No. 21-1333

IN THE

Supreme Court of the United States

REYNALDO GONZALEZ, *et al.*,
Petitioners,

v.

GOOGLE LLC,
Respondent.

**On Writ of Certiorari to
the United States Court of Appeals
for the Ninth Circuit**

AMICUS CURIAE BRIEF OF DEVELOPERS ALLIANCE IN SUPPORT OF RESPONDENT

BRUCE GUSTAFSON
PRESIDENT & CEO
DEVELOPERS ALLIANCE
6218 Georgia Avenue, NW
Suite #1 PMB 3045
Washington, DC 20011
(202) 735-7333
bruce@developersalliance.org

JAMES H. HULME
Counsel of Record
NADIA A. PATEL
ARENTFOX SCHIFF LLP
1717 K Street, NW
Washington, DC 20006
(202) 857-6000
james.hulme@afslaw.com

TABLE OF CONTENTS

TABLE OF AUTHORITIES

Cases **Page(s)**

Statutes

230(f)(4) 10, 27

Other Authorities

Bob Tarantino, *If You Know About It, You're the Publisher – Website Operator Liability for Defamation* (Ent. & Media L. Signal May 15, 2015), http://www.entertainmentmedialawsignal.com/if-you-know-about-it-youre-the-publisher-website-operator-liability-for-defamation/ 24

A Dad Took Photos of His Naked Toddler for the Doctor. Google Flagged Him as a Criminal, NY Times, https://www.nytimes.com/2022/08/21/technology/google-surveillance-toddler-photo.html. 18

EU Regulation, Article 6(1)(b) and Recital 22, No. 2022/2065, of the European Parliament (19 Oct. 2022) 24

Michael L. Rustad & Diane D'Angelo, *The Path of Internet Law: An Annotated Guide to Legal Landmarks*, 10 Duke L. & Tech. Rev. 1 (2011)..12

Marshall McLuhan, Quentin Fiore, *The Medium Is the Massage, An Inventory Of Effects* (Gingko Press, 2001)...........................32

Mohit Singhal et al., *SoK: Content Moderation in Social Media, from Guidelines to Enforcement, and Research to Practice*, (Oct. 27, 2022), https://arxiv.org/pdf/2206.14855.pdf. 10

Restatement (Second) of Torts § 324A (1965)...21

Restatement (Second) of Torts § 578 (1977) 22

Restatement (Second) of Torts § 581 (1977) 21

Sergey Brin & Lawrence Page, *The Anatomy of a Large-Scale Hypertextual Web Search Engine.* http://infolab.stanford.edu/~backrub/google.html (last visited Jan. 17, 2023)....................10

Software Development, https://en.wikipedia.org/wiki/Software_development (last visited Jan. 17, 2023) 27

Symbolics, https://symbolics.com/ (last visited Jan. 17, 2023)...3

INTEREST OF AMICUS CURIAE

The Developers Alliance is a non-profit corporation that advocates for software developers. Our corporate mission is to "[a]dvocate on behalf of developers and the companies that depend on them, support the industry's continued growth, and promote innovation."[1]

Alliance members include industry leaders in consumer, enterprise, industrial, and emerging software, and a global network of more than 75,000 developers.

Amici have no direct financial interest in the outcome of this case but have a strong interest in seeing that the law continues to support innovation in the software industry.

[1] No counsel for any party authored this brief in whole or part, and no person other than *amici* or their counsel made a monetary contribution to the preparation or submission of this brief. All parties have filed a blanket consent with the Court.

SUMMARY OF ARGUMENT

Imagine a world of information and entertainment built just for you, where anything you want or need is instantly available. Searches yield a single, perfect, and comprehensive result, elegantly presented in the most digestible way. Imagine rich and thought-provoking interactions with a global population of thinkers and dreamers, artists, and oracles, without discord, reflective of your values while testing your thinking and expanding your world view.

Now imagine the challenges you would face in building this. The practical limitations of building not 9 billion perfect internets, but a single shared one, filtered and annotated and organized automatically and in real time as new information is flooding in. Imagine trying to decipher what each individual wants or needs from the vast and raging sea of content and interaction. As an engineer, you start with what you have and get to work. And then your lawyer calls.

The fundamental issue presented by this case is that the legal system has not developed a paradigm to address the effects of machine-generated content on the Internet. Traditional legal concepts involving intent, malice, and other measures of human conduct do not work to evaluate verbal content created by computer algorithms. To understand and develop rules to address these issues, it is necessary to understand exactly how these computer systems have developed and how they operate. This brief is intended to do just that.

ARGUMENT

I. Computer Science is a Specialized Field with Its Own Language and Customs Based on Its Historical Roots.

There are many resources that describe the history and early evolution of systems that would one day be labeled the Internet.[2] In order to understand the text and purpose of 47 U.S.C. § 230 — "Protection for private blocking and screening of offensive material," one needs to understand how engineers, software developers and computer scientists have designed the systems that Section 230 is written to address, why common law analogies are inapt, and to understand how the language and culture of computer science is reflected in the text of the law.

