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

BRIEF OF CENTER FOR DEMOCRACY & TECHNOLOGY AND 6 TECHNOLOGISTS AS AMICI CURIAE IN SUPPORT OF RESPONDENT

GREGORY NOJEIM
Counsel of Record
SAMIR JAIN
EMMA LLANSÓ
CAITLIN VOGUS
Center for Democracy & Technology
1401 K St NW, Suite 200
Washington, DC 20005
(202) 637-9800
gnojeim@cdt.org

Counsel for Amici Curiae

TABLE OF CONTENTS

TABLE OF CONTENTS
(continued)

Page

TABLE OF AUTHORITIES

Page(s)

Cases

Statutes

TABLE OF AUTHORITIES
(continued)

Page(s)

TABLE OF AUTHORITIES
(continued)

Page(s)

TABLE OF AUTHORITIES
(continued)

Page(s)

TABLE OF AUTHORITIES
(continued)

Page(s)

TABLE OF AUTHORITIES
(continued)

Page(s)

TABLE OF AUTHORITIES
(continued)

Page(s)

INTERESTS OF AMICI CURIAE[1]

Amici curiae are the Center for Democracy & Technology and six technologists who study and analyze algorithms that recommend content on the Internet.

The **Center for Democracy & Technology** ("CDT") is a non-profit public interest organization. For more than twenty-five years, CDT has represented the public's interest in an open, decentralized Internet and worked to ensure that the constitutional and democratic values of free expression and privacy are protected in the digital age. CDT regularly advocates before legislatures, regulatory agencies, and courts in support of First Amendment rights on the Internet and other protections for online speech, including limits on intermediary liability for user-generated content.

Robin Burke is Professor of Information Science at the University of Colorado, Boulder, and director of That Recommender Systems Lab, a research group studying recommender systems and related technologies.[2]

Matt Cutts is the former Administrator of the United States Digital Service and former Distinguished Engineer at Google. Cutts created the initial version of SafeSearch, Google's family-safe filter, and led Google's efforts to fight search spam as part of Google's Search Quality team.

---

[1] No counsel for a party authored this brief in whole or in part, and no person other than amici or their counsel made a monetary contribution to this brief's preparation and submission.

[2] Individual amici are listed in alphabetical order. Titles are given for identification purposes only.

Dean Eckles is Associate Professor at the MIT Sloan School of Management, where he is also affiliated with the Schwarzman College of Computing. He has expertise in statistics and data science, including as applied to online platforms.

Michael Ekstrand is Associate Professor of Computer Science at Boise State University and co-director of the People and Information Research Team, a research group studying information access systems, including recommender systems and search engines, from a human-centered perspective.

Brandie Nonnecke is Director of the CITRIS Policy Lab at the Center for Information Technology Research in the Interest of Society and the Banatao Institute and Associate Research Professor at the Goldman School of Public Policy, UC Berkeley. Nonnecke has expertise in recommender systems and platform governance.

Jonathan Stray is Senior Scientist at the Center for Human-Compatible AI, UC Berkeley, specializing in the design of recommender systems.

Amici are concerned that the interpretation of Section 230(c)(1) urged by Petitioners relies on a distinction between the "display" and "recommendation" of online content that does not exist as a technical matter. Amici write to provide the Court with technical information about how algorithms that select and rank online content work and to explain that Petitioners' interpretation would have dire consequences for the free expression interests of Internet users and the public.

INTRODUCTION AND SUMMARY OF ARGUMENT

The plain text of Section 230(c)(1) immunizes providers of interactive computer services ("providers") from claims that treat them as the "publisher or speaker" of third-party information. *See* 47 U.S.C. § 230(c)(1).

The Ninth Circuit correctly held that Petitioners' claim—that Respondent Google, Inc. violated the Anti-Terrorism Act (ATA) by allegedly recommending ISIS content posted to YouTube to other YouTube users—seeks to treat Google as the publisher of third-party content. Petitioners and amicus the United States acknowledge that claims based on the display of third-party content, or the failure to block or remove content, treat a defendant as a publisher. Pet'rs' Br. 24–26 (recognizing that a claim treats a provider as a publisher when its gravamen is the dissemination of content); United States' Br. 20. They argue, however, that some or all claims based on the "recommendation" of third-party content do not treat a defendant as a publisher. Pet'rs' Br. 26–29; United States' Br. 26–28.

However, the distinction Petitioners and the United States attempt to draw does not exist. Making choices about what content to display and how to display it is the quintessential activity of traditional publishers. The same is true online. Every interactive computer service provider that displays content must choose what to display from an overwhelming number of available possibilities and order it in some way. Those choices are inherently the provider's "recommendations" as to what content a user should view, typically made using algorithms that rank all possible content according to a set of criteria chosen by the provider, with the highest-ranked items displayed to the user. The distinctions Petitioners suggest between display and recommendation of content—such as whether content is displayed through a system that provides information the user is "actively seeking," or "targeted" to a specific user—are technologically arbitrary and unworkable. Recommendation is functionally indistinguishable from selecting and ordering or ranking items for display, something every provider must do.

Moreover, if the Court holds that claims based on displaying content in ways that allegedly recommend it do not treat providers as publishers, providers will be discouraged from using novel ranking algorithms, which will harm Internet users. Ranking algorithms are necessary to make many services useful. A search engine that did not recommend content by ranking results based on an algorithm designed to identify sites most relevant to a user's query would be close to worthless. Ranking algorithms are also a key component of content moderation for most major services; reducing the visibility of problematic (but not actually illegal) content better protects freedom of expression than deleting it entirely.

