

No. 26-1591

**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-JGB-SSC

Hon. Jesus G. Bernal

**BRIEF OF THE AI COALITION FOR DATA INTEGRITY
AS AMICUS CURIAE IN SUPPORT OF DEFENDANT-APPELLEE
AND AFFIRMANCE**

TOD COHEN
ALYSHA AGARWAL
MANATT, PHELPS & PHILLIPS, LLP
One Embarcadero Center, 30th Floor
San Francisco, California 94111
Telephone: (415) 291-7469
tcohen@manatt.com
Counsel for Amicus Curiae
Additional Amici Curiae listed in the
Statement of Interest

CORPORATE DISCLOSURE STATEMENT

Pursuant to Federal Rules of Appellate Procedure 26.1 and 29(a)(4)(A),
amicus curiae The AI Coalition for Data Integrity states that it has no parent
corporation, it issues no stock, and no publicly held corporation has any ownership
interest in it.

TABLE OF CONTENTS

I.	TRAINING-DATA TRANSPARENCY IS WHAT MAKES THE MARKET FOR LICENSED TRAINING DATA WORK.	10
A.	The Legislature identified copyright transparency as a purpose of AB 2013, and the State asserts that interest here.	10
B.	The market for licensing works as AI training data is real and legally cognizable.....	11
C.	A licensing market cannot function without visibility into what was used.	12
D.	Concealment rewards the developer that acquires content from anywhere over the developer that licenses.....	13
II.	REQUIRING A COMPANY TO DISCLOSE WHERE ITS INPUTS CAME FROM IS ORDINARY REGULATION, AND CALIFORNIA HAS REQUIRED IT OF ELECTRICITY SUPPLIERS FOR NEARLY THREE DECADES.....	13
A.	The Power Source Disclosure program requires the same disclosure, in the same manner, for the same reason.	14
B.	In the same law, the Legislature drew the very line xAI says cannot be drawn.	16
C.	Nearly three decades of practice bears on three questions this Court must answer.....	17
III.	AB 2013 SUPPORTS INNOVATION IN GENERATIVE ARTIFICIAL INTELLIGENCE.....	19
A.	Model quality depends on lawfully obtained human work, which comes from the licensing market.	19
B.	Compliance is feasible, and xAI has already complied.	20
C.	A common disclosure floor rewards developers that license and gives both sides certainty.	21
IV.	A RULING FOR xAI WOULD CONVERT OTHER PEOPLE'S PROPERTY INTO xAI's OWN.....	22
A.	AB 2013 asks about provenance, which is a fact about other people's property.	22
B.	The point extends the State's ingredients-and-recipe distinction from composition to provenance.....	23
C.	For licensed datasets, disclosure reveals no secret at all, which defeats xAI's facial theory.....	23

TABLE OF AUTHORITIES

CASES

<i>Am. Beverage Ass'n v. City & County of San Francisco</i> , 916 F.3d 749 (9th Cir. 2019) (en banc)	17, 21
<i>Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith</i> , 598 U.S. 508 (2023)	11
<i>Campbell v. Acuff-Rose Music, Inc.</i> , 510 U.S. 569 (1994)	11
<i>CDK Global, LLC v. Brnovich</i> , 16 F.4th 1266 (9th Cir. 2021)	21
<i>CTIA – The Wireless Ass'n v. City of Berkeley</i> , 928 F.3d 832 (9th Cir. 2019)	18
<i>InteliClear, LLC v. ETC Glob. Holdings, Inc.</i> , 978 F.3d 653 (9th Cir. 2020)	22
<i>Lucas v. South Carolina Coastal Council</i> , 505 U.S. 1003 (1992)	18
<i>MAI Sys. Corp. v. Peak Comput., Inc.</i> , 991 F.2d 511 (9th Cir. 1993)	22
<i>Nat'l Ass'n of Wheat Growers v. Bonta</i> , 85 F.4th 1263 (9th Cir. 2023)	18
<i>Nationwide Biweekly Admin., Inc. v. Owen</i> , 873 F.3d 716 (9th Cir. 2017)	18
<i>Penn Central Transportation Co. v. New York City</i> , 438 U.S. 104 (1978)	18, 19
<i>Pharmaceutical Rsch. & Mfrs. of Am. v. Stolfi</i> , 153 F.4th 795 (9th Cir. 2025)	13, 18, 23
<i>Ruckelshaus v. Monsanto Co.</i> , 467 U.S. 986 (1984)	18
<i>Thomson Reuters Enter. Ctr. GmbH v. Ross Intell.</i> , 765 F. Supp. 3d 382 (D. Del. 2025)	8, 11, 13
<i>Thomson Reuters Enter. Ctr. GmbH v. Ross Intell.</i> , No. 25-2153 (3d Cir. 2026)	5

