the legal construction of content-generating AI systems.<sup>5</sup> Legal construction of technology looks to the ways in which the law itself contributes to its encounter with, or really its incorporation of, the recently evolving uses of AI systems as “speech engines.”<sup>6</sup> By shifting the

- 
5. Meg first used the term in Meg Leta Jones, *Does Technology Drive Law? The Dilemma of Technological Exceptionalism in Cyberlaw*, 2018 U. ILL. J.L. TECH. & POL’Y 249, 253 (“I argue that technology does not drive law either. Technology is not the locus of legal agency. When testing the theory of technological exceptionalism, no technology has even been exceptional. We must figure out a new way to answer the question, ‘are driverless cars new?’ Because, [*sic*] technological exceptionalism is not up to the task. Instead of analyzing whether technologies are or will be exceptional and in addition to analyzing how the law can and should respond to exceptional or conservative technological advances, this Article argues that cyberlaw research should consider the way in which technologies, practices, and social arrangements are constructed within certain legal contexts: the legal construction of technology.”). Margot discussed the concept in Margot E. Kaminski, *Authorship, Disrupted: AI Authors in Copyright and First Amendment Law*, 51 U.C. DAVIS L. REV. 589, 590-91 (2017) [hereinafter Kaminski, *Authorship, Disrupted*] (“To the extent new technology (or really, the social practice of a new technology) disrupts the law, it does so because of how it encounters existing features of the law, both doctrinal and theoretical. The law, in constructing—that is, building the meaning of—new technological developments and their social uses, takes a central part in its own disruption. Conceiving of technology as some outside force that acts upon the law can lead to a technology-centric approach in which one tries to identify what features of a particular technology are legally disruptive. This kind of disruption narrative gets it wrong. A particular feature of a particular technology disrupts the law only because the law has been structured—doctrinally and theoretically—in a way that makes that feature relevant. The disruptive effects (if any) of a technology become manifest when they encounter, interface with, and are given particular meaning within the law.”). See also Margot E. Kaminski, *Technological “Disruption” of the Law’s Imagined Scene: Some Lessons from Lex Informatica*, 36 BERKELEY TECH. L.J. 883, 893-94 (2021) [hereinafter Kaminski, *Technological Disruption*] (arguing that “the law dynamically constructs technology into its own systems of meaning . . .”).
  6. The term “speech engines” was coined by James Grimmelman in his discussion of search engines, in reasoning that has a lot of resonance for discussions of speech-generating Artificial Intelligence (AI) systems today. James Grimmelman, *Speech Engines*, 98 MINN. L. REV. 868, 873-74 (2014) (explaining that search engines are neither purely neutral conduits nor expressive editors, and are instead best analogized to advisers); see also Jones, *supra* note 5, at 253 (explaining that cyberlaw research should consider how technologies “are constructed within certain legal contexts”); Kaminski, *Authorship, Disrupted*, *supra* note 5, at 591 (examining how AI “encounters and is incorporated into” copyright and First Amendment law); Kaminski, *Technological Disruption*, *supra* note 5, at 886 (“[E]ach law or policy conversation takes place around an understood imagined setting, with technology, or the lack thereof, often playing a central role.”).

main object of inquiry to the law and its features, the method reveals tropes and policy baggage, along with possible new points of intervention. To follow the well-trod path of tech-exceptionalist reasoning and ask, “How does AI speech disrupt the law?” is to blind oneself to more effective ways of discussing how the law encounters, and indeed actively constructs, AI speech.

We proceed as follows: In Part I, we define the method of legal construction of technology. The method examines how the law makes sense of technology, querying the objects, values, and institutions involved in such sensemaking. This is not a purely descriptive move; legal construction makes room for awareness of the pathologies of a particular regulatory path, and the possibility of other interventions.

We then apply the method in Part II. We begin by defining “AI” – not only in the sense of abstract technical specifications, but by placing the development of these systems in social, historical, and legal contexts. We then explore four ways of legally constructing AI speech: (1) as *speech by speakers* (within First Amendment law), (2) as *speech at scale* (within content moderation law), (3) as the output of *risky complex systems* (within risk regulation or tort law), and (4) as market behavior implicating *consumer protection* (within the FTC’s Section 5 authority). For each analysis, we query the values, objects, and institutions of regulation.

Together, these analyses show that no one feature of speech-generating AI systems is driving “disruption” of the legal system. The law, in encountering AI systems, already demonstrates a substantial array of tools, institutions, and orienting values. From analogies to impact assessments, the law is already deeply involved in making legal sense of AI.

Part III outlines the policy implications of this method. We each have clear normative leanings, spelled out in other works, as to which of these and other tools work best for governing AI systems.<sup>7</sup> But our goal here is less to argue for particular tools than to show the things legal construction of technology can empower us to do. A values-first approach to regulating technology would leave

---

7. See, e.g., Meg Leta Jones, *The Ironies of Automation Law: Tying Policy Knots with Fair Automation Practices Principles*, 18 VAND. J. ENT. & TECH. L. 77 (2015); Meg Leta Jones, *The Right to a Human in the Loop: Political Constructions of Computer Automation & Personhood*, 47 SOC. STUD. SCI. 216 (2017) [hereinafter Jones, *Human in the Loop*]; Margot E. Kaminski, *Regulating the Risks of AI*, 103 B.U. L. REV. 1347 (2023) [hereinafter Kaminski, *Regulating the Risks*]; Rebecca Crootof, Margot E. Kaminski & W. Nicholson Price II, *Humans in the Loop*, 76 VAND. L. REV. 429 (2023); Margot E. Kaminski & Jennifer M. Urban, *The Right to Contest AI*, 121 COLUM. L. REV. 1957 (2021); Margot E. Kaminski, *Binary Governance: Lessons from the GDPR’s Approach to Algorithmic Accountability*, 92 S. CAL. L. REV. 1529 (2019) [hereinafter Kaminski, *Binary Governance*]; Margot E. Kaminski & Gianclaudio Malgieri, *Algorithmic Impact Assessments Under the GDPR: Producing Multi-Layered Explanations*, 11 INT’L DATA PRIV. L. 125 (2021); Margot E. Kaminski, *The Right to Explanation, Explained*, 34 BERKELEY TECH. L.J. 189 (2019).

each lawyer, policymaker, or scholar to determine—and to own—what values they want to advance and what harms they want to prevent, and to then choose their legal tools or regulatory abstentions accordingly.

We summarize and conclude with some parting thoughts. Legal constructions of technology can intertwine, bend, evolve, and mutate. What they resolutely do not do, in the face of purportedly new technologies like AI, is passively get acted upon and break. To see law, its institutions, and individual actors in this way is to afford new technologies some special force they do not have. Law and technology are each aspects of, and deeply embedded in, society. As such, they co-construct each other, in the context of deeply political systems that channel or constrain political power.<sup>8</sup>

### I. LEGAL CONSTRUCTION OF TECHNOLOGY

We start with an introduction to our approach to answering the question, “Does AI speech disrupt the law?” We approach the question by rearranging its premise. Technology does not disrupt the law. The law does not follow technology around trying to “keep up”—that is not the nature of the relationship, or so argues what we have called the legal construction of technology. The legal construction of technology method takes the view that the law conducts its own kind of sensemaking of new technological developments and their social uses.<sup>9</sup> The law makes sense of technology in very particular ways, using language, enforcement, abstention, institutional arrangements, and means for garnering nonlegal expertise, that change as legal cultures change. Law makes sense of, and indeed shapes, technology and its trajectories through legal tools and institutions.

---

8. Julie Cohen writes, “[T]echnology is not a monolithic, irresistible force. . . . Information technologies are highly configurable, and their configurability offers multiple points of entry for interested and well-resourced parties to shape their development. . . . Legal institutions too offer multiple points of entry for economic and political power, and as they are enlisted to help produce the profound economic and sociotechnical transformations that we see all around us, they too are being changed.” JULIE E. COHEN, *BETWEEN TRUTH AND POWER: THE LEGAL CONSTRUCTIONS OF INFORMATIONAL CAPITALISM* 1-2 (2019). Cohen argues that legal actors have “asked *how law should respond* to the changes occurring all around it . . . [but have] not asked the broader, reflexive questions about how core legal institutions *are already evolving* in response to the ongoing transformation in our political economy.” *Id.* at 2. Similarly, in science-and-technology studies (STS), Sheila Jasanoff, who started her career as a lawyer, framed biotechnology policy as one of “coproduction.” See Sheila Jasanoff, *Ordering Knowledge, Ordering Society*, in *STATES OF KNOWLEDGE: THE COPRODUCTION OF SCIENCE AND THE SOCIAL ORDER* 13, 13-45 (Sheila Jasanoff ed., 2004).

9. See quotes *supra* note 8.

These tools and institutions are designed to advance particular values, and are aimed at particular people, behavior, or things.<sup>10</sup>

#### A. *Defining Legal Construction of Technology*

Legal construction of technology rests on the theory that technology should be understood through its cultural interpretation – related but not identical to a science-and-technology-studies (STS) approach.<sup>11</sup> It is an analysis of how the long stories of law intersect with the long stories of technology, asking not whether some technology will require legal overhaul, but what choices were made across both social contexts such that a technology developed that would present legal questions worth considering. The law plays as much of a role as technology does: it embeds values, crafts institutions, and even sets up its own fault lines that later may crack. This method emphasizes that the law is worthy of deeper scrutiny; its values and theories, objects of regulation, and institutions all contribute to how technology is brought into contact with, and made meaningful within, the law.

What legal construction of technology does *not* do is ask as an endpoint, “What about this technology disrupts the law?” This framing presumes that technology does the acting, and law the responding. It treats technology as some outside force, perhaps even external to society. The deterministic framing typically also presumes that technological development inherently moves at a pace law is not equipped to follow (the so-called “pacing problem”).<sup>12</sup> Echoes of this framing sound across early cyberlaw debates about whether the Internet was an

---

10. The law can also construct technology in other ways, such as an aspect or disrupting aspect of the “imagined regulatory scene.” See Kaminski, *Technological Disruption*, *supra* note 5, at 886.

11. Legal construction fits within or alongside the foundational work on “coproduction” that Sheila Jasanoff developed in STS over the 1990s and early 2000s, which captures the way epistemic understanding goes hand-in-hand with normative understandings of the world. Jasanoff, *supra* note 8, at 13-45.

12. The pacing problem has been articulated in scholarship (see, for example, Gary E. Marchant, *The Growing Gap Between Emerging Technologies and the Law*, in 7 INT’L LIBR. ETHICS, L. & TECH., THE GROWING GAP BETWEEN EMERGING TECHNOLOGIES AND LEGAL-ETHICAL OVERSIGHT: THE PACING PROBLEM 19, 22-23 (Gary E. Marchant, Braden R. Allenby & Joseph R. Herkert eds., 2011)) and public discourse (see, for example, Vivek Wadhwa, *Laws and Ethics Can’t Keep Pace with Technology*, MIT TECH. REV. (Apr. 15, 2014), <https://www.technologyreview.com/2014/04/15/172377/laws-and-ethics-cant-keep-pace-with-technology> [<https://perma.cc/38KK-VRH8>]).

“exceptional” technology,<sup>13</sup> and whether it presented radically new features that would require radical changes in law.<sup>14</sup>

There are immediate and practical costs to the other way, of technological essentialism. Starting with the technology and asking how the law responds to it locks lawyers and policymakers into preexisting paths, as a matter of both law and rhetoric. As a consequence of politics and history, this often leads to a crippled legal system and strong technological system, in which innovation wins because innovation is best. Path dependencies occur.<sup>15</sup> Legal scholars have detailed path dependencies in judicial use of precedent,<sup>16</sup> administrative expertise and resources,<sup>17</sup> and legislative policy.<sup>18</sup> Part of the path, too, includes institutional expertise, resources, and culture. The path-dependent ways the law makes meaning of technology can hide significant policy choices in design specs,<sup>19</sup> the quantification of risk or harms,<sup>20</sup> and procurement policies,<sup>21</sup> shielding substantive policymaking from democratic processes.<sup>22</sup>

- 
13. See Easterbrook, *supra* note 4, at 215-16; Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 CALIF. L. REV. 513, 514-16 (2015) (“Some early cyberlaw questions have seen a kind of resolution in the twenty years since the rise of cyberlaw. Legislatures and courts have weighed in, and the vigorous debate has continued—around ‘net neutrality,’ for instance, and the impossible puzzle that is privacy. But even here, participants have at least a sense of the basic positions and arguments. Law, in other words, is catching up . . . [R]obotics has a different set of essential qualities than the Internet . . . But the parallels are also strong . . . Ultimately, however, robotics law can and should depart from cyberlaw and form a distinct area of governance and study.”).
  14. This is a common stumbling block for law-and-technology scholars because the essentialism or exceptionalism approach starts with a technical definition and then requires additional steps into cultural contexts (like law and use and economics and infrastructure) just to present their questions.
  15. Gaia Bernstein, *When New Technologies Are Still New: Windows of Opportunity for Privacy Protection*, 51 VILL. L. REV. 921, 921 (2006).
  16. Oona A. Hathaway, *Path Dependence in the Law: The Course and Pattern of Legal Change in a Common Law System*, 86 IOWA L. REV. 601, 650-52 (2001).
  17. Helen Ingram & Leah Fraser, *Path Dependency and Adroit Innovation: The Case of California Water*, in PUNCTUATED EQUILIBRIUM AND THE DYNAMICS OF U.S. ENVIRONMENTAL POLICY 78, 83-84 (Robert Repetto ed., 2006).
  18. Janelle Knox-Hayes, *Negotiating Climate Legislation: Policy Path Dependence and Coalition Stabilization*, 6 REGUL. & GOVERNANCE 545, 559-61 (2012).
  19. WOODROW HARTZOG, *PRIVACY’S BLUEPRINT: THE BATTLE TO CONTROL THE DESIGN OF NEW TECHNOLOGIES* 34-43 (2018).
  20. Kaminski, *Regulating the Risks*, *supra* note 7, at 1397-98.
  21. Deirdre K. Mulligan & Kenneth A. Bamberger, *Procurement as Policy: Administrative Process for Machine Learning*, 34 BERKELEY TECH. L.J. 773, 788-90 (2019); JENNIFER PAHLKA, *RECODING AMERICA: WHY GOVERNMENT IS FAILING IN THE DIGITAL AGE AND HOW WE CAN DO BETTER* 5-8 (2023).
  22. Kaminski, *Regulating the Risks*, *supra* note 7, at 1398-99.

In a more recent conversation, two prominent voices in the field of law and technology again debated technological essentialism.<sup>23</sup> In an article comparing robotics to the Internet, Ryan Calo argued that the “essential qualities” of robots would necessitate “distinct conversation[s]” about the law.<sup>24</sup> Robots’ essential technical qualities, Calo posited, would require changes in the law.<sup>25</sup> Jack M. Balkin responded that drawing out a technology’s essential qualities was not particularly useful when thinking about legal application.<sup>26</sup> Nor is boiling a technology down to essential qualities really even possible, given the way technology in society evolves over time.<sup>27</sup>

Calo certainly recognizes the social aspects of technology.<sup>28</sup> But his initial approach—distilling technology to its essential features that act upon the passive law—reflects technological essentialism. We reject this. Technologies develop in societies with laws that set expectations about how they will work and why. Then, law does the daily work of maintaining those expectations. Law is always producing, maintaining, and “keeping up” with technology. Sometimes this is a matter of structuring how the law incorporates expertise from other fields.<sup>29</sup>

---

23. The debate is fundamental to the philosophy of technology, where technological essentialism, in its most simplistic terms, means that technology has an essence. *See, e.g.*, Andrew Feenberg, *From Essentialism to Constructivism: Philosophy of Technology at the Crossroads*, in *TECHNOLOGY AND THE GOOD LIFE* 294, 294-95 (Eric Higgs, Andrew Light & David Strong eds., 2000) (promoting a continued recognition of technical phenomena but a departure from essentialist theories that define the technical as abstracted from society); Sally Wyatt, *Feminism, Technology, and the Information Society*, 11 *INFO., CMTY. & SOC’Y* 111, 122 (2008) (tying anti-essentialism to feminist technology studies).

24. Calo, *supra* note 13, at 515.

25. *See id.* at 553-58.

26. Jack M. Balkin, *The Path of Robotics Law*, 6 *CALIF. L. REV. CIR.* 45, 45 (2015).

27. *Id.* (“I do not think it is helpful to speak in terms of ‘essential qualities’ of a new technology that we can then apply to law. On the contrary, we should try not to think about characteristics of technology as if these features were independent of how people use technology in their lives and in their social relations with others. Because the use of technology in social life evolves, and because people continually find new ways to employ technology for good or for ill, it may be unhelpful to freeze certain features of use at a particular moment and label them ‘essential.’”).

28. Calo, *supra* note 13, at 516 (“The better we understand how a technology will affect society, the better position we are in to integrate the technology gracefully.”). He later wrote an essay on the importance of STS to law and technology scholarship. Ryan Calo, *The Scale and the Reactor* (Apr. 15, 2022) (manuscript at 16), <https://ssrn.com/abstract=4079851> [<https://perma.cc/4LCW-ATRS>] (“STS is a set of methods, concepts, and assumptions—a way of seeing technology as social fact—that I think of as indispens[a]ble to the study of law and technology.”).

29. The law incorporates expertise through institutional design, decisions about which entity crafts policy (e.g., Congress versus the courts versus administrative agencies), and related



Other times, law “keeps up” with technology the same way it “keeps up” with other evolving social facts: through a combination of fuzzier standards, common law, and new legislation and regulation.<sup>30</sup>

These tasks do not always achieve evolving social goals and values, but that is a matter of politics, not the nature of the relationship between law and technology.<sup>31</sup> So, while Calo could claim that the fact of robots’ embodiment “disrupted” the way U.S. law handles software and liability, the same is not necessarily true of European law.<sup>32</sup> Likewise, the U.S. legislation passed to change the legal category for connected computers (Communication Decency Act Section 230, discussed below), was a political effort made in response to lower-level court cases and only part of the legislative effort to shape content on the Internet. Today, Section 230’s special platform protections are not under threat from new embodied robots, but rather from political actors in Florida and Texas where state legislation attempts to restrict and govern how platforms moderate content.<sup>33</sup>

We find that legal construction benefits from a comparative perspective – whether comparisons within domestic legal systems between different subject-matter areas, or comparisons across legal cultures.<sup>34</sup> A comparative approach, while not required, affords researchers greater perspective on what they are taking as given, versus where there is possibility to imagine things very differently, from analogies to values and institutions. In other words, if a technology appears to be “exceptional” to torts but not to privacy laws, or in the United States but not in Europe, it says more about the political histories of those areas of law and places than about the particular attributes of that technology.<sup>35</sup>

Legal construction of technology requires uncovering more than “what the law says” and “how the technology works”; it demands a deep understanding of the relationship between law and technology in context over time, and of the

---

choices about ex ante versus ex post decision-making. See Rebecca Crootof & BJ Ard, *Structuring Techlaw*, 34 HARV. J.L. & TECH. 347, 380-86, 407 (2021).

