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 the State of California,

Defendant–Appellee.

On appeal from the United States District Court for the
Central District of California — No. 2:25-CV-12295 (Bernal, J.)

BRIEF OF THE KNIGHT FIRST AMENDMENT INSTITUTE AT COLUMBIA UNIVERSITY AS *AMICUS CURIAE* IN SUPPORT OF APPELLEE AND AFFIRMANCE

Stephany Kim
Jake Karr
Xiangnong (George) Wang
Alex Abdo
Knight First Amendment Institute
at Columbia University
475 Riverside Drive, Suite 302
New York, NY 10115
(646) 745-8500
stephany.kim@knightcolumbia.org
jake.karr@knightcolumbia.org

Counsel for Amicus Curiae

Table of Contents

Table of Authorities....ii

Interest of *Amicus Curiae*....1

Introduction and Summary of Argument....2

Argument....4

I. Transparency into generative AI systems empowers the public to make informed decisions about whether and how to use those systems....4

II. The *Zauderer* test appropriately balances the informational interests in compelled commercial disclosures against the potential expressive burdens....10

A. Courts have long acknowledged the public interests in compelled commercial disclosures....10

B. *Zauderer* is the correct standard for evaluating compelled disclosures of factual and uncontroversial information about commercial products or services....11

C. *Zauderer* accommodates competing First Amendment interests....17

III. A.B. 2013 is subject to and, on this record, passes the *Zauderer* test....19

A. A.B. 2013 compels the disclosure of factual and uncontroversial information about generative AI systems available to the public. 19

B. A.B. 2013 is reasonably related to California's interest in helping inform users about the generative AI systems offered to them....22

C. The preliminary injunction record does not demonstrate that A.B. 2013 imposes an undue burden on xAI....25

Conclusion....26

Table of Authorities

Cases

303 Creative LLC v. Elenis, 600 U.S. 570 (2023) 13

Am. Beverage Ass'n v. City & Cnty. of S.F., 916 F.3d 749 (9th Cir. 2019) 17, 18, 19

Am. Hosp. Ass'n v. Azar, 983 F.3d 528 (D.C. Cir. 2020) 25, 26

Am. Meat Inst. v. USDA, 760 F.3d 18 (D.C. Cir. 2014) 26

Beach v. Ocwen Fed. Bank, 523 U.S. 410 (1998) 10

Biden v. Knight First Amend. Inst. at Columbia Univ., 141 S. Ct. 1220 (2021) 4, 5

Buckley v. Valeo, 424 U.S. 1 (1976) 11

Cent. Hudson Gas & Elec. Corp. v. Pub. Serv. Comm'n, 447 U.S. 557 (1980) 16

Chamber of Com. v. SEC, 85 F.4th 760 (5th Cir. 2023) 18, 21

Citizens United v. FEC, 558 U.S. 310 (2010) 11

CompassCare v. Hochul, 125 F.4th 49 (2d Cir. 2025) 19

CTIA – The Wireless Ass'n v. City of Berkeley, 928 F.3d 832 (9th Cir. 2019) 14, 20, 21

Edenfield v. Fane, 507 U.S. 761 (1993) 22

Expressions Hair Design v. Schneiderman, 877 F.3d 99 (2d Cir. 2017) 15

First Nat'l Bank of Bos. v. Bellotti, 435 U.S. 765 (1978) 11

Fresenius Med. Care Orange Cnty., LLC v. Bonta, 172 F.4th 718 (9th Cir. 2026) 14, 15, 16

hiQ Labs, Inc. v. LinkedIn Corp., 31 F.4th 1180 (9th Cir. 2022) 23

Hum. Life of Wash. Inc. v. Brumsickle, 624 F.3d 990 (9th Cir. 2010) 11

Md. Shall Issue, Inc. v. Anne Arundel Cnty. Md., 91 F.4th 238 (4th Cir. 2024) 15, 16

Merck Sharp & Dohme Corp. v. Albrecht, 587 U.S. 299 (2019) 10

Moody v. NetChoice, LLC, 603 U.S. 707 (2024) .....passim

Nat'l Inst. of Fam. & Life Advocs. v. Becerra, 585 U.S. 755 (2018) .....passim

Nationwide Biweekly Admin., Inc. v. Owen, 873 F.3d 716 (9th Cir. 2017) 4, 18, 22, 25

Pac. Gas & Elec. Co. v. Pub. Utilities Comm'n, 475 U.S. 1 (1986).....13, 19

Pharm. Care Mgmt. Ass'n v. Rowe, 429 F.3d 294 (1st Cir. 2005).....15

Pharm. Rsch. & Mfrs. of Am. v. Stolfi, 153 F.4th 795 (9th Cir. 2025)..11, 16, 20, 23

R.J. Reynolds Tobacco Co. v. FDA, 96 F.4th 863 (5th Cir. 2024).....19

S.F. Apartment Ass'n v. City & Cnty. of S.F., 881 F.3d 1169 (9th Cir. 2018) 11, 14, 23

South Dakota v. Wayfair, 585 U.S. 162 (2018).....5

Uber Techs., Inc. v. City of Seattle, 168 F.4th 1202 (9th Cir. 2026).....14

Va. State Bd. of Pharmacy v. Va. Citizens Consumer Council, Inc., 425 U.S. 748 (1976) 10

Volokh v. James, 148 F.4th 71 (2d Cir. 2025).....8, 15, 24

X Corp. v. Bonta, 116 F.4th 888 (9th Cir. 2024).....21

Zauderer v. Off. of Disciplinary Counsel, 471 U.S. 626 (1985) .....passim

Statutes

15 U.S.C. § 1601(a).....10

Cal. Civ. Code § 3110 5

Cal. Civ. Code § 3111 .....passim

Other Authorities

Alex Abdo et al., *A Safe Harbor for Platform Research*, Knight First Amend. Inst. at Columbia Univ. (Jan. 19, 2022) 7