The 1990's saw the migration of complex computer networks from the universities and laboratories, where they were first used, to the wide world outside academia. Some of the early precursors to today's social media and third-party content platforms were bulletin board systems ("BBS") — a computer science term of art that, like many such terms, analogizes to real-world systems and concepts as short-hand for an abstract digital system. The analogy of a physical bulletin board with content posted by other users for everyone to see, loosely organized and ever changing, helped users to grasp how the systems were designed and might be used. Computer science in general has a unique vocabulary

[2] E.g., https://symbolics.com/ (last visited Jan. 17, 2023).

of familiar-looking terms such as “bulletin board,” “posts,” “server,” “client,” “users,” “platform” or “publish” which act as placeholders for more complex and nuanced concepts — just as “partner,” “publisher,” or “reasonable” represent far more than their dictionary definition to legal practitioners. Words like “cache” and “user” in Section 230 alert us that more than Black’s Law Dictionary is required to interpret Congress’s words.

II. The Fundamental Functions Regulated by Section 230 Have Their Roots in Early Computer Networking.

A bulletin board system is a computer server running software that allows remote users to connect to the central server using a terminal program (often called a client). Prior to the modern Internet, users connected their home computers to a hardware modem that dialed the phone number of another modem connected to a server computer capable of accepting many incoming connections. Once connected, client software would run on the user’s computer, allowing it to communicate with the remote server to navigate content, upload and download files and text, and exchange messages with other users connected to the same server.

Computer science terminology often emphasizes the purpose or function of various elements in a system over where or how things are implemented. The result is that terms like “server” can refer to a physical computer, a network function, a software package or process, or a cluster of any of these performing a similar role. An individual device can be both a server and a client at the same time, or from time to time. The result is that the language of computer science and networking can be easily misunderstood by those unfamiliar with the art.

Early BBS users tended to be technical experts and hobbyists, often with backgrounds that included research or higher education where institutional versions of servers and clients inside dedicated networks were common. BBS users could upload and download simple data (eventually including software), read news or bulletins uploaded or posted (created on the system) by others, exchange messages or engage in simple chat. Servers were managed by individuals or small groups who maintained the hardware and software, managed user accounts[3], and decided what software and computing services their system might make available to users. Since everyone involved had some measure of computer science in their background, common concepts and terms like files, directories, login names, passwords, permissions [4] , and system administration were

[3] Users, accounts and clients are related terms with very real differences. A user is most closely analogous to a person (though they are not the same), while accounts are the logical labels a computer system uses to assign system resources and access rights. A single user might have access to several accounts with varying levels of access. The term client indicates the existence of a server, and simply defines one side of the logical client-server relationship.

[4] Permissions or privileges are account attributes that computer systems use to, among other things, catalog what hardware and software services a particular account is allowed to access and what functions it can perform. For example, an

imported naturally from these institutional systems, and they remain foundational to this day.

The system administrator's ("admin") role in academic networks had always included establishing and enforcing the policies that defined appropriate uses for the computer systems they managed. Being technologists, admins established the conditions and limits around how shared processing, storage, and networking resources [5] were used in addition to enforcing the administrative policies of the institutions that owned the assets. Even on these early platforms[6], this often included organizing and structuring any shared data so that content could be found by anyone using a little logic (the human kind) and a lot of browsing (the manual version). Data was just one more shared asset to manage. Early

account may have permission to download content from a shared directory but may not have permission to upload content of its own. Permissions are typically under control of the system administrator whose account will include permissions and privileges allowing it to change the permissions of others.

[5] Resources can be a reference to physical computing elements (hard drives, computer processors, etc.), to the output of physical elements (an amount of storage or processing time, etc.), or to functions, software tools and other logical items. An account's ability to access resources is defined through permissions.

[6] Platforms are a combination of hardware and software that enable higher level services. The name highlights that platforms are a foundation upon which services run. The word is often used in the context of systems that support third-party communications and content sharing. The term "system" is broader, but the terms are often used interchangeably.

platforms had limited search capabilities, so organizing content into some logical structure (alphabetical by file name, topical by content, serially by date or size, user-defined by asking users to supply logical file names for their content, etc.) was critical so content could be easily found.

While early content was being organized, admins established and enforced community policies for the systems they managed. Unlike system-level policies, which focused on technical areas like hardware, software and services, these community policies focused on content and user behavior. A server or a section of a server's storage dedicated to bird watchers might empower the admin to prohibit (and even remove) user-created content about duck hunting. Users that repeatedly violated rules could have their accounts closed, their permissions limited, or their uploaded content removed. Admins thus became moderators, policing content and enforcing content policies, or conscripted moderators from their user communities if the workload was high.[7]

[7] BBS owners recognized early on that content was the key to revenue. Because early systems charged direct fees for access and because the user base was very small at the start, some BBS provided incentives to upload shared content, providing "credits" for unique data uploaded against "debits" for content downloaded.