In addition, a holding excluding "recommendations" made using ranking algorithms from Section 230's liability shield will create strong incentives for providers to limit speech. Because content moderation inevitably results in errors, a provider cannot perfectly remove or block only content that exposes it to liability. Instead, a provider would rationally seek to minimize the risk of liability by taking steps such as imposing categorical limits on the type of content it ranked and displayed or increasing its reliance on automated content moderation tools, the inherent limitations of which will exacerbate the tendency to over-remove innocuous or even beneficial content. As a result, Internet users will be less able to speak freely and everyone will have less access to information.

Even though claims based on algorithmic selection and ranking of content treat providers as publishers, Section 230's liability shield is not absolute. Among other things, it does not apply if the provider even partially creates or develops the information that gives rise to a legal violation. Other areas of law, such as data protection and antitrust law, fall outside the scope of Section 230. Thus, holding that claims based on "recommendations" such as those at issue here treat a provider as a publisher for purposes of Section 230(c)(1) will still leave open avenues to hold providers legally responsible for their actions.

ARGUMENT

I. Claims that seek to hold providers liable for recommending content treat them as publishers.

Section 230(c)(1) provides that "[n]o provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider." 47 U.S.C. § 230(c)(1). Petitioners argue that their claim that Google violated the ATA by recommending certain content to users does not treat Google as a "publisher or speaker" of third-party information. In support of this argument, Petitioners and amicus the United States attempt to distinguish between the "display" of third-party content, which they acknowledge is immunized by Section 230(c)(1), and the "recommendation" of such content, which they assert is not. The Court should reject their plea to differentiate between display and recommendation in determining whether a claim treats a provider as a publisher or speaker for purposes of Section 230(c)(1) for three reasons.

First, just as traditional publishers do, interactive computer service providers must make choices about what content to display and how to display it. Those choices are necessarily a "recommendation" to users that reflects providers' judgments about what content may be most interesting to a user, be most important for a user to view, or meet some other criterion. *See* United States' Br. 27–28 ("The appearance of a video in a user's queue thus communicates the implicit message that YouTube 'thinks you, the [user]—you, specifically—will like this content.'" (quoting *Force v. Facebook, Inc.*, 934 F.3d 53, 82 (2d Cir. 2019) (Katzmann, J., concurring in part and dissenting in part)). Claims based on such recommendations are based on providers' choices of what content to display and how to display it—a quintessential activity of publishers.[3]

Second, many providers across the Internet use sophisticated algorithms to select and order content for display. A holding that attempts to distinguish between claims based on the display of content and

[3] Recommendations inherent in the selection and ordering of content for display are distinct from a provider's own speech explicitly recommending particular content (e.g., a video review). Such speech falls outside the scope of Section 230(c)(1) because it is not information provided by "another information content provider." 47 U.S.C. § 230(c)(1).

claims based on recommendation of content would have sweeping ramifications for many online services beyond social media.

Third, attempts to draw a distinction between recommendations and display of content using factors such as whether the provided results are information that the user did not specifically request or whether the recommendations are "targeted" are neither workable nor descriptive of underlying technical differences.

A. Every provider that displays third-party content must select and order it in some way.

Publishing requires making choices about what content to display and *how* to display it. Those choices are effectively "recommendations" that certain content will be most interesting or appealing to readers. For example, print newspapers choose which news stories to place on "A1" and which to place on "A26" based on editors' judgments about what reports are most important or will be of greatest interest to their readers; the choice to display content on the front page recommends those reports to readers by drawing their attention to them. *See, e.g.*, Suzanne Daley, *Making the Front Page: How All the News Fits in Print*, N.Y. Times (Dec. 23, 2019); Steven Clayman & Ann Reisner, *Gatekeeping in Action: Editorial Conferences and Assessments of Newsworthiness*, 63 Am. Soc. Rev. 178, 178–79 (1998) (explaining that "stories are chosen from the available pool, prioritized in terms of newsworthiness, and arranged within a newspaper or newscast").

Similarly, online providers must make choices about what content to display and how to display it.

Every method of displaying third-party content online involves selecting and ordering content in some way, because there is far too much content available for any individual to consume. Vast amounts of third-party content, from videos to social media posts, tweets, blog posts, and classified advertisements, are posted online every day. Resp.'s Br. 1, 10.

Even on services where the user explicitly "follows" a list of sources, the number of available items can be overwhelming. The typical Facebook user, for example, might have more than 1000 new posts from "friends" that are available for display every time they log in. Akos Lada et al., *How Does News Feed Predict What You Want to See?*, Meta (Jan. 26, 2021), https://tinyurl.com/46bchswt.

Thus, by numerical necessity, providers must make choices about what content to display to a user at any one time, and in what order. To implement these choices on a computer, they rely on algorithms. The algorithmic selection and ordering of content has long been called "ranking" in the context of web search. *See, e.g.*, Lawrence Page et al., *The PageRank Citation Ranking: Bringing Order to the Web,* Stanford InfoLab Publication Server (Jan. 29, 1998), https://tinyurl.com/ydmu2chn (describing Google's initial search ranking algorithm). The term "recommendation" emerged in the early 1990s to describe algorithms used to select and order content for a broad range of online services, such as shopping, news, movies, and music. Zhenhua Dong et al., *A Brief History of Recommender Systems*, arXiv (Sept. 5, 2022), https://tinyurl.com/yx7syn23.

