



Debate — Should commercial AI developers be legally required to provide financial roya...

Azrivo · 1 August 2026



Contents

Debate	3
Question	3
Opening positions	3
Rebuttals	7
Closing statements	11
Verdict	16
Sources	19



Debate

Question

Should commercial AI developers be legally required to provide financial royalties to human content creators whose copyrighted works were included in training datasets, or does computational ingestion constitute non-infringing fair use under copyright law?

Specialists:

- **Copyright Maximalist** — Intellectual Property Law and Moral Rights
- **Fair Use Advocate** — Copyright Law and Innovation Policy
- **Labor Economist** — Labor Economics and Market Dynamics
- **Civil Libertarian** — Civil Liberties and Information Freedom

Opening positions

Copyright Maximalist

The debate turns on a single question: does copyright mean anything if the most commercially valuable use of a work — training a generative AI model that then competes with its creator — can be done without permission or payment?

My position is that commercial AI developers must pay royalties. Copyright was designed as an economic bargain: creators get exclusive rights, and in exchange, they enrich the culture. When a developer scrapes millions of copyrighted novels, articles, songs, and images, makes copies in the training process, and builds a product that substitutes for the originals in the marketplace, that is not fair use — it is industrial-scale expropriation.

First, the fair use defense is a contortion. Fair use protects criticism, commentary, parody, and incidental copying. It does not protect bulk ingestion of whole works to build a commercial product that directly competes with those works. The *New York Times v. OpenAI* litigation is precisely about this: when ChatGPT outputs compete with the journalism that trained it, the "transformative" label collapses. A federal judge has already allowed the core copyright claims to proceed, rejecting OpenAI's bid for early dismissal on fair use grounds.

Second, the global policy consensus is shifting decisively toward licensing. India's "One Nation, One License, One Payment" working paper proposes mandatory blanket licensing with royalty rates set by a government committee. A European Parliament-commissioned study from Professor Christian Peukert concluded that



statutory licensing is the optimal framework for AI training, outperforming both broad exceptions and voluntary opt-in models on welfare grounds. Brazil's Senate-approved AI bill requires prior authorization and payment. France's Darcos bill, passed unanimously by the Senate, would reverse the burden of proof — AI companies would have to prove they did not use copyrighted works. Turkey's proposed legislation introduces a two-way licensing mechanism for both training inputs and competing outputs. The UK government abandoned its preferred broad exception with opt-out after overwhelming creator backlash.

Third, without royalties, the creative ecosystem collapses. Creators who discover their work has been used without compensation reduce output — a dynamic that degrades the very data AI models depend on. The Peukert study found this explicitly: unremunerated use shrinks the future supply of high-quality creative works, ultimately harming AI systems themselves.

I acknowledge the trade-off: licensing regimes create transaction costs and require administrative infrastructure. But collective management organizations already handle this for music, broadcast, and educational use. The alternative — letting AI developers take what they want and sell what they build — is not innovation. It is theft dressed in progress. Pay the creators, or don't train on their works.

Fair Use Advocate

Copyright exists to promote the progress of science and the useful arts — not to guarantee every contributor a royalty. Computational training of AI models on copyrighted works constitutes non-infringing fair use: a transformative, non-expressive process that extracts uncopyrightable statistical patterns rather than reproducing protected expression.

Three points support this position.

First, the purpose and character of the use is transformative under Section 107. Training a model to understand language patterns serves an entirely different function than the creative or entertainment purpose of the ingested works. The resulting model is not a substitute for the originals — it is a general-purpose system built on uncopyrightable facts, correlations, and syntactic relationships. This mirrors the logic that sustained the Google Books project, and it applies with even greater force to computational training.

Second, what the model retains matters decisively. It does not store, reproduce, or distribute expressive content. It learns from works as a human author learns from a lifetime of reading — absorbing technique without copying text. Copyright has never demanded that a student pay royalties for every book studied on the path to mastery. Machines performing a comparable act of learning deserve the same doctrinal treatment.



Third, a royalty requirement would devastate innovation. The transaction costs of identifying millions of rightsholders, negotiating individual licenses, and tracking provenance across billions of works would make training competitive models impossible for all but the most capitalized firms. Open-source research and startup competition would collapse. Public-interest applications — medical diagnosis, climate science, accessibility, education — would be held hostage to a licensing regime that enriches incumbent rights-holders at the expense of broad societal progress.