30. Jones, *supra* note 5, at 277-81; Kaminski, *Technological Disruption*, *supra* note 5, at 891.
31. Jones, *supra* note 5, at 278 (“Not only does law not linearly follow technology, a great deal of legal work shapes technology and the way in which it will be understood in the future.”).
32. *Id.* at 262-63.
33. See *NetChoice, LLC v. Paxton*, 49 F.4th 439 (5th Cir. 2022), *cert. granted in part*, No. 22-555, 2023 WL 6319650 (Sept. 29, 2023).
34. See, e.g., Jones, *supra* note 5, at 260-77 (conducting a comparative analysis of intermediary liability).
35. Kaminski, *Authorship, Disrupted*, *supra* note 5, at 590-91; see also Kaminski, *Technological Disruption*, *supra* note 5, at 885-86 (“The law, in constructing – that is, building the meaning of – new technological developments and their social uses, takes a central part in its own disruption.”).

cultural evolution of governance. Often law-and-technology scholarship is interesting precisely because technological development makes salient something about existing law and governance structures that might not have been visible before.<sup>36</sup> But again, that is not some special function of a technical artifact; it is the result of a process of continual legal meaning-making that sometimes, in constructing technology, highlights a feature or a bug that we failed to notice about the law.

The ways we make legal sense of technologies changes the way we design, integrate, and regulate them. For example, designers, marketers, and lawyers all use analogies to advocate for certain understandings of technologies. A designer might build a portable drive with cameras and microphones to be seen as a “family member,” while a marketer might give it a name like “home assistant,” and a lawyer might analogize it to a “tiny constable.”<sup>37</sup> Is the cloud a fluffy ether, remote storage, or someone else’s computer? Are we thinking of email as letter-like, the way a user experiences it, or as a physical network?<sup>38</sup> These sensemaking choices have value-laden impacts: they drive policy choices, like whether courts will protect an expectation of privacy and require a warrant for a search. There are normative and deeply pragmatic ramifications of understanding a driverless car as a chauffeur, versus group transportation, versus a proxy driver.

As outlined in Part II below, we already see wranglings over analogies for AI-generated speech, both in existing law, and in how commentators have suggested changing the law. A speech-generating AI system could be a speaker,<sup>39</sup> or

---

36. Jack M. Balkin, *Digital Speech and Democratic Culture: A Theory of Freedom of Expression for the Information Society*, 79 N.Y.U. L. REV. 1, 2 (2004) (“This effect—making more central and visible what was already always present to some degree—is important in any study of the Internet and digital technologies. In studying the Internet, to ask ‘What is genuinely new here?’ is to ask the wrong question. If we assume that a technological development is important to law only if it creates something utterly new, and we can find analogues in the past—as we always can—we are likely to conclude that because the development is not new, it changes nothing important. That is the wrong way to think about technological change and public policy, and in particular, it is the wrong way to think about the Internet and digital technologies. Instead of focusing on novelty, we should focus on salience. What elements of the social world does a new technology make particularly salient that went relatively unnoticed before?”).

37. See, e.g., Kevin S. Bankston & Ashkan Soltani, *Tiny Constables and the Cost of Surveillance: Making Cents Out of United States v. Jones*, 123 YALE L.J.F. 335, 339 (2014).

38. Orin Kerr, *The Problem of Perspective in Internet Law*, 91 GEO. L.J. 357, 359-61 (2003).

39. See *infra* Part I.

a platform or tool;<sup>40</sup> a data processor;<sup>41</sup> or a risky complex system.<sup>42</sup> AI-generated speech can be analogized to encyclopedias; or speech-generating AI systems can be analogized to advice-offering physicians.<sup>43</sup> Each of these analogies illustrates the law's attempts at meaning-making and carry policy consequences. The fact that AI systems could, and in fact probably do, fall into multiple legal buckets or no existing legal buckets in a particular area of law, does not tell us that the law is failing to keep up with technology. It tells us that legal systems, or really legal actors, are already using the law's tools of sensemaking to enact underlying values in decisions about what harms and benefits should be afforded, and to whom.

Legal construction of technology, unlike technological-essentialist and pacing-problem-oriented framings, puts legal actors where they accurately sit: not chasing technology, but as potentially powerful interveners in conception and design through to integration and enforcement. The method insists that law not needlessly cede its role in shaping society to technology and its benefactors. As one of us has explained, “[b]y chasing new technologies with legal solutions, law-and-technology scholars, as well as policymakers, unnecessarily accept a degree of irrelevance.”<sup>44</sup>

Legal construction of technology is also unavoidable. We all do it, searching for analogies or trying to match policy outcomes to existing institutional goals. Yet it is very easy to do it badly, by tacitly accepting what is, and turning to technological essentialism as the way to frame a core set of questions. Every law-and-technology scholar, wittingly or not, asks how the law makes meaning of technology. But not every scholar steps back to ask not, “What is special about this artifact?” but, “What are the costs of constructing technology from within existing governance systems, and is there a better way?” Legal constructionists see moments of technological sensemaking as opportunities for intervention, and interventions as part of the sensemaking process. While seemingly descriptive, legal construction as a method provides new possibilities for normative action, some of which we discuss in Part III below.<sup>45</sup>

---

40. See Derek E. Bambauer & Mihai Surdeanu, *Authorbots*, 3 J. FREE SPEECH L. 375, 385 (2023).

41. See Kaminski, *Binary Governance*, *supra* note 7, at 1582-85 (discussing how data protection law constructs AI).

42. See evelyn douek, *Content Moderation as Systems Thinking*, 136 HARV. L. REV. 526, 530-34 (2022); Jack M. Balkin, *Free Speech Versus the First Amendment*, 70 UCLA L. REV. 1206, 1221-24 (2023); see also Kaminski, *Regulating the Risks*, *supra* note 7, at 1350-51.

43. Jane Bambauer, *Negligent AI Speech: Some Thoughts About Duty*, 3 J. FREE SPEECH L. 343, 348-61 (2023).

44. Jones, *supra* note 5, at 256.

45. *But see* Calo, *supra* note 28 (manuscript at 2-3) (arguing that normativity and pragmatism are features of American legal thought, in contrast to an STS approach).

### B. Three Core Elements

Here, we identify three core elements of legal construction: (1) the *values* (theories) behind different legal subfields, (2) the *institutions* involved and the regulatory tools they deploy, and (3) the *objects* or targets of regulation. These elements make up a great deal, but not all, of how law might construct technology.<sup>46</sup> Rather than just performing legal interpretation, effective legal construction of technology seeks to explicitly identify these elements and what they contribute to existing law’s sensemaking. It aims to reveal what we are not seeing: policy baggage, path dependencies, and blind spots. This, in turn, lets us step back to think about other possible interventions: how we might legally construct a particular technology differently.

**Values.** Each area of law has its driving theories, which typically prioritize particular *values*.<sup>47</sup> These theories and their values aim, and indeed typically constrain, the law’s construction. For example, U.S. copyright law is widely understood to have utilitarian underpinnings, prioritizing the production of creative works toward “net social welfare,” often with an economic bent.<sup>48</sup> Scholars have debated the extent to which this framing drives stronger rights for authors versus broader exceptions for remixes and reuse.<sup>49</sup> But the values underlying U.S. copyright law are typically economic, not dignitary. This both explains the Supreme Court’s willingness to carve out special treatment for copyright law under the First Amendment (property rights being given special treatment to promote capitalism) and limits the ways in which copyright exceptionalism might be extended to analysis of regulation that protects dignitary rights, such as privacy law. Foregrounding the values of different parts of the legal system allows legal constructionists to better identify consequent blind spots. It allows us to move

---

46. There are other ways of constructing, too. See, e.g., Kaminski, *Authorship, Disrupted*, *supra* note 5, at 592-93; see also Kaminski, *Technological Disruption*, *supra* note 5, at 895-903 (describing the concept of disruption of the “imagined regulatory scene,” an example of legal construction that “takes place not on the page but in our heads”).

47. For a longer discussion of this, see Kaminski, *Authorship, Disrupted*, *supra* note 5, at 606, 613.

48. See Pamela Samuelson, *Allocating Ownership Rights in Computer-Generated Works*, 47 U. PITT. L. REV. 1185, 1224 (1986) (“If there is no human author of such a work, how can any human be motivated to create it? The copyright system assumes that society awards a set of exclusive rights to authors for limited times in order to *motivate* them to be creative . . .”). Scholars have also debated whether U.S. copyright law should instead or in addition look to dignitary arguments about authorial expression. See, e.g., Justin Hughes, *The Philosophy of Intellectual Property*, 77 GEO. L.J. 287, 330-31 (1988).

49. See, e.g., Pamela Samuelson, *Unbundling Fair Uses*, 77 FORDHAM L. REV. 2537, 2617-18 (2009); Cathay Y.N. Smith, *Creative Destruction: Copyright’s Fair Use Doctrine and the Moral Right of Integrity*, 47 PEPP. L. REV. 601, 662-63 (2020).

back from analogies and legal boxes to query whether we agree with underlying regulatory goals.

**Institutions.** Each area of the law, too, involves differing *institutions* and institutional combinations, which use different regulatory tools and perform different methods of legal construction toward different ends. For instance, without federal data-protection legislation, Americans have come to understand computer technology as consumer services regulated by the Federal Trade Commission (FTC). The FTC is culturally an enforcement agency, enabled and restricted by its charge to protect consumers from unfair and deceptive trade practices. As a consequence of history, it struggles to promulgate regulation; instead, it issues consent decrees.<sup>50</sup> The FTC's oversight has historically been limited to certain commercial entities and does not extend, for example, to nonprofits or to telecommunications companies acting as communications networks regulable by the Federal Communications Commission (FCC). While the FCC, as an independent expert agency, attracts the talents and expertise of economists and telecommunications engineers, the FTC largely attracts justice-oriented attorneys and investigators.<sup>51</sup>

These stories of institutions, their tools, their capacities, and their cultures enable legal constructionists to examine not just the content of substantive law, such as FTC consent decrees, but also the consequences. The consequences of framing affected individuals as “consumers” who have been “deceived,” rather than “data subjects” whose “fundamental rights” have been violated are both rhetorical and immensely practical.<sup>52</sup> Further, these stories allow us to look at

---

50. The Federal Trade Commission (FTC) conducts Magnuson-Moss rulemaking. See Jeffrey S. Lubbers, *It's Time to Remove the "Mossified" Procedures for FTC Rulemaking*, 83 GEO. WASH. L. REV. 1979, 1982-85 (2015). The agency has recently taken the initiative to unburden itself from constraints on its rulemaking authority. See Press Release, Fed. Trade Comm'n, FTC Votes to Update Rulemaking Procedures, Sets Stage for Stronger Deterrence of Corporate Misconduct (July 1, 2021), <https://www.ftc.gov/news-events/news/press-releases/2021/07/ftc-votes-update-rulemaking-procedures-sets-stage-stronger-deterrence-corporate-misconduct> [<https://perma.cc/M2UJ-8LMY>]; see also Aaron L. Nielson, *What Happens If the FTC Becomes a Serious Rulemaker?*, in RULEMAKING AUTHORITY OF THE FEDERAL TRADE COMMISSION 255, 255-57 (2022) (describing how the FTC is beginning to engage in more aggressive rulemaking).

51. The FTC does have economists that are sometimes at odds with its more justice-oriented arms. See CHRIS JAY HOOFNAGLE, FEDERAL TRADE COMMISSION PRIVACY LAW AND POLICY 95 (2016) (observing that at the Bureau of Consumer Protection, “many attorneys, sometimes even those at the top of the bureau, are dissatisfied with the economists’ substantive provisions”).

52. See generally MEG LETA JONES, THE CHARACTER OF CONSENT: THE HISTORY OF COOKIES AND THE FUTURE OF TECHNOLOGY POLICY (forthcoming 2024) (detailing the history of who consents to web cookies by analyzing the legal constructions of people engaged with computer technologies in different areas of law: the consumer in U.S. consumer-protection law, the data

specific legal tools such institutions use and identify what one of us has termed the “policy baggage” of said tools: the pathologies inherent in different practices within the law and its particular ways of practical sensemaking.<sup>53</sup>

**Objects.** The third element we highlight in the examples in Section II.B below is the *object* or target of regulation. Law-and-technology scholars have long discussed whether laws are best crafted as technology-specific or technology-neutral.<sup>54</sup> However, this is just one aspect of what we mean by examining the chosen object of regulation. The object of regulation tells us a lot about how an area of the law constructs technological practices. A law can target facial-recognition systems. It can target drones. It can target personal data, or content-moderation systems, or monopolies, or money. It can target an entire meta-class of “anything that harms consumers.” As we show in the below analyses, not only does identifying the object of regulation often provide a practical answer as to whether speech-generating AI systems are subject to that regulation, it also lets us understand what aspects of said systems are *not* being regulated or even noticed through a particular branch of the law. Identifying the object of regulation could help identify how such systems might be better regulated, or are already regulated, through broad, existing, technology-neutral laws that construct such systems not as an instantiation of the pacing problem, but as the product of readily regulable corporate actors. The recent U.S. (re)turn to antitrust enforcement against Big Tech illustrates this potential move.<sup>55</sup>

In short, legal construction asks: What system of legal meaning-making does technology encounter, toward what ends, using what institutions and tools? The rest of this Essay seeks to illustrate this by way of a closer examination through legal construction of the question: “Do AI speakers disrupt the law?”

---

subject in European data-protection law, and the user in transatlantic-communication-privacy law each results in very different ideas and structures of technical governance).

53. Kaminski, *Regulating the Risks*, *supra* note 7, at 1390 (referring to the “tools, tactics, and troubles already in practice”).
54. See, e.g., Michael Birnhack, *Reverse Engineering Information Privacy Law*, 15 YALE J.L. & TECH. 24, 36-52 (2012); Brad A. Greenberg, *Rethinking Technology Neutrality*, 207 MINN. L. REV. 1495, 1498-1501 (2016); Crootof & Ard, *supra* note 29, at 408-13.
55. See Press Release, U.S. Dep’t of Just., Justice Department Sues Google for Monopolizing Digital Advertising Technologies (Jan. 24, 2023), <https://www.justice.gov/opa/pr/justice-department-sues-google-monopolizing-digital-advertising-technologies> [<https://perma.cc/N6ZN-NFU5>].

## II. CONSTRUCTING AI SPEECH

### A. Definitions

What is AI? The answer is more complex than the abstract technical specifications of an artifact. We begin by “defining AI,” for purposes of legal construction of AI-generated speech, as a technical, socio-historical, and ultimately legal matter. We then progress into how different parts of the legal system already construct AI systems.

When we talk about AI systems in the context of AI-generated speech, we are largely referring to Large Language Models (LLMs). Nina Brown offers a helpful working technical explanation for legal scholars: “Large language models (LLMs) are a type of deep learning algorithm used to model statistical relationships between words and phrases in large bodies of text data in order to generate human-like language.”<sup>56</sup> LLMs are prediction machines: they predict what words or phrases come next, as a matter of (complex) statistics. As technical artifacts, such systems are trained on large databases, often containing personal data and copyrighted data, with substantial and ongoing human involvement.

LLMs, like other AI systems, are always human-machine systems; there is always a human in the loop, in both the narrower and broader senses of the term.<sup>57</sup> Humans and organizations of humans pick goals and craft design parameters for these systems.<sup>58</sup> Humans participate in reinforcement learning to refine and “optimize” system output.<sup>59</sup> Humans and their institutions matter for legal construction of these systems in at least three ways: (1) Depending on how much one broadens the lens on what constitutes an AI system, there is always a human who could be regulated.<sup>60</sup> (2) However, human-machine systems and interfaces are notoriously tricky to get right.<sup>61</sup> (3) As we discuss below, the complexity of human-machine systems means that where law or policy attempts to treat AI as a fungible substitute for individual human actors, the fit is often off.

---

56. Brown, *supra* note 1, at 393.

57. See Crootof, Kaminski, & Price, *supra* note 7, at 443; Jones, *Human in the Loop*, *supra* note 7, at 217; Andrew D. Selbst, Danah Boyd, Sorelle A. Friedler, Suresh Venkatasubramanian & Janet Vertesi, *Fairness and Abstraction in Sociotechnical Systems*, ACM DIGIT. LIBR. (Jan. 29, 2021), <https://dl.acm.org/doi/10.1145/3287560.3287598> [<https://perma.cc/6FJC-VBC9>].

58. See David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653, 694 (2017); Brown, *supra* note 1, at 394.

59. Brown, *supra* note 1, at 394.

60. See Lehr & Ohm, *supra* note 58, at 657-58; Crootof, Kaminski & Price, *supra* note 7, at 443-44.

61. See Crootof, Kaminski & Price, *supra* note 7, at 438; Jones, *Human in the Loop*, *supra* note 7, at xx.

Like Google’s relevance search algorithm, the AI at the center of policy debates today should not be understood as objective math, though of course there is plenty of math involved. Trying to pull understanding “how LLMs work” apart from understanding LLMs as a product or a cultural artifact is a pointless and impossible endeavor. So, we provide some background not just on technical specifications but on LLMs as products or artifacts here.

