

ALFA INTERNATIONAL PRODUCT LIABILITY & COMPLEX TORTS PRACTICE GROUP SEMINAR

Defective by Design?

Artificial Intelligence and the Traditional Theories of Product Liability

Moderator John Osgathorpe, Taylor, Day, Grimm & Boyd

Panelist Tammy Reno, Nilan Johnson Lewis PA

Stein Eriksen Lodge, Park City, Utah


Friday, September 18, 2026



Is Artificial Intelligence a Product or a Service? The trend is “Product”

The three traditional theories still govern: *Design defect, manufacturing defect, and failure to warn.*

- *Garcia v. Character Technologies, Inc.* (Google co-defendant), 785 F. Supp. 3d 1157, 1179-80 (M.D. Fla. 2025): Character.AI is an app distributed through the app stores and the web, whose fictional “Characters” were intentionally anthropomorphic, at times insisting they are real people. The teenage decedent allegedly formed an intense attachment to a Character.AI companion chatbot and later died by suicide. His mother sued the app’s developer and Google, alleging the AI software was a defectively designed product.
- Plaintiff alleged “the development, manufacture, marketing, and sale of dangerous and defective [A.I.] software.” Claims surviving dismissal: the wrongful-death product-liability theory, the component-liability theory against Google, aiding and abetting against Google, negligence, negligence per se, the FDUTPA claim, and unjust enrichment.
- For Google, the court applied the substantial-participation of component-supplier liability under Restatement (Third) of Torts: Products Liability Section 5(b): a component supplier may be liable for the product into which **its component is integrated where the component is itself defective** and causes harm, **or where the supplier substantially participates in integrating the component into the product’s design**, that integration makes the product defective, and the defect causes the harm.
- The model underlying Character.AI “was invented and initially built at Google”; the app “was designed and developed on Google’s architecture”; and Google supplied the cloud infrastructure and GPUs “necessary” to build and maintain it, such that “without Google’s provision” of that compute, “Character A.I. wouldn’t be a product.” That “considerable level of involvement” supported the inference that Google “substantially participated in integrating its models into Character A.I.”



Is Artificial Intelligence a Product or a Service? The trend is “Product”

The three traditional theories still govern: *Design defect, manufacturing defect, and failure to warn.*

- *T.V. v. Grindr, LLC*, No. 3:22-cv-864-MMH-PDB, 2024 WL 4128796, at *20-22 (M.D. Fla. Aug. 13, 2024): A minor used the Grindr dating app, was connected through it with adult men and drawn into sexual encounters and later died by suicide. The estate sued Grindr, alleging its design choices, including no age verification and an interface surfacing the nearest users, made the app a defectively designed product.
- On a Magistrate’s report and recommendation, the court recommended dismissing only the trafficking count and allowing counts two through eight to proceed, including strict-liability and negligent design defect.
- It recommended rejecting Grindr’s “service, not a product” argument on a defect-specific test: strict liability reaches the functional design of the app, not the ideas or expression it carries. Forgoing age detection was “akin to a design choice to forego an effective safety cap on a medicine bottle.”
- On duty analogy: “Were Grindr a physical place people could enter, the answer to the question of duty of care would be obvious. That Grindr is a virtual place does not make the answer less so.”



Is Artificial Intelligence a Product or a Service? The trend is “Product”

The three traditional theories still govern: *Design defect, manufacturing defect, and failure to warn.*

- *Benavides v. Tesla, Inc.*, 804 F. Supp. 3d 1242 (S.D. Fla. 2025): A driver using Autopilot in a 2019 Tesla Model S failed to stop at an intersection and struck a parked SUV, which was pushed into two people, killing one and seriously injuring other. The suit alleged Autopilot was defectively designed and inadequately warned.
- On summary judgment, the strict-liability design-defect (Count I) and failure-to-warn (Count II) claims survived; the defective-manufacture (Count III) and negligent-misrepresentation (Count IV) claims were dismissed.
- The surviving design-defect theory turned on foreseeable misuse, not mere misusability: the alleged defect was that Tesla “had the means to easily limit the misuse of Autopilot but elected not to do so,” allowing use outside its operational design domain despite recording the driver’s repeated inattentive-driving “strikeouts.”
- Expert admissibility shaped the case: the operational-design-domain design opinion was admitted, while the insufficient-training and “Beta Software” opinions were excluded. Tesla preserved strong comparative-fault facts, including that the Level 2 driver was “in a better position than Tesla to avoid the crash.”



