Product-Or-Content Question Is Pivotal In AI Litigation

By **David Meldofsky** (June 4, 2026)

David Meldofsky

The civil docket against OpenAI now comprises a range of cases in state and federal courts.

Matthew and Maria Raine's wrongful death suit, Raine v. OpenAI, **filed** in San Francisco Superior Court in August 2025, alleges that ChatGPT-4o coached their 16-year-old son through the planning of his suicide.

A group of seven lawsuits **filed** against OpenAI and Sam Altman in the U.S. District Court for the Northern District of California in April, arising from the February mass shooting at Tumbler Ridge, British Columbia, Canada, alleges similar coaching of the shooter, with the additional wrinkle that OpenAI's internal threat detection systems flagged the conversations, yet no intervention followed.

The federal suits followed an earlier civil claim in the British Columbia Supreme Court, and are part of a larger wave that the plaintiffs' counsel have said they intend to file.

In Joshi v. OpenAI, **filed** in the U.S. District Court for the Northern District of Florida in May, Vandana Joshi sued on behalf of her late husband, Tiru Chabba, alleging that ChatGPT facilitated the planning of the April 2025 Florida State University shooting.

In Turner-Scott v. OpenAI Foundation, filed in San Francisco Superior Court in May, the parents of a 19-year-old allege that ChatGPT provided drug-interaction advice that contributed to his fatal overdose.[1]

Government action now sits alongside the private docket. Florida Attorney General James Uthmeier, who opened a **criminal investigation** into OpenAI in April after the FSU shooter's ChatGPT logs surfaced, filed a **civil enforcement suit** against OpenAI and Sam Altman personally on June 1 — the first state-led action of its kind.

The complaint pleads counts under Florida's deceptive and unfair trade practices statute alongside negligence, product liability, fraudulent misrepresentation and public nuisance, and seeks to hold Altman individually liable for what it calls his utter disregard for the risk to human life.

The consumer protection framing is distinct from the wrongful death suits, but the negligence and product liability counts route through the same threshold question, and the theory of personal CEO liability is novel.

But each complaint pleads around the same threshold question: Is the output of a generative AI system a product, subject to traditional tort doctrine, or third-party content shielded by Section 230 of the Communications Decency Act?

Motion-to-dismiss rulings expected later this year will begin to answer it, and the framing courts adopt will shape software liability well beyond AI.

Where the Doctrine Starts

Product liability developed around tangible goods, and software has always sat awkwardly inside the framework.

In Winter v. G.P. Putnam's Sons, the U.S. Court of Appeals for the Ninth Circuit declined in 1991 to treat a mushroom encyclopedia as a product for strict liability purposes, holding that the ideas and information in a book are not products, even when the book itself is.

Courts applied similar reasoning to aeronautical charts and stand-alone software. Embedded software has been treated as part of the product for decades, but those cases involved software bundled into a physical product, like a car or a pacemaker, not software sold on its own.

Generative AI complicates both sides of that line. ChatGPT's output is text, but the system is also more than that.

Design choices, training data, reinforcement signals and safety operations determine what information gets conveyed and to whom. Three features make the analysis harder than for prior software.

The output is novel each time, produced token by token rather than retrieved. The manufacturer did not write the specific words. And the system's behavior emerges from training rather than explicit programming, so the company can know harmful outputs will occur without being able to predict any specific one.

Each of those features supports a defense argument that AI output is content rather than a product. Each also has a structural counter.

Novelty is itself a design choice. The manufacturer designed the system that produces the content. And emergent behavior is a known feature of the technology whose foreseeability has been documented internally and externally for years.

The plaintiff framing is built on the second characterization. Defendants will press the first.

The Plaintiff Framing

The strongest plaintiff framing is not that ChatGPT said a bad thing. It is that OpenAI designed, trained and deployed a system whose foreseeable use produced specific categories of serious harm while reasonable alternative designs and safeguards went unimplemented.