III. The Taxonomy of Computer Science Provides the Framework We Need to Place Internet Law and Section 230 in Context.

To understand the language of Section 230, it is imperative to both understand the systems the statute is meant to regulate, and the unique language that connects them to the law. Firstly, all of the online content Section 230 impacts sits on a server owned and managed by somebody. Some servers are easy to identify; for example, a box sitting on someone's desk or in a computer room. Some servers are harder to locate, existing virtually in a vast datacenter filled with hard drives and processors. Some servers are spread across many datacenters, existing as slices of processor time and a vast index of storage fragments constantly being shifted and reassembled by incredibly complex algorithms designed to optimize system performance and user responsiveness. Regardless of how a server is implemented, someone is responsible for what it does, who it serves, and what policies govern its use — perhaps several layers of individuals. In computer science, a server is more a concept than a thing, but it has an owner and an administrator. Importantly, ownership and rights over the content that a server holds are completely independent of the server itself.

The people and corporations accessing servers are users (though the term is broader than that). To maintain order in a multi-user system, system administrators assign permissions or privileges such that each interaction between a user and any definable asset or function in a network can be individually permitted, limited, disallowed, or otherwise controlled. Policies can include the ability to write, read, post, or publish user content. Many of these policies are automated and operate below an individual user's notice. For non-technical users, these systems often fade into the background, leaving users with the illusion of directly interacting with other users and their content directly.

Content and servers have a unique relationship. Like servers, content can exist in one place, or many, or in fragments everywhere, but if it is available on the Internet, content has an owner and is associated with a server. Information about content — its size, type, creator, location, popularity, where it is stored, when it was created, etc. — is not content, but metadata (data about data). Metadata itself has no creative element but may be part of a creative process of selecting and defining descriptors of the data a system needs to manage. Metadata and content relate to one another, but they are separate things, and they are stored and managed differently. Users focus on content as the valuable commodity on the Internet, but computer professionals spend much of their time interacting with metadata and the sophisticated databases that manage it. It is largely through metadata that content is organized, searched, recommended, and even moderated.

Section 230 is meant to encourage online content moderation. Moderation is the function of managing user content to ensure it is consistent with platform policies. It acts on content, though it may use metadata as part of the decision process as to which actions to take. Moderation decisions are dependent on context and judgement and are highly complex and nuanced; robust moderation requires an actual understanding of the content at issue. Although automation can assist in content moderation, its usefulness in this area has limits. One challenge in moderation is that users are often seeking to avoid having their content filtered, and as a result, actively work to overcome and elude moderation systems. The goal of moderation is to filter all content that violates platform policies, but none of the content that does not. Content moderation is a complex and highly studied area with a significant body of research.[8] Ultimately, moderation is as much art as it is science.

The current case is focused on recommendation systems; a term not mentioned in Section 230, but reflected in moderation systems that "… pick, choose … display, … search, … organize" content. *See* 47 U.S.C. § 230(f)(4). Recommendation is based on search. A search engine's task is to find, index, and rank searchable content — predominantly the text that makes up most of the Internet. To do this efficiently, search engines find content and build structured indexes by crawling websites across the Internet and parsing the words that appear on the pages they encounter.[9] Web crawlers note where on a

[8] For a recent overview of moderation practices and challenges and current research, *see, e.g.*, Mohit Singhal et al., *SoK: Content Moderation in Social Media, from Guidelines to Enforcement, and Research to Practice* (Oct. 27, 2022), https://arxiv.org/pdf/2206.14855.pdf.

[9] *See* Sergey Brin & Lawrence Page, *The Anatomy of a Large-Scale Hypertextual Web Search Engine.*

page various words appear, how they are punctuated, what words are nearby, whether they identify links to other sites, and whatever other information and metadata the page contains. Non-text searches rely more heavily on associated metadata, and are indexed independently, but the process is similar. Once the results are stored and organized in an index, search engines will take the list of words from a user query and identify which pages in the index contain all those words, generally returning a very long list. An engine's final task is to rank the matching results by relevance based on the associated metadata and their own algorithms, returning the list to the user. Content on the Internet is too vast, replicative, and nuanced to find a unique result for a generalized query and has been since before Section 230 came into being.

Search differs from moderation in that it relies mostly on metadata and is not reliant on an understanding of the underlying content (though this is evolving alongside the evolution of artificial intelligence and machine learning). In search, users and their content are often seeking to be found — sometimes by pretending they are something they are not. Here too users will actively work to game automated systems. The goal of search is to return content that best satisfies the user query as quickly and efficiently as possible, potentially before the user even asks or has completely defined what they are searching for. Given the vast amount of content

http://infolab.stanford.edu/~backrub/google.html (last visited Jan. 17, 2023).

available, an exhaustive real-time search of the complete Internet for every user query is not feasible, and thus search algorithms use various techniques to limit the search space for any individual request, relying on indexes and metadata to facilitate the task. Search is algorithmic, repeatedly yielding the same result for a fixed set of inputs. Recommendation, at its simplest, is search that predicts probable user interest without an actual query.

