

IN THE UNITED STATES COURT OF APPEALS
FOR THE NINTH CIRCUIT

META PLATFORMS, INC., *Plaintiff-Appellant*,
v.

Case Nos. 26-5120, 26-5121, 26-5132

ROB BONTA, in his official capacity as Attorney General of the State of California, *Defendant-Appellee*.

(Consolidated Appeals)

TIKTOK INC. and TIKTOK USDS JOINT VENTURE, LLC, *Plaintiffs-Appellants*,
v.

ROB BONTA, in his official capacity as Attorney General of the State of California, *Defendant-Appellee*.

GOOGLE LLC; YOUTUBE, LLC, *Plaintiffs-Appellants*,
v.

ROB BONTA, in his official capacity as Attorney General of the State of California, *Defendant-Appellee*.

*On Appeal from the United States District Court for the Northern
District of California
Nos. 5:25-cv-09792-EJD, 5:25-cv-09789-EJD, 5:25-cv-09795-EJD
Hon. Edward J. Davila, District Judge*

**BRIEF OF CITP TECH POLICY CLINIC AS AMICUS CURIAE IN
SUPPORT OF DEFENDANT-APPELLEE AND AFFIRMANCE**

/s/Mihir Kshirsagar

Mihir Kshirsagar
Princeton CITP Tech Policy Clinic
315 Sherrerd Hall
Princeton University
Princeton, NJ 08540
Tel: 609-258-5306
mihir@princeton.edu

Attorney for Amicus Curiae

DATE: October 20, 2026

DISCLOSURE STATEMENT AND STATEMENT OF INTEREST OF AMICUS CURIAE

The CITP Tech Policy Clinic is part of Princeton University’s Center for Information Technology Policy (“CITP”), which works to better understand and improve the relationship between technology and society. The Clinic provides nonpartisan research, analysis, and commentary in the public interest. This brief is the product of a clinic project and reflects the individual views¹ of the following scholars who study the behavior of online platforms and the effect that they have on users and society at large:

Manoel Horta Ribeiro, Assistant Professor in the Department of Computer Science at Princeton University.

Mihir Kshirsagar, Director of the Tech Policy Clinic at Princeton’s Center for Information Technology Policy.

Alissa Cooper, Executive Director, Knight Georgetown Institute.

¹ This brief is principally authored by Manoel Horta Ribiero and Mihir Kshirsagar. We thank Tavi Prem for research assistance. Titles and institutional affiliations are provided for identification purposes only and the views expressed should not be regarded as the position of Princeton University or any other institution.

Pursuant to Fed. R. App. P. 29(a)(2) and 29(a)(4)(E), amicus states that all parties have consented to the filing of this brief, that no party's counsel authored this brief in whole or in part, and that no party, party's counsel, or person other than amicus curiae or its counsel contributed money that was intended to fund preparing or submitting this brief.

Date: October 20, 2026

/s/Mihir Kshirsagar
Mihir Kshirsagar
315 Sherrerd Hall
Princeton University
Princeton, NJ 08540
Tel: 609-258-5306
mihir@princeton.edu
Attorney for Amicus Curiae

TABLE OF CONTENTS

	Page
DISCLOSURE STATEMENT AND STATEMENT OF INTEREST OF AMICUS CURIAE	i
TABLE OF AUTHORITIES.....	iv
SUMMARY OF THE ARGUMENT	1
ARGUMENT	3
I. Platforms do not require persistent implicit behavioral signals to produce functional and useful feeds.	3
II. Limiting reliance on persistent implicit signals for minors is grounded in research.	11
III. Relying on explicit signals works and produces useful feeds.	15
CONCLUSION	26