Alexander Wan et al., *The 2025 Foundation Model Transparency Index* (2025) 7

Arvind Narayanan & Sayash Kapoor, *Generative AI Companies Must Publish Transparency Reports*, Knight First Amend. Inst. at Columbia Univ. (June 26, 2023) 7

Brendan Nyhan et al., *Easy to Produce, Hard to Persuade: The Asymmetric Effects of AI on the Online Information Ecosystem*, *in Artificial Intelligence, Politics, and Political Science* 56 (Nathaniel Persily & Joshua A. Tucker eds., 2026) 5, 6

Daphne Keller, *Platform Transparency and the First Amendment*, 4 J. Free Speech L. 1 (2023) 9

Grok, Tesla Support 22

Helen Toner, *What Are Generative AI, Large Language Models, and Foundation Models?*, Ctr. for Sec. & Emerging Tech. (May 12, 2023) 5

Jennifer Medina, *'Who Should I Vote for?' Voters Turn to A.I. Before Casting Their Ballots*, N.Y. Times (July 8, 2026) 6

Kate Crawford & Trevor Paglen, *Excavating AI*, AI Now Inst., https://excavating.ai 24

Nat'l Telecomms. & Info. Admin., *Artificial Intelligence Accountability Policy Report* (2024) 7, 8, 9, 24

Nathaniel Persily & Joshua A. Tucker, *Introduction*, *in Artificial Intelligence, Politics, and Political Science* (Nathaniel Persily & Joshua A. Tucker eds., 2026), ...9

Sha Sajadieh et al., *The AI Index 2026 Annual Report* (2026).....8

Shayne Longpre et al., *A Safe Harbor for AI Evaluation and Red Teaming*, 235 ICML '24: Procs. 41st Int'l Conf. on Mach. Learning 32,691 (2024).....5

Shayne Longpre et al., *Data Authenticity, Consent, & Provenance for AI Are All Broken*, 235 ICML '24: Procs. 41st Int'l Conf. on Mach. Learning 32,711 (2024).....8, 9

Stratis Tsirtsis et al., *AI-Mediated Communication Can Steer Collective Opinion*, ICML '26: Procs. 43rd Int'l Conf. on Mach. Learning 1 (forthcoming 2026).....6

Tejas N. Narechania & Ganesh Sitaraman, *An Antimonopoly Approach to Governing Artificial Intelligence*, 43 Yale L. & Pol'y Rev. 95 (2024).....6

U.S. Gov't Accountability Off., GAO-25-107172, *Artificial Intelligence: Generative AI's Environmental and Human Effects* (2025) 6

Yoshua Bengio et al., *International AI Safety Report* (2026) 7

Interest of *Amicus Curiae*

The Knight First Amendment Institute at Columbia University ("Knight Institute" or "Institute") is a non-partisan, not-for-profit organization that works to defend the freedoms of speech and the press in the digital age through strategic litigation, research, policy advocacy, and public education. The Institute's aim is to promote a system of free expression that is open and inclusive, that broadens and elevates public discourse, and that fosters creativity, accountability, and effective self-government. The Institute is particularly committed to illuminating the forces that shape our information ecosystem online, and the Institute frequently appears as *amicus curiae* in federal courts, including the Ninth Circuit, to advocate for greater access to and understanding of the privately owned online platforms through which the public increasingly engages in discourse, democracy, and commerce. *See, e.g.*, *Moody v. NetChoice, LLC*, 603 U.S. 707 (2024) (social media platform transparency); *X Corp. v. Ctr. for Countering Digit. Hate, Inc.*, No. 24-2643 (9th Cir.) (platform research and journalism); *Amazon.com Servs. LLC v. Perplexity AI, Inc.*, No. 26-1444 (9th Cir.) (same). This case may have far-reaching implications for the ability of the government to enact legislation essential to ensuring that the digital public sphere serves the public interest.[1]

[1] No counsel for a party authored this brief in whole or in part, and no person other than *amicus curiae* and its counsel contributed money that was intended to fund preparing or submitting the brief. The parties have consented to the filing of this

Introduction and Summary of Argument

The district court was correct to conclude that Plaintiff xAI has not demonstrated a likelihood of success on the merits of its First Amendment claim. ER-13. The court erred, however, in failing to apply the correct standard for evaluating laws that, like A.B. 2013, facilitate the flow of factual and uncontroversial information about commercial products or services. ER-12. *Amicus* submits this brief to make three related points.

First, laws that provide for transparency into generative artificial intelligence ("AI") systems can serve important public interests. Generative AI systems have increasingly come to shape modern public discourse, democracy, and the economy. But the powerful private companies that develop and deploy these systems gatekeep information about how they are built, how they operate, and how they impact us. Laws that require the disclosure of such key information can aid the public in making informed choices about whether and how to use particular generative AI systems, and they can provide researchers and journalists the access necessary to enable independent evaluation of the benefits and risks of these systems.

Second, *Zauderer v. Office of Disciplinary Counsel* provides the appropriate framework for evaluating the compelled disclosure of factual and uncontroversial

amicus brief. The Institute would like to thank legal intern Sylvan Lebrun for her significant contributions to this brief.

information about generative AI systems. 471 U.S. 626 (1985). The threshold question under *Zauderer* is whether a challenged law compels the disclosure of "purely factual and uncontroversial information about the terms under which . . . [products or] services will be available." *Id.* at 651. If so, the law is constitutional so long as it is "reasonably related to the State's interest" and not "unduly burdensome." *Id.* Importantly, *Zauderer* accounts not just for the value of disclosures to the public, but also for the potential burdens on speech that disclosure requirements may impose. And even under the more relaxed review that *Zauderer* prescribes, courts have not hesitated to strike down compelled disclosures that unduly burden protected expression.