IV. As the Internet Emerged, the Law Struggled to Adapt to the Established Norms of a New Social and Commercial Domain.

The predecessors of the Internet (e.g., Arpanet, etc.) specifically prohibited commercial activity (restricted as it was to scholars and researchers). The emergence of the World Wide Web in the 1990's marked the emergence of legal issues alongside commercial and non-academic activity. Digital platforms, (where content replication was almost costless, users were pseudo-anonymous, and where jurisdictional boundaries and definitions were unclear) created novel legal challenges. [10] Commercialization created disputes over trademarks and copyrights, and contract enforcement online began to be litigated.

Starting in the early 1990's, cases involving content dissemination, ownership rights, and

[10] *See* Michael L. Rustad & Diane D'Angelo, *The Path of Internet Law: An Annotated Guide to Legal Landmarks*, 10 Duke L. & Tech. Rev. 1 (2011).

liabilities began to appear, with *Cubby, Inc. v. CompuServe, Inc.*[11], *Playboy Enters., Inc. v. Frena*[12], *MAI Sys. Corp. v. Peak Computer, Inc.*[13], *ProCD v. Zeidenberg*[14], etc. At the heart of several of these early suits was the challenge of adapting common law and intellectual property concepts to this emerging technology and its platform management practices which were anchored in the Internet's academic origins.

While legal professionals have chosen to use the analogy of printed books as a stand-in for the many elements that come together to create a web page, computer professionals use terms borrowed from their academic history such as libraries and journals, logic, and communications networks. This has resulted in overlapping terms and tangled meanings. For instance, in computer science, a user publishes their content when they take whatever steps the platform has defined to change the content's status from private to public (or shared). Whether the unpublished content originates from a user's desktop or is created and edited on a platform server, in

[11] 776 F. Supp. 135 (S.D.N.Y. 1991) (comparing the ISP's role to a newsstand or bookstore).

[12] 839 F. Supp. 1552, 1562 (M.D. Fla. 1993) (finding a computer bulletin board liable for copyright and trademark).

[13] 991 F.2d 518 (9th Cir. 1993) (finding that a software program made infringing copies each time it was installed).

[14] 86 F.3d 1447 (7th Cir. 1996) (confirming the validity of a shrink-wrap software license agreement).

computer science a *user* publishes content, not a platform. On some platforms content can be published, then unpublished, edited, and re-published — all under user control, independent of any platform employee. The content is often accompanied by metadata which enables the platform to automatically categorize it or simplify presentation, or to ensure it conforms to platform policies. From there, content, metadata, and what is presented to other users via the Internet can diverge. If the content is a movie, the metadata might include a thumbnail image, short preview clip, or a text description which might be the only thing displayed to another user as they browse or search for a video online. The actual movie is often stored separately from the clips and images — it may be thousands of miles away and in the hands of unaffiliated servers overseen by administrators and moderators of their own. Publishing in an Internet sense and publishing in common law are not at all synonymous. Section 230 uses publishing in its common law sense, not in its computer science sense; confusing the two can lead to unintended interpretations of Section 230.

Like publishing, "reading" is another term that stumbles at the threshold between law and computer science. Computers routinely read and write data into memory or storage. But this is very different than a human computer operator reading an uploaded file and *understanding* it. In fact, most computer content is only intended to be "machine-readable," with subsequent processes required to produce something a person can understand and appreciate. The software that makes up a computer game likely started out as human-readable programming code, was compiled into something only a computer would easily understand, and only when run by a user was the game content truly accessed and visible. The online content that a search engine provides has thus never been read and understood by a human as part of the search process — the content has been disassembled, analyzed, and organized by algorithms, then searched and returned without human intervention. Likewise, recommendation systems do not rely on the actual content when crafting recommendations, but rather on signals in the metadata that references the content and the users that access it. Only when search and recommendation engines seek to moderate the content they act on do they require human understanding of the actual content in context. Effective moderation requires content knowledge, a fact the text of Section 230 recognizes as well.

What ultimately connects all these processes and policies is the shared understanding in computer science that users are independent actors and that admins manage the machinery and act as the bureaucrats and public safety officers in the digital communities they oversee. The technologists that built and launched the Internet were taken completely off guard when courts assigned them responsibility for the words and actions of the users their systems supported. This seemed particularly cruel given the degree to which admins and moderators had empowered that independence.

V. Recommendation and Moderation are Fundamentally Connected and Predominantly Driven by the Third-Party Data Section 230 was Designed to Enable.

