The Artists Are Winning the AI Copyright War and Losing the Market: 90% Report Lost Income Anyway

“Deeply exploitative” is how Karla Ortiz described AI media models when she sued Stability AI in January 2023, adding there was “no legal precedent to set this right.” Three and a half years on, the precedent still does not exist.

THE SHORT VERSION

Andersen v. Stability AI has survived the motions that mattered and now has a real shot at a jury, but the trial that was supposed to start September 8, 2026 has slipped to April 5, 2027, on a case filed January 13, 2023. In the same window, a May 2026 visual artists survey put diminished income opportunities at 90%. Litigation is running on a four-year clock against a market that repriced in eighteen months. If you build or buy creative tooling, the legal outcome is a tail risk to hedge, and the labor market shift is already on your P&L.

The claim: a jury verdict resets the training-data economy

The position repeated across artist advocacy, trade press and a fair amount of legal commentary is that Andersen is the case that changes the rules. There is a real basis for it. In August 2024, Judge William Orrick let what the plaintiffs call the model theory proceed: the argument that Stable Diffusion itself, holding compressed transformations of the training images, may be an infringing copy independent of anything a user ever generates. Trackers covering the case treat that ruling as the reason it is still alive when so many sibling suits got pared down.

The second reason the claim has teeth is that Orrick declined to resolve fair use at the pleadings stage. He routed around the fight rather than settling it early. Defendants have argued since the earliest filings that training is transformative and analogous to search-engine indexing and non-expressive intermediate copying. That argument has not been rejected. It has been deferred to a record built in discovery, which is exactly where a plaintiff wants it if the model weights turn out to contain anything that looks like storage.

I wrote about the related ruling when it landed, where a federal judge split lawfully acquired books from pirated copies in the Anthropic case.

Four years in, the calendar is the defendant's best asset

The October 31, 2023 joint case management statement set "Trial Begins Tuesday, September 8, 2026." That date got cited for two years as the first US AI copyright jury trial. It is gone. The AI Lawsuit Tracker now lists April 5, 2027 in San Francisco, describing the docket as having slipped roughly seven months as discovery expanded. Summary judgment before Orrick moved from November 4, 2026 to February 17, 2027, and it moved at the plaintiffs' own request.

The secondary-source picture is messy enough that Axis Intelligence declined to publish either trial date as fact, flagging the case as active with the trial date unresolved because the sources conflict. Small detail, and it tells you a lot: the most-watched AI copyright case in the United States currently has no date that trackers agree on.

Slippage is normal in complex federal litigation. Discovery in a case that turns on what is physically inside a diffusion model's weights is genuinely hard, and the plaintiffs asked for part of the delay themselves, which suggests they wanted a stronger record rather than a faster one. That is a defensible choice. It is also a choice with a cost, and the cost shows up in the survey data.

What the market did while the docket moved

The May 2026 visual artists survey reports 54% saying AI diminished their income, 75% saying it diminished their job and clientele security, and 90% saying it diminished their income opportunities. Note the gap between 54% and 90%. Fewer than half of the artists who lost opportunity have yet lost measurable income. That is a leading indicator. The pipeline thinned before the invoices did.

On the employment side, CVL Economics' "Future Unscripted," commissioned by the Animation Guild IATSE 839, CAA and the Human Artistry Campaign, projected 21.4% of US film, TV and animation jobs, roughly 118,500 positions, consolidated or eliminated by 2026. In the same study's survey of about 300 entertainment industry leaders, 75% said generative AI tools had already supported eliminating, reducing or merging jobs in their division. That last figure is the one I would underline. It is executives describing decisions already made, not a forecast.

I covered a concrete instance of this last June, when AMC Networks began generating marketing assets with Runway's Gen-4 across two of its franchises. Runway is a named defendant in Andersen. A studio signed a production deal with a company it could theoretically be enjoined from using. Nobody in that procurement chain waited for April 2027. Copyright litigation adjudicates the past, procurement decisions allocate the future, and only one of those bills monthly.