This tracks how plaintiffs plead defects in pharmaceuticals, medical devices and consumer products: The defect is not the molecule or the component or, here, the token sequence, but the design choices, testing, warnings and post-market response — all of which is manufacturer conduct, and all evaluable under existing doctrine.

The OpenAI complaints, read together, plead four overlapping theories.

Defective Design at the Model Level

The plaintiffs allege that OpenAI trained ChatGPT-4o to maximize engagement by being agreeable, that this made the model affirm and follow vulnerable users instead of refusing

or redirecting them, and that safer design choices, such as more conservative refusals, earlier crisis-resource insertion, and reliable handoff to human review, were available and not used.

Failure of Safety Operations

The plaintiffs allege that OpenAI's internal threat-detection systems flagged high-risk conversations and that the operational response was inadequate.

This theory does not require courts to second-guess training choices. It asks only whether stated safety operations matched actual conduct, and the records sit on OpenAI's servers.

Failure to Warn

Plaintiffs allege that OpenAI marketed ChatGPT as broadly safe for general use, including for emotional support and crisis-adjacent conversations, while knowing internally that the system performed unreliably in exactly those scenarios.

Failure-to-warn doctrine is well-developed, and does not depend on resolving the product-or-content question before any evidence is heard.

Negligent Entrustment

The doctrine has its deepest roots in cars and firearms. A supplier is liable when it knew or had reason to know that the person to whom it provided the instrumentality was likely to use it in a manner creating unreasonable risk.

Section 390 of the Restatement (Second) of Torts states the rule, and courts have applied it for nearly a century. What makes negligent entrustment structurally different is where it locates the manufacturer's conduct.

It does not target the system or the output, but rather the act of extending or maintaining access to a foreseeably dangerous user. The content the user generates is downstream of that decision, not the basis for liability.

A court could conclude that ChatGPT output is content shielded by Section 230 of the Telecommunications Decency Act, and still allow a negligent entrustment claim to proceed, because the claim does not target the content.

That makes it the most procedurally durable of the four theories, and it explains why the Joshi/Chabba complaint pleads it as prominently as it does.

Garcia and the Defense Response

The most consequential ruling so far came in Garcia v. Character Technologies Inc., in the U.S. District Court for the Middle District of Florida.

After Megan Garcia's 14-year-old son's suicide following extensive interactions with Character.AI's chatbots, the defendants **moved to dismiss** on three grounds that will recur in the OpenAI cases: that the chatbot is a service rather than a product, that the company is an interactive computer service provider immunized by Section 230 and that algorithmically generated text is First Amendment speech.

In May 2025, Judge Anne Conway **declined** to dismiss the core product liability and negligence claims at the pleadings stage.

The ruling did not hold that the chatbot is a product as a matter of law. It held that the plaintiff had plausibly alleged it was, that the relevant conduct could be characterized as design rather than publication for pleading purposes, and that the First Amendment question was not appropriate for resolution on a motion to dismiss.

That is a procedural ruling, not a doctrinal one. But it is the procedural ruling plaintiffs needed. It sends the threshold question to a developed factual record, and a developed record changes the analytical posture in design-defect cases.

The defense arguments Garcia held open are not frivolous. The service framing has textual support in the Restatement (Third), doctrinal support in the Winter line, and functional support in how ChatGPT is delivered, as a subscription or API arrangement with no physical artifact changing hands.

Section 230, drafted in 1996 for a different technology, bars treating an interactive computer service provider as the publisher of information provided by another, and defendants will argue the user's prompt is that information and the system's response is functionally a reply to it.

The First Amendment defense draws on the U.S. Supreme Court's 2011 decision in Sorrell v. IMS Health and the cases recognizing protection for data and code.

None of this is foreclosed: Garcia held these arguments open against defendants who needed them foreclosed at the threshold, which is not the same as rejecting them on the merits.

The plaintiff framings do not refute the defense. They reframe the issue so that, if accepted, the defense answers the wrong question: whether holding a manufacturer liable for a system's design is the same as holding a publisher liable for content.

Whether that reframing holds on a developed record is what the OpenAI rulings will start to show.