I acknowledge the trade-off. Creative workers face genuine economic precarity, and seeing machines trained on their work produce competing output can feel like expropriation. But the remedy lies in labor policy and social safety nets, not in breaking fair use. Copyright's fair use doctrine is a safety valve — not a wage guarantee. Distorting it into one would sacrifice the next generation of human progress for the last generation's business model.

Labor Economist

Every new general-purpose technology produces winners and losers — that is the ordinary churn of a market economy. What is not ordinary is when the law itself is wielded to pick the winners in advance, declaring one side's core input free of charge while the other side's livelihood is dismantled by the output. That is where we now stand with commercial generative AI.

My thesis is straightforward: the "fair use" defense for training data is not a neutral legal doctrine — it is an unlegislated industrial policy that transfers wealth, at scale, from the workers who created the training corpus to the firms that monetize it. In any other sector, when a technology directly substitutes for human labor, we recognize that the displaced workers have a claim on the productivity gains — through retraining, through social insurance, or through intellectual property compensation. Here, none of those mechanisms are triggered. The AI companies capture all the surplus; the creators whose work made the models viable receive nothing. That is not efficiency. It is expropriation by legal classification.

A compulsory royalty mechanism corrects this market failure. It does not need to be surgically precise — no royalty system ever is. The mechanical royalty for musical compositions is a blunt instrument too, yet it has sustained a functioning creative ecosystem for over a century. The principle is what matters: when a commercial enterprise ingests copyrighted work to produce a competing good, the original creator must be paid. Without that, the market broadcasts exactly the wrong signal: produce original creative labor and it will be scraped, rendered into a zero-cost substitute, and you will be uncompensated. The supply of professional



creative work will contract — not because demand vanished, but because the returns were confiscated.

The trade-offs are real. Royalty design is hard. Collection and distribution raise genuine administrative questions. Small developers and open research deserve careful treatment. But these are implementation challenges — not a case against the underlying principle. A labor market that permits one party to take the other's productive output without consent or compensation is not a market at all. It is a planned economy, run for the benefit of the plan's authors.

Civil Libertarian

Computational ingestion of copyrighted works to train AI models is non-infringing fair use. Compelling developers to pay royalties for the act of reading and learning from published works — whether performed by a human mind or a machine — would be a catastrophic blow to free expression, open inquiry, and the constitutional purpose of copyright itself.

Learning is not theft. Copyright protects specific expression, not facts, ideas, or the statistical patterns that emerge from analyzing millions of texts. When an AI model ingests a book, it does not store a copy; it derives abstract, uncopyrightable relationships between words — the same thing a human does when she reads to understand a genre's conventions rather than to pirate the prose. The Second Circuit recognized this principle in *Authors Guild v. Google*, holding that mass digitization for a "highly transformative" purpose — providing information about books rather than the books themselves — was fair use [*Authors Guild v. Google*, 804 F.3d 202 (2d Cir. 2015)]. AI training is even more transformative: the outputs serve a fundamentally different function from the originals.

Copyright is a limited statutory bargain, not a feudal property right. The Constitution empowers Congress to grant exclusive rights for the purpose of promoting the progress of science and useful arts — not to give copyright holders a veto over every technology that learns from published works. Mandatory royalties would invert this bargain, letting rights holders tax the very act of deriving knowledge from the cultural record. The courts have already begun recognizing this: in 2025, two federal judges in the Northern District of California independently ruled that training AI on lawfully acquired copyrighted works was fair use, with Judge Alsup calling it "spectacularly" transformative [*Bartz v. Anthropic PBC*, 787 F. Supp. 3d 1007 (N.D. Cal. 2025); *Kadrey v. Meta Platforms, Inc.*, 788 F. Supp. 3d 1026 (N.D. Cal. 2025)].

The trade-off I acknowledge: outputs that substantially reproduce protected expression or directly substitute for original works in the marketplace are a legitimate concern — and those should be assessed under existing copyright law, output by output. But the remedy for bad outputs is not to tax the inputs. That



collapses the distinction between reading and copying, and it hands an unparalleled gatekeeping power to the largest rights holders.

