EXHIBIT A

UNITED STATES DISTRICT COURT
FOR THE SOUTHERN DISTRICT OF NEW YORK

IN RE: OPENAI, INC., COPYRIGHT INFRINGEMENT LITIGATION This Document Relates To: All Matters	25-md-03143 (SHS) (OTW) Hon. Sidney H. Stein Hon. Ona T. Wang

[PROPOSED] BRIEF OF *AMICUS CURIAE* JOSHUA S. GANS ON THE FAIR USE ISSUES PRESENTED BY THE PARTIES' MOTIONS FOR SUMMARY JUDGMENT

Andrew J. Pincus
MAYER BROWN LLP
1999 K Street, N.W.
Washington, D.C. 20006
(202) 263-3000
apincus@mayerbrown.com

Graham White
MAYER BROWN LLP
1221 Avenue of the Americas
New York, NY 10020
(212) 506-2500
gwhite@mayerbrown.com

Counsel for Amicus Curiae

TABLE OF CONTENTS

Page

INTEREST OF AMICUS CURIAE 1
INTRODUCTION AND SUMMARY OF ARGUMENT 3
ARGUMENT 5
I. A REDUCTION IN SALES DOES NOT BY ITSELF ESTABLISH COPYRIGHT HARM 5
A. Learning And Commentary Can Reduce Demand For An Original Work 5
B. Copyright Policy Affects Investment In Both Creation And Subsequent Uses 7
C. Scale Makes The Causal Question More Important, Not The Harm More Certain 9
II. DETERMINING WHETHER AND HOW THE CHALLENGED USE DISPLACES THE ORIGINAL WORK IS CRITICAL TO ASSESSING MARKET IMPACT 11
A. The Market Effects Assessment Examines How The Challenged Use Draws On A Particular Work 11
B. The Amount of Training Material, Standing Alone, Is Not Dispositive of Market Effect 13
III. THE LATER EMERGENCE OF LICENSING MARKETS DOES NOT BY ITSELF ESTABLISH COGNIZABLE HARM 14
A. A Licensing Market Visible Today Does Not Establish That A Workable Market Existed When The Training Occurred 14
B. Later Agreements Require A Comparison With Earlier Contracting Conditions 15
C. AI Licenses Cover Distinct Rights And Services 17
D. Licensing Can Be Easier For Uses That Draw On Identifiable Content 18
CONCLUSION 20
CERTIFICATE OF COMPLIANCE 21

TABLE OF AUTHORITIES

Page(s)

Cases

Am. Geophysical Union v. Texaco Inc.,
60 F.3d 913 (2d Cir. 1994).....4, 16

Andy Warhol Foundation for the Visual Arts, Inc. v. Goldsmith,
598 U.S. 508 (2023).....7

Authors Guild v. Google, Inc.,
804 F.3d 202 (2d Cir. 2015).....5, 7

Bartz v. Anthropic PBC,
787 F. Supp. 3d 1007 (N.D. Cal. 2025).....11

Bill Graham Archives v. Dorling Kindersley Ltd.,
448 F.3d 605 (2d Cir. 2006).....16

Campbell v. Acuff-Rose Music, Inc.,
510 U.S. 569 (1994).....7, 10

Kadrey v. Meta Platforms, Inc.,
788 F. Supp. 3d 1026 (N.D. Cal. 2025).....10, 11

Nihon Keizai Shimbun, Inc. v. Comline Bus. Data, Inc.,
166 F.3d 65 (2d Cir. 1999).....6

NXIVM Corp. v. Ross Institute,
364 F.3d 471 (2d Cir. 2004).....11

Silverman v. OpenAI Inc.,
No. 1:25-cv-03483-SHS-OTW (S.D.N.Y. Sep. 1, 2026).....19

Twin Peaks Prods., Inc. v. Publ'ns Int'l, Ltd.,
996 F.2d 1366 (2d Cir. 1993).....6

Statutes

17 U.S.C. § 102(b).....5

17 U.S.C. § 107.....5, 9

17 U.S.C. § 107(4).....3

Other Authorities

Annie Liang & Jay Lu, *Creative Ownership in the Age of AI* (July 29, 2026).....12

Associated Press, AP, Open AI Agree to Share Select News Content and Technology in New Collaboration (July 13, 2023) 17

Joshua S. Gans, *Copyright Policy Options for Generative Artificial Intelligence*, 69 J.L. & Econ. 1 (2026)..... *passim*