<i>Walker v. Univ. Books, Inc.</i> , 602 F.2d 859 (9th Cir. 1979)	22
<i>Winter v. Natural Res. Def. Council, Inc.</i> , 555 U.S. 7 (2008).....	19
<i>Zauderer v. Office of Disciplinary Counsel</i> , 471 U.S. 626 (1985).....	17

STATUTES

18 U.S.C. § 1839(3)	22
Cal. Civ. Code § 3111(a)(1), (a)(5), (a)(7), (a)(9)	7, 10, 12, 15, 17, 18, 22, 24
Cal. Civ. Code § 3111(b)	15
Cal. Civ. Code § 3426.1(d)	16, 22, 24
Cal. Pub. Util. Code §§ 330-400.3	14
Cal. Pub. Util. Code § 398.1	8, 14
Cal. Pub. Util. Code § 398.4	8, 15, 18
Cal. Pub. Util. Code § 398.5	9, 16
Cal. Pub. Util. Code §§ 398.1-398.5	14

OTHER AUTHORITIES

U.S. Copyright Office, Copyright and Artificial Intelligence, Part 3, Generative AI Training	8, 11
---	-------

RULES

Fed. R. App. P. 29(a)	7
-----------------------------	---

STATEMENT OF INTEREST OF AMICUS CURIAE

The AI Coalition for Data Integrity (“AICDI” or “the Coalition”) is a multi-stakeholder cross-industry organization dedicated to promoting legal and ethical artificial intelligence (“AI”) development by advocating for transparency, proper attribution, and fair licensing in the use of data for AI. AICDI’s membership includes AI developers, technology companies, content creators, academic researchers, trade associations, publishers, aggregators, and public-interest groups, all united by the conviction that a robust, trustworthy AI ecosystem depends on the integrity and responsible stewardship of data. The Coalition works for transparency, attribution, and fair licensing in the data used to build AI systems.

The Coalition's members sit on both sides of the market for licensing works as AI training data. Some license their work to developers. Some are developers who pay for that data. Both depend on the same thing: the ability to know what a model was trained on. A rights holder cannot negotiate over a use it cannot see. A developer cannot demonstrate that it acquired its data lawfully if nothing in the market requires anyone to say where their data came from.

The Coalition filed as amicus curiae in *Thomson Reuters Enterprise Centre GmbH v. Ross Intelligence Inc.*, No. 25-2153 (3d Cir. 2026), addressing the market for licensing works as AI training data. The Coalition appears here because this appeal presents the same question from the other direction. In *Ross*, a developer

argued that copying was permissible because it left no trace in the model's output. Here, a developer argues that it may keep the sources of its training data secret. Both positions leave rights holders and content licensors unable to learn that their works were used.

The following organizations have signed on to this brief in support of the Coalition's position:

- John Wiley & Sons, Inc.
- OpenOrigins, Inc.
- National Association of Voice Actors
- Digital Media Leadership Alliance
- Music Artists Coalition
- Transparency Coalition
- The Association of Medical Illustrators
- American Society for Collective Rights Licensing, Inc.
- Open Markets Institute

These organizations join as amici to underscore the importance of transparency, integrity, and the protection of licensing markets in AI data practices. Their participation should not be construed as agreement with every statement or position expressed in this brief. Each organization reserves its own views on specific issues and supports this filing solely to advance the broader principles outlined herein. This Court should affirm.