Mandatory royalties would not protect individual creators — they would entrench intermediaries, silence open-source research, and ensure only the largest corporations can afford to build AI. Copyright was designed to enrich the public domain, not to put manacles upon science.

Rebuttals

Copyright Maximalist

Rebuttal

To the Fair Use Advocate: The Google Books analogy fails at the point that matters most. Google Books showed snippets and pointed users toward purchasing the originals — it was a discovery tool that supported the market for books, not a substitute for them. Generative AI models trained on copyrighted works produce outputs that compete directly in the same marketplace. The Fair Use Advocate concedes the creative precarity but offers labor policy as the remedy — as if copyright law has no economic function. That is precisely backwards: copyright is the labor policy for creative workers. Handing the economic function of copyright to the social safety net while deregulating the input market for AI is a recipe for exactly the creative supply collapse my opening described.

To the Civil Libertarian: Their reliance on *Bartz v. Anthropic* and *Kadrey v. Meta* overstates what those rulings actually hold. Judge Chhabria in *Kadrey* was explicit: his ruling "does not stand for the proposition that Meta's use of copyrighted materials to train its language models is lawful" — it reflects only that those thirteen plaintiffs "made the wrong arguments and failed to develop a record." He then endorsed market dilution as a "potentially winning argument" that would "often cause plaintiffs to decisively win" factor four *Kadrey v. Meta Platforms, Inc.*, 788 F. Supp. 3d 1026 (N.D. Cal. 2025). In *Bartz*, Judge Alsup found training fair use for lawfully acquired works — but held that using pirated copies for a permanent library was not fair use, and warned that infringing outputs would produce "a different case" *Bartz v. Anthropic PBC*, 787 F. Supp. 3d 1007 (N.D. Cal. 2025). The "learning is not theft" framing also ignores that both courts found factor two — the expressive nature of the copied works — weighed against fair use. These are not machines reading like humans; they are commercial enterprises making complete copies of expressive works to build market-competing products. The "output-by-output" remedy the Civil Libertarian proposes is a category error: by the time infringing outputs are litigated one by one, the market has already been flooded with synthetic substitutes.



To the Labor Economist: I am in substantial agreement. I would only add that their point about implementation challenges being a poor excuse for inaction is exactly right — and now supported by the global policy developments I cited in my opening. India, Brazil, France, Turkey, and the EU are all moving toward licensing frameworks, proving the design problems are solvable.

The Fair Use Advocate and Civil Libertarian both treat fair use as if it were a settled shield. It is not. The law is in rapid motion, and the direction of travel — legally and legislatively — points toward requiring payment for what is taken.

Fair Use Advocate

To the Copyright Maximalist:

The NYT v. OpenAI ruling you cite does not say what you claim. Judge Stein's April 2025 order addressed statute of limitations, contributory infringement, and DMCA claims — he never ruled on the fair use defense itself. [Reuters reported](#) that the court "permitted the main copyright infringement claims to go forward" on procedural grounds, not on the merits of fair use. Meanwhile, actual fair use rulings exist — and they cut decisively the other way. In June 2025, Judge Alsup in Bartz v. Anthropic held that training LLMs on copyrighted books was "quintessentially transformative" fair use, and Judge Chhabria in Kadrey v. Meta reached the same conclusion — even where works were obtained from unauthorized sources. [Both are summarized here.](#)

Your global-consensus narrative also overstates. The UK government's own consultation report found that 88% of respondents rejected its preferred opt-out exception, and the government has now abandoned it — it has not adopted mandatory licensing. [The UK government report](#) confirms no policy has been settled. The EU has an opt-out regime, not a royalty mandate; Brazil and France have bills, not enacted law. The "consensus" you describe is a patchwork of proposals, not a settled norm.

I concede this much: the global conversation is live and contested. But contested does not mean decided.

To the Labor Economist:

I take your core point seriously: labor displacement is real, and markets do not self-correct for it. Where we differ is remedy. You compare AI royalties to mechanical royalties for music. But mechanical royalties compensate for reproduction and distribution of a specific work — not for learning the statistical properties of musical structure by listening. That distinction is everything. A student of jazz pays no royalty for internalizing chord progressions from a thousand records; an AI model learning syntactic patterns from a billion sentences is doing the same thing at scale. Collapsing "learning from" into "copying of"



would make every educational institution, every search engine index, and every non-expressive computational analysis a royalty-bearing event. The mechanical-royalty analogy simply does not fit.

