

No. 26-1591

**IN THE UNITED STATES COURT OF APPEALS
FOR THE NINTH CIRCUIT**

X.AI LLC,

Plaintiff-Appellant,

v.

**ROB BONTA,
in his official capacity as Attorney General of California,**

Defendant-Appellee.

On Appeal from the United States District Court
for the Central District of California
No. 2:25-cv-12295-JGB-SSC
Hon. Jesus G. Bernal

**BRIEF OF AMICI CURIAE
LEGAL ADVOCATES FOR SAFE SCIENCE & TECHNOLOGY
AND COALITION OF 28 SCHOLARS AND CIVIL SOCIETY GROUPS
IN SUPPORT OF DEFENDANT-APPELLEE**

Jason Harrow
Gerstein Harrow, LLP
12100 Wilshire Blvd. Ste. 800
Los Angeles, CA 90025
jason@gerstein-harrow.com
(323) 744-5293
Counsel for Amici Curiae

Benjamin Rashkovich
Legal Advocates for Safe
Science & Technology
125 Park Ave., Floor 25
New York, NY 10017
ben@lasst.org
(914) 330-4672
Counsel for Amici Curiae

TABLE OF CONTENTS

TABLE OF CONTENTS	i
TABLE OF AUTHORITIES	ii
AMICI INTERESTS	1
INTRODUCTION	2
ARGUMENT	4
I. AB 2013 Regulates Commercial Speech	4
A. Commercial speech is not limited to “speech which does no more than propose a commercial transaction.”	4
B. The training data disclosures required by AB 2013 are commercial speech.....	7
C. The Court should reject xAI’s arguments that AB 2013’s mandatory disclosures are not commercial speech.	8
II. The Court Should Apply the <i>Zauderer</i> Test for Compelled Commercial Speech	12
III. AB 2013 Is Constitutional Under <i>Zauderer</i>	14
A. <i>Zauderer</i> applies because AB 2013’s disclosures are factual and uncontroversial.....	15
B. The government has substantial consumer protection interests in AB 2013’s required disclosures.	17
C. The government has other substantial interests in AB 2013’s required disclosures.....	23
D. AB 2013’s disclosures are neither unduly burdensome nor unjustified. .	28
CONCLUSION	30
APPENDIX: LIST OF AMICI	I
CORPORATE DISCLOSURE STATEMENT	VII

TABLE OF AUTHORITIES

Cases

<i>Am. Bev. Ass’n v. City & Cnty. of San Francisco</i> , 916 F.3d 749 (9th Cir. 2019)	5, 13, 14
<i>Am. Meat Inst. v. USDA</i> , 760 F.3d 18 (D.C. Cir. 2014) (en banc)	6, 14
<i>Bolger v. Youngs Drug Prods. Corp.</i> , 463 U.S. 60 (1983)	4, 17
<i>Cal. Chamber of Com. v. Council for Educ. & Rsch. on Toxics</i> , 29 F.4th 468 (9th Cir. 2022), <i>reh’g denied</i> , 51 F.4th 1182 (9th Cir. Oct. 26, 2022)	16
<i>Central Hudson Gas & Electric Corp. v. Pub. Serv. Comm’n of New York</i> , 447 U.S. 557 (1980)	12
<i>Chamber of Commerce v. SEC</i> , 85 F.4th 760 (5th Cir. 2023)	6
<i>CTIA—The Wireless Ass’n v. City of Berkeley</i> , 928 F.3d 832 (9th Cir. 2019)	passim
<i>Disc. Tobacco City & Lottery, Inc. v. United States</i> , 674 F.3d 509 (6th Cir. 2012)	12
<i>Env’t Def. Ctr. v. EPA</i> , 344 F.3d 832 (9th Cir. 2003)	16
<i>Kistler v. Eightfold AI</i> , No. 4:26-cv-01768 (N.D. Cal. Mar. 2, 2026)	21
<i>Md. Shall Issue, Inc. v. Anne Arundel Cnty.</i> , 91 F.4th 238 (4th Cir. 2024), <i>cert. denied</i> , 145 S. Ct. 152 (2024)	6, 24
<i>Milavetz, Gallop & Milavetz v. United States</i> , 559 U.S. 229 (2010)	13, 15, 28
<i>Mobley v. Workday</i> , 740 F. Supp. 3d 796 (N.D. Cal. 2024)	21

<i>Moody v. NetChoice</i> , 603 U.S. 707 (2024)	13, 14
<i>Nat’l Ass’n of Mfrs. v. SEC</i> , 800 F.3d 518 (D.C. Cir. 2015)	17
<i>Nat’l Ass’n of Wheat Growers v. Bonta (NAWG)</i> , 85 F.4th 1263 (9th Cir. 2023)	15, 16
<i>Nat’l Elec. Mfrs. Ass’n v. Sorrell</i> , 272 F.3d 104 (2d Cir. 2001)	6, 14
<i>Nat’l Inst. of Family & Life Advocates v. Becerra (NIFLA)</i> , 585 U.S. 755 (2018)	passim
<i>NetChoice v. AG, Fla.</i> , 34 F.4th 1196 (11th Cir. 2022)	14
<i>NetChoice v. Paxton</i> , 49 F.4th 439 (5th Cir. 2022)	14
<i>Packingham v. North Carolina</i> , 582 U.S. 98 (2017)	25
<i>Pharm. Rsch. & Mfrs. of Am. v. Stolfi</i> , 153 F.4th 795 (9th Cir. 2025)	passim
<i>RJ Reynolds Tobacco Co. v. FDA</i> , 96 F.4th 863 (5th Cir. 2024)	29
<i>Rubin v. Coors Brewing Co.</i> , 514 U.S. 476 (1995)	24
<i>Sorrell v. IMS Health Inc.</i> , 564 U.S. 552 (2011)	26
<i>Thomson Reuters Enter. v. Ross Intel.</i> , 765 F. Supp. 3d 382 (D. Del. 2025)	20
<i>TikTok Inc. v. Garland</i> , 604 U.S. 56 (2025)	26

<i>United States v. United Foods</i> , 533 U.S. 405 (2001)	5, 7, 8
<i>Virginia State Bd. of Pharm. v. Virginia Citizens Cons. Council</i> , 425 U.S. 748 (1976)	17
<i>Volokh v. James</i> , 148 F.4th 71 (2d Cir. 2025)	10
<i>Washington v. Glucksberg</i> , 521 U.S. 702 (1997)	24
<i>Wollschlaeger v. Governor</i> , 848 F.3d 1293 (11th Cir. 2017).....	26
<i>X Corp. v. Bonta</i> , 116 F.4th 888 (9th Cir. 2024)	5, 9, 10, 12
<i>Zauderer v. Office of Disciplinary Counsel</i> , 471 U.S. 626 (1985)	passim

Statutes

15 U.S.C. §§ 1333, 4402(2)(A)	12
21 U.S.C. § 343(q)(1)(D)	11
21 U.S.C. § 352(f)	11
AB 2013	passim
Cal. Civ. Code § 3111(a)	7

Legislative History

Cal. Assem. Comm. on Priv. & Consumer Prot., Hr’g on AB 2013 (Apr. 30, 2024)	23, 26, 27
Cal. Assem. Floor, Concurrence in Senate Amendments, Analysis of Assem. Bill 2013, 2024 Reg. Sess. (2024).....	19
Cal. Sen. Jud. Comm., Analysis of AB 2013, 2023-2024 Reg. Sess. (June 17, 2024)	26

Cal. Sen. Rules Comm., Analysis of AB 1013, 2024 Reg. Sess. (Aug. 20, 2024)	23
--	----