Joshua S. Gans, *Copyright Policy Options for Generative Artificial Intelligence*, 69 J.L. & Econ. 1, 1-2 (2026) 8, 9

Kevin Kelly, *Arguments in Favor of AI Fair Use*, The Technium (Sept. 7, 2026) 6

Martin Beraja & Noam Yuchtman, *Growth at the Resource Frontier* 1–4 (July 4, 2026) 15

U.S. Copyright Office, Off. of the Chief Economist, *Identifying the Economic Implications of Artificial Intelligence for Copyright Policy* 36 (Feb. 2025) 17

INTEREST OF *AMICUS CURIAE*

Joshua S. Gans is a Professor of Strategic Management and holds the Jeffrey S. Skoll Chair in Technical Innovation and Entrepreneurship at the University of Toronto's Rotman School of Management. He is also a Research Associate of the National Bureau of Economic Research. His research examines innovation, competition, intellectual property, and the economics of artificial intelligence (AI). His books on AI include *Prediction Machines* (Harvard Business Review Press, 2018), *Power and Prediction* (Harvard Business Review Press, 2022), and *The Microeconomics of Artificial Intelligence* (MIT Press, 2025). Further research is available at www.joshuagans.com. As an author and teacher, he has an interest in both the incentives to create original work and preserving the opportunities to learn from such works.

His 2026 article, *Copyright Policy Options for Generative Artificial Intelligence*,[1] examines how copyright policy affects two investments: creators' investment in original content and developers' investment in AI models. The article considers a setting in which transaction costs prevent advance licensing. Whether that assumption applies to a particular training process requires evidence. Under the model's conditions, permitting training while compensating creators who later suffer serious harm can serve both investments, subject to financial feasibility and implementation costs.

This brief draws on that analysis and related economic research. That research addresses when a particular work makes an AI output possible, and also addresses how markets for licensing a new kind of use develop after a technology reveals the use's value. The brief provides an economic perspective on the fair-use questions raised by the summary-judgment motions in the

[1] Joshua S. Gans, *Copyright Policy Options for Generative Artificial Intelligence*, 69 J.L. & Econ. 1 (2026).

news cases. It supplies no independent measurements of defendants' systems or plaintiffs' losses and does not address the other elements of the parties' claims.

INTRODUCTION AND SUMMARY OF ARGUMENT

A good economics lecture can save a student the price of a textbook. A review can convey a book's central ideas or persuade a reader that the book is not worth buying. In each case, the professor or reviewer learned from an original work and used that learning in a way that may reduce the author's sales. Such uses can qualify as fair use under the statutory factors. The issue fair use addresses is how to preserve incentives to create original works while allowing others to learn from what has been created.

Generative AI presents a new context for this familiar problem. The parties' summary-judgment motions turn in significant part on the fourth fair-use factor—the effect of the challenged use "upon the potential market for or value of the copyrighted work," 17 U.S.C. § 107(4)—and on the related question of how to identify the markets for the asserted works that are relevant to the fair use analysis. The News Plaintiffs' case for cognizable harm to the asserted works rests on three asserted harms. They contend that Defendants' products can reduce demand for news; that Defendants used the asserted articles to train their models; and that a market for AI training licenses has emerged since Defendants engaged in the training challenged in these actions. This *amicus* brief provides an economic analysis relevant to assessing those claims under the statutory factors.

The News Plaintiffs' first asserted harm—a reduction in the plaintiffs' sales—does not by itself demonstrate cognizable harm. A use that conveys only the facts and ideas contained in a work, as a lecture or a review may, can cost the author a sale without supplying a substitute for the author's protected expression, a distinction relevant to the fourth factor. Copyright supports not only the incentive to create content but also the incentive to build tools that learn from that content, and a rule that treated every displaced sale as actionable can weaken the second without reliably protecting the first.

The second asserted harm—that Defendants used the asserted articles to train their models—also requires interpretation. Training data can be valuable in the aggregate even though no single work is essential to any particular output. The fact that a model was trained on a work therefore does not show that the copyrighted work made a given output possible or produced a given harm to the copyright owner. Assessing market effects requires an account of how using the works produces the asserted displacement.

The third harm—that AI-training licensing agreements now exist—also requires interpretation before it can be used in assessing Defendants' earlier conduct. What is needed is a comparison of the current situation to the conditions that prevailed when that conduct occurred. A new technology can create both a new use for content and also a new market that later forms around that use. But the existence of a market today does not show that a comparable market was "traditional, reasonable, or likely to be developed" at the time of the training. *Am. Geophysical Union v. Texaco Inc.*, 60 F.3d 913, 930 (2d Cir. 1994). Later agreements can be relevant, but their weight turns on that comparison, not on their mere existence.