What we now call search and recommendation have long been understood as closely related functions. Nicolas J. Belkin & Bruce W. Croft, *Information Filtering and Information Retrieval: Two Sides of the Same Coin?*, Comm. of the ACM, Dec. 1992, at 29, https://tinyurl.com/2hvjpn7t. In contemporary technical practice, both "search" and "recommendation" are understood as a type of "ranking" built on general purpose techniques for matching content with context. *See, e.g.*, Sahin Cem Geyek et al., *Fairness-Aware Ranking in Search & Recommendation Systems with Application to LinkedIn Talent Search,* KDD '19: Proc. of the 25th ACM SIGKDD Int'l Conf. on Knowledge Discovery & Data Mining 2221 (July 2019), https://tinyurl.com/y42pkv5x ("Ranking algorithms form the core of search and recommendation systems . . . ."); *see also* Michael D. Ekstrand et al., *Fairness in Information Access Systems*, Found. & Trends in Info. Retrieval, July 2022, at 1 (explaining that "tools to facilitate information access take a variety of forms, including information retrieval and recommendation systems").[4]

The goal of recommendation algorithms is to produce a ranked or ordered list of items. *See* Integrity Institute & AlgoTransparency's Br. 7. Typically, this is done by ranking all candidate items according to some rubric and then displaying to the user only the top scoring items, which means that ranking is also the core method of selecting items for display for most platforms, something explicitly contemplated in

---

[4] Indeed, many services explicitly connect search and recommendation. *See, e.g.*, Mounia Lalmas-Roelleke, *Recommending and Searching @ Spotify*, https://tinyurl.com/5kp7z45j (2019) ("We can conclude that recommender systems and search are also two sides of the same coin at Spotify. They work together to help fans get the music they will enjoy listening." (Emphasis omitted)).

Section 230. 47 U.S.C. § 230(f) (defining "interactive computer service" to include "access software provider" of tools that, *inter alia*, "pick," "choose," "analyze," or "organize" content). For example, a provider may choose to display the top five most viewed posts on the service alongside other content the user views. *See* Badrul Sarwar et al., *Analysis of Recommendation Algorithms for E-Commerce*, EC '00: Proc. of the 2nd ACM Conf. on Electronic Com. 158 (Oct. 2000), https://tinyurl.com/2p8huusv. This "Top-N" method originated several years before Section 230, *see* Paul Resnick et al., *GroupLens: An Open Architecture for Collaborative Filtering of Netnews*, CSCW '94: Proc. of the 1994 ACM Conf. on Computer Supported Cooperative Work 175 (Oct. 1994), https://tinyurl.com/5xjr8m7z, and is used to this day, *see* Luke Thorburn et al., *How Platform Recommenders Work*, Understanding Recommenders (Jan. 20, 2022), https://tinyurl.com/34kd7c9a.

Ranking content necessarily results in a recommendation from a provider to a user; conversely, it is impossible to display content without ordering it in some way. Ranking reflects a provider's judgment of what content the user would want to view. Even chronological ranking of content reflects a provider's recommendation that a user should view content that is recent.[5]

---

[5] Pure chronological ordering should not be presumed a standard or preferred approach because it often compares poorly to other ranking methods. It gives the most exposure to whoever posts most often, which rewards frequent, spammy posting, and

The rubric used to evaluate and order items may be designed to achieve a variety of goals.[6] Providers sometimes rank content by the estimated likelihood of evoking certain user responses, such as giving positive feedback on the item (by rating it, sharing it, or "liking" it) or increasing the amount of time the user spends on the service. These responses are sometimes called "engagement." Increased engagement is often correlated with value to the user, the content creator, and the platform—that is, engagement is not a universally accurate signal, but generally people give more attention to things that are useful or interesting to them. Priyanjana Bengani et al., *What's Right and What's Wrong with Optimizing for Engagement*, Understanding Recommenders (Apr. 27, 2022), https://tinyurl.com/ynk2kmw2.

---

can incentivize the creation of sensationalist content that falls just short of violating a platform's rules. *See* Daphne Keller, *Amplification and Its Discontents*, Knight First Amend. Inst. at Colum. Univ. (June 8, 2021), https://tinyurl.com/4965a4bt. When Facebook tested a purely chronological feed in 2018, it found that users manually marked more posts as irrelevant, and metrics measuring things like spam, misinformation, harassment and hate speech "shot through the roof." Alex Kantrowitz, *Facebook Removed The News Feed Algorithm In An Experiment. Then It Gave Up*, Big Tech. (Oct. 25, 2021), https://tinyurl.com/zfxfd4w9.

[6] In addition, providers may or may not consider the content of the items to be ranked. As discussed below, for example, collaborative filtering can be used to rank items solely on the basis that users who liked item X also tend to like item Y, which does not depend on the provider knowing the content of items X or Y. Thus, contrary to claims by amicus Senator Josh Hawley, Sen. Josh Hawley's Br. 14–15, recommendation of items does not necessarily mean that the provider has knowledge of the content of those items.