Other Authorities

<i>AI Chatbots for Mental Health – What Works, What Harms, and What’s Next</i> , Nat’l Acad. of Med. (May 13, 2026).....	25
Alexander Erlei et al., <i>When Life Gives You AI, Will You Turn It Into A Market for Lemons? Understanding How Information Asymmetries About AI System Capabilities Affect Market Outcomes and Adoption</i> , Ass’n for Computing Machinery (Apr. 2026).....	18
Alexander Wan et al., <i>The 2025 Foundation Model Transparency Index</i> , Stanford Ctr. for Rsch. on Foundation Models (Dec. 2025)	18
Anatol-Fiete Näher et al., <i>Measuring fairness preferences is important for artificial intelligence in health care</i> , 6 The Lancet E302 (May 2024).....	24
Ben Cottier et al., <i>How much does it cost to train frontier AI models?</i> , Epoch AI (June 3, 2024).....	18
Brad Smith, <i>Microsoft announces new Copilot Copyright Commitment for customers</i> , Microsoft (Jan. 5, 2024).....	21
Clare Y. Cho et al., <i>Competition and Antitrust Concerns Related to Generative AI</i> , Cong. Rsch. Serv. (Apr. 16, 2025).....	18
David Thiel, <i>Identifying and Eliminating CSAM in Generative ML Training Data and Models</i> , Stanford Internet Observatory (Dec. 20, 2023)	25
Delaram Rezaeikhonakdar, <i>AI Chatbots and Challenges of HIPAA Compliance for AI Developers and Vendors</i> , 2023 J. L. Med. Ethics 51(4) (2023)	20
Deployment Safety Hub, <i>GPT-5.5 System Card</i> , OpenAI (Apr. 23, 2026).....	29
Dimitra Panteli et al., <i>Artificial Intelligence in Public Health: Promises, Challenges, and An Agenda for Policy Makers and Public Health Institutions</i> , 10(5) The Lancet E428 (Feb. 28, 2025)	24
Emilio J. Castilla, <i>AI is reinventing hiring—with the same old biases. Here’s how to avoid that trap</i> , MIT Mgmt. Sch. (Dec. 15, 2025)	22

Ethan A. Abbott et al., <i>Understanding and Addressing Bias in Artificial Intelligence Systems</i> , J. Am. Coll. Emergency Phys. Open, Feb. 2026.....	22
<i>Ethics & Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models</i> , World Health Org. (2024)	25
<i>Gemini 3.5 Flash Model Card</i> , Google DeepMind (May 2026).....	29
George Mitchell, <i>AI Training Data Comparison: What Models Know</i> , AionX (Feb. 13, 2026)	20
Greg McDonough, <i>Privacy Risks in Training Data and How to Address Them</i> , RSAC Conf. (Oct. 28, 2025)	27
<i>Grok 4.1 Model Card</i> , xAI (Nov. 17, 2025)	29
Internet Watch Found., <i>AI Becoming ‘Child Sexual Abuse Machine,’ Adding to Dangerous Record Levels of Online Abuse</i> , IWF Warns (Jan. 16, 2026).....	25
James L. Cross et al., <i>Bias in medical AI: Implications for clinical decision-making</i> , PLOS Digit Health (Nov. 7, 2024)	22
Julie Rogers & Alexandra Jonker, <i>What Is Data Cleaning?</i> , IBM (2026).....	11
Katharine Miller, <i>Privacy in an AI Era: How Do We Protect Our Personal Information?</i> , Stanford Inst. for Human-Centered AI (Mar. 18, 2024).....	27
Michael M. Grynbaum & Ryan Mac, <i>The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work</i> , N.Y. Times (Dec. 27, 2023)	27
Neal Suggs, <i>Shared fate: Protecting customers with generative AI indemnification</i> , Google Cloud (Oct. 12, 2023).....	21
<i>New models and developer products announced at DevDay</i> , OpenAI (Nov. 6, 2023).....	21
Off. of Tech., <i>AI Companies: Uphold Your Privacy and Confidentiality Commitments</i> , Fed. Trade Comm’n (Jan. 9, 2024).....	27
Off. of the Attorney General, <i>California Attorney General’s Legal Advisory on the Application of Existing California Laws to Artificial Intelligence</i> , Cal. Dep’t of Just. (Jan. 8, 2025).....	21

Robert Post, <i>The Constitutional Status of Commercial Speech</i> , 48 UCLA L. Rev. 1 (2000)	7
Samantha Subin, <i>Musk’s xAI, SpaceX combo is the biggest merger of all time, valued at \$1.25 trillion</i> , CNBC (Feb. 3, 2026)	29
Sha Sajadieh et al., <i>The AI Index 2026 Report</i> 171-99, Stanford Inst. for Human-Centered AI (Apr. 2026)	18
Shawn McGrath, <i>How to mitigate AI discrimination and bias in financial services</i> , Ernst & Young (Mar. 17, 2023)	22
<i>System Card: Claude Opus 4.7</i> , Anthropic (Apr. 16, 2026)	29
The Capital Commitment, <i>Five Considerations for Advisers Implementing AI in Investment Decisions</i> , Proskauer (Mar. 31, 2026)	20
Warren Barkley, <i>The Prompt: What AI model is it anyway?</i> , Google Cloud (Sep. 30, 2024)	19
<i>What Is AI Bias? Causes, Types, & Real-World Impacts</i> , Palo Alto Networks	22
<i>When Algorithms Deny Care: Bias in Healthcare AI</i> , RAIL (Nov. 2, 2025)	22
<i>xAI Frontier Artificial Intelligence Framework</i> , xAI (Dec. 30, 2025)	29

AMICI INTERESTS¹

Legal Advocates for Safe Science & Technology, Inc. (“LASST”) is a nonprofit organization dedicated to making advances in science and technology safer for people and the planet. LASST advocates for legal frameworks governing AI systems that appropriately balance innovation with security and public safety. A list of joining amici is provided in an Appendix.

This case involves a novel challenge to AI transparency mandates. The legal issues in question, however, are clear and well-established. Because this Court’s resolution of xAI’s preliminary injunction motion may influence how courts, regulators, and industry participants approach AI developer transparency going forward, amici respectfully submit this brief to provide context on the public interest considerations that support requiring AI developers to disclose basic facts about their publicly accessible models, as AB 2013 does.

¹ No counsel for any party authored this brief in whole or in part, and no person other than amici curiae, their members, and their counsel made a monetary contribution to the preparation or submission of this brief. Fed. R. App. P. 29(a). All parties consented to the filing of this brief.

INTRODUCTION

California’s Assembly Bill 2013 requires developers of generative AI systems to post a “high-level summary” of training datasets. It does not require developers to produce the datasets. It does not require them to identify proprietary algorithms or reveal source code. It simply asks: what kinds of data did you train on and where did it come from? These are questions a consumer might reasonably ask about any product.

xAI frames this modest transparency requirement as a constitutional crisis. Amici disagree. AB 2013 is a common sense, run-of-the-mill disclosure law—the functional equivalent of an ingredient label—applied to a new and increasingly important category of consumer product. AB 2013 is appropriately tailored to serve the interests of consumers and the public at large. Disclosure requirements have long been a critical feature of commercial regulation, and courts have consistently upheld such laws under the deferential standard of *Zauderer v. Office of Disciplinary Counsel*, 471 U.S. 626 (1985), which governs compelled commercial speech that is purely factual and uncontroversial.

Disclosure requirements help people make informed choices. Financial services, pharmaceuticals, automobiles, food and beverages: disclosure requirements in these and many more industries give consumers the information they need to decide what drugs to take and cars to drive, for example.

Generative AI is no different. Users care whether the AI models they use for relationship advice get trained on personal information. Creatives care whether developers compensate artists for providing training data. Parents care whether the AI chatbots on their kids' phones are trained on datasets known to contain child pornography. Financial advisers care whether their AI-enhanced analytical tools are too opaque to satisfy their fiduciary duties to clients. Procurement decisionmakers in regulated industries care whether an enterprise AI product might expose them to liability. These legitimate concerns can be addressed with *information*. Like the many other disclosure requirements enacted by state and federal legislatures, AB 2013 provides just that—without violating any First Amendment rights.

The Ninth Circuit's reasoning in this case, not merely its result, will shape the constitutional landscape for AI governance. With a technology advancing as rapidly as AI, transparency laws are a reasonable and legitimate way to give the public information required to adapt and make informed decisions. Disclosure requirements are the lightest touch mode of regulation available to the government to address the significant information asymmetries presented by the rise of AI. The Ninth Circuit's jurisprudence permits such laws, including AB 2013.

ARGUMENT

I. AB 2013 Regulates Commercial Speech

AB 2013’s training data disclosures constitute commercial speech. They concern the sources, makeup, and treatment of datasets used by developers for AI training and further the interest of California consumers in the free flow of commercial information regarding the artificial intelligence systems available for their use. AB 2013’s disclosures are akin to the many disclaimers, warnings, labels, ingredient lists, and other types of product information held to be commercial speech: descriptive, factual recitations of a product’s inputs with relevance to those products’ consumers. xAI would have the Court ignore both common sense and clear precedent in holding otherwise.

A. Commercial speech is not limited to “speech which does no more than propose a commercial transaction.”

The “core notion of commercial speech” is “speech which does no more than propose a commercial transaction.” *Bolger v. Youngs Drug Prods. Corp.*, 463 U.S. 60, 66 (1983). Courts initially developed this standard, however, to “apply to laws that regulate *voluntary speech* from a regulated commercial entity to a public audience.” *Pharm. Rsch. & Mfrs. of Am. v. Stolft*, 153 F.4th 795, 816 (9th Cir. 2025) (emphasis added). Applying that same “core notion” of commercial speech to laws that regulate *compelled disclosures* by a commercial entity would challenge courts to fit square pegs into round holes.