None of these observations, by itself, supplies a general answer to the fair-use question. The relevance of a decline in sales, the inclusion of the articles as training data, and the later emergence of licensing deals depends on the economic mechanism behind those effects. The Court therefore must assess the economic mechanism producing those consequences in order to determine their relevance. Such an approach will support fair use protection for general-purpose learning where the statutory factors warrant it, while triggering greater scrutiny when conduct is shown to produce substitutes for protected expression or harm to cognizable licensing markets.

ARGUMENT

I. A REDUCTION IN SALES DOES NOT BY ITSELF ESTABLISH COPYRIGHT HARM

A. Learning And Commentary Can Reduce Demand For An Original Work

The Copyright Act recognizes criticism, comment, teaching, scholarship, and research as potential fair uses, subject to its four-factor inquiry, 17 U.S.C. § 107, and it leaves ideas and methods outside copyright protection altogether, 17 U.S.C. § 102(b). Thus, a lecturer can explain an economic idea in the lecturer's own words without reproducing a textbook's protected expression. A reviewer who quotes a passage raises a different question, which turns on the amount used, the purpose served, and the other statutory considerations. Either activity may cost the publisher sales, but lost sales alone do not establish infringement.

The Second Circuit recognized this distinction in *Authors Guild v. Google, Inc.*, 804 F.3d 202 (2d Cir. 2015). The court considered a copyright claim against Google arising from Google Books, in which Google scanned tens of millions of library books, made their full text searchable, and displayed short "snippets" of text in response to user queries. *Id.* at 207-210. The court held that the search and snippet functions were a fair use and did not infringe. *Id.* at 225.

In assessing the fourth fair use factor, the court acknowledged that a snippet could satisfy a searcher's need and cost the author a sale, but held that such a loss does not establish market harm, because it "will generally occur in relation to interests that are not protected by the copyright," such as a searcher's need for a historical fact. *Id.* at 224 ("[T]he possibility, or even the probability or certainty, of some loss of sales does not suffice to make the copy an effectively competing substitute that would tilt the weighty fourth factor in favor of the rights holder in the original."). A reader seeking a fact may otherwise have bought the book to obtain it, but that possibility does not give the author exclusive control over supplying the fact.

A summary of a work could fall on the other side of that line if it supplies a substitute for the original work's protected expression and not just the facts contained within the work. In *Twin Peaks Prods., Inc. v. Publ'ns Int'l, Ltd.*, 996 F.2d 1366 (2d Cir. 1993), the court considered a claim against the publisher of a book about a television series, where a substantial part of the book recounted the plots of the first eight episodes in elaborate detail and copied dialogue from those episodes. *Id.* at 1377-78. The court held that such a summary could substitute for the original and impair the author's markets, reasoning that "a person who had missed an episode . . . would find reading the book an adequate substitute." *Id.* at 1377. The fair-use defense was not defeated because the book described the series, but rather because the book delivered the protected expression itself. *See also Nihon Keizai Shimbun, Inc. v. Comline Bus. Data, Inc.*, 166 F.3d 65, 72–73 (2d Cir. 1999) (rejecting fair use for infringing news abstracts).

The same economic distinction is relevant to ordinary learning. Professor Gans's 2026 study compares three familiar learning activities to AI tools that perform the same learning function. The activities are a teacher who instructs from a series of textbooks, a service that sells summaries of business books, and a guide that answers questions about a television series.[2]

Each activity can reduce demand for the underlying work. Whether it supplies facts and ideas or a substitute for protected expression requires attention to the use.[3] The comparison concerns the economics of learning; it does not determine the legal status of copying used to enable it.

---

[2] Gans, *supra* note 1, at 2–4.

[3] *Id.* Another commentator makes a related observation: people can acquire useful knowledge through commentary about a work and apply it to tasks quite different from consuming the work itself. *See* Kevin Kelly, *Arguments in Favor of AI Fair Use*, The Technium (Sept. 7, 2026), https://perma.cc/JA54-QHWJ (discussing new uses in point 3 and secondary commentary in point 10).

The economic distinction is also relevant when an AI tool performs a similar task. Assessing market effects requires examining what the use supplies to consumers and how it affects the author's market, alongside the other statutory factors.