TABLE OF AUTHORITIES

	Page(s)
Cases	
<i>NetChoice, LLC v. Bonta</i> , 761 F. Supp. 3d 1202 (N.D. Cal. 2024), <i>aff'd in part, rev'd in part</i> , 152 F.4th 1002 (9th Cir. 2025)	1
Statutes	
Cal. Health & Safety Code § 27000.5	<i>passim</i>
Cal. Health & Safety Code § 27001(a)	13
Legislative Materials	
Assemb. Comm. on Privacy & Consumer Prot., Analysis of S.B. 976, 2023–2024 Reg. Sess. (Cal. July 2, 2024)	19
Foreign Authorities	
Regulation (EU) 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), 2022 O.J. (L 277) 1	23
Other Authorities	
Aisha Malik, <i>Instagram’s New ‘Your Algorithm’ Tool Gives You More Control over the Reels You See</i> , TechCrunch (Dec. 10, 2025)	25
Akaash Kolluri et al., <i>Alexandria: A Library of Pluralistic Values for Realtime Re-Ranking of Social Media Feeds</i> , 20 Proc. Int’l AAAI Conf. on Web & Soc. Media 1309 (2026)	28
Alex Moehring et al., <i>Better Feeds: Algorithms That Put People First</i> (Knight-Georgetown Inst. 2025)	<i>passim</i>
<i>An Update on Fulfilling Our Commitments Under the Digital Services Act</i> , TikTok Newsroom (Aug. 4, 2023)	24

Andrew M. Guess et al., <i>How Do Social Media Feed Algorithms Affect Attitudes and Behavior in an Election Campaign?</i> , 381 Science 398 (2023)	15
Arvind Narayanan, <i>Understanding Social Media Recommendation Algorithms</i> (Knight First Amend. Inst. Mar. 9, 2023)	4, 7, 16, 18
Hunt Allcott et al., <i>The Effect of Deactivating Facebook and Instagram on Users' Emotional State</i> (Nat'l Bureau of Econ. Rsch., Working Paper No. 33697, 2025)	17
Hunt Allcott et al., <i>The Welfare Effects of Social Media</i> , 110 Am. Econ. Rev. 629 (2020)	17
<i>Introducing Instagram Teen Accounts: Built-In Protections for Teens, Peace of Mind for Parents</i> , Meta Newsroom (Sept. 17, 2024)	25
Jon Kleinberg, Sendhil Mullainathan & Manish Raghavan, <i>The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization</i> , 70 Mgmt. Sci. 6336 (2024)	17
Lauren Forristal, <i>YouTube Is Working on a Feature That Will Fix the Messy Home Feed</i> , TechCrunch (Nov. 25, 2025)	26
Lindsay Popowski et al., <i>Social Media Feed Elicitation</i> , in Proc. 2026 CHI Conf. on Hum. Factors in Computing Sys. (2026)	27
<i>More Ways to Discover New Content and Creators You Love</i> , TikTok Newsroom (June 3, 2025)	26
Nick Clegg, <i>New Features and Additional Transparency Measures as the Digital Services Act Comes Into Effect</i> , Meta Newsroom (Aug. 22, 2023)	24
Omar El Malki et al., <i>Bonsai: Intentional and Personalized Social Media Feeds</i> , in Proc. 2026 CHI Conf. on Hum. Factors in Computing Sys. art. 1395 (2026)	27
Paul Covington, Jay Adams & Emre Sargin, <i>Deep Neural Networks for YouTube Recommendations</i> , Proc. 10th ACM Conf. on Recommender Sys. 191 (2016)	29, 31

Tom Cunningham et al., <i>What We Know About Using Non-Engagement Signals in Content Ranking</i> (2024) (unpublished manuscript)	7
Ying Li et al., <i>GenRec: An LLM-Backed Recommendation Ranker at Netflix</i> (Aug. 10, 2026) (unpublished manuscript)	28
<i>Your Algorithm, Your Words: New Ways to Shape What You See on Facebook</i> , Meta Newsroom (Oct. 13, 2026)	25

SUMMARY OF THE ARGUMENT

This brief provides a technical primer on why the provisions of California’s Protecting Our Kids from Social Media Addiction Act (the “Act”) that prohibit the default delivery of addictive feeds to minors do not limit the ability of the major social media platforms to recommend and serve valuable content.

As the district court explained in an earlier challenge to the Act, the Act’s coverage definition and seven statutory exceptions (Cal. Health & Safety Code §§ 27000.5(a)(1)–(7)) draw a clear functional line: the statute restricts feeds that recommend content to minors based on persistent data collected about a user’s previous internet activity, while exempting feeds that recommend content based on a user’s explicit requests or non-persistent session context. *See NetChoice, LLC v. Bonta*, 761 F. Supp. 3d 1202, 1210 (N.D. Cal. 2024), *aff’d in part, rev’d in part*, 152 F.4th 1002 (9th Cir. 2025) (“The seven feed-based exceptions are varied, but they generally serve to exempt feeds from [the Act’s] provisions if those feeds recommend content based on a user’s explicit requests and not based on data collected about that user’s previous internet activity.”). The State reads the statute the same way,

explaining that it regulates only “the mechanism of using an individual user’s harvested data to compile addictive feeds personalized to that user.” Answering Br. 34 (ECF 59).