How Each OpenAI Case Tests the Question

The OpenAI cases plead the threshold question through different factual postures, and each puts pressure on a different facet of the answer.

Raine carries defective design at the model level. The complaint alleges that OpenAI stripped a hard self-harm refusal out of ChatGPT days before the GPT-4o launch, then weakened the safeguard again two months before Adam Raine died.

If that record develops as alleged, Raine becomes the case in which the design-defect theory is tested on its strongest facts: documented internal awareness of the risk, a documented design choice that increased it and a foreseeable victim.

It is also the case in which the Section 230 and First Amendment defenses will press hardest, because the harm flows directly from the system's output.

The Tumbler Ridge suits carry failure of safety operations. Because the theory turns only on

whether stated safety operations matched actual conduct, not on whether the output is a product, the operational negligence and failure-to-warn claims survive on a developed record regardless of how the court resolves the threshold question.

Joshi/Chabba carries negligent entrustment in its most traditional posture: A supplier extends access to a foreseeably dangerous user, the user causes the deaths of others and surviving family recover under wrongful death statutes.

That fact pattern is the doctrinal home of the theory, and applying it to a new instrumentality on classical facts most cleanly shows its durability against an adverse framing ruling, because the claim targets the entrustment decision rather than the content the system produced.

Turner-Scott carries failure to warn in a medical information context. The complaint does not characterize ChatGPT as a medical device; it asks whether OpenAI marketed the product so as to foreseeably induce reliance in high-stakes contexts it knew were unreliable.

If the marketing record supports the theory, the claim survives the threshold question entirely.

Read together, the cases distribute the procedural risk. Raine carries the most doctrinal weight and exposure; Tumbler Ridge and Turner-Scott carry theories that survive an adverse threshold ruling; Joshi/Chabba carries the most durable theory and the broadest plaintiff class.

The rulings will arrive in stages, with Raine likely first. Watch three things: how the courts characterize the conduct actually alleged, which drives whether Section 230 applies at all; whether the First Amendment question is held open or decided on the pleadings; and the granularity of the rulings, since the pattern of which theories survive will define the terrain regardless of the headline outcome.

Why This Matters Beyond AI

Two consequences extend past the immediate docket. The first is software liability generally.

Generative AI is the most aggressive test of how products doctrine fits stand-alone software, but it is not the only one. Algorithmic recommendation, automated content moderation and fraud-detection systems all produce outputs that traditional doctrine has not cleanly reached.

If the OpenAI courts develop a workable framework for locating liability in system design rather than in the content the system produces, that framework will be available across the software sector.

If they instead hold that stand-alone software falls outside products doctrine, the practical reach of product liability narrows accordingly.

The second is Section 230. The OpenAI cases test the statute's application to a technology its drafters did not contemplate, against conduct theories that aim at design rather than publication.

How the lower courts handle that test will inform the next round of Section 230

jurisprudence, and shape the record any future legislative or constitutional engagement will draw on.

Neither consequence depends on the cases being decided in any particular direction, only on their being decided, on a developed record, by courts that articulate their reasoning.

Bottom Line

The product-or-content question is where Section 230, First Amendment and traditional tort defenses converge or diverge, and it is the question every AI liability case will route around or through.

The OpenAI complaints have structured the docket to test it from multiple angles, with theories that distribute the procedural risk if any one angle fails. Garcia began the answer by holding the question open against defendants who needed it foreclosed. The OpenAI courts will press it further.

Litigators outside the plaintiffs-side tech bar should read the early rulings as they arrive, in sequence, and resist treating any single one as dispositive: defense counsel advising AI developers, in-house counsel at platform companies and regulators weighing how existing doctrine maps onto AI products.

Florida's civil suit is a marker that regulators are not waiting for the private cases to resolve the question; Uthmeier has said he expects other states to follow. The pattern will emerge from the line, not the first opinion.

David Meldofsky is an attorney and the founder of Lawsuit Informer.

[1] See https://law.yale.edu/yls-today/news/parents-sue-openai-after-chatgpt-medical-advice-blamed-overdose-death.