Because many creators want their content to be easily discoverable in the cacophony of the online world, they are motivated to participate in the collaborative search process. While quality content is a significant factor in search and recommendation ranking, metadata associated with the content users create is also key. Search engine optimization ("SEO") is the art of placing user-created metadata in relation to content so that, when crawled, search engines will rank the content higher in specific searches. Just as early BBS participants chose descriptive filenames for the work they uploaded to shared hard drives, creators today supply keywords, thumbnail images, short descriptions, or links to credentials and reputable source citations.

Bad actors can use SEO techniques to promote undesirable content and propaganda by tailoring metadata, creating false communities, and otherwise mimicking the trappings of legitimate content. Algorithms are in a constant battle to identify useful metadata and de-weight metadata-driven abuse. Search and recommendation algorithm secrecy is one tool that prevents bad actors from gaming the system and dominating search results or avoiding moderation with irrelevant, dangerous, and misleading links and content. Search and recommendation are ultimately a collaboration between users, search engines, and content creators.

Collecting, organizing, cleaning, and pruning metadata is a significant part of effective search, recommendation and moderation. There is no such thing as a purely neutral search or recommendation, since the algorithm, input factors and their weighting, and context are ultimately the product of human creative processes. Even where the content and metadata are produced by third parties, the algorithm that operates on them is the creative voice of the search provider — every search is the collaborative creation of multiple parties.

Recommendation engines are a variety of search that uses contextual, community and user-specific cues to replace explicit user queries. In most cases, the signals that drive recommendation decisions are themselves third party inputs: public recommendations by experts, peer reviews, user ratings, browsing and content access history of the user or users designated as friends, or new and popular content across a site's community generally. The platform's input is in how algorithms choose and weight the metadata that guides and organizes what content is displayed.

The reliance on both observed and user-generated metadata in search and recommendation systems makes metadata moderation an important ancillary function. For content that resides on third-party servers (which the moderator may lack permissions to erase), moderation by de-weighting metadata associated with risky content, or filtering metadata that doesn't seem to match the content in question, is a critical tool for policy enforcement.

For encrypted or obfuscated content, the ability to moderate or recommend is entirely reliant on metadata as an accurate descriptor of the content itself. Recommendation, moderation, and search algorithms can produce effective results on metadata alone, if it is accurate and complete. However, a content creator's incentive to manipulate metadata to either have their content rank unreasonably high in search or recommendation algorithms or to avoid moderation where their content violates policies, undermines the effectiveness of moderation. Effective moderation will always rely on knowledge of the underlying content, either by the moderator directly, or indirectly through feedback from users that have accessed the content. Metadata tools like user rankings, user blocking, and direct reporting are a key element for moderation in recommendation and search systems.

Even if all the content on the Internet was known and the manpower was in place to block and filter appropriately, perfect moderation is an impossibility. "I know it when I see it," as Justice Potter Stewart said of pornography in *Jacobellis v. Ohio*[15], highlights the subjectivity of content policy enforcement online, where cultural norms vary widely. Beyond local norms, the context surrounding content heavily influences its assessment.[16] Even content that is

[15] 378 U.S. 184, 197 (1964) (Stewart, J., concurring).

[16] *See A Dad Took Photos of His Naked Toddler for the Doctor. Google Flagged Him as a Criminal,* NY Times, https://www.nytimes.com/2022/08/21/technology/google-surveillance-toddler-photo.html.

objectively harmful, and even illegal, can be invaluable for prosecutors, researchers, or criminologists. Content moderation, like search, is thus a complex, evolving, and technically challenging field.

As a practical matter, there is no technical means to completely remove all copies of specific content across the entire Internet. While perfect duplicates can be searched out on active servers using algorithmic techniques, even subtle changes (a handful of pixels in an image) can significantly impact the effort to find replicas. Where content creators are actively working to modify and replicate content, the task is impossible. Likewise, it is impossible in practice to block a specific server and its content on the Internet, though with the cooperation of the international organizations that oversee Internet architecture and protocols, it can be made extremely difficult to keep a site online. Most importantly, there is an inevitable time lag between when content appears online and when moderation systems identify and act on it.

The practical approach to all these challenges has been for each admin and interactive computer service to moderate content on their own servers, based on their own locally defined policies, and to rely heavily on moderation tools embedded in their associated search and recommendation services to identify external content that their policies would not allow locally. Given the many performance advantages of controlling content storage, many content-intensive services mandate that user content be stored on the platform itself, which also improves policy enforcement. Interactive computer services ultimately use multiple approaches — including dedicated human moderators, automated systems and algorithms, and user feedback — to improve their ability to identify and act on content that violates platform policy.

VI. Section 230 Emerged Because Common Law Concepts Threatened to Subvert the Foundational Systems Upon Which the Internet was Built.

The dilemma for the emerging Internet community became clear following *Stratton Oakmont*[17]: the industry's interpretation of the case was that actually inspecting user content as a precursor to organizing or moderating it — or being alerted to harmful content by other users — could open an interactive computer service to liability. This was a radical change to a system which had been evolving for more than a decade from processes proven robust in the lab and classroom.