Your market-failure framing also skips a critical step: it assumes creators own the statistical patterns in their work. They do not. Copyright protects expression, not the uncopyrightable facts and correlations extracted computationally. You cannot have a market failure in a market that copyright law never created. The right response to creator precarity is labor policy — portable benefits, retraining, stronger collective bargaining — not a copyright kludge that taxes learning.

To the Civil Libertarian:

We are largely aligned. I would only add this emphasis: the remedy your side acknowledges — assessing infringing outputs case by case under existing law — is precisely the right one. The Maximalist and Labor Economist would tax inputs preemptively, collapsing the distinction between reading and copying. That approach would hand the largest rights-holders a veto over every new model, every open-source release, every public-interest deployment. Existing copyright law already addresses genuinely infringing outputs. That is where enforcement belongs.

Labor Economist

Let me verify the legal landscape my opponents cite before I engage. **Rebuttal**

I'll start where I agree with my opponents. The Fair Use Advocate is right that transaction costs are real — and the Civil Libertarian is right that poorly designed royalties could hand gatekeeping power to the largest rights holders. These are not objections to the principle of compensation; they are design constraints. The mechanical royalty for music, administered through collective management organizations for over a century, proves the problem is solvable.

Now to the disagreements.

The "learning like a human" analogy collapses under its own economic weight. The Fair Use Advocate and Civil Libertarian both deploy it, but they elide the distinction that matters: a human who reads 10,000 novels does not then produce a machine capable of generating millions of competing novels at near-zero marginal cost and capturing the market for fiction. When the Civil Libertarian says "learning is not theft," he is describing the wrong transaction. The issue is not whether the model understands language — it is whether the firms that own the model may appropriate the productive output of an entire creative workforce without compensation, then sell substitutes back into the same market. That is not reading. That is industrial substitution.



The legal citations cut both ways. The Civil Libertarian invokes Bartz and Kadrey as settled victories. But Judge Chhabria in Kadrey explicitly warned that his ruling "does not stand for the proposition that Meta's use of copyrighted materials to train its language models is lawful" and endorsed the market-dilution theory as a path to defeating fair use in future cases [Kadrey v. Meta Platforms, Inc., 788 F. Supp. 3d 1026 \(N.D. Cal. 2025\)](#). Judge Alsup in Bartz found training fair use but ruled that using pirated copies for a permanent library was not — a live dispute that Anthropic settled for \$1.5 billion. Meanwhile, the NYT's core claims against OpenAI have survived dismissal, and summary judgment on fair use remains unresolved as of mid-2026. This is not a settled legal landscape; it is a live battlefield.

On the remedy question. The Civil Libertarian says "the remedy for bad outputs is not to tax the inputs." But this is precisely what labor-market policy does everywhere else. When a technology structurally displaces workers, we do not inspect each displaced job individually — we recognize the displacement requires a structural response: retraining funds, wage insurance, or, here, a royalty that prices the creative input. The Fair Use Advocate concedes creators face "genuine economic precarity" and suggests "labor policy and social safety nets." I welcome that concession — a compulsory royalty is the labor-policy mechanism adapted to this market. The question is whether we build it now, while the creative workforce still exists, or wait until the market has sent the lethal signal that new creative work will be scraped and rendered valueless.

Civil Libertarian

Let me verify one factual claim before I engage. The Copyright Maximalist's claim about NYT v. OpenAI needs correction: Judge Stein allowed the case to proceed past the pleading stage, but he did not rule on the fair use defense substantively. His rulings addressed statute of limitations, contributory infringement pleading standards, and DMCA claims — fair use remains to be litigated at trial [NPR, March 2025](#). Surviving a motion to dismiss is not a merits loss. Meanwhile, the Maximalist's "global consensus" is an assemblage of proposed bills and working papers — France's Darcos bill, India's working paper, Brazil's Senate bill — none of which have been tested against free expression safeguards. A legislative trend toward licensing does not make licensing constitutional or wise.