Most providers also incorporate other types of information into their ranking algorithms in service of other goals, such as presenting a diversity of information, prioritizing sources judged to be reliable, featuring smaller creators to encourage more people to make content for the platform, or selecting content that is most appealing for advertisers to advertise alongside. *See* Thorburn, *supra*.

Regardless of the provider's goals, ranking reflects choices about how to display content, just as a traditional publisher makes choices about how to display content. These choices necessarily result in a provider displaying certain items over others to a user. Claims based on these "recommendations" seek to treat a provider as a publisher.

B. Ranking systems for recommending content are used across the Internet.

Ranking systems that recommend content were created and adopted before the enactment of Section 230 and are now used across the Internet, including by many services that host third-party content. As a result, a holding that recommendations fall outside the scope of Section 230 would have broad ramifications.

A variety of "algorithmic" ranking systems had already been deployed when Congress adopted Section 230. Recommender systems were widely available to Usenet discussion forum readers with the creation of GroupLens in 1994. Resnick, *supra*. This and other systems used around this time used collaborative filtering, perhaps the earliest modern recommendation algorithm. This technique is based on the logic that users who liked X also liked Y and is still in wide use today, with many variations. Will Hill et al., *Recommending and Evaluating Choices in a Virtual Community of Use*, CHI '95: Proc. of the SIGCHI Conf. on Hum. Factors in Computing Sys. 194 (May 1995), https://tinyurl.com/pv4ek7n5. Net Perceptions, Inc. made this technology available as a commercial product for sale to online retailers in 1996, and one of its first customers was Amazon. Dong, *supra.*

Since then, a wide variety of online services have employed various algorithmic ranking systems to recommend content. The social news aggregator service Reddit, for example, uses several different types of user voting systems to rank users' posts and comments. *See, e.g.*, Amir Salihefendic, *How Reddit Ranking Algorithms Work*, Hacking & Gonzo (Dec. 8, 2015), https://tinyurl.com/29w8k9az. Voting-based ranking requires the use of fairly complex algorithms to work well (e.g., preventing winner-takes-all feedback loops where only already-popular items are displayed). *Id.*

Other services, such as search results, news, online shopping, and advertising, have long displayed algorithmically customized results. Early web portals like Yahoo displayed customized lists of links based on a user's city. *See* Dan Tynan, *The History of Yahoo, and How It Went From Phenom to Has-Been,* Fast Company (Mar. 21, 2018), https://tinyurl.com/2s3cm7bm. Google search has long provided users with personalized results based on past signals of interest. *See, e.g.,* Danny Sullivan, *Google Now Personalizes Everyone's Search Results,* Search Engine Land (Dec. 4, 2009), https://tinyurl.com/47fau86t. Modern search engines use various types of information (such as which results the user previously clicked on) to personalize results.

If the Court concludes that claims based on the recommendation of content do not treat providers as publishers for purposes of Section 230(c)(1), a broad range of algorithmic ranking systems that make customized determinations of what content to display may be excluded from Section 230's liability shield. The potential loss of immunity would affect not only social media companies, but also media streaming platforms, search engines, news services, online marketplaces, and more. As discussed in Section II below, this would harm individuals' ability to access information and speak online.

C. Petitioners' attempts to draw lines between display and "recommendation" of online content fail.

Petitioners posit two ways to distinguish recommendation of online content from the other methods of selecting and ordering content that all providers must use to display content to users, but both ways fail.

First**,** Petitioners attempt to distinguish between the algorithmic ranking performed by a search engine in response to a user query and the ranking performed by other types of services. Pet'rs' Br. 44 ("Whether disseminated material was requested by the recipient affects the availability of the [S]ection 230(c)(1) defense."); *see also* Pet. 31–32 (distinguishing between "a system that provides to a user information that the user is actually seeking (as does a search engine) and a system utilized by an internet company to direct at a user information (such as a recommendation) that the company wants the user to have"). But this distinction does not reflect technical reality.

Ranking systems sometimes categorize the "signals" that may be used to select and order items along several axes. These axes include signals that are more explicit (such as a query) versus more implicit (such as the item the user is currently viewing); signals that are ephemeral (such as current location) versus persistent (such as consumption history); and signals that are submitted by the user (such as past item ratings) or someone else (such as other users' item ratings). Ekstrand, *supra*. In this schema, a user query is an explicit, ephemeral, individual signal. However, this is not the only kind of signal that provides information about what a user is actually seeking, even in the case of search.

To begin with, an explicit query does not fully define a user's information need. If a user types the query "politics podcast" into Spotify, this still requires inference and judgment as to which politics podcasts to select as search results and how to prioritize them. Moreover, as discussed above, many search engines direct information to a user based on more than just the current query. Sullivan, *supra*. Perhaps 5% to 40% of the top Google search results to a given query, depending greatly on what is searched for, are personalized based on information other than the explicit query. *See* Tobias D. Krafft et al., *What Did You See? A Study to Measure Personalization in Google's Search Engine,* 8 EPJ Data Sci. 38 (2019), https://tinyurl.com/yshxcdxn.