Third, A.B. 2013 is subject to and—at least on the preliminary injunction record—passes the *Zauderer* test. The law requires the developer of any generative AI system "made publicly available to Californians for use" to disclose basic information "regarding the data used by the developer to train the system." Cal. Civ. Code § 3111. These disclosures are undisputedly factual, noncontroversial, and related to xAI's commercial products. They also reasonably relate to California's asserted interest in informing users about the generative AI systems on offer in the marketplace. Finally, A.B. 2013's minimal and well-justified burden is not undue. The allegations of harm that xAI has put forth thus far are conclusory, vague, or irrelevant to the *Zauderer* inquiry. Simply put, "[t]he First Amendment does not

generally protect" large technology platforms "from being required to tell prospective customers the truth" about their generative AI systems, *Nationwide Biweekly Admin., Inc. v. Owen*, 873 F.3d 716, 721 (9th Cir. 2017), and xAI has not demonstrated that A.B. 2013 likely violates xAI's First Amendment rights.

This Court should affirm the district court's denial of a preliminary injunction with respect to xAI's First Amendment claim. In doing so, the Court should apply the *Zauderer* test to A.B. 2013 and hold that xAI has failed, on the current record, to demonstrate a likelihood of success on the merits of that claim.

Argument

I. Transparency into generative AI systems empowers the public to make informed decisions about whether and how to use those systems.

Large technology platforms wield tremendous influence over our public discourse, democracy, and economy. Search engines determine what information we are able to access when we browse the Internet. *See Biden v. Knight First Amend. Inst. at Columbia Univ.*, 141 S. Ct. 1220, 1224–25 (2021) (Thomas, J., concurring) ("Google is the gatekeeper between [the] user and the speech of others . . . ."). Social media platforms determine how and with whom we are able to communicate online. *See Moody v. NetChoice, LLC*, 603 U.S. 707, 716 (2024) ("Social-media platforms [are] . . . inescapable. They structure how we relate to family and friends, as well as to businesses, civic organizations, and governments."). And online marketplaces determine which goods we can buy and at what price. *See South Dakota v. Wayfair*,

585 U.S. 162, 184 (2018) ("The Internet's prevalence and power have changed the dynamics of the national economy."). The power that these platforms have to shape our society is made more potent by their concentration; in each of these realms, a handful of dominant companies largely dictates the terms of our engagement. *See Knight First Amend. Inst.*, 141 S. Ct. at 1224 (Thomas, J., concurring) ("[T]his concentration gives some digital platforms enormous control over speech.").

Advances in generative AI are enlarging the platforms' power over our online lives.[2] They have deployed generative AI systems "rapidly in recent years, amassing hundreds of millions of users."[3] Instead of relying on traditional search engines, many of us now "interact[] with an AI intermediary that . . . serves as a gatekeeper to original sources."[4] When users turn to Anthropic's Claude, OpenAI's ChatGPT, or xAI's Grok, these generative AI systems "decide what is relevant, what to

---

[2] In this brief, *amicus* uses the term "generative AI" to refer to "any AI system whose primary function is to generate content," including those that are trained to produce text—also known as "large language models"—and those that are trained to produce other types of media like images or video. Helen Toner, *What Are Generative AI, Large Language Models, and Foundation Models?*, Ctr. for Sec. & Emerging Tech. (May 12, 2023), https://cset.georgetown.edu/article/what-are-generative-ai-large-language-models-and-foundation-models; *see also* Cal. Civ. Code § 3110(c) (defining the term for purposes of the statute).

[3] Shayne Longpre et al., *A Safe Harbor for AI Evaluation and Red Teaming*, 235 ICML '24: Procs. 41st Int'l Conf. on Mach. Learning 32,691, 32,691 (2024) ("*AI Safe Habor*").

[4] Brendan Nyhan et al., *Easy to Produce, Hard to Persuade: The Asymmetric Effects of AI on the Online Information Ecosystem*, *in Artificial Intelligence, Politics, and Political Science* 56, 58 (Nathaniel Persily & Joshua A. Tucker eds., 2026).

summarize, and what to omit from an information environment that these models are increasingly populating."[5] These systems also have become "widely embedded" across the major social media platforms, "effectively mediating communication on the very platforms where humans often exchange and form opinions about contested social and political issues."[6] *See Moody*, 603 U.S. at 795 (Alito, J., concurring) ("[M]any of the biggest platforms are beginning to use AI algorithms to help them moderate content."). Generative AI is growing in availability and capability, right alongside the "outsized influence" of "a relatively small number of companies . . . over the information ecosystem."[7]

But with this great power has not come great transparency. Generative AI systems present "unparalleled opportunities and unprecedented dangers," *Moody*, 603 U.S. at 716, but the public and policymakers know surprisingly little about them.[8] That is because the companies that develop these platforms monopolize

---

[5] *Id*. at 59; *see also, e.g.*, Jennifer Medina, *'Who Should I Vote for?' Voters Turn to A.I. Before Casting Their Ballots*, N.Y. Times (July 8, 2026), https://www.nytimes.com/2026/07/04/us/politics/voters-ai-chatbots-elections.html.

[6] Stratis Tsirtsis et al., *AI-Mediated Communication Can Steer Collective Opinion*, ICML '26: Procs. 43rd Int'l Conf. on Mach. Learning 1, 1, 12 (forthcoming 2026) (emphasis omitted), https://arxiv.org/pdf/2605.16245 ("AI-mediated communication is a novel lever for online platforms to influence opinion formation.").

[7] Tejas N. Narechania & Ganesh Sitaraman, *An Antimonopoly Approach to Governing Artificial Intelligence*, 43 Yale L. & Pol'y Rev. 95, 103 (2024).