The most relevant generative AI system is ChatGPT, which operates both as a standalone app and integrated into other products. As its own service, ChatGPT made headlines. As of September 2023, it touts 52.2 million downloads and \$4.6 million in gross revenue from its subscription service.<sup>62</sup> ChatGPT is also integrated into other systems like Loona the robot dog<sup>63</sup> and Canva’s Magic Studio.<sup>64</sup> It also offers enterprise licensing.<sup>65</sup>

ChatGPT is the product of OpenAI, which originally began as a nonprofit in 2015 but formed a “capped profit” entity with a specified and strange investment structure that early backers and founders like Elon Musk have criticized.<sup>66</sup> Although the details have not been made public, according to reports, Microsoft invested \$13 billion in OpenAI, starting with \$1 billion in 2019. After taking a seventy-five percent share of OpenAI’s profit until the company makes its money back, Microsoft will have a forty-nine percent stake in OpenAI.<sup>67</sup> OpenAI was

- 
62. Sarah Perez, *ChatGPT’s Mobile App Hit Record \$4.58M in Revenue Last Month, But Growth is Slowing*, TECHCRUNCH (Oct. 9, 2023, 3:08 PM EDT), <https://techcrunch.com/2023/10/09/chatgpts-mobile-app-hit-record-4-58m-in-revenue-last-month-but-growth-is-slowing> [<https://perma.cc/SG4M-CHXN>].
63. See *Loona*, KEYI ROBOT, <https://keyirobot.com/pages/loonadetail> [<https://perma.cc/QK75-ZNJG>] (advertising a neural network processor handling five trillion calculations per second that “makes Loona one of the smartest consumer robots in the world”).
64. See, e.g., Jesse Weatherbed, *Canva’s AI Tools Automate Boring, Labor-Intensive Design Tasks*, VERGE (Oct. 4, 2023, 9:00 AM EDT), <https://www.theverge.com/2023/10/4/23902794/canva-magic-studio-ai-design-new-tools> [<https://perma.cc/DD7E-AJ49>] (citing Canva’s description of Magic Studio as “the world’s most comprehensive AI-design platform”).
65. See, e.g., *Introducing ChatGPT Enterprise*, OPENAI: BLOG (Aug. 28, 2023), <https://openai.com/blog/introducing-chatgpt-enterprise> [<https://perma.cc/NE9L-MH5T>] (explaining that ChatGPT is used to “craft clearer communications, accelerate coding tasks, rapidly explore answers to complex business questions, assist with creative work, and much more”).
66. See, e.g., Jordan Novet, *Microsoft’s \$13 Billion Bet on OpenAI Carries Huge Potential Along with Plenty of Uncertainty*, CNBC (Apr. 8, 2023, 9:00 AM EDT), <https://www.cnbc.com/2023/04/08/microsofts-complex-bet-on-openai-brings-potential-and-uncertainty.html> [<https://perma.cc/TR5H-YV2M>] (citing Elon Musk as having tweeted, “OpenAI was created as an open source (which is why I named it ‘Open’ AI), non-profit company to serve as a counterweight to Google, but now it has become a closed source, maximum-profit company effectively controlled by Microsoft”).
67. Sara Morrison, *What Microsoft Gets from Betting Billions on the Maker of ChatGPT*, VOX (Jan. 23, 2023, 3:10 PM EST), <https://www.vox.com/recode/2023/1/23/23567991/microsoft-openai-investment-chatgpt> [<https://perma.cc/E5XK-8EML>]; Charles Duhigg, *The Inside Story of*



also a product of a crypto-stung venture capital scene that dove into AI, investing more than one in every four dollars on AI startups in 2023<sup>68</sup> and driving up salaries for AI labor to close to seven figures.<sup>69</sup> Other generative-AI systems have histories too, each different from the specific history of OpenAI.

None of these histories are predetermined, and all of them take place against an existing regulatory backdrop, whether the law has actively intervened or politics have driven a laissez-faire approach. Imagine how different these prediction machines might have been if U.S. states had amended their constitutions to include human rights to data protection, or if a market for personal data had been developed in the 1990s. Who would have built AI systems and why? What type of data would they have been trained on and what kind of audiences would be consuming their outputs? In an alternate universe somewhere, two law scholars have been asked how AI-generated speech disrupts communications law, wherein the FCC has been strictly regulating inputs to communications produced on or by networked systems since the telegraph – instead of largely taking a back seat to regulating the Internet and subsequent information technologies. What do AI systems look like in that universe?

Legal analyses often include a legal definition of a technological subject. This, too, is legal sensemaking. Typically, the policy debate over legal definitions focuses on how technologically specific a law should be, versus how much regulation should be oriented toward addressing non-technologically specific harms.<sup>70</sup> For example, U.S. copyright law has over time become increasingly technologically neutral, moving from regulating, for example, books and maps, to attempting to encompass any creative work fixed in a tangible medium of expression, including online.<sup>71</sup> Legal definitions can be the site of fervent political battles. They are often the targets of attempts to get out of regulatory coverage through characterizations of a technology as too new and thus challenging to regulate (e.g., we don't understand General Purpose AI yet!<sup>72</sup>) or as not worthy of extra

---

*Microsoft's Partnership with OpenAI*, NEW YORKER (Dec. 1, 2023), <https://www.newyorker.com/magazine/2023/12/11/the-inside-story-of-microsofts-partnership-with-openai> [<https://perma.cc/7RQZ-SPD3>].

68. Joanna Glasner, *AI's Share of US Startup Funding Doubled in 2023*, CRUNCHBASE NEWS (Aug. 29, 2023), <https://news.crunchbase.com/ai-robotics/us-startup-funding-doubled-openai-anthropic-2023> [<https://perma.cc/F2E2-HUHR>].

69. See Chip Cutter, *The \$900,000 AI Job Is Here*, WALL ST. J. (Aug. 14, 2023, 11:46 AM ET), <https://www.wsj.com/articles/artificial-intelligence-jobs-pay-netflix-walmart-23ofc3cb> [<https://perma.cc/V5E6-ZGZC>].

70. See, e.g., Birnhack, *supra* note 54, at 40; Greenberg, *supra* note 54, at 1515.

71. See Greenberg, *supra* note 54, at 1502-03.

72. See Claire Boine & David Rolnick, *General Purpose AI Systems in the AI Act: Trying to Fit a Square Peg into a Round Hole*, Paper Presentation at We Robot 2023 (on file with authors).

attention relative to existing technology (e.g., surveillance drones are no different from cell phone cameras!<sup>73</sup>).

The varied and changing legal definitions of AI systems already evidence these and other tensions.<sup>74</sup> Mark A. Lemley and Bryan Casey claim that while it is tempting to start the regulation of AI or robots with a definition of the technology, “it can’t be done, at least not well.”<sup>75</sup> Instead, they argue for a number of strategies for offsetting pacing issues, including regulating behavior rather than technology, and housing definitional decisions with courts or regulators rather than in legislatures.<sup>76</sup>

Claire Boine and David Rolnick track the evolving and competing definitions of AI systems in the EU AI Act.<sup>77</sup> Boine and Rolnick’s work demonstrates the ways in which legal definitions participate in legal construction of technology,

---

73. See, e.g., Gregory S. McNeal, *Can the “Drone” Industry Compete with the Privacy Lobby?*, FORBES (Aug. 13, 2012, 04:29 PM EDT), <https://www.forbes.com/sites/gregorymcneal/2012/08/13/can-the-drone-industry-compete-with-the-privacy-lobby> [<https://perma.cc/X3UX-HUZ2>] (talking about privacy and technophobia); see also Shimonti Paul, *How Fear-Based Regulations Hurt the Drone Industry: Brendan Schulman, DJI*, GEOSPATIAL WORLD (Aug. 31, 2020), <https://www.geospatialworld.net/blogs/how-fear-based-regulations-are-hurting-the-drone-industry> [<https://perma.cc/2GC4-9L8D>] (describing what is necessary for drones to fly freely and arguing that fear-based regulations are harming the drone industry).

74. See, for example, the Organisation for Economic Co-operation and Development (OECD)’s influential and evolving definition of AI, *Scoping the OECD AI Principles: Deliberations of the Expert Group on Artificial Intelligence at the OECD (AIGO)*, OECD Digital Economy Papers No. 291 (Nov. 2019), <https://www.oecd-ilibrary.org/docserver/d62f618a-en.pdf> [<https://perma.cc/93XZ-9MLP>], at 7 (“An AI system is a machine-based system that is capable of influencing the Environment by making recommendations, predictions or decisions for a given set of Objectives. It does so by utilising machine and/or human-based inputs/data to: *i*) perceive real and/or virtual environments; *ii*) abstract such perceptions into models manually or automatically; and *iii*) use Model Interpretations to formulate options for outcomes.”); and *OECD AI Principles Overview*, OECD.AI, <https://oecd.ai/en/ai-principles> [<https://perma.cc/D3HV-ZYTV>] (defining an AI system as “a machine-based system that is capable of influencing the environment by producing an output (predictions, recommendations or decisions) for a given set of objectives,” that “uses machine and/or human-based data and inputs to (i) perceive real and/or virtual environments; (ii) abstract these perceptions into models through analysis in an automated manner (e.g., with machine learning), or manually; and (iii) use model inference to formulate options for outcomes,” and that is “designed to operate with varying levels of autonomy”). For another definition of AI, see Algorithmic Accountability Act of 2019, H.R. 2231, 116th Cong. § 2(1) (2019) (defining an automated decision system as “a computational process . . . that makes a decision or facilitates human decision making, that impacts consumers”).

75. Bryan Casey & Mark A. Lemley, *You Might be a Robot*, 105 CORNELL L. REV. 287, 293 (2020). The authors add that, “the overlap between people, algorithms, computers, robots, and ordinary machines is sufficiently great that there is no good legal definition.” *Id.*

76. *Id.* at 356-61.

77. See Boine & Rolnick, *supra* note 72, at 7 tbl.1.

and how definitions interact with other aspects of legal construction. They note that the definition of AI systems in the initial 2018 draft of the Act focused on intelligence and agency,<sup>78</sup> while the 2021 text from the EU Commission offered a far more general definition that encompassed nonautonomous statistical software.<sup>79</sup> Later definitions, adopted by the other legislative bodies of the EU, returned to focusing on autonomy,<sup>80</sup> distinguishing autonomous systems from statistical software.<sup>81</sup> Boine and Rolnick also point to the European Parliament's attempts to add a definition of General Purpose AI systems to the Act.<sup>82</sup> They note, however, that the whole structure of the EU AI Act, with its focus on risks in the context of specific uses, and fit to intended purposes, is poorly designed for regulating General Purpose AI systems that are adaptable to many unpredictable uses.<sup>83</sup> Their work demonstrates how definitions interact with other aspects of legal sensemaking—for example, decisions to focus on the developers of systems and on regulating uses in specific contexts.

---

78. *See id.* (“Artificial intelligence (AI) refers to systems that display intelligent behaviour by analysing their environment and taking actions—with some degree of autonomy—to achieve specific goals. AI-based systems can be purely software-based, acting in the virtual world (e.g. voice assistants, image analysis software, search engines, speech and face recognition systems) or AI can be embedded in hardware devices (e.g. advanced robots, autonomous cars, drones or Internet of Things applications).”).

79. *Id.* (“Software that is developed with one or more of the techniques and approaches listed in Annex I [machine-learning approaches, logic- and knowledge-based approaches, and statistical approaches] and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with.”).

80. The AI Act as of December 2022 from the Council of the EU defined AI as “[a] system that is designed to operate with elements of autonomy and that, based on machine and/or human-provided data and inputs, infers how to achieve a given set of objectives using machine learning and/or logic- and knowledge-based approaches, and produces system-generated outputs such as content (generative AI systems), predictions, recommendations or decisions, influencing the environments with which the AI system interacts.” *Id.* The AI Act as of June 2023 (as adopted by the EU Parliament) defined AI as “[a] machine-based system that is designed to operate with varying levels of autonomy and that can, for explicit or implicit objectives, generate outputs such as predictions, recommendations, or decisions, that influence physical or virtual environments.” *Id.*

81. *See id.* at 31.

82. *See id.* at 32 (noting that, in June 2023, the European Parliament defined a General Purpose AI system as one “that can be used in and adapted to a wide range of applications for which it was not intentionally and specifically designed”).

83. *See id.* at 17-19.

### B. *Four Analyses of the Law(s) of AI-Generated Speech*

In the following four analyses of the law (or laws) of AI-generated speech, we identify and query the values, objects of regulation, and institutions that different parts of the legal system bring to bear on the “problem” of AI-generated speech. In doing so, we hopefully illustrate what legal construction of technology brings to the methodological table.

Each analysis takes AI speech in a different legal context and reveals how the prioritized values of each space, which may also be contested and evolving, shape how we understand what the technology is and how it works. Depending on these values, legal actors ask different questions about AI’s structure, arrangements, and functions. Institutions with their own cultures, instruments, and political statuses address sometimes very different conceptions of the objects of regulation. The law, we hope to show, can come to construct AI-generated speech in these four different ways—and many more—with very different policy outcomes.

#### 1. *AI as a Speaker*

Is AI-generated speech covered by the First Amendment? Is it protected? In this Section, we identify where First Amendment sensemaking of AI-generated content is still very much contested and ongoing—and at what levels (doctrinal versus theoretical versus institutional). We largely reason within existing First Amendment doctrine and try to be direct about the scope and feasibility of different possible interventions.<sup>84</sup>

The prototypical regulatory *objects* of First Amendment law are *speech* and *speakers*. That is, the First Amendment regulates speech by speakers (who are only sometimes human). A significant part of the First Amendment puzzle thus constructs the question of AI-generated speech as a “substitution problem”: a question of what happens to the law when you substitute an individual human

---

84. A great example of this practice is Stuart M. Benjamin, *Algorithms and Speech*, 161 *U. PA. L. REV.* 1445, 1447 (2013) (“I conclude that if we accept Supreme Court jurisprudence, the First Amendment encompasses a great swath of algorithm-based decisions—specifically, algorithm-based outputs that entail a substantive communication. We could decide to reject Supreme Court jurisprudence, or read it narrowly in order to limit its application. But for the purposes of this Article, I will not apply that lens to the existing case law. Instead, I will look to broadly accepted sources and forms of legal reasoning—which in the First Amendment context means primarily Supreme Court jurisprudence—and consider whether those sources lead to the conclusion that algorithm-based outputs are speech for First Amendment purposes. I find that the answer is yes for most algorithm-based editing.”).

actor with AI.<sup>85</sup> This is, as we will see, only one of a number of ways one could legally construct speech generated by AI systems—even within current First Amendment law.

First Amendment *values* entail multiple theories of the First Amendment, some of which valorize and even prioritize quantity of speech, while others distinguish between higher- and lower-value speech. Nearly all of these theories, one of us has argued, are likely to drive First Amendment coverage of much AI-generated content, though they diverge in how much of that content might ultimately be constitutionally protected.<sup>86</sup>

The typical *institution* of the First Amendment is courts, particularly the politicized Supreme Court, interpreting the Constitution through two-party legal challenges. That is, First Amendment sensemaking happens through a particularly American form of legal analysis. Legislators have an attenuated role, but a role nonetheless; they craft the laws that courts then analyze.

To the extent that AI-generated content can be understood to “disrupt” First Amendment law, it does so largely because of an oddity about the doctrine’s regulatory objects. Different aspects of First Amendment doctrine differ in whether the target object of legal sensemaking is *speech* versus a human *speaker*. This potentially creates a scenario in which a great deal of AI-generated content is likely covered by the First Amendment as abstract and disembodied *speech*, but it is not clear how regulators can intervene to govern even historically unprotected speech, where the interpretative focus currently is on the intent of human *speakers*.<sup>87</sup>

By way of basic background, First Amendment analysis generally asks questions in two steps. The Court typically asks: (1) whether a genre of communication is covered by the First Amendment; and (2) if it is covered, whether it is constitutionally protected such that a law must be struck down as unconstitutional.<sup>88</sup> Numerous Supreme Court cases ask whether something is “covered” as

---

85. See Balkin, *supra* note 26, at 57 (referring to the substitution effect in his response to Ryan Calo).

86. See Toni M. Massaro, Helen Norton & Margot E. Kaminski, *Siriously 2.0: What Artificial Intelligence Reveals About the First Amendment*, 101 MINN. L. REV. 2481, 2486-91 (2017).

87. See, e.g., *id.*

88. See Frederick Schauer, *Out of Range: On Patently Uncovered Speech*, 128 HARV. L. REV. F. 346, 348 (2015) (“The question of which forms of speech are covered by the First Amendment is . . . distinct from the question of how much protection the speech that is covered will receive.”); see also Frederick Schauer, *The Boundaries of the First Amendment: A Preliminary Exploration of Constitutional Salience*, 117 HARV. L. REV. 1765, 1769 (2004) [hereinafter Schauer, *Boundaries of the First Amendment*] (explaining that evaluating First Amendment coverage precedes applying any test of First Amendment protection).