Both implicit and explicit user signals provide valuable, complementary information about the information a user seeks. *See* Gawesh Jawaheer et al., *Comparison of Implicit and Explicit Feedback from an*

Online Music Recommendation Service, HetRec '10: Proc. of the 1st Int'l Workshop on Info. Heterogeneity & Fusion in Recommender Sys. 47 (Sept. 26, 2010), https://tinyurl.com/yck86vyk. Implicit signals provide irreplaceable information about user needs. It can be difficult to convince users to reliably enter explicit signals such as item ratings, because users incur a cost (to their time and attention) with no immediate perceived benefit. In contrast, implicit data produced during the normal course of interacting with a system (such as the length of time a user views an item) is a rich source of information about what the user wants to see. Certain other signals are neither clearly explicit nor clearly implicit. When a user "likes" a post, they may intend to communicate approval to the post author and/or direct the ranking algorithm. This underscores the difficulty in using explicitness as a legal criterion.

Further, a user can be "actually seeking" the information that they receive via a recommendation algorithm even when they do not provide any explicit information. Certain recommendation systems, such as news aggregators, operate largely in this query-free manner. By accessing the service at all, the user indicates that they are interested in the service's recommendation of content.

Second, Petitioners' question presented suggests a distinction between the recommendation of content that is "targeted" versus untargeted. *See also* Pet. App. 83a-84a (Berzon, J. concurring). But "targeted" has no standard technical definition. To the extent that the term is intended to include customized recommendations that display third-party content to a particular user based on information about or provided by that user, that would encompass the typical practices of numerous services including search engines. *See supra* at 13–14. A holding that claims based on "targeted" recommendations of content do not treat a provider as a publisher for purposes of Section 230(c)(1) would not give providers any guidance as to which of the many different kinds of customized ranking algorithms and signals fall outside of Section 230(c)(1)'s liability shield.

For example, the Facebook News Feed selects and orders items posted by people the user has "friended" and displays other types of customized recommendations, such as a list of groups that many of one's friends have previously joined. Lada, *supra*. When Spotify adds songs or podcast episodes to the end of an existing playlist, it is guided by the items the user has previously placed on that playlist. Likewise, Amazon's algorithm may display a book to a user that has a similar topic to a book a user previously purchased. YouTube itself uses a range of signals to infer a specific user's intent, including some that are explicit (e.g. user search queries and user satisfaction surveys), some that are implicit (e.g. viewing history), and a variety of signals that are customized but not individualized (e.g. geography, type of viewing device, or time of day). *See* Zhe Zhao et al., *Recommending What Video to Watch Next: A Multitask Ranking System*, RecSys '19: Proc. of the 13th ACM Conf, on Recommender Sys. 43 (Sept. 10, 2019), https://tinyurl.com/yc3nstn4. Other services well outside of social media recommend third-party content to particular users based on similar sets of signals.

Petitioners do not explain which of these signals, alone or in combination, result in recommendations that are "targeted." Indeed, no natural dividing line distinguishes "targeted" recommendations from other recommendations. The lack of clarity that would result from an attempt to draw such a line would have the same harmful effects as would a holding that all recommendations are not "publishing" activity.

II. Holding that claims based on recommending content do not treat providers as publishers will harm Internet users and the public.

Internet users benefit from having access to services that rank and moderate content using a variety of approaches. One way that providers compete is through their use of ranking algorithms. *See, e.g.*, Yahoo! Inc., *Yahoo Announces Acquisition of The Factual, Expanding its Commitment to Trusted News and Information* (Sept. 6 2022), https://tinyurl.com/4cr6xnvs (explaining that Yahoo acquired The Factual based on its novel news ranking algorithm). Providers also offer users differing approaches to content moderation including through the use of ranking algorithms, with some significantly moderating content and others taking a more hands-off approach. *See* Nicole Buckley & Joseph S. Schafer, *"Censorship-Free" Platforms: Evaluating Content Moderation Policies and Practices of Alternative Social Media*, For(e)dialogue, Jan. 2022, https://tinyurl.com/3jjrhhc7.

Holding that claims based on providers' recommendation of content do not treat providers as publishers will create incentives for them not to use novel ranking algorithms to manage third-party content and thereby reduce users' ability to select among services that offer different ranking and content moderation approaches. Some services that may have entered the market to offer an improved system for recommending content may no longer do so for fear of potential liability.

Moreover, concerns about liability from recommending content—and the inevitability of errors when moderating huge volumes of content—may cause some providers to categorically limit certain topics or types of speech or to over-remove third-party content. The steps providers take to minimize the risk of liability will harm the free expression interests of users and the public.

A. Internet users benefit from having access to services that take a variety of approaches to content recommendation.

Internet users benefit from having access to services that recommend (i.e., rank) content. Ranking allows users to find useful information online, where there is far too much information for any individual to consume or sort through. *See supra* Section I. In the digital age, the ability to locate and receive customized information online is crucial to participation in democratic self-governance. *See* Natali Hellberger, *On the Democratic Role of News Recommenders,* 7 Digital Journalism 993 (2019), https://tinyurl.com/2p8e5e5h. Systems that do not require explicit queries can identify and deliver the information that most directly concerns the interests of each user, regardless of whether they know to look for it. See Jonathan Stray, *Who Should See What When? Three Principles for Personalized News*, NiemanLab (July 25, 2012), https://tinyurl.com/3nt3ukee. Ranking systems may also help make the news media more responsive to citizens' information needs, as news organizations adapt to provide the type of information that these systems identify as most interesting to users. *See* Hellberger, *supra*.