[8] *See*, *e.g.*, U.S. Gov't Accountability Off., GAO-25-107172, *Artificial Intelligence: Generative AI's Environmental and Human Effects* 19 (2025),

information about how they are built, how they operate, and how they impact our online lives—and they choose to provide only "limited transparency and access into their systems."[9] Companies do not publish standardized transparency reports that would empower users to compare the benefits and risks of using different systems,[10] nor do they grant access to data that would permit independent researchers to evaluate their systems.[11] In fact, the "most capable systems are also the least

---

https://www.gao.gov/assets/gao-25-107172.pdf (noting the difficulty of understanding the risks and challenges of generative AI because the technology "is rapidly evolving, and private developers do not disclose some key technical information"); Yoshua Bengio et al., *International AI Safety Report* 25 (2026), https://internationalaisafetyreport.org/sites/default/files/2026-02/international-ai-safety-report-2026_1.pdf ("[T]here is a limited degree of public and policymaker knowledge about how most advanced models are developed, safeguarded, evaluated, and deployed.").

[9] Longpre et al., *AI Safe Harbor*, at 32,691; *see also* Alexander Wan et al., *The 2025 Foundation Model Transparency Index* 1–3 (2025), https://crfm.stanford.edu/fmti/December-2025/paper.pdf.

[10] Arvind Narayanan & Sayash Kapoor, *Generative AI Companies Must Publish Transparency Reports*, Knight First Amend. Inst. at Columbia Univ. (June 26, 2023), https://knightcolumbia.org/blog/generative-ai-companies-must-publish-transparency-reports; *see also* Nat'l Telecomms. & Info. Admin., *Artificial Intelligence Accountability Policy Report* 30 (2024), https://www.ntia.gov/sites/default/files/publications/ntia-ai-report-final.pdf (noting that the lack of standardization in voluntary disclosures "frustrate[s] meaningful comparison of different models or systems").

[11] *See* Longpre et al., *AI Safe Harbor*, at 32,695 (explaining how researcher access is "tightly[]scoped" and may be revoked at the company's discretion); *cf.* Alex Abdo et al., *A Safe Harbor for Platform Research*, Knight First Amend. Inst. at Columbia Univ. (Jan. 19, 2022), https://knightcolumbia.org/content/a-safe-harbor-for-platform-research (describing how platforms provide social media data that is limited, incomplete, and outdated, and they can withdraw data access on a whim).

transparent, with training code, dataset sizes, and parameter counts increasingly withheld."[12]

Users need greater transparency from these companies to make informed decisions about whether to use particular generative AI systems and how to do so effectively and safely. *Cf. Volokh v. James*, 148 F.4th 71, 91 (2d Cir. 2025) (transparency into "the terms of their engagement with a social media network[] enabl[es users] to make more informed choices about where they spend their screen time and how to interpret the content they find on a given social media network"). Tailored disclosures could help users "assess the trustworthiness of a model and the appropriateness of a given use of that model."[13] Greater data transparency, in particular, could "highlight[] likely gaps in model coverage and abilities."[14]

Researchers also need greater access to these systems to provide the public and policymakers with independent evaluations of their benefits and risks. A researcher who wants to "examine whether an application has produced unlawfully

---

[12] Sha Sajadieh et al., *The AI Index 2026 Annual Report* 12 (2026), https://arxiv.org/pdf/2606.15708; *see also* Longpre et al., *AI Safe Harbor*, at 32,693 ("For external researchers, the models themselves are . . . black boxes, as developers often do not disclose model architectures, sizes, or training data.").

[13] Nat'l Telecomms. & Info. Admin., *supra*, at 27.

[14] Shayne Longpre et al., *Data Authenticity, Consent, & Provenance for AI Are All Broken*, 235 ICML '24: Procs. 41st Int'l Conf. on Mach. Learning 32,711, 32,711–12 (2024) ("*Data Provenance*").

discriminatory outcomes" needs access to the system's input and output data.[15] A researcher who wants to "assess the damage that could result from malign use of advanced AI, such as large language models," may need even more information, such as the model's architecture and weights.[16] With "[a] better understanding of what goes into" generative AI models, researchers can identify "their strengths, weaknesses, and social biases."[17]

Without this information, the public will be unable to meaningfully evaluate the profound impacts of generative AI on our public discourse, democracy, and economy.[18] As citizens, we will lack the information we need to engage in effective self-governance.[19] And as consumers, we will lack the information we need to make informed decisions about where and how we spend our time online.[20]

---

[15] Nat'l Telecomms. & Info. Admin., *supra*, at 36.

[16] *Id*.

[17] Longpre et al., *Data Provenance*, at 32,713.

[18] *See* Nathaniel Persily & Joshua A. Tucker, *Introduction*, *in Artificial Intelligence, Politics, and Political Science*, *supra*, at 1, 4–5.

[19] *See* Daphne Keller, *Platform Transparency and the First Amendment*, 4 J. Free Speech L. 1, 74 (2023).

[20] *See id.* at 54–55.

II. The *Zauderer* test appropriately balances the informational interests in compelled commercial disclosures against the potential expressive burdens.

A. Courts have long acknowledged the public interests in compelled commercial disclosures.

The Supreme Court has recognized that "people will perceive their own best interests if only they are well enough informed." *Va. State Bd. of Pharmacy v. Va. Citizens Consumer Council, Inc.*, 425 U.S. 748, 763, 770 (1976). The "free flow of information" ensures that we can effectively engage as citizens "in the day's most urgent political debate" and make "intelligent and well[-]informed" decisions as consumers in our "predominantly free enterprise economy." *Id.* at 763–65.