Plaintiffs and their amici contend that without personalization their services would be impossible for minors to navigate. They argue that young users will face an “ocean of irrelevant content” (YouTube Br. 2, be forced to sift through an “overwhelming morass” (TikTok Br. 44), or be left with online tools that are “functionally useless” (Chamber of Progress Br. 16 n.35). In this view, the only alternative to profiling a minor’s historical behavior to serve content is an unfiltered chronological feed or a degraded search function.

Plaintiffs’ argument is not supported by the actual design of their systems. *First*, we explain how their industrial pipelines work and why restricting behavioral profiling does not impair a platform’s ability to select, moderate, and order high-quality, age-appropriate items to display on a user-agnostic basis. *Second*, we explain how the behavioral record maintained by a personalized feed is critical to the engagement feedback loop that the Act targets, and explain why empirical research on the costs of relying on passive engagement signals supports

California's choice to limit the use of that record for minors as a reasonable and tailored policy. *Third*, we demonstrate that feeds built on explicit user signals—including search, topic selection, creator subscriptions, or explicit feed controls—are technically viable and are being developed and deployed with success.

ARGUMENT

I. PLATFORMS DO NOT REQUIRE PERSISTENT IMPLICIT BEHAVIORAL SIGNALS TO PRODUCE FUNCTIONAL AND USEFUL FEEDS.

To evaluate the claim that social media services cannot function effectively without persistent behavioral personalization, we need to understand how modern feeds are built. Most large social media platforms no longer organize their feeds chronologically. Instead, platforms such as Facebook, Instagram, TikTok, YouTube, X (formerly Twitter), and Reddit rely on systems that rank content based on predictions of how a user will respond. *See* Arvind Narayanan, *Understanding Social Media Recommendation Algorithms* 7–11 (Knight First Amend. Inst. Mar. 9, 2023),

<https://knightcolumbia.org/content/understanding-social-media-recommendation-algorithms>.

At a high level, the feeds operate by taking information about a candidate piece of content and a user, and estimating outcomes such as the probability that the user will view the video, like the post, or remain watching for a given duration. Those predictions are then combined into a ranking score that determines which content appears near the top of a user's feed. See Alex Moehring et al., *Better Feeds: Algorithms That Put People First* 12–13 (Knight-Georgetown Inst. 2025),

[https://kgi.georgetown.edu/wp-content/uploads/2025/02/Better-Feeds - Algorithms-That-Put-People-First.pdf](https://kgi.georgetown.edu/wp-content/uploads/2025/02/Better-Feeds-Algorithms-That-Put-People-First.pdf) (“*Better Feeds*”).

At the scale of a major platform, it is impractical to run a sophisticated ranking model over every available piece of content whenever a user opens the app. Consequently, a platform's feed is populated through distinct functional processes that are documented in the technical literature published by companies such as Google, Meta, Pinterest, and LinkedIn:

Content Moderation and User-Agnostic Item Scoring. As content is uploaded and indexed, platforms evaluate the items

themselves—independent of any particular viewer—to remove material that violates platform policies (such as unlawful content, nudity, or self-harm promotion), attach objective metadata (such as language, topic category, and media format), and assign viewer-independent item scores to prioritize or deprioritize categories of posts (such as demoting spam or clickbait, or elevating authoritative sources). *See Better Feeds* at 5, 8.

Candidate Generation (Retrieval). When a user opens the service, a retrieval stage narrows a vast corpus of billions of eligible posts down to a plausible subset of candidates—either by matching the user’s explicit requests (such as search queries, followed creators, or selected topics) or by drawing on behavioral profiles. *Id.* at 5.

Behavioral Engagement Prediction. Where a platform uses behavioral personalization, a machine-learning model uses data collected from the user’s past interactions to calculate numeric prediction scores—estimating the probability that the user will click, like, or spend time watching each candidate post—and weights those predictions in a “value model” to order the feed. *See Narayanan, supra*, at 18–24; *Better Feeds* at 12–13; Tom Cunningham et al., *What We*

Know About Using Non-Engagement Signals in Content Ranking (2024) (unpublished manuscript), <https://arxiv.org/abs/2402.06831>. But behavioral predictions are not inherent to the ranking stage: candidates can instead be ordered using signals that do not depend on a persistent record of the particular user’s past behavior, including item-level quality or safety scores, recency, aggregate popularity, or relevance to an explicit query or selected topic.