RULE 29(a)(4)(E) STATEMENT

No party's counsel authored this brief in whole or in part. No party, no party's counsel, and no person other than amicus curiae, its members, and its counsel contributed money intended to fund preparing or submitting this brief. All parties have consented to its filing.

SUMMARY OF ARGUMENT

California Assembly Bill 2013 requires a developer of a generative AI system to post a high-level summary of the datasets it used for training. Among the twelve enumerated items for training dataset disclosure are the sources or owners of the underlying datasets and whether the datasets include material protected by copyright, trademark, or patent. *See* Cal. Civ. Code § 3111(a)(1), (a)(5). xAI says the Constitution forbids California from requiring such disclosures.

The Legislature enacted AB 2013 for several reasons. One of them was copyright. The Senate committee report explains that transparency about training data “shines the light on various other issues, such as privacy and copyright concerns.” ER-146. The State asserts that interest in this Court, identifying among AB 2013's public purposes “making users and copyright owners aware of whether models were trained on copyrighted material.” Answering Br. 52. Amicus writes to develop that interest, which its members hold, and to make four points.

First, the disclosures AB 2013 requires are what allow a market in licensed training data to operate. Courts have recognized that market. *Thomson Reuters Enter. Ctr. GmbH v. Ross Intell.*, 765 F. Supp. 3d 382, 398, 400 (D. Del. 2025), *appeal pending*, No. 25-2153 (3d Cir. argued June 11, 2026). The Copyright Office has documented it. U.S. Copyright Office, Copyright and Artificial Intelligence, Part 3, Generative AI Training 47-48, 55-57 (Pre-Publication Version, May 2025). A market in which the buyer need never disclose what it took is not a market. It rewards the developer who acquires data from anywhere and penalizes the developer who wishes to use only licensed or legally acquired data.

Second, requiring a company to disclose where its inputs came from is an ordinary exercise of legislative power, and California has required exactly that of other industries for nearly three decades. In 1997, the Legislature created such a regime for the energy industry. Since then, every retail electricity supplier selling power consumed in California has published a label identifying its fuel mix by named category, including coal, natural gas, nuclear, large hydroelectric, solar, and wind. *See* Cal. Pub. Util. Code § 398.4(h). The label goes on the supplier's website. *See* § 398.4(b). The Legislature adopted it because consumers need “reliable, accurate, and timely information regarding fuel sources.” § 398.1(a). No supplier has claimed a constitutional right to conceal whether it burns coal.

That statute also answers xAI's central trade secret premise for non-disclosure. In the same portion of the Civil Code, the Legislature exempted granular data suppliers from reporting to the Energy Commission “a trade secret as defined in subdivision (d) of Section 3426.1 of the Civil Code.” § 398.5(b). That is the same definition xAI invokes here. Answering Br. 37. California has run a public source-disclosure requirement and a trade-secret protection side by side, under one definition, for nearly three decades. The Legislature knows how to shield a trade secret. It has never been thought that identifying an input's source is one.

Third, AB 2013 supports the development of generative AI. Models are only as good as the data behind them, and a model trained on only machine-generated output degrades. *See* ER-130. The corrective is lawfully obtained, human-created work, that comes from licensing markets and legal acquisitions. A healthy licensing market runs on provenance. Disclosure lowers its transaction costs and gives buyers and sellers alike a shared factual baseline. Compliance is also demonstrably light: xAI published its AB 2013 disclosures in December 2025 and maintains that they revealed none of its trade secrets. *See* ER-166-76; Opening Br. 13.