Accordingly, courts give governments substantial leeway to compel accurate and useful information from powerful corporate actors. In the commercial context, the Supreme Court has acknowledged that mandatory labels can better inform consumers of the uses and risks of products in the marketplace, *Merck Sharp & Dohme Corp. v. Albrecht*, 587 U.S. 299, 303–04 (2019), and that "meaningful disclosure of credit terms" can enable consumers to "compare more readily the various credit terms available to [the]m and avoid the uninformed use of credit," *Beach v. Ocwen Fed. Bank*, 523 U.S. 410, 412 (1998) (quoting 15 U.S.C. § 1601(a)). This Court has similarly recognized the important interests served by legislative efforts aimed at reducing the "significant informational asymmetries" that exist in the market for pharmaceutical drugs, *Pharm. Rsch. & Mfrs. of Am. v. Stolfi*, 153

F.4th 795, 826 (9th Cir. 2025), and in negotiations between tenants and landlords, *S.F. Apartment Ass'n v. City & Cnty. of S.F.*, 881 F.3d 1169, 1178–80 (9th Cir. 2018).

In recognition of the broad public importance of transparency laws, courts have sanctioned compelled disclosures in other contexts, too. For example, the Supreme Court has upheld campaign finance disclosures that "curb[] the evils of campaign ignorance," *Buckley v. Valeo*, 424 U.S. 1, 68 (1976), and "enabl[e] the electorate to make informed decisions and give proper weight to different speakers and messages," *Citizens United v. FEC*, 558 U.S. 310, 371 (2010). As the Ninth Circuit has stated, these sorts of "[d]isclosure requirements . . . allow the people in our democracy" to "evaluat[e] the relative merits of conflicting arguments" and "consider . . . the source and credibility of the advocate." *Hum. Life of Wash. Inc. v. Brumsickle*, 624 F.3d 990, 1008 (9th Cir. 2010) (quoting *First Nat'l Bank of Bos. v. Bellotti*, 435 U.S. 765, 791–92 (1978)).

B. *Zauderer* is the correct standard for evaluating compelled disclosures of factual and uncontroversial information about commercial products or services.

In *Zauderer*, the Supreme Court established the constitutional test that applies to laws that facilitate the free flow of information about commercial products or services. *Zauderer* applies to laws that compel the disclosure of "purely factual and uncontroversial information about the terms under which . . . [products or] services

will be available." 471. U.S. at 651. If *Zauderer* applies, a law is constitutional so long as it is "reasonably related to the State's interest" and not "unduly burdensome." *Id.*; *see also Nat'l Inst. of Fam. & Life Advocs. v. Becerra* ("*NIFLA*"), 585 U.S. 755, 775 (2018) (reaffirming the "legality" of "purely factual and uncontroversial disclosures about commercial products").

In *Zauderer* itself, the Supreme Court upheld a rule that required lawyers who advertised their services on a contingency-fee basis to disclose that clients could be required to pay fees and costs. 471 U.S. at 650–53. The Court emphasized that the disclosure requirement did "not attempt[] to prevent attorneys from conveying information to the public," but "only required them to provide somewhat more information than they might otherwise be inclined to present"—"purely factual and uncontroversial information about the terms under which [their] services will be available." *Id.* at 650–51; *see also id.* at 651 n.14 ("[T]he First Amendment interests implicated by disclosure requirements are substantially weaker than those at stake when speech is actually suppressed . . . ."). The plaintiff attorney's "constitutionally protected interest in *not* providing any particular factual information in his advertising [was] minimal," the Court reasoned, because the First Amendment's protection for advertising—core "commercial speech"—"is justified principally by the value to consumers of the information such speech provides." *Id.* at 651 (emphasis in original). In other words, "[t]he right of a commercial speaker not to

divulge accurate information regarding his services is not such a fundamental right"—and it is "adequately protected as long as disclosure requirements are reasonably related to the State's interest in preventing deception of consumers." *Id.* 651 & n.14.

The Supreme Court has subsequently made clear that *Zauderer*'s reasoning applies broadly to commercial transparency laws that provide the public with information about the products and services they might consume. Just two terms ago, the Court directed lower courts on remand to apply *Zauderer* to transparency laws that "require a [social media] platform to provide an individualized explanation to a user if it removes or alters her posts." *Moody*, 603 U.S. at 717, 725–26. Even before that decision, however, the Court had made clear many times that "[t]he State, of course, has substantial leeway in determining appropriate information disclosure requirements for business corporations." *Pac. Gas & Elec. Co. v. Pub. Utilities Comm'n*, 475 U.S. 1, 15 n. 12 (1986) (citing *Zauderer*, 471 U.S. at 651); *see also, e.g.*, *NIFLA*, 585 U.S. at 775 ("[W]e do not question the legality of health and safety warnings long considered permissible, or purely factual and uncontroversial disclosures about commercial products."); *303 Creative LLC v. Elenis*, 600 U.S. 570, 596 (2023) ("[O]ur cases have held that the government may sometimes requir[e] the dissemination of purely factual and uncontroversial information, particularly in the context of commercial advertising." (citation and quotation marks omitted)).

The Ninth Circuit has followed suit, applying *Zauderer* to compelled disclosure laws in a wide variety of commercial contexts. Recently, this Court applied *Zauderer* to a California law that compels charitable organizations "to inform patients of their available health care options," *Fresenius Med. Care Orange Cnty., LLC v. Bonta*, 172 F.4th 718, 734 (9th Cir. 2026), and to a Washington law that compels platforms, including rideshare companies, to inform their app-based workers of their account deactivation policies, *Uber Techs., Inc. v. City of Seattle*, 168 F.4th 1202, 1215–16 (9th Cir. 2026); *see also, e.g.*, *S.F. Apartment Ass'n*, 881 F.3d at 1178–80 (form describing tenants' rights and providing contact information for tenants' rights organizations prior to the commencement of buyout negotiations); *CTIA – The Wireless Ass'n v. City of Berkeley*, 928 F.3d 832, 845 (9th Cir. 2019) (cell phone safety warnings).