Thus, an AI response explaining a general economic concept or providing a summary of a particular article may both satisfy a user's demand for information, yet the two draw on protected material for different purposes and compete with the original in different ways. The same use—here, the AI response—may be fair for one purpose and infringing for another. The determination turns on the specific use rather than on the sophistication of the technology. *See Andy Warhol Foundation for the Visual Arts, Inc. v. Goldsmith*, 598 U.S. 508, 532–37 (2023). A documented decline in sales or traffic requires counterfactual analysis; it does not by itself show whether users obtained facts from another source or a substitute for protected expression. A reduction in sales does not by itself establish the harm the fourth factor requires.

B. Copyright Policy Affects Investment In Both Creation And Subsequent Uses

Copyright policy affects incentives both to create works and to use them in valuable ways. Loss of revenue to one creator therefore does not by itself resolve the market-effects inquiry. As the Second Circuit has explained, "[t]he ultimate goal of copyright is to expand public knowledge and understanding," which copyright pursues by giving creators "a financial incentive to create informative, intellectually enriching works for public consumption." *Authors Guild* , 804 F.3d at 212 . The fair-use factors accordingly are "weighed together, in light of the purposes of copyright." *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 578 (1994). Copyright serves those purposes in two ways. It gives authors a return for creating works. It also leaves room for others to build tools that learn from those works.

These two investments—one made by the original author and one made by a third party who builds on the original work—are each valuable and, moreover, depend on one another.

Professor Gans's 2026 article describes an economic model of how copyright rules affect both investments.[4] The model has two kinds of participants: creators deciding how much to invest in the quality of their work, and an AI developer deciding how much to invest in training a general-purpose product.[5] Access to creators' content lowers the cost of training. But the AI product can also reproduce content in ways that reduce demand for the originals.[6] The model also makes two important assumptions. First, no one knows how much harm a given creator will suffer until after the product launches. Second, transaction costs prevent the parties from negotiating licenses in advance.[7]

The model compares three copyright rules. Under no copyright, the developer may train on all content without liability. Under full copyright, creators can bar the use, and the developer trains without their content. Under the third rule—an after-the-fact compensation rule—the developer may train on all content. It must then compensate any creator whose realized harm exceeds a set threshold.[8] Under the model's conditions, this mechanism can increase expected welfare relative to strict exclusion, subject to financial feasibility and implementation costs.[9]

The other two rules involve different investment costs. A full copyright rule shields creators from displacement, but it keeps their content out of training, which harms AI development. The no-copyright rule allows training, but some AI uses could erode creators' return and thus their incentive to invest in quality.[10] The third rule—the *ex post* compensation rule—can mitigate this

---

[4] Gans, *supra* note 1, at 1-2 (2026).

[5] *Id.* at 4,6.

[6] *Id.* at 2,6

[7] *Id.* at 4,6-7.

[8] *Id.* at 5, 9-10.

[9] *Id.* at 1, 5, 11.

[10] *Id.* at 9.

trade-off. By keeping content available, it preserves access to useful training inputs. Because creators expect compensation if serious harm occurs, they invest more in quality than they would under no copyright.[11] The rule works like insurance by preserving access while providing compensation for sufficiently serious realized harm.[12] And no one has to pay the up-front cost of licensing every work used in training.

The model's finding rests on one key assumption: that developers and creators cannot practically negotiate licenses before training begins.[13] Applying that assumption requires evidence about the available contracting arrangements; a model's size or novelty cannot establish it.[14] It may not fit tools that rely on specific, identifiable content, such as a service that retrieves articles from a known publisher. Professor Gans recognizes that licensing may be possible for those uses.[15]

The mechanism is a policy model, not a description of a compensation requirement under the Copyright Act. Fair use is not infringement, and the article supplies no legal basis for compensating every loss caused by a lawful use. See 17 U.S.C. § 107. Its relevance is to show how timing, uncertainty, and contracting conditions affect incentives. The Court should assess the asserted harms under the statutory factors; the model does not replace that inquiry.

C. Scale Makes The Causal Question More Important, Not The Harm More Certain

The scale of AI use makes the cause and extent of harm more consequential; it does not alter the parties' existing legal burdens. The News Plaintiffs argue that Defendants' products can

---

[11] *Id.* at 10-11.

[12] *Id.* at 12.

[13] *Id.* at 4.