Fourth, xAI's trade-secret theory would permit and encourage an AI developer to take copyrighted, confidential, and personal material belonging to

other people and entities, and then hold the fact of that taking as its own concealed asset. The State draws the right line with Coca-Cola: the recipe may be secret; the ingredients on the bottle are not. Answering Br. 40. The sources and ownership of training data are facts about *other people's* property. An AI developer's own secrets lie in its architecture, its weights, and its methods. AB 2013 asks for none of those.

ARGUMENT

I. TRAINING-DATA TRANSPARENCY IS WHAT MAKES THE MARKET FOR LICENSED TRAINING DATA WORK.

A. The Legislature identified copyright transparency as a purpose of AB 2013, and the State asserts that interest here.

The record forecloses any suggestion that copyright transparency is an interest invented for litigation. The Senate committee report states that requiring transparency about training data “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.” ER-146. The Assembly committee report explains that AB 2013 imposes modest requirements “in exchange for providing Californians with a fuller understanding of the state's AI information ecosystem.” ER-131. The statute carries that purpose into its text: a developer must disclose whether its datasets include material protected by copyright, trademark, or patent, *see* § 3111(a)(5), and must identify the sources or owners of those datasets, *see* § 3111(a)(1).

The State asserts the interest in this Court. Answering Br. 52. Amicus's members are among the people that interest protects.

B. The market for licensing works as AI training data is real and legally cognizable.

Copyright law protects markets a rights holder “would in general develop or license others to develop.” *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 592 (1994); *see also Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith*, 598 U.S. 508 (2023) (evaluating the effect of a use on the copyright holder’s licensing markets). Applying that principle to AI training, the district court in *Ross* recognized a cognizable market for licensing works as training data and found that the defendant's copying harmed it. 765 F. Supp. 3d at 398, 400. The Copyright Office reached the same conclusion after extensive study, describing licensing for AI training as a functioning and growing market. Copyright and Artificial Intelligence, Part 3, *supra*, at 47-48.

Amicus's members transact in that market. Publishers license archives. Image libraries license collections. Performers' representatives negotiate over voice and likeness data. Technology companies license their confidential business information and third-party content to AI developers. And independent creators license their works directly within all of the above markets. Developers obtain the rights and pay for all of it. These are commercial dealings in a defined product.

C. A licensing market cannot function without visibility into what was used.

Three facts are indispensable to any license negotiation, and a rights holder or licensor can obtain none of them by only inspecting a model's output. The first is that the work was used at all. The second is the identity of the user of their content and data. The third is enough about scope to set a price. AB 2013 supplies the beginning of each. Section 3111(a)(1) identifies sources and owners of the datasets, which is how a rights holder learns of the use and the identity of the user of their content. *See* § 3111(a)(1). Section 3111(a)(5) and (a)(7) state whether copyrighted material and personal information are present, which frames the legal status of what was taken and the scope of any license to be negotiated. *See* § 3111(a)(5), (a)(7).

The State's brief establishes why output inspection is no substitute. A model's deficiencies are difficult to identify from its outputs alone, and it is unreasonable to expect a consumer to fact-check a model or compare outputs across models. Answering Br. 33. What is true for a consumer assessing quality is true for a rights holder or licensor assessing use. A novel does not appear in a model's weights with a title page attached. Absent disclosure, the author may never learn that her book was ingested, and the developer that ingested it faces no one across a negotiating table.

D. Concealment rewards the developer that acquires content from anywhere over the developer that licenses.

The consequence is not hypothetical. In *Ross*, the developer sought a license from Thomson Reuters, was refused because it was a competitor, and obtained the protected content through an intermediary instead. 765 F. Supp. 3d at 386-88. That sequence became known only through litigation and discovery. A disclosure rule surfaces it at the front end.