First Amendment speech or speech-entwined activity such as movies,<sup>89</sup> parades,<sup>90</sup> video games,<sup>91</sup> the Internet,<sup>92</sup> graphic design for websites,<sup>93</sup> nude dancing,<sup>94</sup> and transfers of doctors' prescription data for purposes of advertising.<sup>95</sup> It is only once something is found to be covered by the First Amendment—that is, constitutionally salient—that the Court will apply some form of scrutiny to the applicable law, including asking whether such covered activity nonetheless falls into a historically unprotected category of speech.<sup>96</sup>

Without getting into the edge cases, most actual speech—what the Court refers to as “pure speech,” whether spoken or written—is covered by the First Amendment. The Court recently referred to “pure speech,” distinct from embodied speech or speech-entwined conduct, as “[a]ll manner of speech—from ‘pictures, films, paintings, drawings, and engravings,’ to ‘oral utterance and the printed word.’”<sup>97</sup> There are historic exceptions of “pure speech” that are nonetheless invisible to and arguably not even covered by the First Amendment.<sup>98</sup> But

- 
89. *Compare Mut. Film Corp. v. Indus. Comm'n of Ohio*, 236 U.S. 230, 244 (1915) (describing movies as “a business, pure and simple . . . [and not] regarded . . . by the Ohio Constitution, we think, as part of the press of the country, or as organs of public opinion”), with *Joseph Burstyn, Inc. v. Wilson*, 343 U.S. 495, 502 (1952) (overruling *Mutual Film Corp.* and describing movies as within the coverage ambit of the First Amendment).
90. *Hurley v. Irish-Am. Gay, Lesbian, & Bisexual Grp. of Bos., Inc.*, 515 U.S. 557, 568-69 (1995).
91. *Brown v. Ent. Merch. Ass'n*, 564 U.S. 786, 790 (2011).
92. *Reno v. ACLU*, 521 U.S. 844, 868-70 (1997).
93. See *303 Creative LLC v. Elenis*, 600 U.S. 570, 586-87 (2023) (finding custom-wedding websites covered by the First Amendment).
94. *Barnes v. Glen Theatre, Inc.*, 501 U.S. 560, 565-66 (1991).
95. *Sorrell v. IMS Health Inc.*, 564 U.S. 552, 567-69 (2011).
96. See Schauer, *Boundaries of the First Amendment*, *supra* note 88, at 1769-70.
97. *303 Creative*, 600 U.S. at 587 (“All manner of speech—from ‘pictures, films, paintings, drawings, and engravings,’ to ‘oral utterance and the printed word’—qualify for the First Amendment’s protections; no less can hold true when it comes to speech like Ms. Smith’s conveyed over the Internet.”).
98. One prominent exception is securities law, which remained not salient to First Amendment analysis until only recently. See Schauer, *Boundaries of the First Amendment*, *supra* note 88, at 1771 (“Securities violations, antitrust violations, criminal solicitation, and many other categories of ‘speech’ remain uncovered by the First Amendment.”). But see Helen Norton, *What Twenty-First-Century Speech Law Means for Securities Regulation*, 99 NOTRE DAME L. REV. 97, 101, 105-06 (2023). Another is speech integral to criminal conduct, whereby somebody uses speech to agree to or facilitate a criminal course of action. There is a question of whether, for example, speech integral to criminal conduct is not covered or not protected. Eugene Volokh characterizes it as one of the categorial exceptions of unprotected speech. See Eugene Volokh, *The Speech Integral to Criminal Conduct Exception*, 101 CORNELL L. REV. 981, 983 (2016) (characterizing it as “a standard item on lists of First Amendment exceptions”). Frederick Schauer, on the other hand, characterizes criminal solicitation as among the “categories of ‘speech’ [that] remain uncovered by the First Amendment,” categories for which the First Amendment

it is likely most AI speech will be found by courts to be covered by the First Amendment.<sup>99</sup> That is, if AI-generated content produces an antitrust violation or facilitates a crime, like a human-generated equivalent, that content probably will not be covered speech, but if AI-generated content constitutes “pictures, films, paintings, drawings . . . oral utterance and the printed word,” its regulation will likely trigger First Amendment analysis of some kind.<sup>100</sup>

But how can this be so? An AI system is not a human speaker with human rights. It doesn’t have emotions; it doesn’t participate in the political process; it doesn’t experience a chilling effect in which it stops speaking for fear of retaliation or shunning. The fact that AI-generated pure speech will probably be covered by the First Amendment tells us as much about First Amendment law as it does about AI systems.<sup>101</sup> As Toni M. Massaro and Helen Norton observe, “[V]ery little in current free speech theory or doctrine makes First Amendment coverage contingent upon a human speaker.”<sup>102</sup>

Why is this the case? The *values* that drive First Amendment law, as interpreted by courts and the Supreme Court as First Amendment *institutions*, allow constitutional coverage and possibly protection of speakerless speech.

First Amendment coverage is explained by multiple theoretical justifications.<sup>103</sup> The three most influential positive arguments for First Amendment

“just does not show up.” Schauer, *Boundaries of the First Amendment*, *supra* note 88, at 1769, 1771. The recent case of *United States v. Hansen* illustrates the confusion, in that the Court refers to solicitation as “unprotected” but treats it as outside of the scope of First Amendment scrutiny. See 599 U.S. 762, 783 (2023) (“To the extent that clause (iv) reaches *any* speech, it stretches no further than speech integral to unlawful conduct. ‘[I]t has never been deemed an abridgement of freedom of speech or press to make a course of conduct illegal merely because the conduct was in part initiated, evidenced, or carried out by means of language, either spoken, written, or printed.’ Speech intended to bring about a particular unlawful act has no social value; therefore, it is unprotected.” (citation omitted)).

99. See Toni M. Massaro & Helen Norton, *Siri-ously? Free Speech Rights and Artificial Intelligence*, 110 NW. U. L. REV. 1169, 1187 (2016) (arguing throughout that AI speech is likely to be covered by the First Amendment because the First Amendment now covers speech, not speakers and distinguishing between speech in the abstract, and conduct or expressive conduct); Massaro et al., *supra* note 86, at 2521 (“Many of the information products that AI produce likely will be characterized not as pure speech, but as expressive conduct or something similar.”).

100. 303 *Creative*, 600 U.S. at 587.

101. See Massaro & Norton, *supra* note 99, at 1170-75; Balkin, *supra* note 36, at 2 (arguing that new technologies can make old features of the law and broader social world newly salient).

102. Massaro & Norton, *supra* note 99, at 1175; see also Massaro et al., *supra* note 86, at 2483 (“[T]he United States Supreme Court now emphasizes listeners’ interests in free speech outputs—rather than speakers’ humanness or humanity—in ways that make it exceedingly difficult to place AI speakers beyond the First Amendment’s reach.”).

103. See HARRY KALVEN, JR., A WORTHY TRADITION: FREEDOM OF SPEECH IN AMERICA 3 (Jamie Kalven ed., 1988) (“The Court has not fashioned a single, general theory which would explain all of its decisions; rather, it has floated different principles for different problems.”).

coverage and protection are that protecting speech promotes democratic self-governance; protecting speech contributes to the marketplace of ideas and thus to the search for enlightenment; and protecting speech preserves individual autonomy and freedom.<sup>104</sup> Under all of these positive theories, First Amendment rights can attach to speech because of the value of that speech to human audiences.<sup>105</sup> That is, AI-generated political diatribes that influence a human reader's vote would be covered by the First Amendment as speech that promotes democratic self-governance, contributes to the reader's search for truth, and contributes to the reader's autonomy. What matters is the speech, not the speaker; the marketplace and its consumers, not the humanity of the vendor.<sup>106</sup>

In recent years, the Court has additionally shown a growing trend toward invoking a negative justification for First Amendment protection, under which speaker identity matters possibly even less. These arguments focus on “the need to constrain the government’s potentially dangerous exercise of control over expression, and are rooted in distrust of the government as regulator.”<sup>107</sup> If the Court’s focus is on what kind of actions the government may or may not take, the humanity of the speaker matters little if at all. If Congress bans AI speech that criticizes members of Congress, courts will and probably should find that law unconstitutional because of censorial government intent.

Thus for purposes of gatekeeping whether there is a First Amendment issue in the first place, the law typically asks, “Is there *speech*?” not, “Is there a human *speaker*?”<sup>108</sup> Absent a radical reconfiguration of the Court’s current reasoning and doctrine, if AI-generated content is found to be speech, courts will typically next ask whether said speech falls into any historic categorical exceptions to First Amendment protection.<sup>109</sup> Otherwise, if there is “pure speech,” a court will typically apply strict scrutiny and often overturn the law.<sup>110</sup> (There are other

---

104. See Massaro et al., *supra* note 86, at 2487–88.

105. See *id.* at 2488–91.

106. See *Citizens United v. FEC*, 558 U.S. 310, 392 (2010) (Scalia, J., concurring) (“The [First] Amendment is written in terms of ‘speech,’ not speakers.”).

107. Massaro et al., *supra* note 86, at 2491 (footnote omitted).

108. There are plenty of exceptions to the “Is it speech?” approach, including, for example, doctrine about newsgathering or other actions antecedent to speech. See, e.g., Ashutosh Bhagwat, *Producing Speech*, 56 WM. & MARY L. REV. 1029, 1032–35 (2015); Margot E. Kaminski, *Privacy and the Right to Record*, 97 B.U. L. REV. 167, 177–82 (2017); see also Neil M. Richards, *Why Data Privacy Law Is (Mostly) Constitutional*, 56 WM. & MARY L. REV. 1501, 1505–08 (2015) (arguing that “regulation of the commercial trade in personal data will be consistent with the First Amendment, at least most of the time”).

109. See, e.g., *United States v. Alvarez*, 567 U.S. 709, 722 (2012).

110. More intermediate levels of scrutiny apply to, for example, expressive conduct. See *United States v. O’Brien*, 391 U.S. 367, 367–77 (1968).



exceptions – for example, for generally applicable laws that nonetheless affect expression.<sup>111</sup>)

The Court has carved out a list of increasingly ossified categorical exceptions to First Amendment protection of “pure speech”: defamation, fighting words, obscenity, and true threats, among others.<sup>112</sup> For many of these historic exceptions, the doctrine asks, “How do we define unprotected speech so as to protect a (human) *speaker*?” not, “How do we carve out objectively unprotected speech?” This sets up a potentially thorny puzzle, wherein the internal logic of current First Amendment doctrine is likely to result in coverage of AI-generated speech but lacks a clear answer as to when such speech might not be protected and thus be regulable.

For example, First Amendment doctrine constrains defamation law such that to show defamation of a public figure, one must show not just that speech is false but that a speaker has actual malice – that is, knowledge of or reckless disregard for the falsity of its assertions.<sup>113</sup> What does it mean to show that an AI speaker – or more probably its developer or user – has knowledge of, or has shown reckless disregard for falsity?<sup>114</sup> This analysis is further complicated by the fact that AI systems are not individual speakers at all but are complex human-machine systems developed using complex supply chains.<sup>115</sup>

---

111. See, e.g., *Giboney v. Empire Storage & Ice Co.*, 336 U.S. 490 (1949).

112. See, e.g., *Counterman v. Colorado*, 600 U.S. 66, 73-74 (2023) (discussing the historic exceptions to First Amendment coverage); see Kathleen M. Sullivan, *Post-Liberal Judging: The Roles of Categorization and Balancing*, 63 U. COLO. L. REV. 293, 295 (1992). But see Joseph Blocher, *Categoricalism and Balancing in First and Second Amendment Analysis*, 84 N.Y.U. L. REV. 375, 397 (2009).

113. See *N.Y. Times Co. v. Sullivan*, 376 U.S. 254, 279-80 (1964); see also Meg Leta Ambrose & Ben M. Ambrose, *When Robots Lie: A Comparison of Auto-Defamation Law* (IEEE Workshop on Advanced Robotics and Its Social Impacts Working Paper, 2014), <https://ssrn.com/abstract=2746250> [<https://perma.cc/66WF-GQ5K>] (discussing the actual-malice standard).

114. See Massaro et al., *supra* note 86, at 2507; see also Brown, *supra* note 1, at 392, 399-401 (arguing that mens rea is the key challenge for regulating AI-generated speech under defamation law, and arguing that neither developers nor other corporate actors would show the requisite mens rea for defamation).

115. See Katherine Lee, Feder A. Cooper & James Grimmelman, *Talkin’ Bout AI Generation: Copyright and the Generative-AI Supply Chain* 4-5 (Sept. 21, 2023) (unpublished manuscript), <https://ssrn.com/abstract=4523551> [<https://perma.cc/MF79-YSQK>]; Jennifer Cobbe, Michael Veale & Jatinder Singh, *Understanding Accountability in Algorithmic Supply Chains* 1, 2023 Assoc. for Computing Mach. Conf. on Fairness, Accountability, & Transparency (May 17, 2023), <https://ssrn.com/abstract=4430778> [<https://perma.cc/8UWS-GYK5>]; see also Kaminski, *Regulating the Risks*, *supra* note 7, at 1368-69 (“Scholars note, too, the challenges raised by the emerging ecology of AI developers, companies, and users, with multiple actors and discrete steps and components, making it potentially hard to ascribe responsibility or identify a place of intervention for governance.”); Matthew U. Scherer, *Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies*, 29 HARV. J.L. & TECH. 353, 369-73

The Court’s selective focus on human speakers and human intent in First Amendment law is not going away anytime soon. In three recent cases, the Court doubled down on the centrality of (human) intent to historic exceptions to First Amendment protection.<sup>116</sup> In *Counterman v. Colorado*, the Court found that the state of Colorado could not use an objective reasonable-person standard for determining whether speech constituted unprotected true threats.<sup>117</sup> The majority held that the First Amendment requires some form of scienter for true threats, as it does for defamation and other categories of unprotected speech, in order to build a buffer around unprotected speech and prevent a chilling effect on (human) speakers.<sup>118</sup> A standard of recklessness, according to the Court, struck the right balance between competing concerns of chilling protected speech and making it harder to prosecute unprotected harmful threats.<sup>119</sup>

In another 2023 case, *United States v. Hansen*, the Court again addressed intent, analyzing the category of “speech integral to criminal conduct.”<sup>120</sup> The Court noted that both criminal solicitation and criminal facilitation contain a fairly strict intent requirement: “an intent to bring about a *particular* unlawful act.”<sup>121</sup> It is unlikely that an AI developer creating a general-purpose content-generating system would ever exhibit such intent.

While *Twitter v. Taamneh* is not, strictly speaking, a First Amendment case, the Court’s discussion of intermediary liability for online terrorist content again indicated its current emphasis on intent as a backstop for expressive values.<sup>122</sup> In *Taamneh*, the Court interpreted a federal statute that imposed liability for facilitating an act of terrorism and found that it required a level of scienter—“truly culpable conduct”—that did not encompass general-purpose online platforms

---

(2016) (describing the diffuseness and opacity of AI systems as “some of the most problematic features of AI” “from a regulatory standpoint”).

116. *Counterman*, 600 U.S. 66; *United States v. Hansen*, 599 U.S. 762 (2023); *Twitter, Inc. v. Taamneh*, 598 U.S. 471 (2023).

117. *Counterman*, 600 U.S. at 77-78. COLO. REV. STAT. § 18-3-602(1)(c) (2023) makes it unlawful to “[r]epeatedly . . . make[] any form of communication with another person . . . in a manner that would cause a reasonable person to suffer serious emotional distress and does cause that person . . . to suffer serious emotional distress.” There was no evidence that Counterman in fact followed or surveilled his target, singer C.W., so the entirety of the case turned on the Facebook messages he repeatedly sent to her. *Counterman*, 600 U.S. at 70-71, 71 n.1.

118. *Counterman*, 600 U.S. at 75-76.

119. *Id.* at 78-80.

120. 599 U.S. at 783; see also *supra* note 112 and accompanying text (discussing and distinguishing between categories of unprotected and uncovered speech).

121. *Hansen*, 599 U.S. at 771 (emphasis added).

122. *Twitter, Inc. v. Taamneh*, 598 U.S. 471, 489 (2023).

and probably would not encompass most AI developers or platforms hosting general-purpose AI systems.<sup>123</sup>

If courts hew to these kinds of strict intent requirements, then AI “speakers” would get off the hook where human speakers would not.<sup>124</sup> This would perversely incentivize more otherwise unlawful speech by AI systems – protecting more speech generated by AI than speech by actual humans. But as is often the case, this problem of First Amendment sensemaking of speech generated by AI systems is not an unsolvable one. It is not “disruption,” in the sense that the use of such systems has somehow broken the law.

There are multiple ways in which scholars, and courts, could intervene. They could intervene – and indeed have already been intervening – at the level of doctrine, arguing that there are more appropriate ways to think about duties of care and liability with respect to content-generating AI systems. For example, Nina Brown invokes products-liability law. Brown suggests that courts analyzing defamation claims arising from AI-generated content should look not to intent but to whether AI developers are making negligent design and development choices, such as training their systems on databases filled with inaccurate content or failing to adequately test their systems.<sup>125</sup> Eugene Volokh, too, invokes a products-liability approach to AI defamation claims and argues that developers should be subject to a “notice-and-blocking” framework, which we discuss further in the next Section.<sup>126</sup> Jane Bambauer explores negligence and duties of care, suggesting that depending on the specific facts of each case, the producers of AI systems might be analogized to doctors, lawyers, encyclopedia publishers, or even parents who negligently supervise their precocious AI children.<sup>127</sup>

Each of these scholarly moves shifts the legal analysis of AI systems from understanding such systems as *speakers* to trying to construe them as something else: products, encyclopedias, advisors, or precocious children. Within some of these more familiar analogies, we see hints, too, of something more interesting:

---

123. *Id.* at 506 (“The point of aiding and abetting is to impose liability on those who consciously and culpably participated in the tort at issue.”). *But see id.* at 502-03 (discussing situations where a platform might nonetheless be held liable).

124. *See* Massaro et al., *supra* note 86, at 2507 (noting it would “create[] a problem if the doctrine were to insulate AI speakers (but not human speakers) from liability because they lack provably culpable mental states”); *see also* Brown, *supra* note 1, at 401-03 (proposing that claims against AI systems be brought under products-liability law instead of defamation law because of the impossibility of proving the requisite mens rea element in a defamation claim against AI).

125. Brown, *supra* note 1, at 421-24.

126. Eugene Volokh, *Large Libel Models? Liability for AI Output*, 3 J. FREE SPEECH L. 489, 514-15, 522-26 (2023).

127. Bambauer, *supra* note 43, at 349-58.

the construction of the problem of AI systems as a problem of *speech at scale*,<sup>128</sup> or as *risky complex systems*,<sup>129</sup> or as a problem for *consumer protection*,<sup>130</sup> each of which we discuss more below. That is, shifting the construction of AI systems away from the constricting box of *speaker* with human intent opens room for other ways of making legal sense of AI systems.

That’s all nice in the rhetorical abstract. But does First Amendment law as it currently exists leave room for these and other shifts in legal sensemaking? The answer is a very strong maybe.

First, while the *Counterman* majority made some strong claims that all pure-speech exceptions require some showing of intent, Justice Barrett’s dissent pointed out that several existing categorical exceptions appear to have objective standards, if they have any standards at all.<sup>131</sup> Second, the *Counterman* majority outlined a recklessness standard that at times appears to leave room for liability for AI developers when they knowingly ignore substantial risks, and could maybe serve as a template for a *sui generis* scienter standard for AI-generated content in other categories of unprotected speech.<sup>132</sup>

Third, both *Counterman* and *Taamneh* described balancing tests or sliding scales of scienter, potentially leaving room for striking a different balance for AI-

---

128. Cf. Volokh, *supra* note 126, at 519 (noting that an AI notice-and-takedown regime would have to “operate at scale with vast numbers of requests being submitted every day,” similar to the “DMCA copyright and trademark regime”).

129. Cf. Brown, *supra* note 1, at 392-97 (detailing the complex mechanisms by which AI chatbots generate output).

130. Cf. Bambauer, *supra* note 43, at 362 (predicting that “consumer advocates” will put pressure on courts to “impose duties of care” on AI developers).

131. *Counterman v. Colorado*, 600 U.S. 66, 81 (2023); *id.* at 108-14 (Barrett, J., dissenting) (considering fighting words, commercial speech, obscenity, defamation of private figures, and incitement).