[14] *Id.*

[15] *Id.* at 4 & n.7

generate enough competing material to impair the continuing production of journalism.[16] The fourth factor does reach harm of that kind, because it asks whether "unrestricted and widespread conduct of the sort engaged in by the defendant . . . would result in a substantially adverse impact on the potential market" for the work—there must be a causal connection between the use and the market harm. *Campbell*, 510 U.S. at 590. It is true that a modest effect from a single use can become consequential when the use is repeated widely. But that is a reason to examine the cause and extent of the harm with care, not to presume the harm from the capacity of AI to produce competing material.

Assessing harm of this kind depends on a comparison with the world that would have existed without the challenged use. A publisher's revenue can rise and still fall short of what it would have earned absent the use, and its traffic can decline for reasons unrelated to copying, including changes in search, distribution, and reader preferences. Proof of displacement therefore requires more than a showing that revenue or traffic moved in the wrong direction. It requires connecting the specific conduct to the specific loss. The relevant causal proof also differs by use. For training, the question is how using the content produced the asserted displacement; for retrieval and delivery, it is what the service supplied in place of the source.

Two decisions confronting this theory in the AI context explain that a plaintiff must offer proof that the challenged use actually caused the claimed market harm, not merely that AI *can be capable* of generating competing works. In *Kadrey v. Meta Platforms, Inc.*, the court thought the dilution theory not only viable but "highly relevant" for large language models, which can generate "countless" competing works, and it observed that the market for some works, "like news articles,"

[16] News Plaintiffs' Combined Summary Judgment Br. (Redacted Version) 63–74, Sep. 4, 2026, Dkt. No. 1709 ("News Pls.' Br.").

might be "even more vulnerable" to that competition. 788 F. Supp. 3d 1026, 1055-56, 1059 (N.D. Cal. 2025). But the court granted summary judgment for the developer, because the plaintiffs offered "no meaningful evidence on market dilution at all," and courts "must decide cases based on the arguments presented and the evidence submitted by the parties," not on a technology's mere capacity to cause harm. *Id.* at 1055, 1060. Among the questions the plaintiffs had left unanswered was the counterfactual one—how the threat to their market in a world where developers may copy their works compares to the threat in a world where they may not. *Id.* at 1056.

Bartz v. Anthropic PBC went further. Assuming that training would "result in an explosion of works" competing with the plaintiffs' books, the court held that this was "not the kind of competitive or creative displacement that concerns the Copyright Act," which "seeks to advance original works of authorship, not to protect authors against competition." 787 F. Supp. 3d 1007, 1031–32 (N.D. Cal. 2025). Whether or not such dilution can ever count, then, it cannot be presumed from the volume of material a model can produce. It must be shown with evidence connecting the challenged use to a substantial effect on the market for the works at issue.

II. DETERMINING WHETHER AND HOW THE CHALLENGED USE DISPLACES THE ORIGINAL WORK IS CRITICAL TO ASSESSING MARKET IMPACT

A. The Market Effects Assessment Examines How The Challenged Use Draws On A Particular Work

The market-effects inquiry requires attention to *how* the challenged use of protected expression produces the asserted harm, rather than treating every lost sale after learning as equivalent. *See NXIVM Corp. v. Ross Institute*, 364 F.3d 471, 482 (2d Cir. 2004).

Plaintiffs try to supply that harm with two facts—that their works were used to train Defendants' models, and that the training data as a whole is collectively valuable to the models. *See* News Pls.' Br. 59-66. Neither fact alone, nor the two together, establishes how using a work produces the asserted displacement. The economic inquiry concerns the contribution of the works

at issue and its connection to the asserted harm. Mere inclusion in valuable training data does not answer that inquiry.

Annie Liang and Jay Lu propose a counterfactual criterion for assessing essential dependence.[17] Under their criterion, an output depends essentially on a work if it could not have been generated without that work.[18] Remove the work (or a collection of works) from the training data, and ask whether the output remains possible. This distinguishes inclusion, contribution, and essential dependence.

A work can improve a model without being indispensable to a particular output. That an output remains possible without the work does not rule out effects on the output's quality, frequency, or cost. Liang and Lu's criterion is an analytical proposal, not the governing test of infringement. Essential dependence should not be added as a legal element, nor should another available source automatically excuse copying.

Suppose a model produces an output that explains a familiar economic principle. Many works in the dataset teach that principle, so removing any single one would not keep the model from producing the explanation, and the output does not depend on any one work. A work that contributes something distinctive stands on different footing. If nothing else in the dataset supplies what that work provides, removing it would leave the model unable to produce the output, and the output then depends on that work.