The Ninth Circuit has already recognized—and cured—this tension: “Our Circuit has characterized speech as commercial ‘even if not a clear fit’ with these legal tests where the speech nonetheless ‘communicates the terms of an actual or potential [commercial] transaction.’” *Stolfi*, 153 F.4th at 820 (quoting *X Corp. v. Bonta*, 116 F.4th 888, 901 (9th Cir. 2024)); *see also United States v. United Foods*, 533 U.S. 405, 409 (2001) (noting commercial speech is only “usually defined as speech that does no more than propose a commercial transaction” (emphasis added)). In other words, the commercial speech inquiry does not end just because the *Bolger* test does not apply. The “fact-driven” commercial speech analysis must instead “give effect to common-sense distinctions between commercial speech and not.” *X Corp.*, 116 F.4th at 900. Commercial speech thus covers transparency requirements and voluntary transaction proposals alike. *See Stolfi*, 153 F.4th at 820-21; *id.* at 815-16 (explaining that “the commercial speech doctrine originated in this advertising context” and is “therefore inapt” where “regulated entities and individuals” produce reports “out of legal obligation, not economic motivation”).²

Commercial speech has long been held in this circuit to encompass

² At the time of filing, a stay application in *Stolfi* is pending before the Supreme Court, No. 25A806. Even if the Court were to grant certiorari, the analysis that follows rests independently on this Circuit’s prior holdings in *CTIA—The Wireless Ass’n v. City of Berkeley*, 928 F.3d 832 (9th Cir. 2019), and *American Beverage Ass’n v. City & County of San Francisco*, 916 F.3d 749 (9th Cir. 2019), both of which predate *Stolfi* and establish the same commercial-speech and *Zauderer* framework on which the district court relied.

compelled disclosures of product information. *See id.* at 820-21 (listing examples). This is especially so where mandated disclosures “communicate product-specific economic information” that “furthers the consumer’s interest in the free flow of commercial information” and “has no independent expressive meaning.” *Id.* at 822. Other circuit courts agree. *See Md. Shall Issue, Inc. v. Anne Arundel Cnty.*, 91 F.4th 238 (4th Cir. 2024) (“[T]he plaintiffs acknowledge that product safety warnings are of a type long considered permissible under the *Zauderer* jurisprudence.” (cleaned up)), *cert. denied*, 145 S. Ct. 152 (2024); *Chamber of Commerce v. SEC*, 85 F.4th 760, 770-72 (5th Cir. 2023) (SEC share buyback disclosures); *Am. Meat Inst. v. USDA*, 760 F.3d 18, 23-24 (D.C. Cir. 2014) (en banc) (country-of-origin label mandates); *Nat’l Elec. Mfrs. Ass’n v. Sorrell*, 272 F.3d 104, 107-08, 113-14 (2d Cir. 2001) (labels disclosing presence of mercury and safe methods of disposal). Courts have repeatedly reaffirmed the fundamental principle that disclosure requirements advance the First Amendment interest of consumers in need of accurate and useful information because speaker autonomy is not at stake in the commercial context.³

³ “Disclosure requirements are permissible within the domain of commercial speech, however, because the autonomy of speakers is not at stake, only the conveyance of information. . . . Within commercial speech . . . the primary constitutional value concerns the circulation of accurate and useful information. For the state to mandate disclosures designed more fully and completely to convey information is thus to advance, rather than to contradict, pertinent constitutional values.” Robert Post, *The Constitutional Status of Commercial Speech*, 48 UCLA L. Rev. 1, 27-28 (2000); *see also United Foods*, 533 U.S. at 425 (Breyer, J., dissenting).

B. The training data disclosures required by AB 2013 are commercial speech.

Treating the transparency disclosures mandated by AB 2013 as commercial speech “makes sense.” *See Stolfi*, 153 F.4th at 822. AB 2013 requires AI developers to provide “high-level summar[ies] of the datasets” used to train models “made publicly available to Californians for use.” Cal. Civ. Code § 3111(a). These disclosures are meant to “communicate product-specific . . . information” that “further[] the consumer’s interest in the free flow of commercial information” and have “no independent expressive meaning.” *Stolfi*, 153 F.4th at 822 (citing the contrast in *Zauderer*, 471 U.S. at 651, between commercial speech regulation and speech regulation “prescrib[ing] what shall be orthodox in politics, nationalism, religion, or other matters of opinion”).

The Court need not look far for analogous commercial speech examples: ingredients and nutritional facts on food packaging, ABV labels on alcoholic drinks, country-of-origin labels for meats and produce, active/inactive ingredients on pharmaceuticals, and many more. High-level summaries of training data serve the same function for consumers. The makeup of an AI model’s training data implicates issues that many consumers care deeply about, from privacy and intellectual property to child safety and healthcare. Such information can and does already affect how consumers assess their options in the marketplace of AI systems. The factual and uncontroversial disclosures mandated by AB 2013 serve

these consumer needs in furtherance of the First Amendment.

C. The Court should reject xAI’s arguments that AB 2013’s mandatory disclosures are not commercial speech.

xAI throws several arguments at the wall in its attempt to convince the Court that AB 2013’s disclosures are not commercial speech. None stick.

First, xAI argues that AB 2013’s disclosures are not commercial speech because they do not “propose a commercial transaction.” Pl.’s Br. at 22-24. Courts have already rejected xAI’s arguments, but it is worth highlighting xAI’s cherry-picking: xAI insists that “the Supreme Court has ‘defined’ the category of ‘commercial speech’ narrowly to reach only ‘speech that *does no more* than propose a commercial transaction,’” citing *United Foods*, 533 U.S. at 409 (emphasis added by xAI). Pl.’s Br. at 22. Had xAI quoted a single word more, the flexibility of this purportedly narrow definition would have undermined xAI’s very point: “A quarter of a century ago, the Court held that commercial speech, *usually* defined as speech that does no more than propose a commercial transaction, is protected by the First Amendment.” *United Foods*, 533 U.S. at 409 (emphasis added). Contrary to xAI’s argument, many courts—including this one—have consistently adopted a broader understanding of commercial speech than the advertising-centric framework laid out in *Bolger*.

Second, xAI argues that the district court erred in relying on *Stolfi*. In particular, xAI focuses on a distinction between that case’s “government reporting

requirement” and this one’s public “disclosure requirement.” Pl.’s Br. at 29-31. xAI omits the relevant context. In *Stolfi*, the act at issue “requires manufacturers to report specific information to a government agency, which then makes that information available to the public unless it is a confidential trade secret.” 153 F.4th at 809. Indeed, the Ninth Circuit described such publications as “notably commonplace.” *Id.* at 810. At any rate, this distinction is not relevant to any of the principles for which the district court cited *Stolfi*: namely, that the *Bolger* factors and the early contours of the commercial speech doctrine arose in the context of advertising as opposed to compelled disclosures and, thus, are less applicable to the latter. *Id.* at 816.

Third, xAI analogizes AB 2013’s run-of-the-mill disclosure requirement to the one rejected in *X Corp. v. Bonta*, 116 F.4th 888 (9th Cir. 2024). But xAI’s comparison here is flawed: as the district court noted, the mandated disclosures in *X Corp.* are easily distinguishable. *X Corp.* concerned AB 587, a law requiring social media companies to disclose their content moderation practices regarding specific types of content disfavored by the government. *Id.* at 894, 901-02. To comply, social media companies would have had to “recast [their] content-moderation practices into language prescribed by the State, implicitly opining on whether and how certain controversial categories of content should be moderated” and thereby “convey the company’s policy views on intensely debated and

politically fraught topics.” *Id.* at 901-02. Notably, the Ninth Circuit held that “a social media platform’s existing TOS [terms of service] and content moderation policies may be commercial speech,” as the content of those policies would simply be factual recitations, but noted that requiring a platform to say whether and how it moderates topics disfavored by the state would reveal the platform’s “opinions about and reasons for those policies,” which are “different in character and kind” from the policies themselves. *Id.* at 901-02; *see id.* at 902-03 (noting that the panel decision did not conflict with 5th and 11th Circuit decisions applying *Zauderer* to general content moderation disclosure policies for this reason); *Volokh v. James*, 148 F.4th 71, 86 (2d Cir. 2025) (holding that “requiring an advertiser to convey purely factual and uncontroversial information about the terms under which it will provide its services” does not implicate controversy).