132. *Id.* at 69 (majority opinion) (“The State must show that the defendant consciously disregarded a substantial risk that his communications would be viewed as threatening violence. The State need not prove any more demanding form of subjective intent to threaten another.”) A developer might be shown to “consciously disregard a substantial risk” in how she designs the system. *Id.* at 79 (quoting *Voisine v. United States*, 579 U.S. 686, 691 (2016)); *see id.* (“That standard involves insufficient concern with risk, rather than awareness of impending harm.”) This arguably imports a risk standard into intent. *But see id.* (“But still, recklessness is morally culpable conduct, involving a ‘deliberate decision to endanger another.’” (quoting *Voisine v. United States*, 579 U.S. 686, 694 (2016))). In the threats context, it means that a speaker is aware “that others could regard his statements as’ threatening violence and ‘delivers them anyway.’” *Id.* (quoting *Elonis v. United States*, 575 U.S. 723, 746 (2015) (Alito, J., concurring in part and dissenting in part)). The latter constraints – that recklessness involves having awareness about particular statements and delivering them anyway – suggest that this would not include behavior by a developer or programmer.

generated speech than human speech.<sup>133</sup> The Court in *Counterman* balanced the harm of not prosecuting unprotected speech against the harm of a chilling effect on human speakers, declaring that a recklessness standard strikes the right balance.<sup>134</sup> However, AI systems do not themselves experience a chilling effect. AI-generated speech, like commercial speech, could be characterized as inherently harder than ordinary speech.<sup>135</sup> There are plenty of economic incentives for AI-speech generation, just as there are with commercial speech, and thus less of a fear of a “chilling effect” on the companies that develop and distribute AI systems. If courts show a willingness to rebalance harms and chilling effects, this might result either in a version of recklessness that could be applied to AI developers across the board, or in the sort of objective product-liability-like standards Brown or Volokh discuss.

There is also room for adaptation at the level of First Amendment *values*. As one of us has suggested with co-authors elsewhere, courts could justify a different doctrinal standard for AI-generated speech by focusing on harms to the human listeners on whose backs coverage of AI-generated content is theoretically justified in the first place.<sup>136</sup> When lawmakers restrict AI speech in the name of protecting the interests of a human audience, courts might mimic areas of law similarly justified by listeners’ interests such as commercial-speech doctrine or professional-speech doctrine by applying less exacting scrutiny and upholding

---

133. The Court in *Twitter, Inc. v. Taamneh* articulates that the scienter requirement for facilitating an unlawful act is a movable standard that shifts based on how much knowledge or how much substantial assistance has been provided. The majority explains that “courts often viewed [the] twin requirements [of knowledge and substantial assistance] as working in tandem, with a lesser showing of one demanding a greater showing of the other.” 598 U.S. 471, 491-92 (2023). The key language for AI-generated content is that “if the assistance were direct and extraordinary, then a court might more readily infer conscious participation in the underlying tort.” *Id.* at 492. Where the internet platforms in *Taamneh* were many steps removed from the terrorist attack at the Reina nightclub in Turkey, they presumably would be more directly responsible for speech produced by their own algorithms. An AI developer might not have knowledge of the exact speech an Large Language Model might produce, but arguably did provide substantial assistance in producing it.

134. For example, the law must require specific intent for incitement, which often sits perilously close to political advocacy, while “the speech on the other side of the true-threats boundary line—as compared with the advocacy addressed in our incitement decisions—is neither so central to the theory of the First Amendment nor so vulnerable to government prosecutions.” *Counterman*, 600 U.S. at 81.

135. Eugene Volokh, Mark Lemley & Peter Henderson, *Freedom of Speech and AI Output*, 3 J. FREE SPEECH L. 651, 657 n.15 (2023). But unlike commercial speech, where the Supreme Court has justified its permissiveness toward the regulation of deception in part by pointing out that commercial speakers have every reason to know what the objective truth actually is (“verifiability”), AI-generated speech may be really hard, if not impossible, for companies to fact-check in real time.

136. Massaro et al., *supra* note 86, at 2491.

regulations that aim to protect human listeners by preventing, for example, deception.<sup>137</sup> How successful this argument might be with the current Court is another matter.

In these ways, law makes room for alternate or changing sensemaking. The law uses language, and language often leaves room for interpretation. The law is justified by values, and values might dictate different results, or might themselves change over time. The law is also a social system set up to handle evolving social facts. Nonetheless, the law is also made up of humans, and humans and their institutions can be deeply political.

This brings us back to *institutions*. The institutions enforcing the First Amendment are primarily the courts, especially the Supreme Court.<sup>138</sup> The Court has recently been described as pursuing First Amendment Lochnerism, using the blunt force of the Constitution to overturn a host of regulations.<sup>139</sup> The Court has been especially solicitous of religious and politically conservative speakers, potentially constraining regulation in all spaces, including those occupied by AI.<sup>140</sup> The issue of platform liability has itself become highly politicized, with the *NetChoice* cases addressing Florida and Texas laws that were enacted to require purportedly liberal internet platforms to carry conservative speech making their way from the Fifth and Eleventh Circuits up to the Court.<sup>141</sup> All this is to say: much of what is in motion on the legal side is spurred more by a shifting political landscape than by the introduction of purportedly fast-paced new technology.

In summary, the First Amendment constructs the puzzle of AI-generated content as largely a problem of intent and individual liability, or lack thereof. This may limit regulatory interventions and, at first glance, back lawmakers into the particular legal struggles of the “substitution” approach: AI does not itself have intent, and responsibility may be extremely hard to determine or apply, given the black-box nature of AI systems, the complexity of human-machine

---

137. *Id.* at 2519.

138. There are other institutions as well, such as Congress. There are also plenty of institutions that make free-speech policy. See, e.g., Mike Ananny, *Probably Speech, Maybe Free: Toward a Probabilistic Understanding of Online Expression and Platform Governance*, KNIGHT FIRST AMEND. INST. (Aug. 21, 2019), <https://knightcolumbia.org/content/probably-speech-maybe-free-toward-a-probabilistic-understanding-of-online-expression-and-platform-governance> [<https://perma.cc/ALT8-MPGL>].

139. Amanda Shanor, *The New Lochner*, 2016 WIS. L. REV. 133, 133.

140. Rebecca Aviel, Margot E. Kaminski, Toni M. Massaro & Andrew Keane Woods, *From Gods to Google: How Religious Speech Cases May Fortify the Deregulatory First Amendment* 3-4 (Oct. 2023) (unpublished manuscript) (on file with authors).

141. Blake E. Reid, *Uncommon Carriage*, 76 STAN. L. REV. (forthcoming 2024), <https://ssrn.com/abstract=4181948> [<https://perma.cc/D9ZN-7VKX>].

systems, and the complexity of AI supply chains.<sup>142</sup> But as we will subsequently explain, other approaches to regulating AI-generated speech touch little if at all upon intent, opening up the possibility of broad regulatory interventions. Scholars of First Amendment law might do well to think about if or how that doctrine could make space for these other approaches, including regulation.<sup>143</sup>

## 2. *AI Speech as a Problem of Speech at Scale*

If the First Amendment focuses on *speech* by individual rights-bearing *speakers*, the closely adjacent field of content-moderation law instead legally constructs AI-generated content as a problem of *speech at speed and at scale*.<sup>144</sup> (To be clear: content-moderation law may itself be backstopped by the First Amendment; that is not our focus here.<sup>145</sup>) For our purposes, we refer to this as “speech-at-scale” construction. It differs from the First Amendment construction not just in the substantive law — statutory schemes rather than constitutional interpretation — but also in its objects, values, and institutions.

Here is the argument for speech-at-scale construction: AI systems will quickly and at ostensibly low cost produce an enormous quantity of content. Some of that content will be high-value content (e.g., political speech), some of it will be low-value but legal content (e.g., hate speech and porn),<sup>146</sup> and some of it will be illegal content (e.g., defamation). That content will, like human-

142. See Lee et al., *supra* note 115, at 1-2; Cobbe et al., *supra* note 115, at 1.

143. We address this a bit in Part III.

144. Kaminski & Urban, *supra* note 7, at 1964 (comparing privatized processes at speed and at scale). See also Mark P. McKenna & Woodrow Hartzog, *Taking Scale Seriously in Robotics and A.I. Law*, B.U. SCH. L. 3 (2023), <https://www.bu.edu/law/files/2023/09/McKenna-Hartzog-Scale-v7.pdf> [<https://perma.cc/3S8D-K754>] (arguing that a shift in scale can mean a difference not just in quantity but in kind); Balkin, *supra* note 42, at 1238. Jack M. Balkin refers to this as “the Algorithmic Society.” Balkin, *supra* note 42, at 1236. We differ in that we distinguish speech-at-scale construction from risky systems construction (what Balkin terms hygienic, epidemiological, environmental, and probabilistic (HEEP) construction). See discussion *infra* Part III.

145. For discussion of whether the First Amendment requires shielding intermediaries from liability, see, for example, Felix T. Wu, *Collateral Censorship and the Limits of Intermediary Immunity*, 87 NOTRE DAME L. REV. 293 (2011); Christina Mulligan, *Technological Intermediaries and Freedom of the Press*, 66 SMU L. REV. 157 (2013); Benjamin, *supra* note 84; Frank Pasquale & Oren Bracha, *Federal Search Commission? Access, Fairness and Accountability in the Law of Search*, 93 CORNELL L. REV. 1149 (proposing that the First Amendment does not prevent the regulation of search engines); Grimmelmann, *supra* note 6, at 893-911 (offering a middle way based in listener rights).

146. What Daphne Keller calls “lawful but awful” speech. Daphne Keller, *The EU’s New Digital Services Act and the Rest of the World*, VERFASSUNGSBLOG (Nov. 7, 2022), <https://verfassungsblog.de/dsa-rest-of-world> [<https://perma.cc/TNB2-SLJE>].

generated content, be distributed both to individual users as the result of individual queries and to the general public online. The puzzle of AI speech at scale both involves and resembles existing discussions of online-platform liability and content moderation.<sup>147</sup> Online platforms will have to moderate AI-generated speech, just as they do human speech. And AI systems may trigger questions about whether they should themselves be treated as speakers or information intermediaries.

Unlike First Amendment construction, which frames AI as an individual speaker and backs lawyers into a substitution problem focused on speaker intent, speech-at-scale construction focuses less on individual speakers and more on the problems of fairly moderating quickly distributed networked content at scale. Thus, the *objects* of speech-at-scale construction are typically communications platforms; content at scale, rather than speech from a rights-bearing individual; and individual platform “users,” rather than speakers in the Constitutional sense.

Speech-at-scale construction asks what kind of law establishes the best incentives and processes for ensuring that content at scale is on balance not too harmful, while also avoiding an undue chilling effect on both speech and innovation (whatever that means).<sup>148</sup> The *value* central to AI as speech at scale is something like efficient fairness. And its sensemaking *institutions* are markedly different as well. Legislators serve as the primary, albeit not exclusive, policy-makers—in many instances deferring to private platforms’ decision-making. Courts serve as the primary institution of interpretation of legislation.

While First Amendment rights and values exist in the backdrop of speech-at-scale construction, the central policy tools look little to nothing like First Amendment law. Under speech-at-scale construction, the law of AI-generated speech online is not much different from the law of human-generated speech online. Both are the law of intermediary liability and content moderation.

There have been, until very recently, two approaches to intermediary liability and content moderation in the United States. The first approach, exemplified by Section 230 of the Communications and Decency Act (CDA 230), immunizes online intermediaries from liability for either hosting or removing most user-generated content.<sup>149</sup> The much-debated rationale behind CDA 230 is that if online platforms are on the hook for user-generated content, they will be

---

147. See Grimmelmann, *supra* note 6, at 944-45; Volokh, *supra* note 126, at 491-93; Bambauer & Surdeanu, *supra* note 40, at 375-86 (taking different sides as to whether Section 230 of the Communications and Decency Act applies directly to AI-generated content).

148. See, e.g., *Zeran v. Am. Online, Inc.*, 129 F.3d 327 (4th Cir. 1997).

149. See Eric Goldman, *Why Section 230 Is Better than the First Amendment*, 95 NOTRE DAME L. REV. REFLECTION 33 (2019); JEFF KOSSEFF, *THE TWENTY-SIX WORDS THAT CREATED THE INTERNET* (2019); Danielle K. Citron & Mary Anne Franks, *The Internet as a Speech Machine and Other Myths Confounding Section 230 Reform*, 2020 U. CHI. LEGAL F. 45.



incentivized to overcensor, and innovation and free speech will suffer.<sup>150</sup> This rationale is paired with recognizing platform hosts' need to moderate content to create domains in which people want to participate, without being held liable as publishers of any legally troublesome user-posted content. In the legal gap created by CDA 230, private governance of online speech has flourished, with companies "voluntarily" creating extensive private content-moderation regimes largely to ward off government scrutiny or additional law, or in response to collateral governmental pressures.<sup>151</sup>

The second approach to online liability and content moderation in the United States is the "notice-and-takedown" regime established in the Digital Millennium Copyright Act (DMCA), which creates a different system for copyright law: one that directly and in great detail dictates how companies should moderate online content.<sup>152</sup> The DMCA regime is granular, highly procedural, and filled with minutiae aimed at balancing the legitimate interests of copyright holders and the potential chilling effects on those who post online content. The DMCA regime divides intermediaries into several different classes, requiring platforms that cache, host, or link to content to respond to a formal notification of copyright infringement by taking down the identified content.<sup>153</sup> The user who posted said content has the opportunity to state that it is not infringing and request that it be reinstated (a "put back").<sup>154</sup> In recent years, companies such as Google have created a "para DMCA" or "DMCA-plus" regime that bypasses the DMCA's requirements and safeguards and attempts to automate and monetize

---

150. Anupam Chander & Uyen P. Le, *Free Speech* 501 (U.C. Davis Legal Stud. Rsch. Paper No. 351, Sept. 2013), <https://ssrn.com/abstract=2320124> [<https://perma.cc/J99K-76QH>]. *But see* Danielle Keats Citron & Benjamin Wittes, *The Internet Will Not Break: Denying Bad Samaritans Internet Immunity*, 86 *FORDHAM L. REV.* 401 (2017).

151. *See, e.g.*, Kate Klonick, *The New Governors: The People, Rules, and Processes Governing Online Speech*, 131 *HARV. L. REV.* 1598, 1603-04 (2018); Hannah Bloch-Wehba, *Global Platform Governance: Private Power in the Shadow of the State*, 72 *SMU L. REV.* 27, 37-39 (2019); Jack M. Balkin, *Free Speech is a Triangle*, 118 *COLUM. L. REV.* 2011, 2021-25 (2018); Genevieve Lakier, *Jawboning at a Problem of Constitutional Evasion*, *KNIGHT FIRST AMEND. INST.* (Oct. 13, 2023), <https://knightcolumbia.org/blog/jawboning-as-a-problem-of-constitutional-evasion-or-why-the-significant-encouragement-test-is-not-so-bad> [<https://perma.cc/QW3A-64PD>].

152. Digital Millennium Copyright Act (DMCA) § 202, 17 U.S.C. § 512 (2018).

153. *See, e.g.*, Jennifer M. Urban, Joe Karaganis & Brianna L. Schofield, *Notice and Takedown in Everyday Practice* 1 (UC Berkley Pub. L. Rsch. Paper No. 2755628, 2017), <https://ssrn.com/abstract=2755628> [<https://perma.cc/DU6P-VUZN>].

154. *See* Wendy Seltzer, *Free Speech Unmoored in Copyright's Safe Harbor: Chilling Effects of the DMCA on the First Amendment*, 24 *HARV. J.L. & TECH.* 171, 187-200 (2010) (detailing the steps in the DMCA takedown and appeal system and analyzing its chilling effects).

copyright content moderation.<sup>155</sup> The takedown system is designed to achieve the same speech-at-scale value of efficient fairness, to avoid overburdening platforms with monitoring efforts, but also to address harms to copyright holders.

Until only very recently, these two approaches represented the two main ways of legally constructing speech at scale in the United States: the hands-off privatized CDA 230 approach and the highly proceduralized DMCA approach. However, in the past few years, several states (e.g., Florida and Texas) have imposed detailed transparency and must-carry requirements on online intermediaries that would interfere with privatized content-moderation policies—what might be characterized as a developing “third way” of platform liability.<sup>156</sup> Aspects of these laws are soon to be challenged at the Supreme Court in *NetChoice, LLC v. Paxton*.<sup>157</sup>

There is no perfect content-moderation system. Under any approach, some legal content will be blocked, and some illegal content will slip through the guardrails. The policy *value* is thus to foster an overall system that works well enough and appears to be fair and effective—hence, the focus on procedure and transparency.<sup>158</sup> In an ideal world, current content-moderation systems are calibrated to target undesired and sometimes illegal content efficiently and accurately.

There is also no inevitable content-moderation system. Content-moderation policy is the result of politics, and politics play out through *institutions* and people. As one of us has written:

Other countries have extended liability to platforms once the operator has knowledge of legally actionable content. Platforms are not

---

155. See Urban et al., *supra* note 153, at 57–59; see also Matthew Sag, *Internet Safe Harbors and the Transformation of Copyright Law*, 93 NOTRE DAME L. REV. 499, 500 (2017) (describing these as “DMCA-plus”).

156. Reid, *supra* note 141 (manuscript at 1); see also Daphne Keller, *The EU’s New Digital Services Act and the Rest of the World*, VERFASSUNGSBLOG (Nov. 7, 2022), <https://verfassungsblog.de/dsa-rest-of-world> [<https://perma.cc/TNB2-SLJE>] (“The resulting litigation has sent an epoch-defining First Amendment question hurtling toward America’s newly reckless, conservative-dominated Supreme Court. Other countries’ incremental creep toward carriage mandates for major platforms may abruptly be bypassed by tremendous changes in the U.S.”).