These examples illustrate the framework rather than describe Defendants' systems. They show why aggregate value and an individual work's small share do not settle the contribution question. Nor does determining that a single article can be removed without affecting an output

---

[17] Annie Liang & Jay Lu, *Creative Ownership in the Age of AI*, 13-15 (July 29, 2026) (unpublished manuscript).

[18] *Id.*

imply that a publisher's entire archive can be removed without that effect. The other works that make one article unnecessary may be the publisher's own articles. Commentary about a work may also allow a model to produce an output without the work itself. But that commentary may itself derive from the original work, so its availability does not show that the original made no contribution to the output.

The critical point is that the economic consequences of using a work cannot be assessed without determining how that use affects the particular copyright owner. The value of the training data as a whole does not answer that question. Nor does the mere inclusion of a work in that data.

B. The Amount of Training Material, Standing Alone, Is Not Dispositive of Market Effect

The amount of relevant training data can affect essential dependence under this framework. Several works may support the same capability. Accumulating relevant works can therefore reduce dependence on a single source, without establishing that using a replaceable work has no economic effect.

Liang and Lu's analysis identifies conditions for this result.[19] In their model of gradual, light-tailed innovation, and under their conditions on the generation technology, the share of possible outputs that remain generable after removing any single work approaches one.[20] Other works increasingly support capabilities that once depended on an individual contribution.

That conclusion changes under their conditions for sufficiently heavy-tailed innovation. Breakthroughs recur, and that share of output space repeatedly falls arbitrarily close to zero despite the growing dataset. Contributions that open new creative dimensions can likewise remain

[19] *Id.* at 22.

[20] *Id.* at 24-25.

pivotal.[21] Accumulated knowledge makes some contributions easier to replace, but a distinctive contribution can expand what existing works support.

One limit keeps these results in proportion. Liang and Lu measure how many possible outputs depend on a given work, not how much those outputs are worth. A small set of outputs can account for most of the value—so a conclusion that most outputs do not depend on any one work does not tell the Court how much money is at stake for the outputs dependent on a single work.[22] These output-space results do not estimate actual output frequencies, lost revenue, or welfare. Establishing which conditions describe a particular model requires evidence.

Scale alone therefore establishes neither the absence of dependence nor the absence of harm. The framework helps assess economic claims about contributions and capabilities; direct copying during training and substantial similarity in an output remain distinct legal questions.

III. THE LATER EMERGENCE OF LICENSING MARKETS DOES NOT BY ITSELF ESTABLISH COGNIZABLE HARM

A. A Licensing Market Visible Today Does Not Establish That A Workable Market Existed When The Training Occurred

The News Plaintiffs treat the content-licensing deals that have appeared since this dispute began, and the intermediaries now brokering them, as evidence of a market in which they suffered harm when the challenged training occurred years earlier. *See* News Pls.' Br. 71. But a new use can become valuable before any arrangements for licensing that use exist, so a licensing market visible today does not show that a workable one existed when the training occurred.

Economic research on innovation describes this sequence. Martin Beraja and Noam Yuchtman examine how institutions governing a newly valuable resource can emerge as a

[21] *Id.* at 25-30.

[22] *Id.* at 27-28.

consequence of innovation rather than as a precondition to it.[23] They explain that when an advance reveals that a resource has value, the institutions that define and price access to it—contracts, intermediaries, and the litigation and regulation that settle ownership—tend to form afterward, once that value is recognized.[24] Beraja and Yuchtman treat the training of AI models as a leading example, alongside earlier episodes such as the radio spectrum, oil, and digital advertising, in which use ran ahead of settled institutional arrangements.[25] Copyrighted works already had owners; the question here is how existing rights apply to a new use and how parties transact over that new use. Determining whether licensing was feasible for a particular use therefore requires evidence.

Transaction costs explain why that sequence is particularly relevant to licensing. Assembling permissions across a large and dispersed dataset can be prohibitively costly before experience reveals which inputs matter and before intermediaries organize to be able to clear rights at scale.[26] A use can therefore be valuable yet lack a workable means of licensing. A royalty is principally a transfer; reluctance to pay it does not establish a transaction cost. Costly bilateral negotiations also do not rule out collective arrangements.