Two developers building comparable models face a choice. One negotiates and obtains licenses and pays. The other takes what it wants without asking. Without any obligation to say where the data came from, the second developer's cost advantage is invisible, and its conduct is unreviewable by the people it affects. Amicus's members that choose to license are undercut by a rule of silence. California's interest in "reducing those asymmetries, facilitating informed commercial transactions, and improving the efficiency of the [generative AI] market," *Pharmaceutical Rsch. & Mfrs. of Am. v. Stolfi*, 153 F.4th 795, 826 (9th Cir. 2025), reaches this market too.

II. REQUIRING A COMPANY TO DISCLOSE WHERE ITS INPUTS CAME FROM IS ORDINARY REGULATION, AND CALIFORNIA HAS REQUIRED IT OF ELECTRICITY SUPPLIERS FOR NEARLY THREE DECADES.

The State describes AB 2013 as a law compelling "disclosure of the inputs into a product," comparable to laws requiring disclosure of ingredients in food or

chemicals in consumer products. Answering Br. 19. The comparison is apt and it does not stop at food and chemicals. California already requires an entire industry to publish the provenance of its inputs, in the same manner, for the same reason, in more detail than AB 2013 demands.

A. The Power Source Disclosure program requires the same disclosure, in the same manner, for the same reason.

California Senate Bill 1305 (Sher, ch. 796, 1997 Cal. Stat.) established the Power Source Disclosure program, codified at Cal. Pub. Util. Code §§ 398.1-398.5. Article 14 sits in Chapter 2.3 of the Public Utilities Act, titled “Electrical Restructuring.” Cal. Pub. Util. Code §§ 330-400.3. The Legislature enacted it as California opened the retail electricity market to competition, so that consumers could compare competing suppliers on the merits of what those suppliers actually sold. It found “a need for reliable, accurate, and timely information regarding fuel sources for electric generation offered for retail sale in California,” § 398.1(a), and declared the program's purpose to be disclosure of “accurate, reliable, and simple to understand information on the sources of energy that are used to provide electric services,” § 398.1(b).

AB 2013 rests on the same premise. Its author explained that to build consumer confidence in AI, “we need to start with the foundations, and for AI that is the selection of training data,” and that the required disclosures would let a consumer “compare competing systems and services.” ER-131. Both statutes

answer the same problem: a buyer choosing among competing sellers cannot evaluate a product whose inputs are concealed.

The mechanics track AB 2013 closely. Every retail supplier offering electricity for consumption in California must disclose its electricity sources. § 398.4(a). The disclosure goes to potential end-use consumers in product-specific materials, “including the retail supplier's internet website.” *Compare* § 398.4(b) *with* § 3111 (an AI developer must post documentation on its website). The electricity supplier must express its sources as a percentage of annual sales and identify them within enumerated categories named in the statute: coal, large hydroelectric, natural gas, nuclear, and eligible renewables including biomass and biowaste, geothermal, eligible hydroelectric, solar, and wind, along with any further categories the Energy Commission designates. *Compare* § 398.4(g), (h) *with* § 3111(a) (enumerating twelve required training-data disclosure categories). Where an electricity supplier cannot trace its power to a generator, it discloses the share drawn from “unspecified sources.” § 398.4(g)(1)(A). A supplier that does not know its source says that it does not know. And the program reaches retail suppliers generally, exempting only on-site generation and over-the-fence transactions, *see* § 398.4(l), much as AB 2013 exempts categories of AI systems consumers are unlikely to encounter, § 3111(b).

The electricity program has operated for more than a quarter century. The resulting Power Content Label is a familiar document. No retail supplier has claimed that identifying its fuel mix compels it to speak in violation of the First Amendment, or that publishing the percentage of its portfolio drawn from coal is a taking of its property.

B. In the same law, the Legislature drew the very line xAI says cannot be drawn.

Section 398.5 completes the design and refutes xAI's takings claim premise. Suppliers that disclose specific purchases must report to the Energy Commission the kilowatt-hours they purchased “by generator and fuel type,” consistent with metered data. § 398.5(a)(1). That granular, counterparty-level information is commercially sensitive, and the Legislature protected it: information submitted under that section “that is a trade secret as defined in subdivision (d) of Section 3426.1 of the Civil Code shall not be released except in an aggregated form such that trade secrets cannot be discerned.” § 398.5(b).