Other federal courts of appeals have likewise applied *Zauderer* to all manner of compelled commercial disclosures, including platform transparency laws. The Fifth and the Eleventh Circuits agreed in the *NetChoice* cases (in decisions that were later vacated on other grounds) that *Zauderer* applies to Texas and Florida laws requiring platforms to disclose information and data about their content moderation policies and practices (even as they came to opposing results in applying the test to the disclosure laws at issue). *See Moody*, 603 U.S. at 722. And the Second Circuit explained that a law "merely requir[ing] social media networks to publicly disclose

their content moderation policies" would be properly subject to *Zauderer* review. *Volokh*, 148 F.4th at 89; *see also, e.g.*, *Expressions Hair Design v. Schneiderman*, 877 F.3d 99, 103 (2d Cir. 2017) ("*Zauderer*'s less-exacting standard applies whenever a law mandates that a merchant disclose specific facts relating to a product it is offering for sale . . . ."); *Md. Shall Issue, Inc. v. Anne Arundel Cnty. Md.*, 91 F.4th 238, 248–50 (4th Cir. 2024) (suicide prevention pamphlets at gun dealerships); *Pharm. Care Mgmt. Ass'n v. Rowe*, 429 F.3d 294, 309–10 (1st Cir. 2005) (pharmacy benefit managers' conflicts of interest and financial arrangements with third parties).

The parties suggest that *Zauderer* has another threshold requirement: namely, that the law in question compel the disclosure of "commercial speech." *See* Appellant's Br. 32 n.1, ECF No. 15.1; Appellee's Br. 26–27, ECF No. 20.1. But that is not so. The Supreme Court justified *Zauderer* based on the interests of consumers in knowing more about the products and services that they might consume. 471 U.S. at 651. That is why the only threshold question under *Zauderer* is whether the law compels the disclosure of factual and uncontroversial information about commercial products or services. *See id.*; *NIFLA*, 585 U.S. at 775. And it is why this Court, in *Fresenius*, applied *Zauderer* without engaging in a commercial speech inquiry, explaining that "[t]he State may compel disclosure of factual and uncontroversial information to consumers, . . . , so long as that disclosure does not burden expressive association." 172 F.4th at 734 (citations and quotation marks omitted). As this Court

put it in that context, *Zauderer* is only "inapplicable when the compelled speech is 'in no way relate[d] to the services that [the regulated entity] provide[d].'" *Id.* (quoting *NIFLA*, 585 U.S. at 769)).[21]

In any event, this Court has adopted an expansive definition of what constitutes "commercial speech." Under this Court's precedent, where a disclosure "is closely tethered to the sale of a product and 'assists consumers and furthers the societal interest in the fullest possible dissemination of information,' it is properly categorized as commercial speech." *Stolfi*, 153 F.4th at 822 (quoting *Cent. Hudson Gas & Elec. Corp. v. Pub. Serv. Comm'n*, 447 U.S. 557, 561–62 (1980)); *see also, e.g.*, *Md. Shall Issue, Inc.*, 91 F.4th at 248 ("[S]peech connected with the sale of a good or a service—promoting the product or service, explaining it, or giving warnings about it—is commercial . . . ."). Thus, whether interpreted according to its plain terms or through the lens of commercial speech as defined by this Circuit, the *Zauderer* test applies broadly to laws that compel the disclosure of purely factual and uncontroversial information about commercial products or services.

[21] The court below also appeared to believe that it first needed to analyze whether the law in question is content based. ER-8–10. But that inquiry, too, is irrelevant to *Zauderer*'s threshold test. Indeed, compelled commercial disclosures are generally based on the content of the disclosure.

C. *Zauderer* accommodates competing First Amendment interests.

Importantly, *Zauderer* accounts not just for the value of disclosures to the public, but also for the potential burdens on speech that disclosure requirements may impose. Finding that *Zauderer* applies to a given law does not end the inquiry. As noted above, courts still must determine whether the disclosure requirement at issue is "unduly burdensome"—that is, whether the "requirement is not justified when balanced against its likely burden" on protected expression. *Am. Beverage Ass'n v. City & Cnty. of S.F.*, 916 F.3d 749 (9th Cir. 2019). This inquiry calls for a sliding-scale balancing of a disclosure's informational benefits against its expressive burdens. The greater the burden a disclosure requirement imposes, the greater the state's interest must be to justify it.

Even under the more relaxed review that *Zauderer* prescribes, this Court and sister circuits have not hesitated to strike down compelled disclosures that unduly burdened protected activity. In *American Beverage Association*, for example, the Court reversed the denial of a preliminary injunction against a San Francisco ordinance that required health warnings on advertisements for sugar-sweetened beverages to "occupy at least 20% of the advertisement." 916 F.3d at 754, 758. The Court concluded that, based on the record before it, the 20% requirement was "unduly burdensome" because it risked "drown[ing] out" the advertisers' message where "a smaller warning—half the size—would accomplish [San Francisco's]

stated goals." *Id.* at 757 (citing *NIFLA*, 585 U.S. at 778). And in the Eleventh Circuit's now-vacated *NetChoice* decision, the court partially affirmed a preliminary injunction against a Florida law to the extent that it would have required social media platforms to provide a "thorough rationale" every time they "remove[d] or alter[ed]" a user's posted content because, the court reasoned, that requirement would have unduly burdened the platforms' protected speech. *Moody*, 603 U.S. at 720 (citation omitted).