Section 230 and its operative language in subsections (c)(1) and (c)(2) together re-enabled content moderation by insulating it from the burden of the inappropriate common law precedents the courts were applying. Section 230(c)(2) was understood by the Internet community to immunize an interactive computer service from civil liability where they enable good faith moderation. Section 230(c)(2)

[17] *Stratton Oakmont, Inc. v. Prodigy Servs. Co.*, No. 31063/94, 1995 WL 323710, at *4-5 (N.Y. Sup. Ct. May 24, 1995).

addressed the common law punishment of a clumsy "Good Samaritan" intervention.[18] It also clarified that an interactive computer service has discretion in how it makes moderation decisions, in keeping with the First Amendment. This has the practical significance of emphasizing that policy control belongs to the platform provider and not to the state. Unresolved in Section 230(c)(2) however, was the fact that moderation often necessitates direct exploration of third-party content — the issue at play in *Stratton* — which can attract civil liability based not on action, but on knowledge.[19] Section 230(c)(1) is therefore understood to immunize an interactive computer service irrespective of their knowledge or awareness of the actual content from another information content provider, effectively making content knowledge irrelevant and removing the final impediment to moderation without civil liability. A traditional distributor with notice of actual content is

[18] Restatement (Second) of Torts § 324A (1965), *Liability to Third Person for Negligent Performance of Undertaking*, ("One who undertakes, gratuitously or for consideration, to render services to another which he should recognize as necessary for the protection of a third person or his things, is subject to liability to the third person for physical harm resulting from his failure to exercise reasonable care to protect his undertaking.").

[19] Restatement (Second) of Torts § 581 (1977). Transmission of Defamation Published by Third Person.

"(1) . . . one who only delivers or transmits defamatory matter published by a third person is subject to liability if, but only if, he knows or has reason to know of its defamatory character."; *see also Smith v. California*, 361 U.S. 147 (1959) (regarding illegal content).

treated as a publisher at common law.[20] Section 230(c)(1) prohibits treating an interactive computer service as a publisher of another's content. Section 230 places third-party content into a black box civil law cannot open, relegating the interactive computer service to the common law definition of third-party content distributor exclusively, without the ability for common-law publisher liability to attach. It is worth noting that in all cases, civil liability does not disappear; it remains with the original content creator. Knowledge of illegal (versus merely harmful) content and criminal law are also unaffected. The practical importance of Section 230(c)(1) is to enable moderation of harmful content (the law's goal) based on awareness of the actual content itself.

Without both parts of Section 230, information service providers would face increased costs and liability risk for moderating content. Liability risk would increase where platforms acquire *knowledge* of harmful content through inspection associated with moderation. Risk would also increase where platforms acquire *notice* of harmful content through user complaints. And finally, liability might increase where someone claimed moderation practices caused unreasonable harm. Costs would increase as providers add resources to improve moderation effectiveness in order to reduce these liability risks. This would be true despite investments in automation to reduce the level of human moderation required,

[20] Restatement (Second) of Torts § 578 (1977) ("one who repeats or otherwise republishes defamatory matter is subject to liability as if he had originally published it").

since human moderation is indispensable to effective moderation. Litigation risk would increase since users could look for compensation for moderation practices or failures they found harmful, and because cases based on First Amendment protections alone would be longer and more complex. Providers would need to weigh these costs against the commercial benefits of moderation. In enacting the statute, Congress decided that social policy was also fundamentally important, as is their constitutional role, and so they reweighed these risks to favor voluntary content moderation.

Without Section 230, interactive service providers have four practical options when setting third-party content moderation polices: 1) accept the incremental cost and liability risk and the obligation of improving and defending the effectiveness of their processes, or 2) avoid common-law publisher liability by moderating based only on metadata and other indirect signals and reacting promptly when given notice of harmful content by third parties, or 3) choose not to proactively moderate and simply respond to third party complaints, or 4) aggressively pre-moderate, such that nothing harmful appears on their public servers. A fifth option — to only host their own original content — would redefine what the Internet is and does as an information and communication network.

Avoiding content inspection and relying on metadata and other signals as a proxy for actual content has proven technically ineffective. Although metadata is often used to assess the probability that any given content violates policies, outcomes based on this process will always be best guesses and error rates will be high because many moderation decisions depend on context and judgement. Human intervention remains indispensable to content moderation (despite the high cost)[21], and in any case, direct content knowledge is virtually essential to reach reasonable outcomes. In accepting notice of harmful content from third parties, moderators are required to either adjudicate flagged content and accept the risk of under-reacting or act on all user complaints and accept the risk of over-reacting.[22] There is no practical way to ignore user complaints without attracting liability (as *Stratton* teaches us), and that option is at least commercially unreasonable for obvious reasons. The most effective direct moderation approach would be pre-moderation — reviewing all content before it is placed online. Moderators could review social media posts before