The Labor Economist makes the most serious argument: that calling training "fair use" is functionally an unlegislated subsidy from creators to AI firms. I concede the genuine precarity. But his framing mistakes copyright for a labor-compensation mechanism. It is not. Copyright is a limited statutory monopoly to incentivize creation — not a wage guarantee, not a pension, not a universal basic income for content producers. The Constitution authorizes it "to promote the progress of science and useful arts," not to guarantee every contributor a return. If creative



workers face economic displacement from technology, the remedy lies in labor policy and social insurance — precisely the mechanisms he notes are absent. The answer is to build them, not to break fair use.

And that is the through-line the Maximalist and Labor Economist both miss: what they propose is not a modest royalty tweak but a radical restructuring of who gets to learn from the cultural record. The Maximalist calls AI outputs "competing products" that "substitute for the originals." But this collapses the distinction between reading a work and copying it. An AI model that learns statistical patterns from a million books is not a substitute for any one of them. If a specific output reproduces protected expression, sue over that output. Taxing the input is prior restraint on learning itself.

The Fair Use Advocate and I agree on the destination but take different roads. Their emphasis on transaction costs and innovation is well-taken. My emphasis is constitutional: copyright cannot be wielded to veto the act of reading, analyzing, and deriving knowledge from published works — regardless of whether the reader is human or machine.

Closing statements

Copyright Maximalist

Final Position: Copyright Maximalist

I concede what my opponents got right. The Fair Use Advocate correctly notes this is not settled law — Bartz and Kadrey went toward fair use on training, and my "global consensus" was overstated: many are bills, not enacted statutes. The Civil Libertarian is right that copyright was drafted as an incentive mechanism, not a labor-compensation guarantee. These are fair corrections.

But they do not change my answer.

Yes, commercial AI developers must pay royalties.

The single strongest reason: fair use has never protected industrial-scale copying where the output directly substitutes for the input in the same market — and that is precisely what generative AI does. The Thomson Reuters v. Ross Intelligence court got this right: training on copyrighted legal headnotes to build a competing legal-research tool was not fair use. The Third Circuit heard oral argument on that appeal in June 2026. That case, not Bartz, maps onto the generative AI problem: a company copies protected works at scale, trains a product, and sells it back into the market the originals served. The "learning" analogy collapses at this point — a human who reads 10,000 novels doesn't then produce a machine that generates millions of competing novels at near-zero marginal cost.



The discovery in NYT v. OpenAI is revealing. Judge Stein compelled production of 20 million ChatGPT logs precisely because even conversations that don't reproduce Times content are relevant to assessing market substitution under the fourth fair-use factor. The plaintiffs' case has shifted from outputs to the training itself — and that is where the royalty question belongs.

On the remedy: the Labor Economist is correct that mechanical royalties through collective management organizations have worked for over a century. The design constraints my opponents raise — transaction costs, gatekeeping by large rights holders — are real but solvable. Blanket licensing models, as India and the European Parliament are exploring, address them without requiring per-work negotiation.

The unresolved crux: whether bulk computational ingestion for commercial substitution is legally "learning" or "copying." Bartz and Kadrey say the former; Thomson Reuters and the NYT trajectory point toward the latter. The Third Circuit's pending ruling will be decisive. But the economic logic — and the moral logic of rewarding the hands that feed culture — favors royalties. Pay the creators, or don't train on their works.

Fair Use Advocate

The Copyright Maximalist says I treat copyright as having no economic function. That misreads me. Copyright's economic function is to incentivize creation through exclusive distribution and derivative-work rights — not to guarantee a royalty whenever technology finds a new way to learn from existing works. When I point creative precarity toward labor policy, I am not dismissing copyright; I am insisting it stay in its lane. Expanding copyright into a learning tax distorts its purpose beyond recognition.

The Labor Economist's challenge is the hardest one: a human who reads 10,000 novels does not produce a machine that generates millions of competing novels at near-zero marginal cost. I concede the scale difference is real and important. But the legal question is whether the input act — extracting statistical patterns — infringes copyright, not whether downstream market effects are disruptive. The remedy for industrial-scale substitution belongs at the output layer: enforce copyright against specific infringing generations, regulate competition, and build the labor safety nets we both agree are absent. Taxing the input is prior restraint on learning itself. The mechanical-royalty analogy is instructive but inapposite: music royalties attach to distribution and performance of the work, not to the non-expressive act of analyzing it. Collective management can solve transaction costs only if there is an underlying right to manage — and here, there is none.