Plaintiffs’ briefs blur these distinct functions and the different signals that can be used to perform them, suggesting that restricting one particular input—persistent behavioral data—will impair platforms’ ability to maintain a safe, relevant, and high-quality service. YouTube argues, for example, that “[g]iven the sheer volume of content on YouTube, personalization is critical to ensuring that users receive content that is high quality for them,” and that “YouTube cannot maintain its content-quality standards without personalization.” YouTube Br. 38. YouTube further suggests that without persistent personalization, a minor could be served “content in a language she does not speak” (*id.*), or be exposed to “spam and low-quality content” (*id.* at 52; *see also id.* at 42). TikTok similarly states that “[a]bsent

personalization, TikTok’s recommendation system could not identify users for which such content should be excluded or disfavored.” TikTok Br. 13.

As a technical matter, however, viewer-dependent behavioral predictions are distinct from a platform’s user-agnostic content-moderation and quality standards. They are also distinct from non-persistent session information and device settings available when a user opens the service. We can think of these as different classes of signals:

- **Engagement Predictions (Viewer-Dependent).** These estimate the likelihood that a specific user will take an action on an item—such as clicking, watching, or lingering—based on a persistent profile of that user’s past behavior. *Better Feeds* at 12.
- **Item-Level Scores (User-Agnostic Quality and Safety Evaluations).** An item-level score is “a score assigned to an item independent of the viewer with a clear valence of value, that is, where a high score is more desirable and a low score is less desirable in some way.” *Id.* at 8. As the report notes, “[i]tem-level scores could include an item’s informativeness, toxicity, or

likelihood to be spam.” *Id.* Platforms can score age-appropriateness and source authority the same way. Because these scores evaluate the item itself rather than a user’s behavioral history, they do not depend on any user’s data.

- **Content Properties and Session Context.** These reflect objective properties of the content (such as its language, topic, or freshness) and non-persistent context from the user’s current session or device settings (such as device language locale, accessibility settings, or an active search query). Cal. Health & Safety Code §§ 27000.5(a)(1), (2), (3).

Below is a table illustrating how a platform’s functions in selecting and ordering items to display operate independently of persistent behavioral profiling:

Feed Component	Function	Signals Used	Restricted?
Item-Level Quality & Safety Evaluations	Enforces user-agnostic platform standards by removing or demoting spam, clickbait, and harmful content, and elevating informative or age-appropriate items	Viewer-independent item scores (spam classifiers, toxicity ratings, source authority, age ratings)	No. Operates user-agnostically.
Basic Delivery & Device Parameters	Ensures displayed items match the user's spoken language, device capabilities, and accessibility needs	Technical device information (e.g., OS/browser language locale), user-selected accessibility settings, or age-assurance signals	No. Exempted under §§ 27000.5(a)(1) and (a)(3).

User-Directed Selection & Ordering	Retrieves and orders items that fulfill the user's deliberate requests and topic interests	Search queries, followed creators, selected topic channels, and explicit "show less" commands	No. Exempted under §§ 27000.5(a)(2) and (a)(4).
Passive Behavioral Prediction	Predicts which items will maximize a specific minor's viewing time based on past behavior	Persistent record of past watch history, millisecond dwell time, scroll speed, and clickstreams	Yes. Restricted for minors absent verifiable parental consent (§§ 27000.5(a), 27001(a)).

As the table shows, restricting persistent behavioral profiling does not impair a platform's ability to determine which items are eligible, safe, or high-quality to display. A platform does not need a longitudinal behavioral profile of a fourteen-year-old's past viewing habits to filter out spam, demote borderline or clickbait content, or elevate authoritative, age-appropriate videos. Those evaluations reside in item-level scores assigned to the videos themselves. Notably, platforms like Reddit and Twitter built large user bases without passive behavioral

prediction using explicit user input and aggregate-level information about the popularity of content.

Regardless of whether a minor's feed is populated through a search query, a topic selection, or a subscription list, the platform can still evaluate, score, and order every candidate item for quality and safety before displaying it. And platforms can use aggregate data from user interactions to inform that judgment. Plaintiffs' own declarants describe moderation systems that act on posts user-agnostically before any personalization runs. Answering Br. 23 (citing 3-MetaER-450 ¶ 7; 4-TikTokER-620–22 ¶¶ 3–6; 7-YouTubeER-1341 ¶¶ 24–25).