On the other hand, courts routinely uphold compelled disclosures that do *not* impose an undue burden on First Amendment rights. This Court previously rejected a mortgage servicer's challenge to a California law requiring it to disclose its lack of authorization from lenders in mail solicitations. *Owen*, 873 F.3d at 731, 735. "All mandatory disclosures impose *some* burden," the Court observed, but "[t]he First Amendment does not generally protect corporations from being required to tell prospective customers the truth." *Id.* at 721, 734 (emphasis added). Because the "relatively brief disclosures in letters that span one to two full pages of text" were "not disproportionate" to the state's interest in "protect[ing] against consumer confusion," the law was not unduly burdensome. *Id.* (quotation marks omitted); *see also, e.g.*, *Chamber of Com. v. SEC*, 85 F.4th 760, 766, 772 (5th Cir. 2023) (rule compelling stock issuer to explain reasons for share repurchases was not unduly burdensome because the "issuer is free to speak (or not) however and whenever it

wishes apart from a privately crafted explanation of its reasons for repurchasing shares"); *R.J. Reynolds Tobacco Co. v. FDA*, 96 F.4th 863, 867, 886 (5th Cir. 2024) (law compelling warnings on tobacco products was not unduly burdensome because of "limited" infringement of tobacco company's "minimal [expressive] interest"); *CompassCare v. Hochul*, 125 F.4th 49, 67 (2d Cir. 2025) (law compelling workplace disclosures was not unduly burdensome because "[i]t does not impair Plaintiffs' ability to convey their own beliefs, nor does it directly contradict Plaintiffs' core mission and message").

The *Zauderer* test is not a rubber stamp. It anticipates that courts can and will engage in thoughtful consideration of a disclosure law's benefits and burdens, providing the government with enough slack to determine "appropriate information disclosure requirements for business corporations," *Pac. Gas & Elec. Co.*, 475 U.S. at 15 n.12, while reining in the government when the informational benefits of a disclosure requirement fail to justify its expressive burdens, *see Am. Beverage Ass'n*, 916 F.3d at 757.

III. A.B. 2013 is subject to and, on this record, passes the *Zauderer* test.

A. A.B. 2013 compels the disclosure of factual and uncontroversial information about generative AI systems available to the public.

A.B. 2013 requires the developer of a generative AI "system or service" that is "made publicly available to Californians for use" to disclose basic information "regarding the data used by the developer to train the generative [AI] system or

service." Cal. Civ. Code § 3111. The law compels only a "high-level summary of the datasets used in the development of" the system, including the sources, types, function, and estimated quantity of the data in the datasets, the time period during which the data were collected, and simple yes-or-no responses to questions such as whether the datasets contain copyrighted materials or personally identifiable information. *Id.* § 3111(a)(1)–(7), (10). These disclosures consist of purely factual and noncontroversial information about commercially available generative AI products or services—a notion that xAI does not meaningfully contest.[22]

First, A.B. 2013 compels purely factual information because it requires disclosure of "literally true" and accurate information about the data that a developer uses to train its generative AI products or services. *CTIA*, 928 F.3d at 846–47. In requiring developers to provide the public with aggregate information about the "presence and amount of certain ingredients"—*i.e.*, inputs—in their generative AI systems, A.B. 2013 resembles conventional consumer protection laws relating to products and services in the marketplace. *Stolfi*, 153 F.4th at 835 (discussing product labels mandated by the Pure Food and Drugs Act of 1906); *see* Cal. Civ. Code

[22] In its opening brief, xAI cites only one reason why *Zauderer* should not apply to A.B. 2013: "there is no advertising at issue here." Appellant's Br. 32 n.1. But as explained above, that argument is plainly inconsistent with Supreme Court and Ninth Court precedent. *See supra* Section II.B.

§ 3111(a)(1), (a)(3)–(12). In requiring developers to summarize the "intended purpose" of their systems or training processes, or the "functional need" of certain data, A.B. 2013 does no more than "forc[e] a company to explain the reason for its actions"—"a purely factual disclosure." *Chamber of Com.*, 85 F.4th at 769–70 (citation and quotation marks omitted); *see* Cal. Civ. Code § 3111(a)(2), (a)(9), (12). Indeed, xAI itself acknowledges that A.B. 2013's disclosures consist of "statements of fact." Appellant's Br. 18, 26.

Second, A.B. 2013 compels noncontroversial information because it does not require a developer to take "sides in a heated political controversy" or "convey a message fundamentally at odds with its mission." *CTIA*, 928 F.3d at 845, 848. In *X Corp. v. Bonta*, this Court distinguished a prior California law that required "social media companies to report whether and how they believe particular, controversial categories of content should be defined and regulated on their platforms"—which crossed the line—from Florida and Texas platform transparency laws that required companies to disclose "high-level statistics about their moderation efforts" and "information about their content-moderation standards and rule changes"—which did not. 116 F.4th 888, 902–03 (9th Cir. 2024) (citations and quotation marks omitted). Like these latter laws, A.B. 2013 does no more than compel "high-level statistics" and summary information about xAI's generative AI data collection and training "efforts." *Id.* And xAI admits as much. *See* Appellant's Br. 18, 26

(acknowledging that A.B. 2013 does not compel "expressions of value, opinion, or endorsement" (citation omitted)).

Finally, A.B. 2013 compels information about commercial products or services because it applies to generative AI systems "made publicly available to Californians for use." Cal. Civ. Code § 3111. xAI is in the business of offering generative AI systems for public consumption, either directly for use by consumers or embedded in other applications or platforms that consumers use. *See* Appellant's Br. 8–9.[23] In its opening brief, xAI repeatedly acknowledges that its systems are "products" in the "AI marketplace." Appellant's Br. 1, 3, 14, 16, 25, 31, 34.

B. A.B. 2013 is reasonably related to California's interest in helping inform users about the generative AI systems offered to them.