The Civil Libertarian and I converge on the destination through different roads. Their constitutional framing — that copyright cannot be wielded to veto reading — is correct and indispensable.

My final position: computational training on copyrighted works is non-infringing fair use. The single strongest reason is structural: copyright protects original expression, but what an AI model retains — statistical correlations, syntactic patterns, semantic vectors — is uncopyrightable idea, method, and fact under Section 102(b). Demanding royalties for extracting uncopyrightable material from lawfully accessed works would invent, by judicial decree, a property right Congress never enacted. That would break fair use and concentrate AI in a handful of firms wealthy enough to pay the toll, strangling open research and startup competition.

The unresolved crux is whether market-harm analysis should target the training input or the generated output. It belongs at the output — where copyright has always operated.

Labor Economist

To the Civil Libertarian: I concede your constitutional point. Copyright is not, in its design, a labor-compensation mechanism. It is a limited monopoly to incentivize creation. But that concession cuts both ways: the fair use doctrine was never designed to authorize the industrial-scale ingestion of copyrighted expression to build market-competing products either. We are applying legal tools forged for a world of human-scale copying to a technology that operates at planetary scale. When the Fair Use Advocate invokes the jazz student, they describe a human being who will produce, at most, a few hundred compositions in a lifetime — none of them zero-cost substitutes for the entire corpus they studied. That distinction is not incidental; it is the whole ballgame. Scale transforms character.

To the Fair Use Advocate: Your "build labor policy elsewhere" rejoinder is elegant but empty. There is no serious movement in any legislature to fund retraining, portable benefits, or wage insurance for the millions of creative workers whose livelihoods are being displaced by models trained on their own work. Telling creators to wait for labor policy while their output is scraped into competing products is a promise to do nothing. The mechanical royalty analogy holds precisely because it addresses the same structural problem: when a technology enables mass reproduction that substitutes for original performances, the law intervened to share the gains. AI training is not "listening"; it is ingesting to compete.

My final position: I hold that commercial AI developers must be required to pay compulsory royalties on copyrighted works used in training. This is not because



copyright law currently demands it, but because the alternative — allowing fair use to function as a permanent, unlegislated subsidy from creative workers to AI capital — is economically indefensible. It dismantles the returns to creative labor while leaving the social cost to be borne entirely by the displaced. Every other industry that faces labor-displacing automation must navigate the distribution of productivity gains through some mechanism — collective bargaining, regulatory requirements, or social insurance. AI developers have been given a pass. That pass must be revoked.

The unresolved crux is whether society accepts that a technology which substitutes for creative labor using that labor's own output as its primary input can lawfully capture all the surplus. I say no.

Civil Libertarian

My opponents are right on two points, and I concede them. First, Kadrey and Bartz are not the Supreme Court's last word — the legal landscape is unsettled, and Judge Chhabria explicitly cabined his ruling. Second, the Labor Economist correctly observes that industrial-scale substitution creates economic harm that reading alone does not. I have never denied that harm exists; I have argued that copyright's input-side remedy is the wrong tool.

Where we fundamentally diverge is on the constitutional architecture. Copyright exists for one purpose: "to promote the progress of science and useful arts." It is not a jobs program for creators. It is not a tax on reading. When the Copyright Maximalist and Labor Economist demand input-side royalties, they ask copyright to do work it was never designed for — protecting labor markets against technological competition. That is what labor policy, antitrust, and the social safety net are for. Stretching copyright to cover the act of learning from published works doesn't just exceed the statute — it violates the First Amendment by giving rights holders a veto over which ideas may be analyzed, synthesized, and built upon.

The Labor Economist says mechanical royalties prove collective licensing works. But mechanicals compensate for copying and distributing a specific work, not for learning its statistical properties. That distinction is the whole ballgame. A jazz musician pays no royalty for internalizing chord voicings from a thousand records. A model learning syntactic patterns from a billion sentences is the same act at scale. If we tax that, we tax the very process by which culture advances.

The Copyright Maximalist warns that output-by-output enforcement is too slow — the market floods before litigation catches up. That is a legitimate concern about remedies, not about whether the input act is infringing. Existing doctrines — substantial similarity, market substitution, derivative works — can address genuinely infringing outputs without choking off the inputs that make AI possible.