II. LIMITING RELIANCE ON PERSISTENT IMPLICIT SIGNALS FOR MINORS IS GROUNDED IN RESEARCH.

California's decision to turn off personalization based on implicit signals by default for minors is a reasonable policy choice that is supported by the evidence. Because a persistent behavioral record helps power the engagement feedback loop, limiting the default collection and use of that record for minors is a reasonable and tailored way to interrupt that loop. *See, e.g.,* Andrew M. Guess et al., *How Do Social Media Feed Algorithms Affect Attitudes and Behavior in an Election*

Campaign?, 381 Science 398 (2023),

<https://www.science.org/doi/10.1126/science.abp9364> (finding that moving adult users out of algorithmic feeds and into chronological feeds substantially decreased the time they spent on the platforms and their activity).

Over the past fifteen years, we have seen that the content users consume on major social media platforms has become progressively less intentional, particularly as platforms have pivoted toward short-form video feeds such as TikTok, Instagram Reels, and YouTube Shorts. Unlike earlier platforms that were organized around interpersonal connections and explicit subscriptions, these short-form video feeds do not depend on a user-constructed social graph. Earlier social media asked users to explicitly construct a network of people and topics they wanted to hear from. TikTok-style feeds instead infer what users are interested in from their passive behavior and construct an information environment on their behalf. *See* Narayanan, *supra*, at 21–23.

As platforms have moved toward less intentional, behaviorally inferred feeds, researchers have documented growing concerns regarding compulsive use and user well-being. Large-scale randomized

deactivation experiments find that stepping away from modern, behaviorally ranked social media platforms measurably improves users' emotional well-being. *See, e.g.,* Hunt Allcott et al., *The Welfare Effects of Social Media*, 110 Am. Econ. Rev. 629 (2020); Hunt Allcott et al., *The Effect of Deactivating Facebook and Instagram on Users' Emotional State* (Nat'l Bureau of Econ. Rsch., Working Paper No. 33697, 2025), <https://www.nber.org/papers/w33697> (finding in a randomized study that deactivating Facebook and Instagram improved emotional state indices, with effects on Instagram particularly pronounced among women under 25).

The computer science literature explains why optimizing recommender systems around passive behavioral signals raises concerns. The result is a mismatch: recommender systems have become good at learning from what users do, while giving users comparatively limited ways of telling those systems what they want. *See* Jon Kleinberg, Sendhil Mullainathan & Manish Raghavan, *The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization*, 70 Mgmt. Sci. 6336 (2024), <https://doi.org/10.1287/mnsc.2022.03683> (modeling how, when a user's

in-the-moment choices diverge from what she wants, optimizing a feed for engagement can leave her worse off).

When a ranking model is trained to maximize predicted engagement using implicit behavioral telemetry—such as how many milliseconds a user pauses while scrolling, whether a video loops in the background, or how long a session lasts—it treats these passive behavioral signals as inputs to predictions of engagement. Yet as the Knight-Georgetown Institute’s report observes, “not all signals are created equal.” *Better Feeds* at 10. Passive behavioral signals can diverge from a user’s deliberate preferences; for example, “users in a disorganized mental state may take a long time to look through different videos as they struggle to make a decision,” or may linger on provocative or distressing content that they do not affirmatively wish to see. *Id.* at 10 & n.24; *see also* Narayanan, *supra*, at 25–27 (explaining how engagement optimization can create feedback loops that amplify borderline or emotionally charged content).

In light of this research, California’s decision to require parental consent before platforms rely on persistent behavioral profiling for minors—while leaving platforms free to recommend content based on a

minor’s explicit requests—is a reasonable, research-backed policy choice. *See Better Feeds* at 37.

III. RELYING ON EXPLICIT SIGNALS WORKS AND PRODUCES USEFUL FEEDS.

The practical question is whether platforms can build functionally useful feeds based on explicit user signals that fulfill the legislative aim to allow a minor “to search for or follow whatever content or content creators they want.” Assemb. Comm. on Privacy & Consumer Prot., Analysis of S.B. 976, 2023–2024 Reg. Sess., at 12 (Cal. July 2, 2024) (3-MetaER-489). We explain how platforms can do so using many of their existing practices.