Ranking has also become an important content moderation tool. Providers may deprioritize or "downrank" undesirable content to reduce its visibility or reach. Emma Llansó et al., *Artificial Intelligence, Content Moderation, and Freedom of Expression*, Transatlantic Working Group on Content Moderation Online & Freedom of Expression (Feb. 26, 2020), https://tinyurl.com/49purumm; Tarleton Gillespie, *Do Not Recommend? Reduction as a Form of Content Moderation*, Soc. Media + Soc'y, July-Sept. 2022, https://tinyurl.com/d82m3j4d. Providers use downranking, for example, to curtail the spread of "borderline" content that nearly violates their rules but does not actually do so. Gillespie, *supra*. Downranking allows providers to moderate such content without removing it entirely, meaning it remains available for a user to find directly. In this way, downranking can preserve a measure of freedom of expression by keeping content accessible to those who specifically seek it out.

The benefits to users from providers' use of ranking algorithms for recommendations and content moderation will be lost if providers face the risk of liability from reliance on such algorithms and therefore eliminate or substantially curtail their use.

B. Free expression interests of users and the public will be harmed if the Court holds that claims based on recommendation of content do not treat providers as publishers.

If the Court adopts Petitioners' interpretation that Section 230(c)(1) does not apply to claims based on recommendations, providers are likely to over-remove content for fear of liability. Because "recommendations" made using ranking algorithms are inherent to the process of displaying information, providers' ranking activities could expose them to liability for the content they display. As a result, they will have a strong incentive to block or remove content to avoid the risk of liability—exactly the dynamic Section 230 was designed to prevent.[7]

That incentive will apply most clearly to content that may itself give rise to potential liability, particularly controversial speech. As a result, users will be limited in their ability to express themselves, and online information available to the public will be skewed. A restaurant review site, for example, that might otherwise rank highly a negative review (e.g., because it is recent or contains specific details) may instead block or remove the review, particularly in the face of a threat of a defamation claim from the restaurant's owner. By contrast, the site will have no such incentive for positive reviews.

But the free expression harms will not stop there. Because content moderation is necessarily imperfect, providers cannot reliably target for removal only content that itself gives rise to liability. *See* Evelyn Douek, *Governing Online Speech: From "Posts-As-Trumps" to Proportionality and Probability*, 121 Colum. L. Rev. 759, 792 (2021). To minimize the risk of liability, they will inevitably over-remove innocuous

---

[7] The distinction the United States attempts to draw between claims based on the failure to block or remove content and claims based on the recommendation of content collapses upon itself. Claims that seek to hold providers liable for ranking content necessarily seek to hold them liable for failing to block or remove it from the pool of content from which they are generating rankings to determine what content to display. Thus, claims based on the recommendation of content are necessarily based on the failure to block or remove that content.

and even beneficial speech, or content that they would have previously merely downranked, and thereby negatively impact the ability of users and the public to receive information.

Some providers may respond to the risk of liability by eliminating categories of content users may post to avoid the risk they would otherwise fail to block or remove specific content in that category that could result in liability. For example, some providers may prohibit all content related to terrorism, including news reports, documentation of terrorists' human rights violations, or even anti-indoctrination materials. Smaller providers with fewer resources to spend on content moderation may be especially likely to adopt this approach. This result limits people's ability to discuss or seek out information about controversial topics.

Providers have used categorical topic bans in the past when faced with the threat of liability for third-party content. Following enactment of the Allow States and Victims to Fight Online Sex Trafficking Act/Stop Enabling Sex Traffickers Act, Pub. L. 115-164, 132 Stat. 1253 (2018), for example, some platforms responded by prohibiting content related to sex altogether. *See, e.g.*, Shannon Liao, *Tumblr Will Ban All Adult Content on December 17th*, The Verge (Dec. 3, 2018), https://bit.ly/2SmoC5A.

Further, given the scale of user-generated content online, providers will likely have to increase their reliance on automated content moderation tools. *See* Hannah Bloch-Wehba, *Automation in Moderation*, 53 Cornell Int'l L.J. 41, 55 (2020) (describing the rising use of automated tools to moderate content due in part to providers' "focus on scale"). The inherent limits of automated content moderation technologies will cause providers to over-remove innocuous or even beneficial content.[8] *See generally* Carey Shenkman et al., *Do You See What I See? Capabilities and Limits of Automated Multimedia Content Analysis*, Ctr. for Democracy & Tech. (May 2021), https://tinyurl.com/47m2nbyd.

One of the chief limitations of automated content moderation technologies is that they struggle to discern context. Natasha Duarte et al., *Mixed Messages? The Limits of Automated Social Media Content Analysis*, Ctr. for Democracy & Tech. 9, 19 (Nov. 2017), https://tinyurl.com/2p829azn. Tools that rely on matching previously identified violative content with content newly uploaded by users cannot accurately moderate content that may be objectionable in one context but acceptable in another. For example, a matching-based tool could not determine whether an image of known terrorist propaganda was posted to recruit new adherents or to be debunked or analyzed. *See, e.g.*, *Countering Daesh Propaganda: Action-Oriented Research for Practical Policy Outcomes*, The Carter Ctr. (Feb. 2016), https://tinyurl.com/msjaf6f2 (including images created by ISIS or Daesh in an interdisciplinary guide to "counter Daesh propaganda"). In contrast, the matching tool PhotoDNA is successful in part because the context in which known child sexual abuse

---

[8] These limitations also demonstrate the continued need for Section 230, which encouraged providers to engage in content moderation by immunizing them from claims based on their failure to remove certain content. Contrary to the suggestion of some amici, *see* Sen. Josh Hawley's Br. 16; Counter Extremism Project (CEP) & Hany Farid's Br. 21–24, and the Ninth Circuit below, Pet. App. at 42a–43a, advances in technology do not allow providers to accurately detect, flag, and remove "dangerous content" at scale.

material (CSAM) is posted is irrelevant, since the posting of CSAM is always illegal.[9] Shenkman, *supra*, at 15.