157. See 49 F.4th 439 (5th Cir. 2022), *cert. granted in part*, No. 22-555, 2023 WL 6319650 (Sept. 29, 2023).

158. See, e.g., MARTIN HUSOVEC & IRENE ROCHE LAGUNA, *Digital Services Act: A Short Primer*, in PRINCIPLES OF THE DIGITAL SERVICES ACT (forthcoming 2024) (manuscript at 1), <https://ssrn.com/abstract=4153796> [<https://perma.cc/732M-VFLM>] (emphasizing “fairness, trust, and safety”).

considered neutral, automated systems and accountability is effectuated through human involvement and design choices.<sup>159</sup>

While CDA 230 may feel inevitable in the United States, countries in the European Union have long done content moderation differently. In Europe, until very recently, the E-Commerce Directive established a DMCA-like notice-and-takedown regime that applied to *all* potential online liability, not just to alleged copyright infringement.<sup>160</sup> The system is calibrated differently. Platforms have higher monitoring burdens, but they publicly host less harmful content as a consequence. As we discuss in greater detail in the next Section, the recent Digital Services Act (DSA) supplants the old regime while also retaining its basic notice-and-takedown framework.<sup>161</sup>

For the most part, then, speech-at-scale construction will treat AI-generated content just as it treats human-generated content, as a problem of content at scale, solved by systems aiming toward efficient fairness. Perhaps policymakers will intervene in any number of ways: to require that AI-generated content be identified as such,<sup>162</sup> or to recalibrate notice-and-takedown policies to adapt to faster content generation by AI systems. But these changes are neither inevitable nor required by the development of AI technology. The law may adapt, but once again, it does not break.

Then there is what we call the legal construction of AI systems as “splatforms” (speaker-platform hybrids). This analysis asks whether companies that produce AI-generated content themselves might be directly protected from liability by CDA 230. Scholars are split on this question.<sup>163</sup> AI systems arguably loosely resemble search engines in that they both simultaneously create content

---

159. Jones, *supra* note 5, at 262.

160. See Directive 2000/31/EC of the European Parliament and of the Council of 8 June 2000, on Certain Legal Aspects of Information Society Services, in Particular Electronic Commerce, in the Internal Market (Directive on Electronic Commerce), 2000 O.J. (L 178) 15.

161. The Digital Services Act (DSA) adds on increased process protections for individuals whose content has been taken down, including the possibility of resorting to Alternative Dispute Resolution (ADR). It also requires transparency reports that show how the system is or is not working, which is long sought after by researchers in this space. See Daphne Keller, *What Does the DSA Say?*, CTR. INTERNET & SOC’Y (Apr. 25, 2022, 7:04 PM), <https://cyberlaw.stanford.edu/blog/2022/04/what-does-dsa-say-o> [<https://perma.cc/7NWJ-KLS5>]; Urban et al., *supra* note 153, at 31-34 (citing transparency reports Google and Reddit).

162. See, e.g., Madeline Lamo & Ryan Calo, *Regulating Bot Speech*, 66 UCLA L. REV. 988 (2019).

163. Compare Volokh, *supra* note 126, at 494 (arguing that “§ 230 likely doesn’t provide AI companies with immunity for material composed and communicated by their AI programs”), and Matt Perault, *Section 230 Won’t Protect ChatGPT*, J. FREE SPEECH L. 363, 364 (2023) (arguing that “current speech liability protections do not apply to certain generative AI use cases”), with Bambauer & Surdeanu, *supra* note 40, at 376 (arguing that § 230 “is likely to protect the creators, distributors, and hosts of online services that include ChatGPT in many cases”).

(as speakers or platform users) and point users to information they did not generate themselves (as platforms or conduits).<sup>164</sup> The majority of scholars have concluded that CDA 230 does not protect companies that offer algorithms that directly produce content, characterizing those algorithms as more like users/speakers than platforms/search engines. But Derek E. Bambauer and Mihai Surdeanu argue, in an interesting moment of legal construction, that the answer will be highly fact dependent.<sup>165</sup> They argue that LLMs can be understood under CDA 230 as both a platform (for existing content within training databases) and a user (for newly generated content).<sup>166</sup> The user of LLMs, they argue, is even more actively involved in the creation of content than the user of, say, a search engine.<sup>167</sup> This characterization of LLMs recalls James Grimmelmann's work on search engines as "speech engines," wherein he highlighted the ways in which search engines operate as both conduits and speakers.<sup>168</sup> We take no position here, but note the ways in which the platform debate hinges not just on attributes of the technology but on existing categories in the law.

As Volokh's recent work on generative AI systems illustrates, speech-at-scale framing can also be deployed in discussions of the direct, rather than intermediary, liability of the companies that develop and distribute such systems.<sup>169</sup> Volokh reasons that OpenAI and Google are not covered by CDA 230, as the output of their AI systems is content their software has produced, not hosted content generated by others.<sup>170</sup> He nonetheless incorporates speech-at-scale reasoning into what is ostensibly First Amendment-driven analysis.<sup>171</sup> Volokh concludes that companies that create and distribute AI systems should be required under libel law to establish a "notice-and-blocking" regime similar to the DMCA—and that a failure to do so would constitute negligence and even "actual malice."<sup>172</sup>

Volokh thus shifts from the substitution construction of First Amendment law (what, exactly, does it mean to ask if AI systems show intent of "actual malice"?) to speech-at-scale construction of content moderation (how does one

---

164. See Lee et al., *supra* note 115, at 91; Bambauer & Surdeanu, *supra* note 40, at 384.

165. Bambauer & Surdeanu, *supra* note 40, at 386-87.

166. *Id.* at 384.

167. *Id.* at 377-78.

168. Grimmelmann, *supra* note 6, at 889.

169. See Volokh, *supra* note 126, at 493.

170. *Id.* But see Bambauer & Surdeanu, *supra* note 40, at 382 (arguing that OpenAI is the host, not the speaker, of information).

171. Volokh, *supra* note 126, at 542-44 (discussing policy reasons for using notice-and-blocking, including balancing harm to innovation and the listener rights of AI-system users with speech harms to affected libeled individuals).

172. *Id.* at 519.

balance the benefits of AI system output against their harms, through an efficient but procedurally fair system?).

To summarize, legally constructing AI-generated speech as a problem of speech at scale invokes particular legal practices, tools, and policy conversations. In the context of content moderation, it treats AI speech more or less as it does human-generated speech, with the *object* of regulation being networks and networked content: the large corporate entities that serve as online intermediaries or platforms, and their users. Speech-at-scale construction attempts to handle content moderation either by handing it off to private governance or by treating it as a problem of properly calibrating ex post individual processes and systemic transparency. The *institutions* relevant to the private management of these network systems include courts occasionally maintaining otherwise private management, though state legislators have started to try to restructure the network or recalibrate the balance in service of shifting values. Constructing AI speech as speech at scale is less concerned with respecting individual speakers' rights or policing the substantive boundaries of protected and unprotected speech. Its primary *value* is calibrating efficient fairness between platforms and users, with some free-speech and innovation values sprinkled in.

### 3. *AI as a Risky Complex System*

A third way to construct speech by AI systems is through risk regulation.<sup>173</sup> This approach constructs AI not as a speaker, nor as the generator of speech at scale, but as a *risky complex system*. The *value* this approach seeks to achieve is the ex ante mitigation of widespread, significant, and typically quantifiable harms: some quite concrete, and others harder to measure and more contestable.

The risky-complex-system approach appears in at least two places. One construction, currently largely taking place in Europe, is highly regulatory in nature. It uses regulatory tools borrowed from other areas of the law, such as data-protection law, environmental law, and product-safety regulation. That approach focuses primarily on ex ante risk assessment and risk mitigation, either in soft law or with centralized agencies as *institutions* exerting light-touch to moderate top-down control.<sup>174</sup> Another construction, which may be more familiar to U.S. readers, constructs AI as risky complex systems through American tort law.

The *object* of a risky-complex-system construction is not just the AI system itself but also the humans and organizations around it. This construction is more like the AI-human systems of content moderation than the individual AI speaker

---

173. See generally Kaminski, *Regulating the Risks*, *supra* note 7 (discussing the advantages, drawbacks, and different models of AI risk regulation).

174. *Id.* at 1379-80.

constructed through First Amendment doctrine, but unique in important ways that present opportunities for powerful forms of intervention.<sup>175</sup> To construct AI as a risky complex system, one shifts the conception of AI away from viral social media posts and toward nuclear plants, cars, and medical devices, where regulation aims to mitigate the risk of failure and bolster resilience through a number of regulatory tools. Whether through mandated design specs or more general recording and reporting requirements, legal construction of risky systems typically accepts that such complex systems will inevitably crash, but, with varying degrees of firmness, tries to task the developers and users of such systems with lowering risks to the population at large.

The risky-complex-system construction is increasingly popular, including outside of the context of AI-generated speech. Light-touch risk regulation has rapidly proven to be the go-to tool for governing AI systems in general.<sup>176</sup> For example, both the EU's General Data Protection Regulation (GDPR) and the EU AI Act use *ex ante* risk assessment and risk mitigation to address perceived problems with algorithms and AI.<sup>177</sup> Each law requires the developers of "risky" AI systems to conduct an impact assessment (GDPR) or conformity assessment (AI Act) and to mitigate discovered risks. While neither the GDPR nor AI Act are, strictly speaking, targeted at AI systems as speakers, they both target AI as risky systems: the GDPR applies to systems developed using personal data, and the AI Act applies to systems it categorizes as "high risk" or "unacceptably risky."<sup>178</sup> The *value* this approach seeks to achieve is the prevention of

---

175. For examples of complex systems in the copyright context, see Maayan Perel & Niva Elkin-Koren, *Accountability in Algorithmic Copyright Enforcement*, 19 STAN. TECH. L. REV. 473 (2016); Maayan Perel & Niva Elkin-Koren, *Black Box Tinkering: Beyond Disclosure in Algorithmic Enforcement*, 69 FLA. L. REV. 181 (2017).

176. See Kaminski, *Regulating the Risks*, *supra* note 7, at 1408; Margot E. Kaminski, *The Developing Law of AI Regulation: A Turn to Risk Regulation*, LAWFARE (Apr. 21, 2023, 1:23 PM), <https://www.lawfaremedia.org/article/the-developing-law-of-ai-regulation-a-turn-to-risk-regulation> [<https://perma.cc/575J-TBAR>].

177. See Kaminski, *Regulating the Risks*, *supra* note 7, at 1408.

178. The European Commission is to determine which systems are high-risk systems. Currently, the Commission has proposed high-risk classifications for systems including the use of AI in law enforcement, border control, employment, and other circumstances. See *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts*, annex 3, COM (2021) 206 final (Apr. 21, 2021). Unacceptably risky systems are also defined in Title II of the Act itself. *Id.* at tit. II, art. 5. As of the 2021 draft of the Act, unacceptably risky systems include some forms of manipulation that may be executed through speech, for example, the use of AI in advertising:

[T]he placing on the market, putting into service or use of an AI system that deploys subliminal techniques beyond a person's consciousness in order to materially distort a person's behaviour in a manner that causes or is likely to cause that person

widespread, significant, and quantifiable damage. Many other regulations and guidelines center around risk regulation, including: the National Telecommunications and Information Administration AI Risk Management Framework, the Singapore Model AI Governance Framework, and several proposed U.S. laws both at the state and federal level.<sup>179</sup>

A similar move to risk regulation is now occurring in the context of speech and content moderation. There has been a recent shift in legal scholarship from talking about individual First Amendment rights, or intermediary liability and efficient process, to talking about speech systems and population-level risks.<sup>180</sup> This kind of sensemaking funnels lawmakers to a different kind of regulatory instrument: constructing the old online speech-at-scale problem as a problem of risk regulation. It moves the conversation from being about process or substantive individual rights to one about designing a regulatory system that involves *ex ante* mitigation and *ex post* resilience, and often centralized regulators.<sup>181</sup> This approach shifts the object of study to system design and the institutions to administrative agencies and potentially state courts.

This move to risk regulation of online content systems is echoed and implemented in the EU's new DSA. The DSA in part preserves the EU's historic 2000s framework for intermediary liability discussed above; that is, it constructs online content as a problem of speech at scale, focusing on due-process protections and systemic transparency. But the DSA also shifts toward legally constructing

---

or another person physical or psychological harm; (b) the placing on the market, putting into service or use of an AI system that exploits any of the vulnerabilities of a specific group of persons due to their age, physical or mental disability, in order to materially distort the behaviour of a person pertaining to that group in a manner that causes or is likely to cause that person or another person physical or psychological harm . . . .

*Id.* In initial drafts of the AI Act, few if any AI “speakers” clearly fell into high-risk categories. However, the European Parliament’s most recent version of the Act added “AI systems intended to be used by social media platforms that have been designated as very large online platforms” and “AI systems intended to be used for influencing the outcome of an election or referendum or the voting behaviour of natural persons.” *See id.* at annex 3.

179. *Id.*; Kaminski, *supra* note 176 (discussing the cyclical and complimentary relationship between risk regulation and tort law).

180. *See, e.g.,* douek, *supra* note 42, at 531-32; Balkin, *supra* note 42, at 1246-47 (“Because of the sheer scale and speed of digital communication, many commentators – and digital companies themselves – have begun to use a different language and a distinctive set of metaphors to describe the effects of online speech and the goals of platform governance. This way of thinking is hygienic, epidemiological, environmental, and probabilistic (HEEP).”); Massaro et al., *supra* note 86, at 1183; COHEN, *supra* note 8, at 242, 246; Ananny, *supra* note 138. *See generally* Omri Ben-Shahar, *Data Pollution* (Dec. 18, 2020), <https://ssrn.com/abstract=3191231> [<https://perma.cc/8TTD-BQZR>] (developing the “data pollution” framework to counter the notion that the injuries from digital-data enterprise are solely private).

181. *See* sources cited *supra* note 180.

content moderation as a problem of large-scale risky systems. The DSA establishes for the first time risk-assessment and mitigation requirements for Very Large Online Platforms and Very Large Online Search Engines.<sup>182</sup> The DSA’s legal toolkit includes: formal risk assessments, risk-mitigation plans, third-party audits,<sup>183</sup> the appointment of compliance officers,<sup>184</sup> oversight by accredited researchers who get access to internal data,<sup>185</sup> and a slew of public transparency measures, including publicly available repositories of information about advertisements, and public transparency reports that include information about risk assessments, mitigation, and audits.<sup>186</sup> The DSA’s risk regulation involves ex ante intervention, including in product design.<sup>187</sup>

Once again, this is an example of how the law does not just react to technology. The EU could have stuck to speech-at-scale construction for online content moderation, pursuing the value of efficient fairness. Instead, it returned to the problem and supplemented speech-at-scale construction with a risky-complex-systems approach.

Stateside, there is no centralized risk regulation—yet. But scholars have pushed to establish risk mitigation for AI systems by making sense of them through tort law. One could imagine courts assessing tort claims eventually establishing a requirement of ex ante risk regulation for libelous AI systems, just as tort litigation eventually led to risk-mitigation requirements for car manufacturers.<sup>188</sup> For example, Bryan H. Choi details the long history of courts

---

182. HUSOVEC & LAGUNA, *supra* note 158, at 4. (“The real novelty of the DSA lies in the creation of fully-fledged tiers of due diligence obligations. Instead of ‘liability’[] the DSA insists on the ‘responsibility’[] that grows with the service’s size and societal impact.”); *id.* at 11 (illustrating the responsibilities of different intermediaries). See also Daphne Keller, *What Does the DSA Say*, STAN. CTR. INTERNET & SOC’Y BLOG (Apr. 25, 2022), <https://cyberlaw.stanford.edu/blog/2022/04/what-does-dsa-say-o> [<https://perma.cc/7NWJ-KLS5>]. Beyond its risk regulation framework for Very Large Online Platforms and Very Large Online Search Engines, the DSA also requires all platforms to indicate how and when they use automation, including “a legal requirement that they assess the impact on human rights of their content-related decision-making.” HUSOVEC & LAGUNA, *supra* note 158, at 4.

183. Regulation 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services and amending Directive 2000/31/EC (Digital Services Act), art. 34-37, 2022 O.J. (L 277) 1, 64-69.

184. Digital Services Act, art. 41.

185. Digital Services Act, art. 40.

186. Digital Services Act, arts. 42, 44.

187. Digital Services Act, art. 16.

188. See Bryan H. Choi, *Crashworthy Code*, 94 WASH. L. REV. 39, 87 (2019) (arguing for a change in focus from “prevention” to “mitigation” of code crashes by analogizing to tort liability for manual car accidents in the 1960s); Kaminski, *supra* note 176; Volokh, *supra* note 126, at 514-18 (discussing how the responsibility elements of libel law could be used to shoehorn a notice-and-blocking requirement into libel law).



dismissing claims of buggy software, because they rarely involve a physical injury and it is difficult to determine causality in complex computational systems.<sup>189</sup> Choi explains that even if a physical injury is presented, software errors “cannot be reasonably prevented via *ex ante* design or *ex post* testing.”<sup>190</sup> Like software before it, AI systems of today certainly cannot be made error proof. Choi argues that we should shift from a prevention to mitigation pathology, arguing for “crashworthy code” that would create a duty to minimize injuries when the inevitable error occurs.<sup>191</sup>

Nina Brown, too, discusses risk regulation through tort law. As discussed, Brown analyzes the specific issue of chatbots generating defamatory content.<sup>192</sup> After recounting the challenges AI chatbots present to assigning liability under traditional defamation law, namely the tenuous link between developers and the actual defamatory statement and the difficulty of proving the requisite mental state, she suggests products-liability law as an alternative.<sup>193</sup> For defective design, Brown would construct AI as a consumer technology with expected uses and require that developers mitigate foreseeable harms.<sup>194</sup>

Calls to treat software as a product for products-liability purposes stretch back at least to the 1980s.<sup>195</sup> The reasons those calls have been unsuccessful have little to do with the difference between software from decades ago and AI today. Nonetheless, a strong sense of novelty, societal pressure, and institutional drive by the right actors, may yet bring this construction to the fore.

The risky-complex-systems construction thus brings different regulatory tools to the table. It expands regulatory imagination—or, at least, moves to

---

189. Choi, *supra* note 188, at 69-70.

190. *Id.* at 44.

191. Choi, *supra* note 188, at 86-100. Choi uses the important shift that occurred in automobile design and litigation to situate his argument. The courts struggled to distribute liability effectively while the pathology focused on preventing crashes, but after the shift to mitigation, largely pushed by Ralph Nader, the approach to car safety changed across society including car manufacturers and courtrooms.

192. *See generally* Brown, *supra* note 1 (evaluating the challenges of applying defamation law to chatbots).

193. *Id.* at 401-03.

194. *Id.* at 421-24.