Plaintiffs and their amici suggest that without personalization, users are left with unworkable alternatives—either an unranked chronological feed that buries relevant posts beneath unwanted content (such as subjecting a teenager who dislikes sports to “an endless stream of NFL early-season highlights,” *TikTok Br. 6*), or a degraded search function where an aspiring young artist searching for “painting tips” might receive only videos about home renovation (*YouTube Br. 43*) and

where users cannot discover content they do not already know to search for (NetChoice Br. 30; CDT Br. 8–10).

In practice, removing persistent behavioral inference does not force a platform to present an unfiltered firehose of global uploads or a broken search bar. While a chronological feed is one lawful option under the statute, it is far from the only one. Under the Act’s exceptions for explicit user requests and non-persistent session context (Cal. Health & Safety Code §§ 27000.5(a)(1)–(4)), platforms can deliver rich content discovery using well-established information retrieval mechanisms as well as modern explicit-preference controls.

Standard information retrieval and explicit-preference mechanisms directly address each of the navigation challenges identified in Plaintiffs’ briefs:

- **Topic Selection and Subscription Graphs.** A platform complying with the Act does not need to place a teenager into an undifferentiated chronological feed of all uploads on the platform. The platform can curate topic feeds without using any individual user’s behavioral data and allow the user to choose which feeds to view and follow. The Act leaves platforms free to curate feeds

around any content or subject in the same way that Spotify gives users a choice of different playlists. When a user selects categories of interest during onboarding or in a topic menu—such as “music,” “chemistry,” and “basketball”—and excludes others, the platform retrieves candidates from those chosen categories and ranks them using the viewer-independent item-level quality scores described in Section I, *supra*. Within those user-selected topics, the user discovers new creators and videos without needing to know their names in advance, while unrequested topics (such as NFL highlights for a user who did not select football) do not appear.

- **Query Disambiguation and Non-Persistent Search.**

Similarly, when a user searches for an ambiguous term like “painting tips” (exempt under § 27000.5(a)(2)), modern search systems do not require a longitudinal profile of the user’s past behavior to return helpful results. Standard search interfaces routinely resolve ambiguity by displaying diverse top-ranked results across common meanings and presenting immediate topic-refinement options (such as “Watercolor,” “Acrylic,” or “Home

DIY”) that allow the user to clarify her intent with a single tap (§§ 27000.5(a)(1), (2), (4)).

Major platforms already operate non-profiled feeds in practice. To comply with Article 38 of the European Union’s Digital Services Act—which requires very large online platforms to offer at least one recommender system option “which is not based on profiling,” Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services and Amending Directive 2000/31/EC (Digital Services Act), art. 38, 2022 O.J. (L 277) 1—TikTok and Meta have engineered and deployed non-profiled feed options in Europe. TikTok’s non-profiled feed displays popular videos from the user’s region and around the world, while Meta’s displays content from accounts the user follows. These deployments demonstrate that operating non-profiled recommendation and feed systems at platform scale is technically feasible. *See An Update on Fulfilling Our Commitments Under the Digital Services Act*, TikTok Newsroom (Aug. 4, 2023), <https://newsroom.tiktok.com/fulfilling-commitments-dsa-update>; Nick Clegg, *New Features and Additional Transparency Measures as the*

Digital Services Act Comes Into Effect, Meta Newsroom (Aug. 22, 2023), <https://about.fb.com/news/2023/08/new-features-and-additional-transparency-measures-as-the-digital-services-act-comes-into-effect/>.

Beyond traditional topic menus and search bars, major platforms have recently deployed a variety of commercial features that allow users to state explicitly what topics or types of content they want to see in their feeds:

- **Meta’s Topic and Prompt Controls.** On Instagram, Teen Accounts already allow minors to “select topics they want to see more of in Explore and their recommendations,” *Introducing Instagram Teen Accounts: Built-In Protections for Teens, Peace of Mind for Parents*, Meta Newsroom (Sept. 17, 2024), <https://about.fb.com/news/2024/09/instagram-teen-accounts/>, and Instagram’s *Your Algorithm* interface lets users inspect a list of topics and explicitly add or remove topics to shape their recommendations, see Aisha Malik, *Instagram’s New ‘Your Algorithm’ Tool Gives You More Control over the Reels You See*, TechCrunch (Dec. 10, 2025), <https://techcrunch.com/2025/12/10/instagrams-new-your->

[algorithm-tool-gives-you-more-control-over-the-reels-you-see.](#)