Automated tools that go beyond merely matching previous violative content and try to make predictions about whether novel content violates a service's rules may also fail to understand context. Duarte, *supra*, at 12–13. For example, even a sophisticated text classifier may be unable to distinguish between the use of the term "slant" as a slur to insult people of Asian descent from the "reclaiming" of the slur by a member of the Asian community. *See, e.g.*, *Matal v. Tam*, 137 S. Ct. 1744, 1754 (2017) (describing how the lead singer of the band "The Slants" "chose this moniker in order to 'reclaim' and 'take ownership' of stereotypes about people of Asian ethnicity"). Similarly, an image classifier "may be able to identify nudity, but not make a judgment about whether that nudity is occurring in the context of artistic expression or abuse." Shenkman, *supra*, at 29.

Providers may also respond to the risk of liability for user-generated content they rank by configuring machine learning tools to block or flag content even when comparatively less sure that the content is prohibited. *See* Nadia Chowdhury, *Automated Content Moderation: A Primer*, Stanford Cyber Pol'y Ctr. 5 (2022), https://fsi-live.s3.us-west-1.amazonaws.com/s3fs-

---

[9] Thus, the fact that algorithmic content moderation tools are relatively successful in removing certain kinds of content, such as CSAM, does not show that they can "control their platforms" and accurately moderate other kinds of user-generated content, contrary to the claims of some amici. *See, e.g.*, CEP & Hany Farid's Br. 24.

public/automated_content_moderation_a_primer.pdf (describing how "a company sets a threshold for what level of confidence is required before removing or taking other action on a piece of content"). As a result, providers will remove even more innocuous content than they already do.

These concerns are not hypothetical. Providers relying on automated content moderation tools routinely misidentify benign content as forbidden, including specifically in the enforcement of providers' policies against content that promotes terrorism. For example, an "enforcement error" caused Instagram to remove a series of user-generated posts about the Al-Aqsa Mosque—one of Islam's holiest sites—because the term "al-Aqsa" also appears in the name of a designated terrorist organization. Jon Porter, *Instagram Blames "Enforcement Error' for Removal of Posts About Al-Aqsa Mosque*, The Verge (May 31, 2021), https://tinyurl.com/yn8kujej. Facebook reportedly suspended dozens of Middle Eastern journalists after potentially "miscategorizing their accounts as having links to terrorism." Olivia Solon, "*Facebook Doesn't Care": Activists Say Accounts Removed Despite Zuckerberg's Free-Speech Stance*, NBC News (June 15, 2020), https://tinyurl.com/4ve5nypm.

The uncertainty created by a holding that claims based on recommendation of content do not treat providers as publishers may cause some providers to rely even more on automated content moderation tools notwithstanding the known limitations of those tools. This, in turn, will exacerbate the existing problem of over-removing harmless and even beneficial speech, meaning that Internet users will be less able to speak freely online, and the public will lose access to valuable information.

III. Threats to individuals' rights caused by providers can be addressed in other ways.

Amici acknowledge that providers may act in ways that threaten individuals' rights and our democracy. *See, e.g.*, Jonathan Stray et al., *Building Human Values into Recommender Systems: An Interdisciplinary Synthesis*, arXiv (July 20, 2022), https://tinyurl.com/5bwdncv2; Gabriel Nicholas, *Shedding Light on Shadowbanning*, Ctr. for Democracy & Tech. (Apr. 2022), https://tinyurl.com/yc2czmjw; Brandie Nonnecke et al., *Harass, Mislead, & Polarize: An Analysis of Twitter Political Bots' Tactics in Targeting the Immigration Debate for the 2018 US Midterm Election*, 19 J. of Info. Tech. & Pol. 423 (2020), https://tinyurl.com/4b764amn. While these threats cannot be coherently addressed by drawing artificial distinctions between display and recommendation of content and declaring the latter to be outside the scope of "publishing" activity, that does not mean providers are free from legal responsibility. Section 230 does not give providers blanket immunity, and other approaches in law and policy also can hold providers accountable.

A. Section 230 does not shield providers from liability when they are responsible in whole or in part for the development of the information at issue.

Some claims against providers fall outside the scope of Section 230(c)(1)'s liability shield even when the claim treats the provider as a "publisher." Congress also separately required that the provider be publishing "information provided by *another* information content provider." 47 U.S.C. § 230(c)(1) (emphasis added). Section 230 states that a provider may be deemed an "information content provider" when it "is responsible, in whole or in part, for the creation or development of information." *Id.* § 230(f)(3).