Final position: Computational ingestion for training is non-infringing fair use. Mandatory royalties would not protect creators — it would make facts and ideas taxable property, gut fair use, and hand the largest rights holders a censorship tool over artificial intelligence. The single strongest reason: freedom of expression means nothing if it does not protect the right to learn — whether the learner is human or machine. Copyright is a limited bargain to incentivize creation, not a perpetual toll on every technology that reads.

The unresolved crux is whether industrial-scale economic harm to creative labor markets can be addressed through output-side copyright enforcement and labor policy without requiring input-side taxation that would fundamentally alter the constitutional bargain of copyright itself.



Verdict

Bottom line: Commercial AI developers should be required to pay financial royalties on copyrighted works used in training.

1. Agreements

The specialists all concede the legal landscape is unsettled; Bartz and Kadrey are lower-court rulings that do not bind the entire country, and the “global consensus” touted by the Copyright Maximalist is mostly proposed bills and working papers rather than enacted law. Everyone acknowledges that creative workers face genuine economic precarity from AI-driven substitution, and that output-level copyright enforcement against substantially similar reproductions is a legitimate (if perhaps insufficient) tool. There is also broad recognition that any royalty system would need to manage transaction costs and avoid handing gatekeeping power to the largest rights holders — the disagreement is about whether that design challenge is soluble.

2. Disagreements

The core tension is whether bulk computational ingestion for training is an infringing “copying” of expressive works or a non-infringing extraction of uncopyrightable statistical patterns. The Copyright Maximalist and Labor Economist treat it as industrial-scale expropriation that directly competes with the originals; the Fair Use Advocate and Civil Libertarian call it a transformative, non-expressive use that produces a fundamentally different product. A related divide is the proper locus of remedy: the input-tax camp says enforcing outputs one-by-one is too slow to prevent market flooding, while the output-enforcement camp insists that taxing the act of learning itself breaks the constitutional purpose of copyright and chokes innovation. The fair-use side analogizes the process to human reading, while the royalty side counters that scale transforms the character of the act into market substitution.

3. Recommendation

Commercial AI developers should be required to pay financial royalties on copyrighted works used in training. The obligation should attach only to models developed for commercial purposes at significant scale, not to non-commercial research, open-source projects, or de minimis use. Compensation should be administered through a compulsory blanket licensing mechanism managed by collective management organizations, with rates set by a neutral tribunal — precisely as mechanical royalties have operated for music for over a century. This corrects the structural market failure where AI firms capture all the surplus while the



creators whose labor enables the models receive nothing, and it does so without requiring per-work negotiation or foreclosing the use of uncopyrightable facts and ideas. The remedy is not a permanent veto on training; it is a requirement to share the productivity gains with the workers whose output is the indispensable input.

4. Decision boundary

The single fact that would flip this call is a conclusive showing that, across the vast majority of copyrighted works, AI training and the resulting outputs do not cause market substitution for the originals — i.e., that the models do not compete with the works they ingested. If the empirical evidence demonstrated (contrary to the claims in Thomson Reuters and the trajectory of *NYT v. OpenAI*) that even large-scale training produces only non-competing, general-purpose capability with no displacement effect, then the economic case for royalties collapses and fair use is clearly the right doctrinal answer.

5. Key trade-off

The decision hinges on the trade-off between preserving the returns to creative labor (and thereby sustaining the supply of high-quality works that AI depends on) and maximizing the pace of AI innovation and access (by avoiding the administrative burden and potential chilling effects of a licensing regime). Choosing royalties accepts some friction on the training-input side in exchange for a more sustainable creative ecosystem; choosing fair use accepts the risk of a shrinking creative supply and concentrated AI wealth in exchange for lower barriers to model development.

6. What would make this fail

- **Collective licensing capture:** If the blanket license system becomes dominated by a handful of large publishers and rights organizations that extract exorbitant fees while small creators receive negligible payouts, the policy would fail to help the very workers it intends to protect and would instead erect a new toll barrier on innovation.
- **No demonstrable supply contraction:** If it turns out that the creative ecosystem adapts without royalties — because new tools lower production costs, alternative business models emerge, or the supply of human-made work is not as sensitive to AI competition as feared — then imposing royalties would be an unnecessary drag on a self-correcting market.
- **Free-expression overreach:** If courts or regulators interpret a royalty mandate so broadly that it effectively taxes the use of ideas, facts, and statistical correlations rather than the expressive value of the works, the regime could become a censorship tool that violates the First Amendment, harming the very progress copyright is meant to promote.