With new features Meta introduced on Facebook on October 13, 2026, users can express their specific preferences after selecting “Interested” or “Not Interested” on both Reels and posts, or directly instruct the recommender system in natural language what they want to see more or less of. *See Your Algorithm, Your Words: New Ways to Shape What You See on Facebook*, Meta Newsroom (Oct. 13, 2026),

[https://about.fb.com/news/2026/10/introducing-your-algorithm-facebook/.](https://about.fb.com/news/2026/10/introducing-your-algorithm-facebook/)

- **TikTok’s *Manage Topics*.** TikTok similarly provides *Manage Topics*, which enables users to specify broad content categories they would like to see more or less frequently in their feed. *More Ways to Discover New Content and Creators You Love*, TikTok Newsroom (June 3, 2025), [https://newsroom.tiktok.com/more-ways-to-discover-new-content-and-creators-you-love.](https://newsroom.tiktok.com/more-ways-to-discover-new-content-and-creators-you-love)
- **YouTube’s *Custom Feeds*.** YouTube is likewise developing *Custom Feeds*, which allow users to describe in plain language what they want to see and generate a dedicated feed around that

explicit request. *See Dive in and discover what's next on YouTube with our latest viewer features*, YouTube Official Blog (September 23, 2025) <https://blog.youtube/news-and-events/made-on-youtube-viewer-custom-feeds-shopping-tools/>.

Although these commercial controls currently operate alongside behavioral systems on Plaintiffs' platforms, their deployment demonstrates that platforms already possess the engineering capability to parse explicit user preferences—whether expressed through structured topic sliders or natural-language prompts—and translate them into feed recommendations at scale.

Recent computer science research confirms that recommendation pipelines can be driven directly by such explicit user instructions, interactive preference elicitation, or user-selected values rather than passive behavioral tracking. *See, e.g., Omar El Malki et al., Bonsai: Intentional and Personalized Social Media Feeds*, in Proc. 2026 CHI Conf. on Hum. Factors in Computing Sys. art. 1395 (2026), <https://doi.org/10.1145/3772318.3791855> (translating natural-language user instructions into a sourcing and ranking pipeline); Lindsay Popowski et al., *Social Media Feed Elicitation*, in Proc. 2026 CHI Conf.

on Hum. Factors in Computing Sys. (2026), <https://doi.org/10.1145/3772318.3791569> (eliciting feed preferences through interactive dialogue); Akaash Kolluri et al., *Alexandria: A Library of Pluralistic Values for Realtime Re-Ranking of Social Media Feeds*, 20 Proc. Int’l AAAI Conf. on Web & Soc. Media 1309 (2026) (ranking candidates according to user-articulated goals); *see also* Ying Li et al., *GenRec: An LLM-Backed Recommendation Ranker at Netflix* (Aug. 10, 2026) (unpublished manuscript), <https://arxiv.org/abs/2608.10257> (describing an industrial LLM-based ranker operating over text representations rather than traditional engineered behavioral features).

An example from YouTube’s brief illustrates how the explicit signal pipeline differs from an implicit signal pipeline in practice. YouTube notes that on its service, “a teen interested in robotics might find videos about mechanical engineering.” YouTube Br. 4. Both pipelines draw from the same corpus of videos, and both apply the same user-agnostic content-moderation filters and item-level quality scores. They differ only in whether the system relies on a persistent record of the teenager’s past behavior.

The Behavioral Pipeline. In a traditional behavioral architecture, as described by YouTube’s engineers, candidate generation draws on the user’s persistent watch history, search history, and demographic profile to retrieve candidate videos. *See* Paul Covington, Jay Adams & Emre Sargin, *Deep Neural Networks for YouTube Recommendations*, Proc. 10th ACM Conf. on Recommender Sys. 191, §§ 3.2–3.3 (2016). The ranking model then scores each candidate by predicting “expected watch time per impression.” *Id.* § 4. Under this design, the system infers the teenager’s interests from which videos she watched and for how long—and it infers other interests the same way, including topics she never affirmatively sought out. Each video she watches adds to that record and refines the next watch-time prediction. The system need not represent the teenager’s interests as an explicit, user-endorsed set of topics; instead, it can infer which candidate videos are likely to generate engagement from patterns in her behavior and patterns learned across users.