The substance of AB 2013’s disclosures, while important, is quotidian: objective information and high-level summaries about training data. Nor does AB 2013 compel xAI to couch its disclosures in language prescribed by the government. xAI focuses on two specific requirements of AB 2013—that developers disclose how training data sets “further the intended purpose” of the AI system and the “intended purpose” of any data cleaning efforts—but there is no controversy around why developers clean data or what their systems are intended

to do.⁴ AB 2013 mandates the disclosure of straightforward product information, not political expression. In other words, AB 2013 mandates the disclosure of commercial speech.

Fourth, xAI protests that its training data disclosures are not commercial speech because “xAI has no economic motivation to make the A.B.2013 disclosures” and, in fact, would prefer not to do so. Pl.’s Br. at 25. But this argument proves too much. Many compelled commercial disclosures are made by speakers who would prefer silence. Therein lies their value. “Indeed, it is entirely common for courts to conclude that regulatory disclosures related to a specific product are commercial speech, even where the specific information disclosed is directly *contrary* to the speaker’s economic interest and therefore would not be disclosed absent regulation.” *Stolfi*, 153 F.4th at 820 n.15 (citing cases). A pharmaceutical manufacturer has no affirmative economic motivation to disclose adverse drug interactions; a tobacco retailer has no economic motivation to disclose potential health risks; a food producer has no economic motivation to disclose that its product contains saturated fat.⁵ The point of compelled disclosure

⁴ See, e.g., Julie Rogers & Alexandra Jonker, *What Is Data Cleaning?*, IBM (2026), <https://www.ibm.com/think/topics/data-cleaning>.

⁵ See, e.g., 21 U.S.C. § 352(f) (Federal Food, Drug, and Cosmetic Act disclosure mandate of “unsafe dosage or methods or duration of administration or application”); 21 U.S.C. § 343(q)(1)(D) (Nutrition Labeling and Education Act disclosure mandate of saturated fat and other nutrients); 15 U.S.C. §§ 1333, 4402(2)(A) (Family Smoking Prevention and Tobacco Control Act disclosure

laws is that the speaker’s preference for silence does not end the inquiry.

xAI asks the Court to ignore a “common-sense distinction between commercial speech and other varieties of speech,” *X Corp.*, 116 F.4th at 900, and view AB 2013’s disclosures as fundamentally different from, say, ingredient labels. At oral argument, though, xAI had no trouble analogizing training data to ingredients in the context of its trade secrets claims. *See* Hr’g Tr. 20:15-22:19, Feb. 23, 2026. xAI cannot have it both ways.

II. The Court Should Apply the *Zauderer* Test for Compelled Commercial Speech

The Court must next determine which standard of First Amendment scrutiny to apply to the disclosures required by AB 2013. Where, as here, the First Amendment concern at issue deals with compelled rather than restricted commercial speech, courts in this circuit apply the more permissive standard of scrutiny articulated in *Zauderer v. Office of Disciplinary Counsel*, 471 U.S. 626 (1985). Courts only apply intermediate scrutiny when assessing restrictions on commercial speech under *Central Hudson Gas & Electric Corp. v. Public Service Commission of New York*, 447 U.S. 557 (1980), or if the compelled disclosures do not fit *Zauderer*. *See Milavetz, Gallop & Milavetz v. United States*, 559 U.S. 229,

mandate of health warnings on cigarette packaging); *see also Disc. Tobacco City & Lottery, Inc. v. United States*, 674 F.3d 509, 525-27, 530-31 (6th Cir. 2012) (applying *Zauderer* to the Family Smoking Prevention and Tobacco Control Act).

249-51 (2010); *see also* *Am. Bev. Ass’n v. City & Cnty. of San Francisco*, 916 F.3d 749, 762 (9th Cir. 2019) (“Because *Zauderer* does not apply here, *NIFLA* directs us to consider whether the ordinance survives heightened scrutiny [under *Central Hudson*].”). The Court should apply the more permissive standard set out in *Zauderer*, as it has done in other cases challenging similar disclosure requirements.

To state the obvious, AB 2013 is a *disclosure* requirement. It compels speech. As xAI cannot contest this, it instead argues that *Zauderer* is inapposite because the Supreme Court has only applied *Zauderer* in the context of misleading advertising. As an initial matter, xAI is incorrect: the Supreme Court did not question the lower courts’ determination that *Zauderer* was the appropriate standard for reviewing the disclosure laws at issue in *Moody v. NetChoice*, 603 U.S. 707, 725 (2024) (“The next order of business is to decide which of the laws’ applications violate the First Amendment . . . [that] means asking, again as to each thing covered, whether the required disclosures unduly burden expression,” citing *Zauderer*).

Further, this Court should emphasize what it has previously and repeatedly articulated: *Zauderer* reaches beyond the narrow domain of misleading advertising and applies to First Amendment analysis of compelled commercial speech generally. In *CTIA*, the Ninth Circuit “squarely address[ed] the question whether, in the absence of a prevention-of-deception rationale, the *Zauderer* compelled-

disclosure test applies.” 928 F.3d at 843. This Circuit joined several others in concluding that it does. *Id.* (“Our sister circuits have thus held under *Zauderer* that the prevention of consumer deception is not the only governmental interest that may permissibly be furthered by compelled commercial speech. The Supreme Court also signaled its agreement with this reading of *Zauderer* [in *NIFLA*].”); *Am. Bev. Ass’n*, 916 F.3d at 755-56 (“We also rejected the argument that *Zauderer* applies only to situations in which the government requires disclosures to prevent consumer deception, pointing out that we were joining the holdings of several of our sister circuits.”); *see also Am. Meat Inst.*, 760 F.3d at 22; *Nat’l Elec. Mfrs. Ass’n*, 272 F.3d at 115; *NetChoice v. AG, Fla.*, 34 F.4th 1196, 1230 (11th Cir. 2022); *NetChoice v. Paxton*, 49 F.4th 439, 485 (5th Cir. 2022).⁶ Despite xAI’s attempt to minimize the applicability of *Zauderer*, precedent on this point is unambiguous. This Court should apply *Zauderer* in assessing the constitutionality of AB 2013’s mandated disclosures.

III. AB 2013 Is Constitutional Under *Zauderer*

For *Zauderer* to apply, a disclosure requirement must compel commercial speech that contains “purely factual and uncontroversial information.” 471 U.S. at 651. Courts then apply rational basis review to compelled disclosure laws that are “purely factual and uncontroversial” under *Zauderer*. The disclosure requirement

⁶ Both *NetChoice* cases were remanded on other grounds by *Moody*, 603 U.S. 707.

must thus be “reasonably related to a substantial governmental interest.” *Nat’l Ass’n of Wheat Growers v. Bonta (NAWG)*, 85 F.4th 1263, 1275 (9th Cir. 2023); *see also Nat’l Inst. of Family & Life Advocates v. Becerra (NIFLA)*, 585 U.S. 755, 768 (2018). The disclosure law at issue must also not be so “[u]njustified or unduly burdensome” as to “offend the First Amendment by chilling protected speech.” *Milavetz*, 559 U.S. at 249-50; *see also NIFLA*, 585 U.S. at 776 (noting that the burden prong focuses on the threat of chilled speech).

A. *Zauderer* applies because AB 2013’s disclosures are factual and uncontroversial.

AB 2013’s mandated disclosures are factual on their face. The “factual” analysis under *Zauderer* is straightforward: disclosure requirements pass muster if they compel information that is “literally true,” with a narrow exception for requirements compelling speech that “may be literally true but nonetheless misleading and, in that sense, untrue.” *CTIA*, 928 F.3d at 847-48.

AB 2013 requires generative AI system developers to share high-level summaries of inarguably factual information: numbers and descriptions of training data sources, numbers and descriptions of data points within those datasets, whether datasets contain personal information or sources outside the public domain, whether any of the training data was generated for the purpose of model development (known as “synthetic data”) rather than collected from the real world, and so on. Unless xAI chooses to report such information inaccurately, the contents

of its disclosures must be “literally true.”

The “literally true but nonetheless misleading” exception does not apply. Unlike cases where the compelled speech at issue failed the “factual” threshold of *Zauderer* for being misleading, it is xAI writing its own disclosures, rather than the government supplying contested language. If xAI believes that any facts or figures it discloses under AB 2013 may somehow mislead a consumer, it is entitled to supplement those disclosures with clarifying information or narrative.