Many lower courts have examined the issue of what it means for a provider to "develop" information so that it becomes an "information content provider." As other amici have discussed, *see, e.g.*, Lawyers Comm. for Civil Rights Under Law et al.'s Br. 10–14, the leading interpretation stems from the Ninth Circuit's ruling in *Fair Housing Council v. Roommates.com,* 521 F.3d 1157 (9th Cir. 2008). In that case, the Ninth Circuit held that the defendant website Roommate was not immunized by Section 230(c)(1) from claims that it had violated the Fair Housing Act and California law by requiring its users to provide preferences for the race, sex, sexual orientation, and family status of potential roommates, because Roommate was the developer, at least in part, of the allegedly illegal content. *Id.* at 1166. The court held that "a website helps to develop unlawful content, and thus falls within the exception to [S]ection 230, if it contributes materially to the alleged illegality of the conduct." *Id.* at 1168. This "material contribution" test has since been applied by nearly every Circuit court. *See Kimzey v. Yelp!, Inc.*, 836 F.3d 1263, 1269 n.4 (9th Cir. 2016).

Whether the material contribution test is the appropriate way to determine when providers are responsible, in whole or in part, for the creation or development of information is not properly before the Court in this case, however. The question presented for review on certiorari focuses solely on a separate element of Section 230 immunity—whether "targeted recommendations" of third-party content are distinct from "traditional editorial functions" of a publisher.[10] Pet. i. The scope of what it means to treat a provider as a publisher of third-party content is distinct from the question whether a provider is itself partially responsible for that content's creation or development. *See Izumi Seimitsu Kogyo Kabushiki Kaisha v. U.S. Philips Corp.*, 510 U.S. 27, 31–32 (1993) (per curiam) (stating that "[a] question which is merely 'complementary' or 'related' to the question presented in the petition for certiorari is not 'fairly included therein'" (quoting *Yee v. Escondido*, 503 U.S. 519, 537 (1992)). Accordingly, the Court should not address in this case the myriad issues that stem from the definition of "information content provider" and the meaning of "developing" content *See* Sup. Ct. R. 14.1(a).

But holding here that claims based on recommendation of content treat providers as publishers would nevertheless leave open for this Court to determine in an appropriate future case whether and when recommendations or the use of ranking algorithms may mean a provider is

[10] Petitioners later raised the prospect, in their opening brief, that recommendations may be content provided by the information content provider itself, due to the website's generation of URLs and notifications to users about third-party content. Pet'rs' Br. 34–40. This argument is undeveloped and technically unsound. As the United States explains, "the creation of navigational hyperlinks is inherent in the provision of an online platform. . . . A website does not act as an information content provider by taking the technical steps necessary to render user-generated online content visible to others." United States' Br. 33 (citation omitted).

"responsible, in whole or in part, for the creation or development of information," 47 U.S.C. § 230(f)(3), and therefore is not entitled to the protection of Section 230(c)(1).

B. Other areas of law outside of Section 230 also can address unlawful actions by providers.

Providers' business practices, and their consequences, are governed by many areas of law unaffected by Section 230. Laws that neither directly restrict speech nor create incentives for providers to do so are better suited to address specific concerns about providers' actions.

For example, comprehensive privacy or data protection laws are the best way to limit online services' ability to collect, use, package, and sell personal information about users and to limit the use of such information for targeting content. Several states have passed privacy laws that limit what online services may do with individuals' personal information and that provide people with greater control over their data. *See, e.g.*, California Consumer Privacy Act of 2018, Cal. Civ. Code §§ 1798.100–.199.100; Conn. Data Privacy Act, Conn. Pub. Act 22-15; Colorado Privacy Act, Colo. Rev. Stat. §§ 6-1-1301–1313. Congress also has considered data privacy legislation; most recently, it debated (but failed to pass) the bipartisan American Data Privacy and Protection Act, H.R. 8152, 117th Cong. (2022). Comprehensive federal privacy legislation could address some of the harms caused by content ranking systems and providers' collection of users' personal information.

Competition laws can address the risks to competition, innovation, and consumer choice from a concentration of power within a few major firms in the online ecosystem. Appropriate enforcement of the antitrust laws is consistent with free expression interests. *See Associated Press v. United States*, 326 U.S. 1, 20 (1945) ("The First Amendment affords not the slightest support for the contention that a combination to restrain trade in news and views has any constitutional immunity."). Here too Congress has debated enacting additional laws focused on large online platforms. *E.g.* American Innovation and Competition Online Act, S. 2992, 117th Cong. (2022); Open App Markets Act, S. 2710, 117th Cong. (2022). Concerns about the power and influence these firms wield in the marketplace, and how they may curtail potential competitors from entering markets that they control, are the appropriate focus of competition law.

Further, any online service that engages in unfair and deceptive trade practices, including in representations and omissions it may make about the operation of its ranking algorithms and recommendations, may appropriately fall into the scope of the Federal Trade Commission's existing enforcement authority. *See* 15 U.S.C. § 45(a).

In short, there are multiple avenues for policymakers to address concerns over providers' data collection, concentration of power, and potential unfair or deceptive practices. These alternative paths, coupled with the limits in Section 230 itself, mean that providers can be held accountable for unlawful practices without resort to Petitioners' technically baseless interpretation of the term "publisher" under Section 230.

CONCLUSION

For the foregoing reasons, this Court should affirm the Ninth Circuit's decision below.

Respectfully submitted,

GREGORY NOJEIM
Counsel of Record
SAMIR JAIN
EMMA LLANSÓ
CAITLIN VOGUS
Center for Democracy & Technology
1401 K St NW, Suite 200
Washington, DC 20005
(202) 637-9800
gnojeim@cdt.org

Counsel for Amici Curiae

January 18, 2023