B. Later Agreements Require A Comparison With Earlier Contracting Conditions

Later licensing deals show that some content can be licensed under later conditions. They may also illuminate what was reasonably likely to develop; but inferring earlier feasibility requires comparison with the conditions then prevailing. Indeed, the fourth fair use factor considers only

---

[23] Martin Beraja & Noam Yuchtman, *Growth at the Resource Frontier* 1–4 (July 4, 2026) (unpublished manuscript).

[24] *Id.* at 22.

[25] *Id.* at 1-4. Beraja and Yuchtman's model allows contracting and does not establish that firms lacked alternatives to proceeding without agreement. Their welfare analysis also omits some costs, including creating original content.

[26] Gans, *supra* note 1, at 4–5.

harm to markets that are "traditional, reasonable, or likely to be developed," *Texaco*, 60 F.3d at 930, and it does not treat the mere absence of a paid license as cognizable harm.

Thus, in *Bill Graham Archives v. Dorling Kindersley Ltd.*, 448 F.3d 605 (2d Cir. 2006), the Second Circuit held that a copyright owner cannot create a cognizable market simply by demanding payment for, or offering to license, a transformative use; a rights holder's willingness to license does not by itself establish harm to a traditional market. *Id.* at 614-615. And in *Texaco*, the Second Circuit found cognizable harm from a company's unlicensed photocopying of journal articles where a functioning mechanism for licensing that copying already existed and made permissions readily available; the court expressly declined to decide whether cognizable harm would exist if no such mechanism were available. 60 F.3d at 930-31.

A later-in-time contract demonstrates that a buyer and seller could reach agreement under the conditions then prevailing. But the critical next step of inferring earlier feasibility requires a comparison: whether the conditions at that earlier time would have allowed comparable rights to be assembled and licensed for the earlier training process and the data it used. That approach gives the later evidence an appropriate role, by requiring a basis for drawing an inference about the earlier-in-time conditions.

Whether the earlier conditions supported comparable licensing requires evidence. That is because several concrete forces can make contracting easier over time. Experience reveals useful applications and the value of particular inputs. Publishers organize rights and delivery services, and intermediaries aggregate permissions and reduce the cost of finding counterparties. Standard terms and ways to identify and measure uses further reduce the burden of reaching agreement. A license to archive from one identifiable owner can also resolve a transaction that is much simpler than clearing permissions across a dispersed training dataset.

The comparison must also move forward as conditions change. New training runs and continuing services call for analysis of their own circumstances, including relevant potential markets. Evidence about licensing should be based on the conditions prevailing at the time when the infringing use is alleged to have occurred.

C. AI Licenses Cover Distinct Rights And Services

The agreements must be examined before they can establish the availability or price of a comparable training license, because a single "AI license" can bundle distinct rights and services. The term can describe permission to train, to retrieve material, to display text, receive current updates, to use a brand, or to obtain a curated archive. A contract can also bundle those permissions with technical services or settlement of past claims.

The publicly reported deals demonstrate just that, with some agreements licensing content for use in a training dataset and others supplying current content for a consumer-facing feature—a different use on different terms.[27] A single price for such a bundle does not reveal the price or value of each component, or establish that each required a copyright license.

The Associated Press's 2023 arrangement with OpenAI illustrates the pitfalls of relying on "AI licensing agreements" to demonstrate comparability. OpenAI licensed part of AP's text archive, while AP received access to OpenAI's technology and product expertise.[28] The announcement supplies an example of archive licensing paired with reciprocal services, and—importantly—it does not disclose all terms or establish comparability to each challenged use. Content agreements did not begin only recently, and some cover training.

---

[27] U.S. Copyright Office, Off. of the Chief Economist, *Identifying the Economic Implications of Artificial Intelligence for Copyright Policy* 36 (Feb. 2025), https://perma.cc/H6RU-4BTE.

[28] Associated Press, AP, Open AI Agree to Share Select News Content and Technology in New Collaboration (July 13, 2023), https://perma.cc/7RFP-K7CV.

An assessment of the parties' agreements should, therefore, identify the material and rights supplied, the activity and period covered, and the arrangements for aggregating permissions and delivering content. Transaction costs matter alongside the agreed price. Payments settling past disputes also require care: they can reflect litigation costs and risk as well as the value of prospective access. Careful assessment of all of these features is necessary to determine what, if anything, can be learned from an agreement about licensing another use.

D. Licensing Can Be Easier For Uses That Draw On Identifiable Content

Whether a use can be licensed depends on what the use is. Some uses draw on identified content when a model answers a user's request. Retrieving current articles from a known publisher to answer a query is one example. Obtaining permission for such a use can be easier than assembling permissions across a dispersed general-purpose training dataset. Those agreements do not by themselves establish comparable licensing feasibility for the training at issue.