Nor are AB 2013’s mandated disclosures controversial. AB 2013 “involves no compelled recitation of a message and no affirmation or belief,” nor does it prohibit an AI system developer from supplementing the mandatory disclosures with additional information or its own views. *See Env’t Def. Ctr. v. EPA*, 344 F.3d 832, 850 (9th Cir. 2003) (upholding safe toxin disposal disclosures under *Zauderer*). AB 2013 does not require xAI to “convey a message fundamentally at odds with its mission” or wade into a “robust disagreement by reputable scientific sources.” *See NAWG*, 85 F.4th at 1277-78 (citing *CTIA*, 928 F.3d at 845, and *Cal. Chamber of Com. v. Council for Educ. & Rsch. on Toxics*, 29 F.4th 468, 478 (9th Cir. 2022), *reh’g denied*, 51 F.4th 1182 (9th Cir. Oct. 26, 2022)). Nor does AB 2013 require xAI to “take sides in a heated political controversy.” *CTIA*, 928 F.3d at 848. And to the extent xAI argues that AB 2013’s subject matter is inherently controversial, both the Ninth Circuit and the Supreme Court have already rejected

such an expansive understanding of controversy under *Zauderer*. See *Stolfi*, 153 F.4th at 824; *Bolger*, 463 U.S. at 68 (“We have made clear that advertising which ‘links a product to a current public debate’ is not thereby entitled to the constitutional protection afforded noncommercial speech.”).

xAI dropped its arguments that AB 2013 mandates non-factual and controversial disclosures for an obvious reason: AB 2013 is facially constitutional under *Zauderer*. As with product labeling, safety warnings, and country-of-origin identifications, AB 2013 promotes the dissemination of purely factual information that is significant to consumers.

B. The government has substantial consumer protection interests in AB 2013’s required disclosures.

California has a “more than trivial” governmental interest in fostering transparency in the marketplace of AI systems available to Californian consumers. See *CTIA*, 928 F.3d at 844; see also *NIFLA*, 585 U.S. at 776. The government certainly has a non-trivial interest in promoting market efficiency, consumer protection, investor awareness, and product transparency in the fast-moving, high-disruption, and deeply asymmetric new industry of generative artificial intelligence. See *Virginia State Bd. of Pharm. v. Virginia Citizens Cons. Council*, 425 U.S. 748, 763-64 (1976); *Nat’l Ass’n of Mfrs. v. SEC*, 800 F.3d 518, 534 (D.C. Cir. 2015).

California’s interest in AI training data transparency is especially acute given

the speed at which generative AI has spread across all domains of public and private life,⁷ the breakneck pace of its technological development,⁸ and the significant consolidation of market power in an industry characterized by nearly unfathomable barriers to entry.⁹ Consumers in the generative AI marketplace could hardly be facing more “severe buyer-supplier market asymmetries.”¹⁰ It is precisely these circumstances where the state’s substantial interest in “reducing those

⁷ Sha Sajadieh et al., *The AI Index 2026 Report* 171-99, Stanford Inst. for Human-Centered AI (Apr. 2026),

https://hai.stanford.edu/assets/files/ai_index_report_2026.pdf (discussing breadth and depth of generative AI adoption and noting that “Generative AI adoption reached 53% in three years, faster than the personal computer or the internet”).

⁸ *Id.* at 68-80 (discussing pace of AI development).

⁹ See Clare Y. Cho et al., *Competition and Antitrust Concerns Related to Generative AI*, Cong. Rsch. Serv. (Apr. 16, 2025), <https://www.congress.gov/crs-product/IF12968> (describing concentrated control over key inputs and the risk that large incumbents entrench market power in AI); Ben Cottier et al., *How much does it cost to train frontier AI models?*, Epoch AI (June 3, 2024), <https://epoch.ai/publications/how-much-does-it-cost-to-train-frontier-ai-models> (“If the trend of growing training costs continues, the largest training runs will cost more than a billion dollars by 2027, suggesting that frontier AI model training will be too expensive for all but the most well-funded organizations.”).

¹⁰ See Alexander Erlei et al., *When Life Gives You AI, Will You Turn It Into A Market for Lemons? Understanding How Information Asymmetries About AI System Capabilities Affect Market Outcomes and Adoption*, Ass’n for Computing Machinery (Apr. 2026), <https://dl.acm.org/doi/10.1145/3772318.3791420> (“Complex AI systems can appear highly accurate while making costly errors or embedding hidden defects. While there have been regulatory efforts surrounding different forms of disclosure, large information gaps remain.”); see also Alexander Wan et al., *The 2025 Foundation Model Transparency Index*, Stanford Ctr. for Rsch. on Foundation Models (Dec. 2025), <https://crfm.stanford.edu/fmti/December-2025/paper.pdf> (“The upstream domain, which covers topics like training data and training compute, is the least transparent yet most heterogeneous.”).

asymmetries, facilitating informed commercial transactions, and improving the efficiency” of a highly information-asymmetric market reaches its apex. *See Stolfi*, 153 F.4th at 826.

The disclosures required by AB 2013 are “reasonably related” to these substantial government interests. AB 2013 “furthers consumer protection in California by granting the state’s residents insight into how the GenAI systems and services they engage with are trained.”¹¹ According to the bill’s author, “[c]onsumers may use this knowledge to better evaluate if they have confidence in the AI system or service, compare competing systems and services, or put into place mitigation measures to address any shortcomings of the particular system or service.”¹²

These are far from hypothetical concerns or idle curiosities for consumers. Understanding training dataset makeup and collection methodology is increasingly critical to how individual and enterprise consumers of generative AI systems decide which products to use.¹³ For instance, some consumers need a “multimodal”

¹¹ Cal. Assem. Floor, Concurrence in Senate Amendments, Analysis of Assem. Bill 2013, 2024 Reg. Sess., at 2 (2024).

¹² *Id.*

¹³ *See, e.g.,* Warren Barkley, *The Prompt: What AI model is it anyway?*, Google Cloud (Sep. 30, 2024), <https://cloud.google.com/transform/choosing-right-gen-ai-model-trade-offs-benefits-the-prompt> (listing “[t]he availability of labeled data, domain specific data, data quality, and more” as important considerations for “customizing” models).

AI system that has been trained on different types of information, from text and images to audio and video, while other consumers require systems trained on highly particular datasets.¹⁴ Financial advisers and other professionals with fiduciary duties need insight into an AI system’s training data to comply with their client obligations, like explainability.¹⁵ Enterprise consumers in regulated industries require access to information about AI systems’ training data in order to protect themselves against potential liability for infringing copyright,¹⁶ breaching HIPAA due diligence obligations,¹⁷ and violating privacy laws,¹⁸ for example. In

¹⁴ George Mitchell, *AI Training Data Comparison: What Models Know*, AionX (Feb. 13, 2026), <https://aionx.co/productivity/ai-training-data-comparison/> (“AI models work with different data types depending on their function. Understanding these categories helps you choose appropriate datasets for your use case.”).

¹⁵ The Capital Commitment, *Five Considerations for Advisers Implementing AI in Investment Decisions*, Proskauer (Mar. 31, 2026), <https://www.proskauer.com/blog/risk-3-five-considerations-for-advisers-implementing-ai-in-investment-decisions> (“Training data is another critical element. If the adviser does not know whether material non-public information or other inappropriate data has been ingested (whether inadvertently or intentionally), this could raise concerns about the adviser’s compliance with Section 204A of the Advisers Act . . .”).

¹⁶ See, e.g., *Thomson Reuters Enter. v. Ross Intel.*, 765 F. Supp. 3d 382, 397-401 (D. Del. 2025).

¹⁷ See, e.g., Delaram Rezaeikhonakdar, *AI Chatbots and Challenges of HIPAA Compliance for AI Developers and Vendors*, 2023 J. L. Med. Ethics 51(4) (2023), <https://pmc.ncbi.nlm.nih.gov/articles/PMC10937180/>.

¹⁸ Off. of the Attorney General, *California Attorney General’s Legal Advisory on the Application of Existing California Laws to Artificial Intelligence* 5, Cal. Dep’t of Just. (Jan. 8, 2025), <https://oag.ca.gov/system/files/attachments/press-docs/Legal%20Advisory%20-%20Application%20of%20Existing%20CA%20Laws%20to%20Artificial%20Intelligence.pdf>.

these circumstances, enterprise consumers of AI systems may be liable to their own downstream customers¹⁹—and, recognizing this tension, some AI system developers offer intellectual property indemnification clauses.²⁰ Transparency into training data is the floor, not the ceiling, for the sorts of information customers need to navigate the dynamic and asymmetric market for artificial intelligence.

xAI decries that AB 2013 is partially meant to address “biased” models, *see, e.g.*, Pl.’s Br. at 33, but identifying and mitigating biases in AI system datasets is a technical matter rather than one of “political correctness” or editorial freedom, as xAI describes it. In the context of AI development—and data generally—bias is a *functional* problem representing several kinds of systemic flaws that could hamper an AI system (whether in its training data or elsewhere) and result in undesirable

¹⁹ Several suits filed in the Northern District of California are premised on arguments that the use of AI systems in hiring may violate the Fair Credit Reporting Act and Title VII, for instance. *See, e.g., Mobley v. Workday*, 740 F. Supp. 3d 796, 804-06 (N.D. Cal. 2024) (“Nothing in the language of the federal anti-discrimination statutes or the case law interpreting those statutes distinguishes between delegating functions to an automated agent versus a live human one.”); Not. of Removal at 35-37, *Kistler v. Eightfold AI*, No. 4:26-cv-01768 (N.D. Cal. Mar. 2, 2026) (claiming the defendant violated the Fair Credit Reporting Act).