[21] *See* n.8, *supra.*

[22] This is the operative standard in some jurisdictions outside the United States. *See, e.g.*, Bob Tarantino, *If You Know About It, You're the Publisher – Website Operator Liability for Defamation* (Ent. & Media L. Signal May 15, 2015), http://www.entertainmentmedialawsignal.com/if-you-know-about-it-youre-the-publisher-website-operator-liability-for-defamation/. In some jurisdictions, moderation of illegal content based on notice is mandated, but exempt from liability, by statute. *See, e.g.*, Article 6(1)(b) and Recital 22 of Regulation (EU) No. 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).

they were displayed,[23] with every video and picture reviewed before it is posted. China, the country closest to this approach, uses a vast army of moderators and still cannot effectively pre-filter what users choose to share. Pre-moderation of content is also likely to only pass content that is particularly benign and be hostile to content that is outside the mainstream or in other ways controversial.

While an Internet based on pre-moderation of all content or the elimination of third-party content is easy to imagine, it is not as clear how removing proactive moderation would work. To avoid civil liability, presumably providers would need to strictly adhere to the common law role of third-party content distributors. It is unclear whether they could enact policies that prohibited certain content categories with the caveat that they would only be removed should the provider become aware of them. For any content brought to the provider's attention by a user or otherwise, the decision on whether to take it down would be weighted by the common-law publisher liability that could attach. Computer scientists are not lawyers, and certainly not judges, and the algorithms that would emerge would likely remove anything that drew a complaint and leave anything that did not. Easier still, providers could give users the ability to remove third-party content directly, just as a pedestrian passing a bulletin board could freely remove anything they object to that others have

[23] Many online comment systems support this capability, especially where the number of posts is small relative to the capacity to moderate.

posted. Presumably litigation over abuses of user-driven moderation would occur between users and without interactive service providers being involved, though this is not clearly the case. Another alternative might be for courts to arbitrate user content battles and for service providers to moderate only after a court decision. In any case, a new body of law would need to emerge alongside a new Internet, and the online policy precedents inherited from academic networks and BBS would need to be reworked.

VII. Analogies to Common Law that Go Beyond the Text of Section 230 are Inapt and Distort the Law's Purpose.

Traditional publishing is a weak and potentially dangerous analogy for a platform that enables billions of users to create and share content with billions of others with limited intermediary intervention. For many users, the experience is one of interacting directly with content creators; the intermediary fades into the background. At another level, bad actors see intermediaries as watchmen to be tricked and evaded as part of their criminal or anti-establishment enterprise. No traditional author was out to trick his publisher, and certainly the traditional publishing process created significant impedance to the rapid and scattershot dissemination of unfiltered, artless, and sometimes empty speech. Where users are their own publishers, and distribution is automatic and unrestricted, book seller analogies simply do not fit. A straight-forward reading of Section 230(c)(1) might be that an interactive service provider is simply not liable for content created by someone else. Likewise, an allegory that encourages people to help those in need gives the system administrator or software developer little concrete guidance in managing the content and community policies on a modern network. Section 230 seems to tell them it is good to try and remove content they believe the community will find objectionable, or to enable others to do the same.

Given that the language of computer science is often functional and symbolic, as opposed to literal, any analogy anchored purely in the language of the common law must be treated with suspicion. Section 230 is written using terms from both domains and should be interpreted appropriately. The functions listed in Section 230(f)(4) are clearly from computer science, as is "user" throughout the statute. "Development" in Section 230(f)(3) is also a computer science reference, meaning the "process of conceiving, specifying, designing, programming, documenting, testing, and bug fixing involved in creating and maintaining applications, frameworks, or other software components."[24] Because the language of computer science looks like ordinary English, as opposed to a foreign language, laws that import its terms must be treated carefully. Perhaps future legal scholars will embrace the use of italics, as they do with concepts anchored in Latin, to highlight where they stray from one language to the other.

[24] Software development, https://en.wikipedia.org/wiki/Software_development (last visited Jan. 17, 2023).

Of particular importance is the word “publisher” in Section 230(c)(1), a reference to common law and not to computer science. In the online world, publication, distribution, and presentation are often purely automated and under user control. If the proscription against applying common law encumbrance remains, the mismatch in analogy and practice is moot. Where the courts seek to apply common law analogies to the Internet, however, the lack of an equivalent to the publisher and distributor roles leads to erroneous and illogical outcomes. A traditional publisher does not moderate content. A traditional bookseller does not leave the door open for authors to place their works on the shelf independently. Section 230 states that interactive service providers should not be treated as (common law) publishers or speakers, which remains good law particularly because trying to do so would highlight how inappropriate such labels are on modern networks.