Section 3426.1(d) of the Civil Code is the same definition on which xAI's takings claim depends, and the same definition the State applies. *See* Cal. Civ. Code § 3426.1(d); Answering Br. 37. So, California has operated, within a single article, a public disclosure of input sources and a trade-secret shield for granular input data, under one statutory definition, for nearly three decades. The Legislature

knows how to protect a trade secret in this setting. It protected the generator-level contract data. It published the fuel mix.

The line tracks the familiar one between a recipe and an ingredient list. A supplier's portfolio-level fuel mix corresponds to the ingredient list. Its generator-by-generator purchasing, and how it assembles a portfolio, corresponds to the recipe. AB 2013 sits on the ingredient side of that line. It requires a high-level summary of dataset sources and characteristics. § 3111(a). It does not require a developer to produce its datasets, its weights, its architecture, or its methods, and it does not require it to explain what any processing entailed or how it was performed. *See* § 3111(a)(9) (requiring disclosure only of whether datasets were cleaned, processed, or modified, and the intended purpose of those efforts, not the methods used).

C. Nearly three decades of practice bears on three questions this Court must answer.

The Power Source Disclosure program is legislative practice rather than judicial holding, and amicus offers it as context. It bears on three issues before this Court.

Uncontroversial. Under *Zauderer v. Office of Disciplinary Counsel*, 471 U.S. 626, 651 (1985), compelled commercial disclosure must be purely factual and uncontroversial. *Am. Beverage Ass'n v. City & County of San Francisco*, 916 F.3d 749, 756 (9th Cir. 2019) (en banc). A disclosure is controversial when it forces a

speaker into a heated political dispute or an unresolved scientific debate. *CTIA – The Wireless Ass'n v. City of Berkeley*, 928 F.3d 832, 845 (9th Cir. 2019); *Nat'l Ass'n of Wheat Growers v. Bonta*, 85 F.4th 1263, 1278 (9th Cir. 2023). Identifying the origin of an input is neither. California consumers have read a label reporting fuel origin for a generation. Saying that a dataset came from a licensed archive, or from a web crawl, or from a source the developer cannot specify, is a statement of the same order as saying that a kilowatt-hour came from natural gas.

Not unduly burdensome. A disclosure is unduly burdensome when it effectively rules out the speech it accompanies. *Nationwide Biweekly Admin., Inc. v. Owen*, 873 F.3d 716, 734 (9th Cir. 2017). Section 398.4 requires percentage breakdowns across a named list of fuels, computed annually across an entire portfolio, and posted publicly. *See* § 398.4. Section 3111(a) requires a high-level summary. *See* § 3111(a). AB 2013 asks less than a regime California has administered without constitutional objection for nearly three decades.

Investment-backed expectations. An alleged taking of trade secrets is analyzed as a regulatory taking under *Penn Central Transportation Co. v. New York City*, 438 U.S. 104, 124 (1978). *Stolfi*, 153 F.4th at 833. Under the first *Penn Central* factor, expectations are “necessarily tempered” in fields long subject to public concern and government regulation. *Stolfi*, 153 F.4th at 834 (quoting *Ruckelshaus v. Monsanto Co.*, 467 U.S. 986, 1007 (1984)); *see also Lucas v. South*

Carolina Coastal Council, 505 U.S. 1003, 1027-28 (1992). xAI says it invested on the understanding that its dataset information would stay confidential. Answering Br. 47. But California has required firms in another input-driven industry to publish the provenance of their inputs since before xAI existed, and the Legislature said in AB 2013's own history that copyright was among its concerns. ER-146. No firm that builds a product out of material obtained from other people can reasonably expect to be the single industry excused from saying where that material came from.