California asserts an interest in providing "consumers with useful information regarding generative AI products." Appellee's Br. 29 (citing ER-148–49). There can be no real doubt that this interest is sufficiently strong. *See, e.g.*, *Zauderer*, 471 U.S. at 651 (recognizing "the State's interest in preventing deception of consumers"); *Edenfield v. Fane*, 507 U.S. 761, 769 (1993) (recognizing "Florida's interest in ensuring the accuracy of commercial information in the marketplace"); *Owen*, 873 F.3d at 721 ("The required disclosures are meant to protect against consumer

[23] *See, e.g.*, *Grok*, Tesla Support, www.tesla.com/support/grok ("You can talk to Grok, your AI companion built by xAI, hands-free in your Tesla vehicle."), www.tesla.com/support/grok.

confusion, and are therefore permissible under *Zauderer* . . . ." (citation omitted)). Nor can there be any doubt that the harm that California seeks to alleviate is "potentially real, not purely hypothetical." *NIFLA*, 585 U.S. at 776 (citation omitted); *see supra* Part I. As this Court has acknowledged, "information monopolies" in the digital public sphere "disserve the public interest." *hiQ Labs, Inc. v. LinkedIn Corp.*, 31 F.4th 1180, 1202 (9th Cir. 2022). In *Stolfi*, this Court held that the state's interest in "reducing [informational] asymmetries, facilitating informed commercial transactions, and improving the efficiency of the pharmaceutical market" justified a requirement requiring pharmaceutical companies to disclose information related to the costs, revenues, and prices of certain prescription drugs. 153 F.4th at 827; *see also S.F. Apartment Ass'n*, 881 F.3d at 1178 (recognizing that the state has an "interest in increasing the fairness of buyout negotiations" between landlords and tenants). Similarly here, California seeks to reduce information asymmetries in the marketplace for generative AI products and services. Appellee's Br. 30–33.

A.B. 2013's disclosures are reasonably related to California's asserted interest because they will provide consumers with helpful information to better understand the generative AI systems on offer in the marketplace. For example, information about the "sources" and "types" of data used to train a generative AI system, Cal. Civ. Code § 3111(a)(1)–(3), can help consumers "assess the trustworthiness of a

model and the appropriateness of a given use of that model."[24] Information about the "time period during which the data . . . were collected," *id.* § 3111(a)(10), can help consumers weigh the timeliness and accuracy of the outputs they receive from a generative AI system. Information about whether a system was trained on unlicensed or personally identifiable data, *id.* § 3111(a)(5)–(7), can help consumers "make more informed choices about where they spend their screen time," *Volokh*, 148 F.4th at 91. Each of these disclosures, though modest, will assist consumers in comparing the various generative AI systems on the market and deciding which to use, if any, and how.

xAI argues that A.B. 2013 "does not require AI companies to disclose information users actually find useful—*e.g.*, how well a model performs real-world tasks." Appellant's Br. 11. But this assertion embraces a far too narrow conception of what is "useful" to the public, and it obfuscates the importance of training datasets, which "are the foundation on which contemporary [generative AI] systems are built" and "are an essential part of understanding" them.[25] In any event, "*Zauderer*'s reasonably related analysis" does not require a compelled disclosure to "be immediately or directly useful for many . . . consumers" or demand "evidentiary

[24] Nat'l Telecomms. & Info. Admin., *supra*, at 27.

[25] Kate Crawford & Trevor Paglen, *Excavating AI*, AI Now Inst., https://excavating.ai; *see also, e.g.*, Br. of *Amici Curiae* Ctr. for AI & Digit. Pol'y et al. 7–17, ECF No. 23.1.

parsing where . . . the government uses a disclosure mandate to achieve a goal of informing consumers about a particular product trait." *Am. Hosp. Ass'n v. Azar*, 983 F.3d 528, 540–41 (D.C. Cir. 2020) (citations and quotation marks omitted). Researchers and experts can often translate commercial disclosures into forms more useful to ordinary consumers, especially in the context of complex and opaque technologies.

C. The preliminary injunction record does not demonstrate that A.B. 2013 imposes an undue burden on xAI.

xAI claims that A.B. 2013 "substantially burden[s] protected speech," Appellant's Br. 34, but it "offers little more than conclusory statements to that effect," *Owen*, 873 F.3d at 734. The company merely asserts that the law "forces xAI to speak a message it does not want to convey." Appellant's Br. 20. Yet as this Court has already observed, "[a]ll mandatory disclosures impose *some* burden," and a disclosure is "unduly burdensome" only when it "effectively rules out" protected speech. *Owen*, 873 F.3d at 721, 734 (emphasis added). xAI has not even attempted to make any such showing.

The only other potential "burden" that xAI asserts relies on equally conclusory and vague allegations that A.B. 2013 might cause it to suffer some competitive or financial injury. Appellant's Br. 24–25. Even if these allegations were well-plead and substantiated, they would not amount to a "burden" in any sense relevant to the *Zauderer* test, which asks whether the disclosures "unduly burden *expressive*

activity." *Moody*, 603 U.S. at 727 n.3 (emphasis added). xAI's asserted financial costs would be relevant only if they were so great that the law's disclosure scheme "essentially operates as a restriction on constitutionally protected speech." *Am. Meat Inst. v. USDA*, 760 F.3d 18, 27 (D.C. Cir. 2014); *see also Azar*, 983 F.3d at 541 (making clear that a petitioner "must demonstrate a burden on speech," rather than a financial one). In *American Meat Institute*, for example, the D.C. Circuit noted that a country-of-origin labeling rule was "of great concern to plaintiffs because of its cost implications," 760 F.3d at 21—the plaintiffs would have to segregate livestock in order to comply with the rule. In its undue-burden analysis, however, the court did not mention these costs and considered only whether there was evidence that the requirement had chilled plaintiffs' protected speech. *Id.* at 27.

Conclusion

For the foregoing reasons, the Court should hold that A.B. 2013 is subject to and—on the preliminary injunction record—passes the *Zauderer* test and affirm the district court's denial of a preliminary injunction with respect to xAI's First Amendment claim.

July 22, 2026

Respectfully submitted,

/s/ Stephany Kim

Stephany Kim
Jake Karr
Xiangnong (George) Wang
Alex Abdo
Knight First Amendment Institute at
Columbia University
475 Riverside Drive, Suite 302
New York, NY 10115
(646) 745-8500
stephany.kim@knightcolumbia.org
jake.karr@knightcolumbia.org

Counsel for Amicus Curiae