Finally, the output from a user search or platform recommendation does not constitute the speech of one or the other, but of the user in question, other users, the content creator, and platforms acting collaboratively. The metadata that underlies search and recommendation systems virtually always includes user-created information. More than that, the specific recommendations or search results that a single user sees will be the collaborative speech of a much larger community of users, often one that they themselves have curated, combined with the algorithmic decisions of the hosting platform. Unlike a newspaper, whose employees arrange the page, search and recommendation systems weigh the collaborative input of readers and experts alongside their own decision processes to provide results tailored for and by their users — often without direct knowledge of the actual content being recommended.

VIII. Commentary on the Future of Search, Recommendation and Moderation in the context of Section 230.

As this is being written, decades of computer science based on programming and logic are being challenged by systems based on training and algorithms that mimic the ways that humans are thought to process information. The ability to understand how these AI systems work is as limited as the ability to explain how human beings make decisions. These systems mimic intelligence without *being* intelligent, but increasingly this is enough to radically change how information is accessed and presented. An algorithm capable of reading all the text available on the Internet — every book and legal brief and social media post — and then responding to queries as if it understands the question and is not simply assessing the billions of patterns and probabilities invisibly embedded in its training data, is now available. It will seem primitive and limited within a year.

At the same time, tools to create artificial content — movies of historic figures saying and doing things never before said and done, for example — are in the field and advancing rapidly. From a world where seeing was believing but words could lie, we are entering a world where digital content of any type is capable of manipulation. In fact, much of this content emerges from only the smallest of creative seeds — a few words prompting an algorithm to generate an image or a movie that never previously existed, based not on copies of existing data but from the patterns hidden within existing data. Courts will soon wrestle with applying the laws that bind and promote human creativity and innovation in the context of algorithmic inventors and creators (generative AI).

The emerging challenges for search, moderation and recommendation will not be simply more content, but *different* content and different purposes. Search algorithms may no longer be tasked with finding the answer to a query, but instead to *create* the answer based on the sum total of the information they can reach. Moderation might evolve to encompass not just demoting what is undesired, but to identifying and weighting philosophies and epistemologies that emerge naturally from the totality of the content that they can access. Recommendation using all the Internet's knowledge could amplify or replace traditional education and tuition. In every case, innovation will rely on free access to as much of the world's knowledge and creativity as possible, coupled with the means to qualify, weight, and authenticate the content that it finds.

Section 230 is as important to the future Internet as it is to the Internet today and to the Internet it enabled when it became law. It focuses civil liability on Internet users for the content they create, leaving in place platform liability for its own creative content. These are appropriate, future proof, foundational principles which seem fully capable of maintaining order and innovation online for the foreseeable future.

CONCLUSION

Section 230 continues to fulfill a critical purpose as a bulwark against the temptation of imposing 19th century law to a system beyond contemplation when those precedents were born. At the same time, it provides Americans with the fundamental tools necessary to maintain order on the Internet and encourages their use. It is Congress' role to promote public policy through enacting the laws and to correct the court's errors where they misapply precedent and risk damaging the innovative and economic health of the country. Section 230 filled a vacuum that courts were forced to attempt to fill on their own.

There is little doubt that the Internet would have evolved differently had interactive service providers been held liable for harmful user content. At a minimum, the scope and diversity of voices online, and the diversity of digital creativity, would have been reduced. In turn, institutional and otherwise powerful voices would be more dominant. At its core, Section 230 removed some of the barriers that kept user content from flourishing online.

Section 230 re-established the processes and functions which computer science had developed for managing multi-user systems. It did this by removing the weight of common law principles that were poorly suited to a new domain based on distorted analogies that simply did not fit well. It is Congress' role to re-orient the courts when precedent goes astray.

Section 230 remains good law, with a purpose and scope that is as relevant and beneficial today as when it was enacted. It is a barrier to inappropriate legal analogy that protects the underlying technical and management systems upon which the Internet is built, while promoting socially beneficial practices. Its text is clear and specific and focused on allocating liability in a practical and logical way. If the medium is indeed the message,[25] then the Internet collapses space such that all speakers and content are in direct contact with each other. Rules that blunt traditional concepts of intermediary liability seem appropriate to that idea.

Section 230 does just enough, and not too much, to allow the machinery of the Internet to maintain order without suppressing participation. It lifts the burden of inapt common law and First Amendment analogies, leaving room for the law to develop new rules for a new medium. Above all, it makes clear that an interactive service provider is not liable for third-party content, and equivalently, for third-party inputs in content recommendations.

The judgment of the court of appeals should be affirmed.

[25] Marshall McLuhan, Quentin Fiore, *The Medium Is the Massage, An Inventory Of Effects* (Gingko Press, 2001).

Respectfully submitted,

James H. Hulme
Counsel of Record
Nadia A. Patel
ARENTFOX SCHIFF LLP
1717 K Street, NW
Washington, DC 20006
(202) 857-6000
james.hulme@afslaw.com

Bruce Gustafson
President & CEO
DEVELOPERS ALLIANCE
6218 Georgia Avenue NW
Suite #1 PMB 3045
Washington, DC 20011
bruce@developersalliance.org

January 19, 2023