Where I land: the remedy artists need is priced, not adjudicated

MY READ

Even a full plaintiff win in Andersen pays for training that already happened. It does not restore the commission pipeline, and it does not make a studio hire an illustrator in 2028. The asset artists actually hold is the one Columbia measured: identical work is valued 62% lower when labeled AI-generated. That is a demand-side premium available today, and it decays every month it stays unlabeled and unenforced. My position is that mandatory provenance and disclosure move more money to working artists over the next three years than any damages award in this case will.

The Columbia Business School research from June 2025 is the most commercially useful number in this whole file. Participants valued art labeled as AI-generated 62% lower than identical work labeled human-made, while conceding the skill involved was comparable. They were paying for origin, not judging quality.

A 62% valuation premium is a market that exists. It requires one thing to function: a credible, cheap, enforceable signal of who made the thing. Without labeling, that premium is unclaimable, because the buyer cannot tell which item is which and defaults to assuming the cheaper provenance.

Which is why EU Parliament document A-10-2026-0019 interests me more than the San Francisco docket. It requires generative AI providers to disclose copyright-protected training content as an itemised list identifying each item, regardless of where the training occurred. Itemised, per item, extraterritorial. If that survives implementation in anything close to its current form, it does two things at once: it builds the evidentiary record that copyright plaintiffs currently have to buy through years of discovery, and it creates the infrastructure that makes the 62% premium chargeable.

Unconfirmed, and I want to be plain about this: I have not seen data on how that itemised-list obligation is being enforced in practice, or whether any provider has produced a compliant list. The text is the text. Implementation is a separate question and I do not have evidence on it.

The strongest argument against me runs through the

model theory

The case for litigation over market structure is better than I have made it sound. Courts create leverage that markets do not. If Andersen reaches a jury and the model theory holds, the settlement value of every pending image-model case reprices overnight, and that repricing pushes licensing deals that create durable revenue for artists rather than one-time payouts.

There is a second point I take seriously. Voluntary disclosure has a free-rider problem that only a court or a regulator solves. If labeling is optional, the firms with the most AI content in their pipeline simply do not label, and the premium never materialises. Orrick's decision to let the model theory proceed is the kind of ruling that makes disclosure legally compulsory rather than commercially optional.

So the honest version of my position is narrower than the contrarian framing suggests: litigation and disclosure are complements, and the sequencing matters. Damages arrive once, in 2027 at the earliest, for training that finished years ago. Disclosure, if it works, compounds. An artist who can prove human authorship at point of sale charges the premium on every piece, every year, without a lawyer.

What would move me off this position

I expect the April 5, 2027 date to move again. That is a prediction, not a fact, and I would be glad to be wrong. My reasoning: the February 17, 2027 summary judgment hearing sits seven weeks before trial, which is tight for a ruling that could reshape what goes to the jury, and the plaintiffs have already shown willingness to trade speed for a better record. Discovery on model internals has no obvious natural end.

Things that would falsify my argument:

If summary judgment on February 17, 2027 produces a clear ruling that the model theory survives and fair use fails, settlement pressure across every pending image-model case jumps immediately, and litigation becomes the faster remedy. I would revise.

If the next visual artists survey shows the 90% opportunity figure holding while the 54%

income figure closes the gap toward it, the damage is compounding faster than any 2027 remedy can address, and the argument for procurement-side pressure gets stronger.

If a provider publishes an itemised A-10-2026-0019 training list and nothing changes commercially (no premium, no licensing deals, no enforcement), then disclosure is a compliance cost rather than a market mechanism, and I would have to concede that damages awards are the only lever that actually moves money.

For anyone running a creative pipeline right now, the practical read is unglamorous. The 75% of entertainment leaders who already cut or merged roles did so without waiting for a court, and the vendors they signed with include a named defendant. **Provenance metadata costs almost nothing to add today** and is the only thing that makes a 62% premium legible to a buyer. Whatever a jury decides in San Francisco, it will decide it about training that happened years before the verdict.

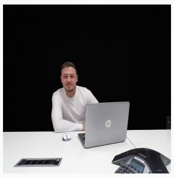

WRITTEN BY

Artur Markus

Software engineer in Basel building AI infrastructure and agentic systems, with a background in pharmaceutical automation, GxP compliance, and validation. I help teams turn AI from a chatbot into a reliable operational layer.

Book a meeting →