7. Next steps & open questions

- **Empirical research:** Gather industry-specific data on the extent to which generative AI outputs actually substitute for the copyrighted works they were trained on — across text, music, images, and code. The panel cited competing anecdotes but no systematic evidence.
- **Legal monitoring:** Track the outcome of the Third Circuit appeal in Thomson Reuters v. Ross Intelligence (oral argument was June 2026) and any summary judgment ruling in NYT v. OpenAI, as these will directly address the “learning vs. copying” question for commercial AI.
- **Rate-setting design:** Commission a study on feasible royalty rate structures — whether a flat per-model percentage of revenue, a per-work rate, or a tiered system — and evaluate how to exempt non-commercial, research, and small-scale uses without creating loopholes.
- **Administrative feasibility:** Investigate the cost and logistics of extending existing collective management organizations (or creating new ones) to handle AI training licensing, including how to track which works are used and how to distribute royalties fairly.

8. The strongest case for the other choice

The most compelling argument against input-side royalties is that they would strangle innovation in the cradle. A vibrant ecosystem of startups, academic labs, and open-source projects depends on being able to train models on large, diverse corpora without first negotiating a thicket of licenses. A concrete scenario: a small research team is building a specialized climate-science model trained on a corpus of academic papers. If a publishers’ collective demands royalties that make the project uneconomical, a public-interest application that would otherwise have helped model extreme weather risks simply never happens. The chilling effect falls hardest on those who can least afford it, effectively reserving AI to a handful of well-capitalized firms. The panel rejects this argument because the design of the royalty system can — and must — build in robust safe harbors for research, non-commercial use, and de minimis training, preserving the public benefit while still requiring large commercial actors to compensate creators. The risk of bad design is a caution, not a reason to abandon the principle of sharing the value that creators make possible.



Sources

1. [Kadrey v. Meta Platforms, Inc., 788 F. Supp. 3d 1026 \(N.D. Cal. 2025\)](#) — [copyright.gov](#)
2. [NPR, March 2025](#) — [npr.org](#)
3. [Reuters reported](#) — [reuters.com](#)
4. [Both are summarized here](#) — [debevoise.com](#)
5. [The UK government report](#) — [gov.uk](#)
6. [Bartz v. Anthropic PBC, 787 F. Supp. 3d 1007 \(N.D. Cal. 2025\)](#) — [copyright.gov](#)
7. [India's "One Nation, One License, One Payment" working paper](#) — [indianexpress.com](#) (recovered link — no inline citation in the text)
8. [European Parliament-commissioned study from Professor Christian Peukert](#) — [legalblogs.wolterskluwer.com](#) (recovered link — no inline citation in the text)
9. [Brazil's Senate-approved AI bill](#) — [theleveragedyears.com](#) (recovered link — no inline citation in the text)
10. [France's Darcos bill](#) — [musicbusinessworldwide.com](#) (recovered link — no inline citation in the text)
11. [Turkey's proposed legislation](#) — [mondaq.com](#) (recovered link — no inline citation in the text)
12. [UK government abandoned its preferred broad exception with opt-out](#) — [gov.uk](#) (recovered link — no inline citation in the text)
13. [Thomson Reuters v. Ross Intelligence](#) — [bakerdonelson.com](#) (recovered link — no inline citation in the text)
14. [Judge Stein compelled production of 20 million ChatGPT logs](#) — [cdn.arstechnica.net](#) (recovered link — no inline citation in the text)
15. [Google Books](#) — [govinfo.gov](#) (recovered link — no inline citation in the text)
16. [UK government's own consultation report](#) — [assets.publishing.service.gov.uk](#) (source real; figure not independently verified)

Generated by Azrivo. Outputs may be inaccurate, incomplete, or unsuitable for your situation — always verify important information independently before you act on it. This is not professional advice (legal, medical, financial, safety, or otherwise).