Consider a service that repeatedly obtains current articles from an identified publisher. The parties can specify what will be delivered, how often, and on what terms. They can negotiate attribution, permitted display, update speed, and reliable access. A service using a specified archive can offer similarly concrete terms. Because the content and the recurring use are identifiable, agreement is easier, even if assembling comprehensive permissions for an earlier general-purpose training dataset posed a materially different problem.[29] Such agreements can support continued production of content. A service that needs timely and reliable information has an interest in the publisher's investment in producing it, and contracting can connect that interest to the publisher's return when the contribution and use are identifiable.

---

[29] Gans, *supra* note 1, at 4–5.

The United States' submission recognizes specialized-access agreements while expressly taking no position on the financial or logistical feasibility of a general licensing regime.[30] The distinction follows naturally from the different contracting problems involved. Feasibility alone does not establish a legal entitlement to payment. Some retrieval uses may qualify as fair use, while some training uses may be straightforward to license.

Defendants separately defend two retrieval features.[31] OpenAI's Browse service searches the web in response to a user's request and draws on the retrieved pages to answer it. Microsoft's Copilot "web grounding" retrieves webpage content from the Bing search index to inform Copilot's responses. Those features raise their own fair-use questions, and the Court should evaluate them on their own terms. The economic case for fair use of general-purpose training, where the statutory factors warrant it, is compatible with agreements for identifiable uses of content.

Separate analysis also allows the Court to trace connections between training, retrieval, and harmful outputs where the evidence supports them. Professor Gans's economic model, discussed in Part I.B, also accounts for the risk that AI outputs displace demand for original works.[32] Recognizing the value of general-purpose learning therefore leaves room to assess related conduct. The Court can still ask whether particular conduct, such as a specific output or retrieval feature, substitutes for protected expression. It can also ask whether a licensing market for that conduct was workable or likely to develop.

---

[30] Statement of Interest of the United States, *Silverman v. OpenAI Inc.,* No. 1:25-cv-03483-SHS-OTW (S.D.N.Y. Sep. 1, 2026), Dkt. No. 316.

[31] Defendants' Mem. in Supp. of Mot. for Summ. J., Sep. 4, 2026, Dkt. No. 1753; Microsoft Corp.'s Mem. in Supp. of Mot. for Summ. J. in News Plaintiffs' Consolidated Cases, Sep. 4, 2026, Dkt. No. 1698.

[32] Gans, *supra* note 1.

CONCLUSION

Copyright policy affects both the economic return for creating original work and the economic return for building on such works. Any economic analysis therefore should identify how the challenged use changes each of those incentives. A lost sale requires an account of what displaced it. A claim that a new use depends on a work *requires* more than showing its inclusion in a valuable dataset. A later license *requires* comparison with the rights, uses, and contracting conditions relevant to the challenged conduct.

The Court should therefore reject the News Plaintiffs' attempt to categorically preclude a finding of fair use for general-purpose training based on those broad assertions and instead assess supported harm and cognizable markets based on the statutory factors. Where identifiable content and recurring uses make agreement workable, licensing can connect the value of access to incentives for continued creation. The development of those opportunities should inform the analysis of the uses to which they apply, with due attention to when and how they became feasible.

Date: October 30, 2026

Respectfully submitted,

Andrew J. Pincus
MAYER BROWN LLP
1999 K Street, N.W.
Washington, D.C. 20006
(202) 263-3000
apincus@mayerbrown.com

Graham White
MAYER BROWN LLP
1221 Avenue of the Americas
New York, NY 10020
(212) 506-2500
gwhite@mayerbrown.com

Counsel for Amicus Curiae

CERTIFICATE OF COMPLIANCE

I certify that this brief complies with the 6,000-word limitation in the Court's August 7, 2026 Order, ECF No. 1663. The brief contains 5,707 words, including footnotes and excluding the caption, tables of contents and authorities, signature blocks, and this certificate, as permitted by Local Civil Rule 7.1(c). This count was generated by the software used to prepare the document.

The brief is prepared in 12-point Times New Roman type, with 10-point footnotes, one-inch margins, and double-spaced body text, consistent with Local Civil Rule 7.1(b).

Dated: October 30, 2026

/s/ Andrew J. Pincus
Andrew J. Pincus