²⁰ *See, e.g.*, Brad Smith, *Microsoft announces new Copilot Copyright Commitment for customers*, Microsoft (Jan. 5, 2024), <https://blogs.microsoft.com/on-the-issues/2023/09/07/copilot-copyright-commitment-ai-legal-concerns/>; *New models and developer products announced at DevDay*, OpenAI (Nov. 6, 2023), <https://openai.com/index/new-models-and-developer-products-announced-at-devday/>; Neal Suggs, *Shared fate: Protecting customers with generative AI indemnification*, Google Cloud (Oct. 12, 2023), <https://cloud.google.com/blog/products/ai-machine-learning/protecting-customers-with-generative-ai-indemnification>.

outcomes, from errors and malfunctions to actively harmful outputs.²¹ For one example, a traffic prediction model trained only on weekday commute data will badly misrepresent traffic conditions on weekends and holidays—due to “biased” or unrepresentative training data. For another, consider a medical imaging diagnostic model only trained on certain populations, like patients over a certain age, that therefore disproportionately overlooks life-threatening symptoms in younger patients. Bias mitigation is crucial to professional and enterprise consumers in the fields of medicine,²² finance,²³ cybersecurity,²⁴ and recruiting,²⁵ to name just a few. The technical biases that AI training data or model weights can contain are different in kind from the constitutionally protected “editorial judgments” that xAI describes.

²¹ See Ethan A. Abbott et al., *Understanding and Addressing Bias in Artificial Intelligence Systems*, J. Am. Coll. Emergency Phys. Open, Feb. 2026, at 2-5.

²² *When Algorithms Deny Care: Bias in Healthcare AI*, RAIL (Nov. 2, 2025), <https://responsibleailabs.ai/knowledge-hub/articles/healthcare-ai-bias>; James L. Cross et al., *Bias in medical AI: Implications for clinical decision-making*, PLOS Digit Health (Nov. 7, 2024), <https://journals.plos.org/digitalhealth/article?id=10.1371/journal.pdig.0000651>.

²³ Shawn McGrath, *How to mitigate AI discrimination and bias in financial services*, Ernst & Young (Mar. 17, 2023), https://www.ey.com/en_us/insights/forensic-integrity-services/ai-discrimination-and-bias-in-financial-services.

²⁴ *What Is AI Bias? Causes, Types, & Real-World Impacts*, Palo Alto Networks (last accessed May 14, 2026), <https://www.paloaltonetworks.com/cyberpedia/what-is-ai-bias>.

²⁵ Emilio J. Castilla, *AI is reinventing hiring—with the same old biases. Here’s how to avoid that trap*, MIT Mgmt. Sch. (Dec. 15, 2025), <https://mitsloan.mit.edu/ideas-made-to-matter/ai-reinventing-hiring-same-old-biases-heres-how-to-avoid-trap>.

xAI’s argument that California has no interest in the subject of AB 2013’s disclosures beyond “mere consumer curiosity” is misplaced: AB 2013’s disclosures about the training datasets used by AI systems available to Californian individuals and enterprises empower these consumers to make informed decisions in the marketplace for AI products and services.

C. The government has other substantial interests in AB 2013’s required disclosures.

California’s substantial interests in AB 2013’s disclosures go beyond consumer protection and market transparency. Legislative history reflects the state’s broader set of interests. *See, e.g.*, Cal. Sen. Rules Comm., Analysis of AB 2013, 2024 Reg. Sess., at 3 (Aug. 20, 2024) (“Requiring transparency about the training data used for AI systems helps identify and mitigate biases, addressing hallucinations and other problematic outputs, and shines the light on various other issues, such as privacy and copyright concerns.”); Cal. Assem. Comm. on Priv. & Consumer Prot., Hr’g on AB 2013, at 5-6 (Apr. 30, 2024) (“In their race to obtain vast quantities of training data, major AI developers have not hesitated to move fast and break things,” such as by scraping websites for data including “images containing child pornography” and “introduc[ing] biases into a system’s output”).

While any one interest alone is sufficient to defeat xAI’s First Amendment challenge, it is worth describing several—namely public health, child safety, and privacy—to highlight the breadth of the benefits that Californians may receive

from AB 2013’s required disclosures.

To begin with, AB 2013’s disclosures serve California’s interest in promoting public health. It is uncontestable that California has a “significant interest in protecting the health, safety, and welfare of its citizens.” *Rubin v. Coors Brewing Co.*, 514 U.S. 476, 485 (1995); *see also Md. Shall Issue, Inc.*, 91 F.4th at 250 (“It is elemental that government . . . has an interest in the health and safety of its citizens and, in particular, an ‘interest in preventing suicide, and in studying, identifying, and treating its causes.’” (quoting *Washington v. Glucksberg*, 521 U.S. 702, 730 (1997))).

Insight into whether AI systems developers have trained their models on datasets correlated with “AI psychosis” or suicidal ideation versus datasets curated to prevent such incidents is critical to that substantial interest, for instance. So too is transparency into how AI systems cure their training data of biases that, when deployed in public health surveillance or behavioral epidemiology, can “exacerbate existing health disparities.”²⁶ These are just two ways of many that AI intersects with public health issues that professional and regulatory organizations alike have

²⁶ Dimitra Panteli et al., *Artificial Intelligence in Public Health: Promises, Challenges, and An Agenda for Policy Makers and Public Health Institutions*, 10(5) *The Lancet* E428 (Feb. 28, 2025), [https://www.thelancet.com/journals/lanpub/article/PIIS2468-2667\(25\)00036-2/fulltext](https://www.thelancet.com/journals/lanpub/article/PIIS2468-2667(25)00036-2/fulltext); *see also* Anatol-Fiete Näher et al., *Measuring fairness preferences is important for artificial intelligence in health care*, 6 *The Lancet* E302 (May 2024), [https://www.thelancet.com/pdfs/journals/landig/PIIS2589-7500\(24\)00059-1.pdf](https://www.thelancet.com/pdfs/journals/landig/PIIS2589-7500(24)00059-1.pdf).

recognized and called for disclosures around, including AI training data disclosures.²⁷

It is similarly inarguable that “the prevention of sexual exploitation and abuse of children constitutes a government objective of surpassing importance.” *Packingham v. North Carolina*, 582 U.S. 98, 111 (2017) (cleaned up). The boom in AI development has led to a stunning increase in AI-generated child sexual abuse material (CSAM). For instance, analysts at the Internet Watch Foundation uncovered 3,440 photorealistic AI videos of child sexual abuse, often including real and recognizable child victims—a 26,362% increase from 2024.²⁸ Studies have demonstrated the presence of CSAM in “several prominent ML [machine learning] training datasets.”²⁹ The risks of training AI models on datasets containing CSAM are neither minor nor speculative, as researchers have explained

²⁷ See, e.g., *Ethics & Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models* 43, World Health Org. (2024), <https://iris.who.int/server/api/core/bitstreams/e9e62c65-6045-481e-bd04-20e206bc5039/content>; see also *AI Chatbots for Mental Health – What Works, What Harms, and What’s Next*, Nat’l Acad. of Med. (May 13, 2026), <https://nam.edu/news-and-insights/ai-chatbots-for-mental-health-what-works-what-harms-and-whats-next/>.

²⁸ Internet Watch Found., *AI Becoming ‘Child Sexual Abuse Machine,’ Adding to Dangerous Record Levels of Online Abuse, IWF Warns* (Jan. 16, 2026), <https://www.iwf.org.uk/news-media/news/ai-becoming-child-sexual-abuse-machine-adding-to-dangerous-record-levels-of-online-abuse-iwf-warns/>.