195. Michael C. Gemignani, *Product Liability and Software*, 8 RUTGERS COMPUT. & TECH. L.J. 173, 203 (1981); Nancy L. Birnbaum, *Strict Products Liability and Computer Software*, 8 COMPUT. L.J. 135, 144 (1988); Donald R. Ballman, *Software Tort: Evaluating Software Harm by Duty of Function and Form*, 3 CONN. INS. L.J. 417, 421 (1996); Frances E. Zollers, Andrew McMullin, Sandra N. Hurd, & Peter Shears, *No More Soft Landings for Software: Liability for Defects in an Industry that Has Come of Age*, 21 SANTA CLARA COMPUT. & HIGH TECH. L.J. 745, 768-69 (2005); Michael D. Scott, *Tort Liability for Vendors of Insecure Software: Has the Time Finally Come?*, 67 MD. L. REV. 425, 467 (2008).

transplant tools from other legal areas such as environmental law – demonstrating that there are more possible points of legal intervention than notice-and-takedown regimes.<sup>196</sup> But as one of us has noted elsewhere, risk regulation, like any form of regulation, is not a neutral tool. It brings with it “policy baggage” that includes a tendency to frame what should be policy discussions as obscure mathematical or scientific questions for technical experts, and a tendency to undervalue harms that are harder to quantify. Moreover, it reflects a willingness to accept risks, rather than prevent harms, and to deploy a technology regardless of the risks it brings or the other policy interventions its use eschews.<sup>197</sup>

In the speech context, aspects of the risky-complex-system construction raise important questions about freedom of speech and the extent to which regulators can pursue *ex ante*, wholesale interventions that downplay individual rights.<sup>198</sup> Understanding AI through the *value* of risk mitigation, the *institutions* of expert agencies and state courts, and the *object* of system design, directs us to investigate, define, and demand different things about the technology than constructions of it as speaker, or as the producer of networked content or a network itself.

#### 4. *AI and Consumer Protection*

In our fourth and final example, AI can be understood as speech that messes with the marketplace, either by deceiving consumers or by providing inadequate notice of how such systems work, under consumer-protection law. The regulatory *objects* of consumer protection law are consumers, businesses, products, and more broadly speaking, the marketplace. The consumer-protection construction understands AI not as a speaker, but as a product or a manipulative salesperson. AI may fail to adequately disclose information pertinent to consumer choice, or may generate information that manipulates consumer choice in a marketplace. The consumer-protection approach is situated within consumer-protection agencies (*institutions*), contextualized by experts in investigation and enforcement. Its *values* are market fairness and efficiency – in contrast to content-moderation law’s narrower goal of procedural efficiency.

It would be hard to overstate the power the consumer-protection construction has had over the way computers are understood in American law and beyond. The United States created an area of laws, rights, and institutions for

---

196. Keller, *supra* note 146; douek, *supra* note 42, at 554.

197. Kaminski, *Regulating the Risks*, *supra* note 7, at 1379; Kaminski, *supra* note 176. Jessica Eaglin has termed this “techno-correctionism.” Jessica M. Eaglin, *When Critical Race Theory Enters the Law & Technology Frame*, 26 MICH. J. RACE & L. 151, 155-57 (2021).

198. See Balkin, *supra* note 42, at 1257. See also *NetChoice, LLC v. Paxton*, 49 F.4th 439 (5th Cir. 2022), *cert. granted in part*, No. 22-555, 2023 WL 6319650 (Sept. 29, 2023).

computers through consumer-protection regimes the way Europeans did over the course of the latter half of the 20th century through data protection. Enthusiastic and important political actors carved out a path through the Department of Commerce and the FTC (the relevant *institutions* in the consumer-protection framing), paying less attention to alternatives like the FCC, courts, or changes to state constitutions to recognize data-privacy rights as Europe does.<sup>199</sup> As such, the FTC was relatively well-positioned to move into the AI arena, providing a legal understanding of AI as a factor in markets and defining its *object* of regulation as the consumer experience in a fair marketplace.<sup>200</sup>

The FTC frames AI as data and a model to scrutinize,<sup>201</sup> though its most recent actions have also prioritized scrutiny of talent and computational resources.<sup>202</sup> This framing capitalizes on the FTC's mandate in Section 5 of the Federal Trade Commission Act to prevent unfair or deceptive practices. As such, AI garners the attention of the FTC when used to “influence people’s beliefs, emotions, and behavior.”<sup>203</sup> Broadly, the FTC will investigate models that cause more harm than good, “that is, in Section 5 parlance, if it causes or is likely to cause substantial injury to consumers that is not reasonably avoidable by consumers and not outweighed by countervailing benefits to consumers or to competition.”<sup>204</sup> According to the agency, this includes discriminatory outcomes like the case of digital redlining against Facebook brought by the Department of Housing and Urban Development in 2019.<sup>205</sup>

---

199. JONES, *supra* note 52 (manuscript at 113-18) (on file with authors).

200. *Id.*

201. Elisa Jillson, *Aiming for Truth, Fairness, and Equity in Your Company's Use of AI*, FED. TRADE COMM'N (Apr. 19, 2021), <https://www.ftc.gov/business-guidance/blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai> [<https://perma.cc/QB6C-8UMT>]; see also Andrew D. Selbst & Solon Barocas, *Unfair Artificial Intelligence: How FTC Intervention Can Overcome the Limitations of Discrimination Law*, 171 U. PA. L. REV. 1023, 1077-78 (2023) (describing a data-security approach to regulating discriminatory AI).

202. Staff in the Bureau of Competition & Off. of Tech., *Generative AI Raises Competition Concerns*, FED. TRADE COMM'N (June 29, 2023), <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2023/06/generative-ai-raises-competition-concerns> [<https://perma.cc/9WJG-Q3CH>].

203. Michael Atleson, *The Luring Test: AI and the Engineering of Consumer Trust*, FED. TRADE COMM'N (May 1, 2023), <https://www.ftc.gov/business-guidance/blog/2023/05/luring-test-ai-engineering-consumer-trust> [<https://perma.cc/4TTB-4FWW>].

204. Jillson, *supra* note 201.

205. Press Release, Dep't of Just., Justice Department Secures Groundbreaking Settlement Agreement with Meta Platforms, Formerly Known as Facebook, to Resolve Allegations of Discriminatory Advertising (June 21, 2022), <https://www.justice.gov/opa/pr/justice-department-secures-groundbreaking-settlement-agreement-meta-platforms-formerly-known> [<https://perma.cc/2TEN-6FEL>] (note the housing case was about using actual race categories in roommate searches, not inferred algorithmic outcomes).

Maintaining an efficient and fair market involves correcting information asymmetries. Thus, transparency and duties to warn are essential policy tools for the consumer-protection objective of keeping consumers informed. There are a range of potential notice requirements that could become the dominant legal approach to AI, similar to the way an overemphasis on notices swallowed the legal approach to consumer-data privacy in the United States. There is the potential to label AI simply and appropriately as potentially dangerous products, leaving it to consumers to accept risks, abstain, or go shopping.<sup>206</sup>

The consumer-protection construction of AI systems may be gaining momentum in the United States. In July 2023, the FTC announced its investigation into OpenAI, the high-profile startup that created ChatGPT. In particular, the agency is investigating unfair practices related to “reputational harm.”<sup>207</sup> It wants details of all complaints OpenAI received about “false, misleading, or disparaging” statements generated by its products.<sup>208</sup> The agency is also looking into the way that ChatGPT stores and uses interactions with users after an incident in March 2023 that exposed some users’ conversations with the system—and whether proper notice about recording and training are effectively presented to consumers.<sup>209</sup>

The FTC has also focused on deceptive notices. This could lead to an agency emphasis on requiring companies to effectively articulate the way AI systems collect and use data as well as checking bold statements about what the systems can do. For example, the FTC brought enforcement actions against Facebook in 2019<sup>210</sup> and Everalbum (a photo app) in 2017<sup>211</sup> for misrepresenting their uses of facial-recognition technology. Regulatory actors might also take issue with statements about AI systems’ capacity. “AI snake oil” is a term popularized by

---

206. Brown, *supra* note 1, at 416–21 (describing a duty-to-warn approach to AI liability).

207. Cat Zakrzewski, *FTC Investigates OpenAI over Data Leak and ChatGPT’s Inaccuracy*, WASH. POST (July 13, 2023, 6:00 a.m. EST), <https://www.washingtonpost.com/technology/2023/07/13/ftc-openai-chatgpt-sam-altman-lina-khan> [<https://perma.cc/A775-MPJC>].

208. *Federal Trade Commission Civil Investigative Demand Schedule, FTC File No. 232-3044*, FED. TRADE COMM’N 11-12 (2023), <https://www.washingtonpost.com/documents/67a7081c-c770-4f05-a39e-9d02117e50e8.pdf> [<https://perma.cc/74R4-FE7Q>].

209. According to OpenAI, this was caused by “a bug in an open-source library which allowed some users to see titles from another active user’s chat history.” *March 20 ChatGPT Outage: Here’s What Happened*, OPENAI (Mar. 24, 2023), <https://openai.com/blog/march-20-chatgpt-outage> [<https://perma.cc/ELD6-GXZF>].

210. Complaint at 6, *United States v. Facebook*, No. 19-cv-2184 (D.D.C. July 24, 2019), [https://www.ftc.gov/system/files/documents/cases/182\\_3109\\_facebook\\_complaint\\_filed\\_7-24-19.pdf](https://www.ftc.gov/system/files/documents/cases/182_3109_facebook_complaint_filed_7-24-19.pdf) [<https://perma.cc/6WM6-L2TF>].

211. Complaint at 2, *Everalbum, Inc.*, FTC File No. 192-3172 (Jan. 11, 2021), [https://www.ftc.gov/system/files/documents/cases/everalbum\\_complaint.pdf](https://www.ftc.gov/system/files/documents/cases/everalbum_complaint.pdf) [<https://perma.cc/2GG4-F855>].

Arvind Narayanan, who shared presentation slides with the title on Twitter.<sup>212</sup> Commentators use the term to emphasize the hype around AI and focus the public's attention on the limitations of AI systems. In consumer-protection construction, rather than constructing AI systems as producing dangerous public discourse, critics and regulators construct them as the objects of false commercial speech intended to direct consumers to products and services.<sup>213</sup>

Preventing consumer manipulation is a central mandate of the FTC. Thus, another aspect of the OpenAI investigation is how difficult it is to opt-out or change defaults, in particular how difficult it is to disable chat history and whether those steps constitute a “dark pattern.” “Dark pattern” is a term that refers to computer-interface designs and strategies that trick, nudge, or mislead consumers to act contrary to their intended goals. While tricks, nudges, and misdirection are not new to marketplaces, “dark patterns” was coined by a researcher in 2010 specifically for digital interaction.<sup>214</sup> The FTC's 2022 report, *Bringing Dark Patterns to Light*, provides extensive examples of design changes, like font sizes and colors, and nudges and tricks, like confusing language and hidden settings.<sup>215</sup> Although the report makes no mention of AI, the FTC has since warned companies not to overstate claims about AI products<sup>216</sup> and not to skimp on notifying consumers about how data may be used to train AI systems.<sup>217</sup> The agency is on the lookout for dark patterns in chatbots, targeting, and embedded ads.<sup>218</sup> This narrow set of relatively aggressive activities represent the institutional tools and culture of law enforcers promoting consumer empowerment that enables purportedly free choices in an open market.

---

212. Arvind Narayanan (@random\_walker), TWITTER (Nov. 19, 2019, 2:18 PM), [https://twitter.com/random\\_walker/status/1196870349574623232](https://twitter.com/random_walker/status/1196870349574623232) [<https://perma.cc/5ME3-GDXB>].

213. See also Jane Bambauer, *Snake Oil Speech*, 93 WASH. L. REV. 73, 80-83 (2018) (arguing that First Amendment protections should turn on risk analysis, rather than a distinction between truth and falsity).

214. Dr. Harry Brignull coined the term “dark patterns.” See Harry Brignull et al., *Deceptive Patterns – User Interfaces Designed to Trick You*, TESTIMONIUM LTD. (Apr. 25, 2023), <https://www.deceptive.design/about-us> [<https://perma.cc/S3S7-9W6C>].

215. Bureau of Consumer Prot., *Bringing Dark Patterns to Light*, FED. TRADE COMM'N 4 (Sept. 2022), [https://www.ftc.gov/system/files/ftc\\_gov/pdf/P214800%20Dark%20Patterns%20Report%209.14.2022%20-%20FINAL.pdf](https://www.ftc.gov/system/files/ftc_gov/pdf/P214800%20Dark%20Patterns%20Report%209.14.2022%20-%20FINAL.pdf) [<https://perma.cc/MKG6-KJ6J>].

216. Michael Atleson, *Keep Your AI Claims in Check*, FED. TRADE COMM'N (Feb. 27, 2023), <https://www.ftc.gov/business-guidance/blog/2023/02/keep-your-ai-claims-check> [<https://perma.cc/TJ5B-D8YS>].

217. Andrew Smith, *Using Artificial Intelligence and Algorithms*, FED. TRADE COMM'N (Apr. 8, 2020), <https://www.ftc.gov/business-guidance/blog/2020/04/using-artificial-intelligence-and-algorithms> [<https://perma.cc/75BS-N7E9>].

218. Atleson, *supra* note 203.

The *values* of the FTC all relate to consumer protection that serves efficient markets. As a consequence, the FTC has been limited to only a tangential authority over commercial computers and data as objects of regulation, with broad gaps in its authority (for example, over nonprofits, air carriers, banks, and telecommunication companies).<sup>219</sup> Yet the FTC’s political prominence, resources, and tools have made it the de facto governance force for AI developers to consider in the United States.

One positive feature of the consumer-protection construction is that the FTC, with its legislative mandate to go after the broad and evolving standards of deception and “unfairness,” has a long history of being unphased by technological newness. FTC Chair Lina Khan has made it clear that she is not interested in AI exceptionalism.<sup>220</sup> However the FTC’s efforts shape AI in the future, it is important to consider that the FTC has already been fundamental to the existence of the AI we have now. The information economy and the data, algorithms, and computing power that make AI possible were fostered by an FTC understanding of technology and governance. So, as the FTC makes sense of AI, it does so as a coproducer of AI in the U.S. context.

## II. LEGAL CONSTRUCTION AS A POLICY CHOICE

This brings us to the “so what?” section, in which we address the normative valence of legal construction of technology. We have walked through four different legal constructions of AI-generated speech: as speech or speaker, as speech at scale, as the product of risky complex systems, and as a faulty digital product subject to snake-oil sales tactics under consumer-protection law. We have examined the meaning-making modes, tools, and institutions of the law. Assuming all of this has seemed descriptively accurate enough, what is a lawyer, scholar, or policymaker to do with it?

We each have thoughts, spelled out at length elsewhere, as to which of these legal constructions might be more descriptively accurate as to the technology. Additionally, we each have thoughts as to which institutions and policy tools, or

---

219. *What the FTC Does*, FED. TRADE COMM’N, <https://www.ftc.gov/news-events/media-resources/what-ftc-does> [https://perma.cc/3774-PDYJ].

220. See Press Release, Fed. Trade Comm’n, FTC Chair Khan and Officials from DOJ, CFPB and EEOC Release Joint Statement on AI (Apr. 25, 2023), <https://www.ftc.gov/news-events/news/press-releases/2023/04/ftc-chair-khan-officials-doj-cfpb-eoc-release-joint-statement-ai> [https://perma.cc/PL9Z-H9KY] (“We already see how AI tools can turbocharge fraud and automate discrimination, and we won’t hesitate to use the full scope of our legal authorities to protect Americans from these threats . . . Technological advances can deliver critical innovation – but claims of innovation must not be cover for lawbreaking. There is no AI exemption to the laws on the books, and the FTC will vigorously enforce the law to combat unfair or deceptive practices or unfair methods of competition.”).

combinations thereof, might best achieve regulatory goals.<sup>221</sup> Our aim here is not to win the reader over to what we've each already argued about regulating AI systems. Nor does this Essay's length afford us the space to once and for all solve the puzzles and problems raised by AI-generated speech. Rather, we want to win the reader over to the legal-construction-of-technology method and its concrete normative policy benefits, by showing its significant implications for conversations about AI-generated speech.

As ongoing practitioners of the method, we are still uncovering what it affords us. We hope to convince you that while one key benefit of legal construction of technology is its descriptive accuracy, the method isn't merely descriptive. Sometimes, it is revelatory. And often, it is deeply empowering.

Legal construction of technology is *descriptively accurate*. We don't want to shy away from the importance of the fact that legal construction of technology accurately describes what we think the law is doing. The law conducts sense-making of technology, through language and institutions, toward particular values or goals. Each of the above constructions involves different institutions with different histories, performing meaning-making toward sometimes wildly different normative ends. The First Amendment attempts to legally construct through litigated doctrine, AI-generated content as speech or AI systems as a speaker, toward values of democratic participation, pluralism, and the protection of individual speaker and listener autonomy. The FTC attempts to legally construct AI-generated content as consumer deception or a problem of inadequate disclosure or bad data governance, with the goal of maintaining a functional marketplace in which information asymmetries are addressed and consumers are adequately protected from self-interested company behavior. Both of these analyses are descriptively accurate, if simplified, versions of how U.S. law handles AI speech.

If this observation — that U.S. law constructs AI speech through different institutions in different ways, at the same time — is all that legal construction of technology afforded us, it would still, we think, be a meaningful step in the right direction. It may be descriptively accurate to say in a narrow sense that AI systems create an interesting problem for First Amendment doctrine because of the law's lopsided focus on the intent of human speakers. It is not descriptively accurate to say that AI-generated content "disrupts the law" in the sense of being generally hard for the law to regulate, just because construction within First Amendment law takes a hard look at human intent sometimes.

Legal construction of technology can also be *revelatory*. As a method, it leads us to identify salient aspects of the law and allows us to more accurately pinpoint

---

221. See sources cited *supra* note 7.

where and how puzzles arise. This then lets us come up with creative and appropriately targeted interventions.

For example, it is possible to use straightforward legal analysis to arrive at the observation that running AI speech through First Amendment doctrine gets sticky around exceptions and speaker intent. Several scholars have already done so.<sup>222</sup> It is also possible, using the kind of analogous reasoning that resounds throughout the law, to suggest solutions for this problem: treat AI speech as an encyclopedia, the advice of a doctor, the output of a product, or the result of faulty parenting.<sup>223</sup> This reasoning, we believe, in fact *is* performing legal construction of technology: using available tools of legal sensemaking to let the law get to the “right” normative outcome – whatever that may be for an individual lawyer, policymaker, or scholar. It is not, however, taking advantage of what a more self-aware and deliberate approach to legal construction of technology can afford.