III. AB 2013 SUPPORTS INNOVATION IN GENERATIVE ARTIFICIAL INTELLIGENCE.

Whether a law benefits an industry is not a constitutional test. It bears on the substantiality of the State's interest, on tailoring and burden, on the reasonableness of a regulated party's expectations under *Penn Central*, and on the balance of equities and the public interest that govern injunctive relief. *Winter v. Natural Res. Def. Council, Inc.*, 555 U.S. 7, 20, 24 (2008). On each of those, the record shows that AB 2013's disclosures help the enterprise xAI says they threaten.

A. Model quality depends on lawfully obtained human work, which comes from the licensing market.

A generative model draws on patterns in its training data, and is “only as accurate, reliable, or unbiased as the training data from which [it is] built.”

Answering Br. 5 (quoting record sources). Feed a model machine-generated text

and it degrades: training on AI output produces progressively worse output, and the legislative record reflects that synthetic data can cause a model to break down over time. ER-130.

The corrective is human-created work of known origin, acquired lawfully. That supply comes from the market amicus's members build. Its raw material is the archive, the photograph, the recording, the performance. And that market runs on provenance. A licensor cannot price a use it cannot observe. A licensee cannot demonstrate clean title without records of what it acquired and from whom. Enterprise purchasers increasingly demand assurances about the origin of the data behind the models they deploy, and a developer that has kept no account of its sources cannot give them.

Amicus is positioned to say this because its membership spans the transaction. The State can describe the degradation problem. A coalition of developers and creators can connect it to the market that answers it. Transparency about training data is the infrastructure that market stands on.

B. Compliance is feasible, and xAI has already complied.

The best evidence of AB 2013's burden is in the record. xAI published disclosures under the statute in December 2025, ER-166-76, and maintains in this Court that those disclosures reveal none of its trade secrets. Opening Br. 13. The law took effect on January 1, 2026, and developers have been disclosing under it

since. AB 2013 requires a single public posting. It does not interfere with a developer's advertising or drown out its own messaging. *Cf. Am. Beverage Ass'n*, 916 F.3d at 757.

A developer that has complied, that says compliance cost it no secrets, and that faces no enforcement action is not a developer whose ability to innovate hangs on an injunction.

C. A common disclosure floor rewards developers that license and gives both sides certainty.

Concealment lets the developer that scraped price its product below the developer that licensed, while hiding the reason for the gap. A shared baseline of provenance corrects that. It lets licensors find counterparties, lets licensees prove diligence, and lets courts and buyers distinguish between them. This Court has recognized that a legislature may adjust economic relations to advance “consumer privacy and competition” without effecting a taking. *CDK Global, LLC v. Brnovich*, 16 F.4th 1266, 1283 (9th Cir. 2021).

The equities follow. Enjoining AB 2013 would deny consumers information the Legislature judged they need and would also remove the foundation of a lawful market for training data that is still being built. Amicus's members would bear that loss directly.

IV. A RULING FOR xAI WOULD CONVERT OTHER PEOPLE'S PROPERTY INTO xAI's OWN.

A. AB 2013 asks about provenance, which is a fact about other people's property.

A trade secret is information that derives value from not being generally known and that its owner has taken reasonable steps to keep secret. Cal. Civ. Code § 3426.1(d); 18 U.S.C. § 1839(3). Information publicly available or widely known in an industry is not a trade secret. *Walker v. Univ. Books, Inc.*, 602 F.2d 859, 865 (9th Cir. 1979). The party asserting a trade secret must identify it with particularity and bears the burden of establishing that it exists. *MAI Sys. Corp. v. Peak Comput., Inc.*, 991 F.2d 511, 522 (9th Cir. 1993); *InteliClear, LLC v. ETC Glob. Holdings, Inc.*, 978 F.3d 653, 658-59 (9th Cir. 2020). The district court found that xAI had not done so, resting instead on “generalizations and hypotheticals.” ER-7-8.