²⁹ David Thiel, *Identifying and Eliminating CSAM in Generative ML Training Data and Models* 10, Stanford Internet Observatory (Dec. 20, 2023), <https://purl.stanford.edu/kh752sm9123>.

that it plausibly enables the reproduction of specific victims’ likenesses.³⁰ As the California Assembly’s Committee on Privacy and Consumer Protection noted in discussing AB 2013, these mandatory disclosures would help the public to assess whether AI systems train models on datasets known to contain CSAM or correlate with AI-generated CSAM, thereby serving California’s substantial interest in preventing the sexual exploitation of children.³¹

Finally, courts have also long recognized privacy as an important state interest. *See Sorrell v. IMS Health Inc.*, 564 U.S. 552, 596 (2011) (Breyer, J., dissenting) (listing cases); *see also TikTok Inc. v. Garland*, 604 U.S. 56, 74-76 (2025) (concerning data privacy specifically); *Wollschlaeger v. Governor*, 848 F.3d 1293, 1314 (11th Cir. 2017) (“We recognize that protection of individual privacy is a substantial government interest.”). The California legislature noted this state interest in its discussions of AB 2013’s aims and scope.³²

³⁰ *Id.* (“The presence of repeated identical instances of CSAM is also problematic, particularly due to its reinforcement of images of specific victims.”).

³¹ *See* Cal. Assem. Comm. on Priv. & Consumer Prot., Hr’g on AB 2013, at 6 (Apr. 30, 2024).

³² *See, e.g.*, Cal. Sen. Jud. Comm., Analysis of AB 2013, 2023-2024 Reg. Sess., at 1, 5 (June 17, 2024) (noting that AI has “been used to constrain personal autonomy, compromise privacy and security, foment social upheaval, exacerbate inequality, spread misinformation, and subvert democracy” and explaining that “[r]equiring transparency about the training data used for AI systems helps identify and mitigate biases, addressing hallucinations and other problematic outputs, and shines the light on various other issues, such as privacy and copyright concerns”); Cal. Assem. Comm. on Priv. & Consumer Prot., Hr’g on AB 2013, at 2 (Apr. 30, 2024).

As the legislature articulated and AI developers have themselves acknowledged, AI training datasets can—and often do—contain personally identifiable information.³³ AI system developers and AI researchers likewise recognize that this personal information swept up in their training data should be—but often isn’t—properly anonymized.³⁴ Such privacy harms are not merely academic, since AI systems can sometimes be prompted, either unwittingly or by bad actors, to divulge the verbatim contents of their training data, from a private email address or phone number to entire passages of confidential text.³⁵ AB 2013’s mandatory disclosures are properly tailored to serve this substantial state interest,

³³ See Katharine Miller, *Privacy in an AI Era: How Do We Protect Our Personal Information?*, Stanford Inst. for Human-Centered AI (Mar. 18, 2024), <https://hai.stanford.edu/news/privacy-ai-era-how-do-we-protect-our-personal-information>; Off. of Tech., *AI Companies: Uphold Your Privacy and Confidentiality Commitments*, Fed. Trade Comm’n (Jan. 9, 2024), <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/01/ai-companies-uphold-your-privacy-confidentiality-commitments> (“[The] business incentive to constantly ingest additional data can be at odds with a company’s obligations to protect users’ data, undermining peoples’ privacy or resulting in the appropriation of a firm’s competitively significant data.”).

³⁴ Greg McDonough, *Privacy Risks in Training Data and How to Address Them*, RSAC Conf. (Oct. 28, 2025), <https://www.rsaconference.com/library/blog/privacy-risks-in-training-data-and-how-to-address-them>.

³⁵ See Michael M. Grynbaum & Ryan Mac, *The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work*, N.Y. Times (Dec. 27, 2023), <https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html> (“The complaint cites several examples when a chatbot **provided users with near-verbatim excerpts** from Times articles [included in training datasets] that would otherwise require a paid subscription to view.” (emphasis added)).

as AI developers must disclose whether their training datasets include personal and aggregate consumer information, as well as whether those datasets were cleaned, processed, or modified in any way (including for privacy harm mitigation).

These are but a few of the substantial state interests the state contemplated serving with AB 2013. xAI's argument that "consumer curiosity" is the sole interest advanced by AB 2013 misrepresents what is obvious from the text of the statute as well as its legislative history.

D. AB 2013's disclosures are neither unduly burdensome nor unjustified.

Finally, AB 2013's mandatory disclosures do not hinder xAI's ability to engage in protected speech. The Supreme Court has emphasized that "[u]njustified or unduly burdensome disclosure requirements offend the First Amendment by chilling protected speech." *Milavetz*, 559 U.S. at 249-50; *see also NIFLA*, 585 U.S. at 776. Despite xAI's far-fetched and unsupported assertion that AB 2013 imposes an "onerous and economically devastating" burden, Compl. ¶ 12, xAI has failed to allege what that burden actually *is*, let alone how any such burden threatens to chill xAI's First Amendment rights to engage in protected speech. AB 2013 also does not prohibit xAI from conveying any additional information alongside its factual disclosures, which further cures any potential burden concerns. *See Milavetz*, 559 U.S. at 250-52 (noting the statutory "flexibility to tailor the disclosures to [plaintiff's] individual circumstances" and "convey[] any additional information").

And as for xAI’s claim that it will suffer “economic devastation” from a narrow mandate to disclose high-level summaries of information that AI developers already track and publish in the form of public “model cards,”³⁶ neither evidence in the record nor common sense suggests that this is accurate for a company recently valued at \$250 billion.³⁷ xAI has already made a disclosure pursuant to AB 2013, after all, and it has not yet shuttered its doors.³⁸

To the extent AB 2013 does impose any nominal burden on xAI’s ability to engage in protected speech, such a burden is not unjustified. AB 2013’s mandated disclosures do not “impose a burden excessive or disproportionate to the benefits gained.” *See RJ Reynolds Tobacco Co. v. FDA*, 96 F.4th 863, 885-86 (5th Cir. 2024). The benefits gained by California’s consumers from AB 2013’s disclosures are clear and meaningful, while the alleged burdens to covered entities are neither.

³⁶ See, e.g., *Grok 4.1 Model Card*, xAI (Nov. 17, 2025), <https://data.x.ai/2025-11-17-grok-4-1-model-card.pdf>; Deployment Safety Hub, *GPT-5.5 System Card*, OpenAI (Apr. 23, 2026), <https://deploymentsafety.openai.com/gpt-5-5>; *System Card: Claude Opus 4.7*, Anthropic (Apr. 16, 2026), <https://cdn.sanity.io/files/4zrzovbb/website/037f06850df7fbe871e206dad004c3db5fd50340.pdf>; *Gemini 3.5 Flash Model Card*, Google DeepMind (May 2026), <https://storage.googleapis.com/deepmind-media/Model-Cards/Gemini-3-5-Flash-Model-Card.pdf>.

³⁷ Samantha Subin, *Musk’s xAI, SpaceX combo is the biggest merger of all time, valued at \$1.25 trillion*, CNBC (Feb. 3, 2026), <https://www.cnbc.com/2026/02/03/musk-xai-spacex-biggest-merger-ever.html>

³⁸ *xAI Frontier Artificial Intelligence Framework* 10-11, xAI (Dec. 30, 2025), <https://data.x.ai/2025-12-31-xai-frontier-artificial-intelligence-framework.pdf>.

CONCLUSION

For the foregoing reasons, the Court should affirm the district court's denial of xAI's Motion for Preliminary Injunction and clarify that AB 2013 survives First Amendment scrutiny under *Zauderer*, which is the appropriate standard to apply.

Date: July 22, 2026

/s/ Benjamin Rashkovich

Benjamin Rashkovich

Legal Advocates for Safe Science & Technology

125 Park Ave., Floor 25

New York, NY 10017

ben@lasst.org

(914) 330-4672

Counsel for Amici Curiae

/s/ Jason Harrow

Jason Harrow

Gerstein Harrow, LLP

12100 Wilshire Blvd. Ste. 800

Los Angeles, CA 90025

jason@gerstein-harrow.com

(323) 744-5293

Counsel for Amici Curiae

APPENDIX: LIST OF AMICI

I. Organizational Amici

Alliance for Secure AI

The Alliance for Secure AI is a nonprofit organization that educates the public about the implications of advanced AI.

Americans for Responsible Innovation (ARI)

Americans for Responsible Innovation (ARI) is a nonprofit organization dedicated to policy advocacy in the public interest, focused on emerging technologies like artificial intelligence (AI). ARI advocates for the U.S. government to take a thoughtful and proactive approach to AI governance that protects the public while maintaining our country's competitive edge.

California Initiative for Technology and Democracy (CITED)

The California Initiative for Technology and Democracy (CITED), a project of nonprofit organization California Common Cause, seeks state-level solutions to the threats that AI, disinformation, social media, and other new and evolving technologies pose to our democracy, our communities, and our families.