The Explicit-Signal Pipeline. When the teenager creates an account or customizes her feed, she explicitly selects “robotics” and “engineering” from a topic menu, follows creators she likes, and mutes

topics she wishes to avoid. Candidate generation retrieves recent videos tagged with her selected topics alongside uploads from creators she follows. The ranking stage then orders those candidates using viewer-independent item-level scores—such as production quality, freshness, age-appropriateness, and overall popularity within the “robotics” category—which allows the teenager to discover robotics creators she has not seen before. When she searches for a specific concept like “servo motor tutorial,” the search engine ranks results based on the query itself and item-level quality scores, without consulting her watch history or dwell time.

The Practical Difference for Minors. Under both pipelines, the teenager can view robotics and engineering videos, discover new creators, and be protected by the platform’s content-moderation standards. The difference lies in the optimization objective. Because the behavioral pipeline is typically trained to maximize engagement-related objectives across all possible topics, it can move the teenager toward unrelated content because she paused on a video. As YouTube’s engineers explained, behavioral models rely on implicit watch-time feedback to make predictions “deep in the tail where explicit feedback is

extremely sparse.” Covington et al., *supra*, § 3.1. While explicit signals are sufficient to identify the topics a teenager wants to explore, they do not allow the platform to predict, video by video, which item is most likely to prolong her session. Removing behavioral targeting may reduce the system’s ability to predict which individual item will maximize engagement and may reduce some forms of serendipitous discovery. But that reduction in behavioral optimization is part of the mechanism by which the Act limits behavioral profiling and the feedback loop it enables. Answering Br. at 31. By relying on explicit signals only, it is possible for the teenager to keep access to the robotics and engineering content she came for. The difference is that her feed is no longer designed to serve the platform’s optimization goals of keeping her online.

In sum, the engineering tools required to build engaging, navigable, and safe feeds from explicit user signals already exist. Requiring verifiable parental consent before subjecting minors to persistent behavioral profiling does not leave young users stranded in an unranked morass; it simply shifts default youth curation toward

explicit topic selection, search, and intentional feeds that seek to respect the user's deliberate choices.

CONCLUSION

For the foregoing reasons, the district court's order denying a preliminary injunction as to the Act's addictive feed provisions should be affirmed.

Dated: October 20, 2026

Respectfully submitted,

/s/Mihir Kshirsagar

Mihir Kshirsagar
Princeton CTP Tech Policy Clinic
315 Sherrerd Hall
Princeton, NJ 08540
609-258-5306
Attorney for Amicus Curiae

UNITED STATES COURT OF APPEALS
FOR THE NINTH CIRCUIT

Form 8. Certificate of Compliance for Briefs

Instructions for this form:

<http://www.ca9.uscourts.gov/forms/form08instructions.pdf>

9th Cir. Case Number(s) 26-5120, 26-5121, 26-5132

I am the attorney or self-represented party.

This brief contains 4380 words, excluding the items exempted by Fed. R. App. P. 32(f). The brief's type size and typeface comply with Fed. R. App. P. 32(a)(5) and (6).

I certify that this brief (*select only one*):

☐ complies with the word limit of Cir. R. 32-1.

☐ is a **cross-appeal** brief and complies with the word limit of Cir. R. 28.1-1.

☒ is an **amicus** brief and complies with the word limit of Fed. R. App. P. 29(a)(5), Cir. R. 29-2(c)(2), or Cir. R. 29-2(c)(3).

☐ is for a **death penalty** case and complies with the word limit of Cir. R. 32-4.

☐ complies with the longer length limit permitted by Cir. R. 32-2(b) because (*select only one*):

☐ it is a joint brief submitted by separately represented parties;

☐ a party or parties are filing a single brief in response to multiple briefs; or

☐ a party or parties are filing a single brief in response to a longer joint brief.

[] complies with the length limit designated by court order dated _____.

[] is accompanied by a motion to file a longer brief pursuant to Cir. R. 32-2(a).

Signature /s/Mihir Kshirsagar **Date** October 20, 2026

CERTIFICATE OF SERVICE

I certify that on October 20, 2026, this brief was e-filed through the ACMS System of the U.S. Court of Appeals for the Ninth Circuit. I certify that all participants in the case are registered ACMS users and that service will be accomplished by the ACMS system.

Dated: October 20, 2026

/s/Mihir Kshirsagar

Mihir Kshirsagar
Princeton CITP Tech Policy Clinic
315 Sherrerd Hall
Princeton, NJ 08540
609-258-5306
mihir@princeton.edu
Attorney for Amicus Curiae