The legal-construction-of-technology method lets us identify that First Amendment doctrine constructs content differently, at different moments: as speech, when it comes to coverage, and as speech by a human speaker when it comes to historic exceptions to protection. That is itself not just a description, but a minor revelation. That’s surprising, and a bit weird! So, why is it happening – this bifurcation of caring about abstract speech versus caring about human speakers? Because (leaving Lochnerian skepticism aside for just a moment) in service of protecting public-sphere pluralism, contemporary First Amendment doctrine has been structured to pull more and more laws more readily under protective and purportedly neutral judicial scrutiny, and to be highly careful about allowing partisan actors to regulate speech. The big-picture story motivating contemporary First Amendment doctrine over at least the past half-century has been that courts help to structure a pluralistic public sphere, where we all learn to live with speech and opinions that may deeply offend or even hurt us – and with courts protecting minority human speakers whose speech might readily be chilled.<sup>224</sup> Widening the threshold for coverage by constructing more and more laws as touching on “speech” gives judges more oversight and power; interpreting exceptions from constitutional protection as applying only to “speakers with adequate (human) intent” is meant to protect minority (human) speakers and their (human) audiences from chilling effects in the face of majority, partisan rule.

---

222. See, e.g., Massaro et al., *supra* note 86, at 2514; Brown, *supra* note 1, at 402; Volokh, *supra* note 126, at 513-14.

223. See, e.g., Bambauer, *supra* note 43, at 349-58; Brown, *supra* note 1, at 391-93.

224. See, e.g., ROBERT POST, *DEMOCRACY, EXPERTISE, AND ACADEMIC FREEDOM* (2013).



Thus, examining First Amendment law's construction of AI speech reveals the following about the area of law: First Amendment law is the situated product of courts creating doctrine in a liberal pluralistic democracy, in response to clever lawyers who know how to trigger that liberal pluralistic framing by invoking a specter, sometimes true, of partisan interference in public debates, in front of a Supreme Court increasingly solicitous to religious and conservative speakers. A core value running through this project is that quantity of speech is a good thing: the mode by which we hope "good speech" rises to the top is by relying on sheer quantity of and competition between differing information and opinions, because we cannot trust elected political leaders with polemical views to prevent harms while also preserving deeply valued pluralism.<sup>225</sup> But if quantity of speech is actually part of the problem – of misinformation, consumer manipulation, the circulation of deep fakes, and election fraud – then this pluralistic-individualistic First Amendment framing may have to shift.<sup>226</sup> Or at least, courts will have to make room in the doctrine for handling the massive amount of communicative output generated at low cost by nonhuman speech engines, distinguishing in some way between them and the opinion-holding human speakers and listeners that the liberal pluralistic framing is designed to protect.

Legal construction of technology thus makes evident particular oddities in legal sensemaking.<sup>227</sup> It makes us explore what, exactly, those particular oddities might in turn evidence about the broader area of law, its institutions, and its underlying values. It reveals what might happen when existing values are pursued to illogical or inconsistent ends, or when the makeup of interpreting institutions changes.

That's all very big picture, though. More practically, the method yields some very concrete policy revelations, too. It reveals, for example, the limits of resorting only to analogies for construction. Sure, an AI system is not a malicious human speaker. But it's also not an advice-offering doctor, a precocious child of a negligent parent, or a physical product like a chair. It's a complex human-machine prediction system with a befuddling supply chain, often deployed as software-as-a-service. It's not that legal construction of technology is the first to identify the significance or limitations of analogies in legal reasoning. But it enables one to do interesting things with analogies: be appropriately self-aware of their significance and the constructions they trigger, recognize their limitations,

---

225. See Toni M. Massaro & Helen Norton, *Free Speech and Democracy: A Primer for Twenty-First Century Reformers*, 54 U.C. DAVIS L. REV. 1631, 1640-42 (2021); Tim Wu, *Is the First Amendment Obsolete?*, KNIGHT FIRST AMEND. INST. (Sept. 1, 2017), <https://knightcolumbia.org/content/tim-wu-first-amendment-obsolete> [<https://perma.cc/AV4J-5T5T>].

226. See Massaro & Norton, *supra* note 225, at 1640-42.

227. See generally Balkin, *supra* note 36 (revealing these oddities as applied to different areas of robotics).

recognize their embedded normative goals, and recognize the extent to which technological design is itself nudging you toward a particular analogy (like understanding email as a letter).<sup>228</sup>

Legal construction of technology, as noted above, frequently yields the insight that the developing uses of a technology will likely be constructed differently through multiple different institutions at the same time. Awareness of these differing constructions can be revelatory in a number of ways. It can lead to effective counterfactuals: Why compare AI to a chair, as one part of the legal system might do, if it's more accurate to talk about it as a complex risky system, as another part does? It can lead to crucial explorations of how different areas of the law interact: Does one mode of construction leave important things out? Does another mode adequately fill identified gaps? Do the different modes preclude each other? Leave room for each other? Create helpful amplifications, redundancies, even resilience with each other?<sup>229</sup>

Let's again make this more concrete with examples drawn from our four analyses above. If we think that speech-at-scale construction, risky-complex-system construction, or consumer-protection construction are in some ways desirable, does First Amendment doctrine leave space for, or even amplify, those approaches? Right now, at least, the Supreme Court's First Amendment caselaw hasn't had to deal directly with the constitutionality of speech-at-scale regulations, such as notice-and-takedown regimes, or for that matter, with risk regulation of speech. (It has, however, potentially started dismantling aspects of consumer protection, much to the dismay of many.<sup>230</sup>) Congress's policy decision to immunize online platforms in CDA 230 actually punted First Amendment analysis of speech-at-scale regulation by, for the most part, privatizing content moderation instead of drafting command-and-control laws that would be subject to First Amendment analysis. (Blake E. Reid has referred to this as the First Amendment's "interpretative debt," attributable to CDA 230.<sup>231</sup>)

---

228. Thanks to Tiana Wang in Paul Schwartz's class for helpful comments that led to this observation. See also Orin S. Kerr, *The Problem of Perspective in Internet Law*, 91 GEO. L.J. 357, 357 (2003); Stephanie A. Gore, "A Rose by Any Other Name": *Judicial Use of Metaphors for New Technologies*, 2003 U. ILL. J.L. TECH. & POL'Y 403, 416 (2003); Lyria Bennett Moses, *Recurring Dilemmas: The Law's Race to Keep up with Technological Change*, 2007 U. ILL. J.L. TECH. & POL'Y 239, 255-56 (2007).

229. See Crootof & Ard, *supra* note 29, at 376-79 (thinking through interactions between parts of the legal system).

230. See, e.g., *NIFLA v. Becerra*, 138 S. Ct. 2361 (2018) (Breyer, J., dissenting); Christopher T. Robinson, *The Tip of the Iceberg: A First Amendment Right to Promote Drugs Off-Label*, 78 OHIO ST. L.J. 1019 (2017).

231. Gonzalez, Taamneh, and *Section 230's Interpretive Debt*, BLAKE E. REID (Feb. 23, 2023), <https://blakereid.org/gonzalez-taamneh-and-section-230s-interpretive-debt> [<https://perma.cc/B3W4-PGW7>] ("[W]e have almost no idea how huge swaths of law – e.g.,

Thus, although First Amendment questions around AI speech are exciting, the First Amendment probably doesn't matter much, at least not right now. This reality has nothing to do with how AI works or what it can do. U.S. legal choices have left considerable space for the European Union's DSA risk-regulation construction to be important and influential, even to U.S.-based AI players. While the United States has abstained, the European Union has proactively regulated in this space, openly pursuing global influence through the so-called Brussels Effect, whereby global companies subject to EU regulation de facto export EU laws by abiding by them elsewhere.<sup>232</sup> That is, global private companies may start doing a DSA-style mix of due process and risk mitigation all over the world, and they have the space to do that in the United States because of the regulatory gap CDA 230 creates and preserves.

But legal construction is historically contingent, politics are ever changing, and a shift to institutions and U.S. jurisprudence could change everything. If the Supreme Court decides to let states regulate platforms with must-carry laws, which construct content moderation as something like "common carriage," those laws will compete with the DSA's construction.<sup>233</sup> A decision by the Court that platforms' moderation decisions aren't covered by the First Amendment may leave room, too, for stateside risk regulation in this space. If alternatively, the Court decides there is a First Amendment right of some kind for platforms to conduct content moderation, the ability of the U.S. government to regulate platforms and design features in the future will be significantly limited. On the one hand, this may block stateside DSA-style risk regulation and government-driven notice-and-takedown, but on the other, it will preserve the CDA 230 gap and thus may perversely strengthen the Brussels Effect of the DSA's construction.

The interaction of judicial interpretation of CDA 230 and the First Amendment with the European Union's risk regulation is just one way these constructions create interlocking pieces. Legal construction of technology encourages its

---

tort, contract, civil rights, state criminal law, etc.—might (or might not) be applied to user-generated content platforms as a threshold matter. We have almost no idea how the First Amendment might apply. This means that rugging Section 230, whatever your policy preferences, is going to create *an enormous amount of work for courts*."). See generally Blake E. Reid, *Section 230's Debts*, FIRST AMEND. L. REV. (forthcoming), <https://ssrn.com/abstract=4624865> [<https://perma.cc/8RPR-SLW2>] ("Section 230 has accumulated *interpretive debt*—in addition to the Supreme Court's lack of First Amendment jurisprudence, the absence of a common law that applies existing substantive legal regimes to regulate platforms' carriage and moderation practices.").

232. Anu Bradford, *The Brussels Effect*, 107 NW. U. L. REV. 1, 23 (2012).

233. See Reid, *supra* note 141 (manuscript at 8) (arguing that these are not common carriage laws but draw on the rhetoric of common carriage). See also Christopher S. Yoo, *Common Carriage's Domain*, 35 YALE J. REGUL. 991 (2018); Christopher S. Yoo, *The First Amendment, Common Carriers, and Public Accommodations: Net Neutrality, Digital Platforms, and Privacy*, 1 J. FREE SPEECH L. 463, 468 (2021).

practitioners to think seriously about design coherence between the systems – and to think carefully about how one ill-fit system of construction might impede better-fitting or more effective systems.

Practicing legal construction of technology thus takes us out of narrow frameworks. It removes blinders. It makes clear what aspects of the law the technology is running into, *so that you can fix them*. This brings us to empowerment.

We find legal construction of technology to be immensely *empowering* in the policy realm. That is because it inherently pushes lawyers away from surrendering to the idea that the technology is in charge of legal changes, and instead asks them to be creative about imagining how things might be another way. Here are just a few examples.

Legal construction of technology can lead to sometimes helpful legal transplants. A legal transplant occurs when the tools or constructions of one type of existing law are brought to bear on a problem arising in another.<sup>234</sup> Transplants can be transnational: for example, when a country copies or imports the U.S. approach to notice-and-takedown in copyright law,<sup>235</sup> or when U.S. states copy aspects of the GDPR.<sup>236</sup> Transplants can be domestic, such as when states copy each other’s privacy laws.<sup>237</sup> Or transplants can occur across different policy fields, such as when policymakers use metaphors and regulatory approaches from natural-resource law to address the allocation and governance of digital assets.<sup>238</sup>

Legal construction of technology empowers us to discover transplantable constructions, select more appropriate transplants, and be wary of the troubles a particular transplanted construction might bring. For example, AI risk regulation transplants a lot from other areas of law, including data-protection law, environmental law, and EU product-safety regulation.<sup>239</sup> Boine and Rolnick observe that the EU AI Act stitches together a particular version of risk regulation with EU products-safety law. They claim that this particular transplanted way of legally constructing AI systems – with its ultimate focus on *ex ante* mitigation,

---

234. Vanessa Casado Pérez & Yael R. Lifshitz, *Natural Transplants*, 97 N.Y.U. L. REV. 933, 935 (2022) (explaining that a legal transplant is “a transfer of a legal regime or rule from one jurisdiction to another”).

235. See Margot E. Kaminski, *The Capture of International Intellectual Property Law Through the U.S. Trade Regime*, 87 S. CAL. L. REV. 977, 1050 (2014).

236. See Anupam Chander, Margot E. Kaminski, & William McGeeveran, *Catalyzing Privacy Law*, 105 MINN. L. REV. 1733, 1787 (2021).

237. See *id.*

238. See Pérez & Lifshitz, *supra* note 234, at 940.

239. See Kaminski, *Regulating the Risks*, *supra* note 7, at 1352, 1404-05.

product design, intended purpose, and fit for purpose – failed to contemplate or prepare for General Purpose AI.<sup>240</sup>

Legal construction of technology enables us to bring different, maybe better, modes of construction to the table, within an existing meaning-making mode. For example, if one is a First Amendment scholar, one could get trapped in substitution debates (e.g., what to do about intent?). Or, one could use analogies to transplant another mode of construction. We see, for example, Brown’s suggestion that courts analyze defamation claims through a products-liability lens as an attempt to port risk mitigation into the tort-speech interface.<sup>241</sup> We similarly understand Volokh’s suggestion that First Amendment defamation law leaves room for “notice-and-blocking” requirements as an attempt to transplant speech-at-scale construction into the First Amendment.<sup>242</sup>

Legal construction of technology, too, helps us discover when the necessary intervention must occur at the level of values. If we can identify that the problem is that the motivating forces behind the minutia of laws and regulation are mismatched to the problems we’re trying to solve, then we know to look elsewhere for solutions. If one’s concern over AI-generated speech is that it will destroy democracy, the right policy answer is probably not going to be a market-driven policy solution such as copyright law.<sup>243</sup> Or, to return to a previous example, if the problem one is trying to solve is caused by scale, then one needs to be thinking in terms of scale-directed solutions, driven by values that don’t overemphasize the scale of speech production as an inherently positive thing.

---

240. See Boine & Rolnick, *supra* note 72, at 4, 41-42.

241. See Brown, *supra* note 1, at 421-24; see also text accompanying notes 192-193 (discussing Brown’s suggestion to use products-liability law to address the problem of chatbots generating defamatory content).

242. Volokh, *supra* note 126, at 542-44 (discussing policy reasons for using notice-and-blocking, including balancing harm to innovation and the listener rights of AI-system users with speech harms to affected libeled individuals); see also text accompanying notes 169-172 (discussing Volokh’s work on generative AI systems).

243. Oren Bracha, Remarks at the Silicon Flatirons Conference 43 (Oct. 6, 2023), (transcript available at <https://perma.cc/ZYB6-QYNV>) (observing that a “copyright solution is market-based . . . . And there are known advantages and disadvantages to that strategy, but that’s what copyright is about. Now enter generative AI in the field of expression . . . the broad legal arguments, when you really examine them, are not about the specific policy problem that copyright is designed to solve. Well not all of them . . . concerns about the disappearance of sources of livelihood and human income in the creative industries. Second, concerns about opportunities and potentials for people to engage and enjoy the inherent value of creativity. Not the market value, the inherent value of creativity. Third, concerns about the sources and potentials for disruptive innovation, the kind of innovation that doesn’t only deliver very well within given patterns, but rather, break[s] the patterns. Those are very important, legit[imate] concerns, I think. Those are exactly the concerns that copyright is not built to deal with, it was not designed to deal with”).

Finally, legal construction of technology enables lawyers to figure out the shape of the particular issue they're trying to solve and draw on appropriate tools. If we're in an expertise-gap situation, we have plenty of tools within the law for getting more technical expertise into policymaking. Not every problem is the pacing problem, but when the concern is that companies developing technologies are attempting to engage in regulatory arbitrage by falling between legal boxes and into legal cracks, the law has plenty of tools (with tradeoffs) for trying to get around that problem.<sup>244</sup>

## CONCLUSION

A stable legal construction of AI has not yet set in, and negotiations are occurring with all the histories of global and national technology governance as the setting. We have presented only four potential constructions of AI to show a values-first approach, as opposed to a technology-first approach, to technology governance. From copyright law to tax law to competition law, there are plenty of others to choose from. When we skip to asking how a technology will disrupt some law, we don't do the normative work of recognizing these variations and choices.

The law makes meaning of the social uses of technology—it constructs them. That's not to say that's all the law does. Law is construction with consequences. It conducts sensemaking through language, through institutions, through policy tools, and through motivating theories. Legal actors are not passive, even when they fail to enact new law. They have substantial agency in how AI is and will be legally constructed. A technological-exceptionalist approach to technology law can inaccurately characterize technology policymakers as unduly incapacitated and reactive. At the same time, technological exceptionalism can afford technology-law scholars and policymakers a degree of exceptionalism that detracts from what they can learn and sometimes transplant from practices across a wide array of other areas of law.

While these interactions can be hard to predict, as can politics, the legal construction of AI should not be considered neutral or objective. AI, like any technology, cannot be understood outside of its social context—and legal actors should fight for that approach, because it positions them to fight for technologies that further or protect certain democratic values. Legal scholars and technology law experts should not pretend to be caught off guard or fall into the disruption rhetoric. When legal actors, institutions, and ideals make meaning of technology they make paths for powerful democratic participation in technology. Without taking that path, we choose to follow technology.

---

244. See Greenberg, *supra* note 54, at 1500-01, 1508-11.

*Meg Leta Jones, Provost's Distinguished Associate Professor in Communication, Culture, & Technology; Faculty Advisor of Institute for Technology Law & Policy; Core Faculty in Science, Technology, & International Affairs; and Faculty Fellow in Ethics Lab at Georgetown University.*

*Margot Kaminski, Professor of Law, Colorado Law School; Fulbright-Schuman Grantee 2023-2024; Fernand-Braudel Senior Fellow, European University Institute 2024; Affiliated Fellow, Information Society Project at Yale Law School; Faculty Associate, Harvard Berkman Klein Center 2023-2024.*

*Thanks to Nicholson Price, Paul Schwartz, and Ari Waldman and their extraordinary students for workshopping drafts of this Essay. Thanks also to Jack Balkin, participants in the Yale ISP Ideas Lunch, Blake Reid, the Rocky Mountain First Amendment Scholars workshop, and the We Robot and Privacy Law Scholars Conference communities for important and helpful feedback at various stages of these ideas.*