Center for Humane Technology

Center for Humane Technology is a nonprofit dedicated to ensuring that today's most consequential technologies, such as AI and social media, actually serve humanity. We bring clarity to how the tech ecosystem works in order to shift the incentives that drive it.

Cybersafety Research Center (CRC)

Cybersafety Research Center (CRC) is a research-based, nonpartisan, and independent multi-university center for problem-driven research and research-driven policy. CRC conducts cutting-edge computer science research to better understand the distorting effects of algorithms and AI tools on large online networks and how they affect safety, public health, and democracy. CRC works with platforms and regulators to help all parties understand the implications of their findings and develop solutions.

Digital Childhood Institute

The Digital Childhood Institute (DCI) is a 501(c)(3) nonprofit research and education organization founded by longtime child safety advocates, parents, and strategists. Through original research and relentless accountability work, DCI pushes the world's most powerful tech companies to adopt practices that put children ahead of profit.

Electronic Privacy Information Center (EPIC)

The Electronic Privacy Information Center ("EPIC") is a public interest research center in Washington, D.C., established in 1994 with a focus on emerging privacy and civil liberties issues. EPIC advocates for platform and AI governance structures that protect users' privacy, safety, and speech rights.

Encode

Encode AI Corporation ("Encode") is a nonprofit organization dedicated to advancing AI policy that protects the public interest. Encode advocates for AI governance that prioritizes safety, transparency, and alignment with human values.

Fairplay

Founded 25 years ago, Fairplay is a national, independent nonprofit organization dedicated to protecting children from unfair commercial practices in media and digital environments, where it is a nationally recognized expert. Fairplay engages in research, regulatory advocacy, coalition building, and public education addressing social media design, children's privacy, algorithmic amplification, AI deployment, and related harms.

Heat Initiative

Heat Initiative is a nonprofit organization dedicated to pressuring the technology industry to protect children online, holding tech companies accountable for enabling and profiting from child sexual abuse. Heat Initiative builds the evidentiary record through research and impact litigation, equips survivors, advocates, and policymakers with the knowledge and coordination to drive change, and applies sustained pressure on the companies where even small shifts can have an outsized impact on kids' safety.

Legal Advocates for Safe Science & Technology (LASST)

Legal Advocates for Safe Science and Technology, Inc. ("LASST") is a nonprofit organization dedicated to making advances in science and technology safer for people and the planet. LASST advocates for legal frameworks governing AI systems that appropriately balance innovation with security and public safety.

National Association of Voice Actors (NAVA)

NAVA is a social impact non-profit association created to advocate and promote the advancement of the voice acting industry through action, education, inclusion, and benefits. As advocates for voice, image, name, and likeness protections, NAVA has been an integral part of the global discussion around personality rights in the age of AI and has helped shape legislation and protections in the US and around the world.

The National Center on Sexual Exploitation (NCOSE)

Founded in 1962, the National Center on Sexual Exploitation (NCOSE) is the leading national non-profit organization exposing the links between all forms of sexual exploitation such as child sexual abuse, prostitution, sex trafficking, and the public health harms of pornography.

Secure AI Project

The Secure AI Project develops and advocates for pragmatic policies to reduce risks of severe harm from advanced AI.

Surveillance Technology Oversight Project (STOP)

The Surveillance Technology Oversight Project (“S.T.O.P.”) is a non-profit civil rights organization based in New York. It was founded in 2019 to promote community freedom by ending systems of mass surveillance. S.T.O.P. advocates in courts and legislatures across the country to ensure that technology serves the public without sacrificing privacy, with particular focus on how personal information is collected, used, and exposed by AI systems.

TechEquity

TechEquity is a non-profit organization dedicated to raising public consciousness about economic equity issues that result from the tech industry’s products and practices, and advocating for change that ensures tech’s evolution benefits everyone. TechEquity envisions a world where the tech industry is responsible for building widespread economic prosperity and is held accountable for the economic harms it creates in our communities.

Tech Justice Law (TJL)

Tech Justice Law (“TJL”) is a nonpartisan nonprofit impact litigation and advocacy organization bringing justice to individuals and communities harmed by tech products. Through direct and amicus curiae participation in litigation presenting AI safety concerns, TJL works to ensure artificial intelligence systems remain accountable to the rule of law through safety by design, human oversight, and transparency.

Tech Oversight Project

The Tech Oversight Project is a Washington, D.C.-based nonprofit advocacy organization that campaigns to hold large technology companies accountable. It advocates for stronger antitrust enforcement, privacy and data protections, online safety (including protections for children and families), and regulatory oversight of artificial intelligence and platform practices, using research, communications campaigns, and direct engagement with lawmakers.

Transparency Coalition

The Transparency Coalition is a non-profit organization committed to helping initiate legislation to regulate Artificial Intelligence to protect kids and families. We are a non-partisan, action-oriented organization with a voice for an otherwise overwhelmed general public, and a general demand for AI safety.

II. Individual Amici*

David Atkinson

David Atkinson is Legal Strategist at the Future of Life Institute (FLI). FLI is the world's oldest and largest AI think-tank, advocating for policies and approaches that mitigate AI catastrophic risks while advancing positive futures for humanity.

Prof. Helen Norton

Helen Norton is University Distinguished Professor and Rothgerber Chair in Constitutional Law at the University of Colorado School of Law. Her scholarly and teaching interests include constitutional law (especially First Amendment and equal protection law), and antidiscrimination law—and their intersection with technology.

Prof. Kyle Langvardt

Kyle Langvardt is an Associate Professor of Law at the University of Nebraska College of Law. He has written extensively on First Amendment issues raised by the regulation of software and the internet.

* These amici file this brief in their individual capacities; they provide institutional affiliations solely for purposes of identification. The positions taken in this brief are those of amici alone and should not be attributed to any institution with which amici are affiliated.

Prof. Lawrence Lessig

Lawrence Lessig is the Roy L. Furman Professor of Law and Leadership at Harvard Law School. Prior to returning to Harvard, he taught at Stanford Law School, where he founded the Center for Internet and Society, and at the University of Chicago. Once cited by *The New Yorker* as “the most important thinker on intellectual property in the Internet era,” Lessig has turned his focus from law and technology to “institutional corruption”—relationships which, while legal, weaken public trust in an institution—especially as that affects democracy.

Prof. Margot Kaminski

Margot Kaminski, the Moses Lasky Professor of Law at the University of Colorado, is an expert in AI law (co-authoring a forthcoming casebook with Foundation Press), algorithmic accountability, data privacy, and the First Amendment. She is the two-time recipient of Fulbright grants in 2018 and 2024, and her work has been published among other places in the *Yale Law Journal* and *Columbia Law Review*. Her work was cited by the Ninth Circuit sitting en banc in *Project Veritas v. Schmidt*.

Prof. Olivier Sylvain

Olivier Sylvain is a Professor of Law at Fordham University. His research and writing are on information and communications law and policy.

Prof. Peter Ormerod

Peter Ormerod is an Associate Professor of Law at Villanova University’s Charles Widger School of Law. His research interests include data governance, artificial intelligence, and the First Amendment.

Prof. Stuart Russell

Stuart Russell, OBE, FRS, is a Distinguished Professor of Computer Science at Berkeley and President of the International Association for Safe and Ethical AI. He is a leading researcher in artificial intelligence, a member of the National Academy of Engineering, a Fellow of the Royal Society, and the author (with Peter Norvig) of the standard text in the field. He has been active in arms control for nuclear and autonomous weapons. His latest book, *Human Compatible*, addresses the long-term impact of AI on humanity.

Prof. Vivek Krishnamurthy

Vivek Krishnamurthy is an Associate Professor at the University of Colorado Law School, where he directs the Samuelson-Glushko Technology Law & Policy Clinic.

Prof. Zephyr Teachout

Zephyr Teachout is a Professor of Law at Fordham Law School where she focuses on the intersection of corporate power and political power. In 2021, she took a leave to work as Special Advisor and Senior Counsel for Economic Justice at the New York Attorney General's Office.

CORPORATE DISCLOSURE STATEMENT

Pursuant to Federal Rule of Civil Procedure 26.1(a), the undersigned counsel states that the nonprofit, non-governmental organizational amici listed above have no parent corporations and that no publicly held corporation owns 10% or more of their stock.

Date: July 22, 2026

/s/ Benjamin Rashkovich

Benjamin Rashkovich

Legal Advocates for Safe Science & Technology

125 Park Ave., Floor 25

New York, NY 10017

ben@lasst.org

(914) 330-4672

Counsel for Amici Curiae