Sections 3111(a)(1) and (a)(5) ask who owned the data and what its legal status is. Those are facts about property belonging to someone else. An author's ownership of her novel is not xAI's confidential business information. A photographer's copyright in his archive is not xAI's competitive edge. xAI's own secrets lie elsewhere, in its architecture, its weights, its training methods, and its judgments about how to assemble and process a corpus. AB 2013 reaches none of these, and the statute by its terms does not require a developer to disclose what any processing entailed or how it was performed. *See* § 3111(a)(9).

The label a party attaches to information does not determine whether it is a secret. Were it otherwise, a developer could take copyrighted works, declare the fact of the taking proprietary, and place the taking beyond the reach of the people whose works were taken. Trade-secret law has never worked that way: secrecy protects what a company creates, not what it conceals about what it took. That is the rule xAI asks this Court to announce.

B. The point extends the State's ingredients-and-recipe distinction from composition to provenance.

The State's Coca-Cola comparison answers what is in the product.

Answering Br. 40. Amicus's point answers where it came from and who owned it. Even a beverage maker with a secret formula holds no secret, as against its own suppliers, in the fact that it buys sugar from them. Its suppliers know. So does every rights holder whose work an AI developer has licensed.

C. For licensed datasets, disclosure reveals no secret at all, which defeats xAI's facial theory.

xAI mounts a pre-enforcement challenge, so it must show that AB 2013 operates unlawfully in every application. *Stolfi*, 153 F.4th at 834, 839. It cannot, and amicus's members illustrate why.

A developer that licenses a corpus from a publisher has a counterparty. That counterparty knows the agreement exists, knows what it conveyed, and is free to say so. Disclosing that a model was trained on licensed material from an identified

category of sources tells that counterparty nothing it does not already know, and it reveals no information that derives value from being unknown. Cal. Civ. Code § 3426.1(d). For every dataset a developer acquired by agreement, and for every dataset in the public domain, § 3111(a)(1) and (a)(5) compel a fact that is already known to someone with every right to know it. A statute that requires disclosure of such facts is not a taking in every application, and xAI's claim fails on its own chosen terms.

The developers that suffer no injury from these disclosures are the ones that obtained their data lawfully. That is the point.

CONCLUSION

The district court's order should be affirmed.

Dated: July 22, 2026

Respectfully submitted,

/s/ Tod Cohen

TOD COHEN

ALYSHA AGARWAL

MANATT, PHELPS & PHILLIPS, LLP

Counsel for Amicus Curiae

AI Coalition for Data Integrity

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

Form 8. Certificate of Compliance for Briefs

Instructions for this form: <http://www.ca9.uscourts.gov/forms/form08instructions.pdf>

9th Cir. Case Number(s) No. 26 1591

I am the attorney or self-represented party.

This brief contains 4,365 **words, including** 0 **words**

manually counted in any visual images, and excluding the items exempted by FRAP 32(f). The brief's type size and typeface comply with FRAP 32(a)(5) and (6).

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

- ☐ complies with the word limit of Cir. R. 32-1.
- ☐ is a **cross-appeal** brief and complies with the word limit of Cir. R. 28.1-1.
- ☒ is an **amicus** brief and complies with the word limit of FRAP 29(a)(5), Cir. R. 29-2(c)(2), or Cir. R. 29-2(c)(3).
- ☐ is for a **death penalty** case and complies with the word limit of Cir. R. 32-4.
- ☐ complies with the longer length limit permitted by Cir. R. 32-2(b) because (*select only one*):
- ☐ it is a joint brief submitted by separately represented parties.
- ☐ a party or parties are filing a single brief in response to multiple briefs.
- ☐ a party or parties are filing a single brief in response to a longer joint brief.
- ☐ complies with the length limit designated by court order dated .
- ☐ is accompanied by a motion to file a longer brief pursuant to Cir. R. 32-2(a).

Signature s/Tod Cohen

Date 7/22/2026

(use "s/[typed name]" to sign electronically-filed documents)

Feedback or questions about this form? Email us at forms@ca9.uscourts.gov