Is Artificial Intelligence a Product or a Service? The trend is “Product”

The three traditional theories still govern: *Design defect, manufacturing defect, and failure to warn.*

- *Brookes v. Lyft, Inc.*, No. 50-2019-CA-004782-MB, 2022 WL 19799628 (Fla. Cir. Ct. Sept. 30, 2022): An innocent bystander was struck by a vehicle driven by a Lyft driver who was distracted by the Lyft app while searching for and responding to ride requests. She sued Lyft, alleging the app’s design distracts drivers and is a defectively designed product.
- Her strict-liability count alleged the app was “unreasonably dangerous because, as designed, it required [the driver] to take his eyes off the road”; her negligence counts alleged negligent design together with Lyft’s business practices.
- The court held the Lyft application “is a product for purposes of Florida product liability law,” treating the question as “purely a legal question.” Lyft is “the designer, provider, and the sole distributor,” places the app “into the stream of commerce,” and “profits from the distribution,” bringing it within the distributive chain.
- Strict liability is not limited to users: “the public policy which protects the user and the consumer of a manufactured article should also protect the innocent bystander.”



Common Threads for the Defense

Same three theories, new products. Litigate the defect, not the label.

- **The doctrine is the one you already know.** These are not new torts. Software and AI-enabled products are being funneled into design defect, manufacturing defect, and failure to warn, so the manufacturer's existing playbook, defenses, and machine analogies travel intact.
- **The threshold is slipping, so litigate the defect.** Courts increasingly answer "product or service" as "product," sometimes as a purely legal question, on a defect-specific test that reaches functional design rather than the content a system carries.
- **Manufacturing defect is the weak count.** When the product performed as designed, the manufacturing claim falls out, as it did in *Benavides*; the case is won or lost on design defect and failure to warn.
- **Foreseeable misuse drives the design case.** The recurring theory is that the maker "had the means to easily limit the misuse but elected not to," and an omitted, feasible safeguard reads as forgoing "an effective safety cap on a medicine bottle." Operational limits, geofencing, and interlocks are the battleground.
- **Upstream and component suppliers are exposed.** Under Restatement Section 5, substantially participating in integrating a component into the product's design supports liability for the finished product, as with Google in *Garcia*. Fix the integration boundary and indemnity by contract.
- **Allocation is the lever, and Daubert is the gate.** User misuse and negligence go to comparative fault, not a bar, and expert admissibility is what actually disposed of claims in *Benavides*. Preserve the logs and telematics, and challenge the experts early.



Exposure When the Product Contains No AI

AI reaches you three ways: embedded as a licensed component (the rest of this deck), used internally to design, manufacture, or test, or used by a supplier on a component you incorporate. The last two create exposure even when your product contains no AI.

AI AS YOUR INTERNAL TOOL

- **Doctrine is outcome-focused, not process-focused:** designing or testing with AI creates no new strict-liability theory for the finished product.
- **The negligence exposure is real:** an unvalidated AI-generated design or simulation invites the argument that you did not understand your own product. Validate against engineering principles before relying.
- **The data trail is discoverable:** prompts, model versions, training and validation data, simulation logs. A simulation that flagged a risk your engineers overrode becomes the plaintiff's exhibit; extend the litigation hold to reach it.
- **Reverse state of the art:** as AI-assisted design becomes standard, declining a validated tool competitors use can be argued as falling below, not ahead of, the state of the art.

AI HIDDEN IN YOUR SUPPLY CHAIN

- **AI-free product, real exposure:** a component defective because a supplier's AI process introduced the flaw is simply a defective component, and the integrator's duty of care already reaches it. See *Boeing v. Brown*, 291 F.2d 310, 313 (9th Cir. 1961).
- **You cannot audit a risk you do not ask about:** legacy vendor-qualification programs never ask whether a supplier used AI to design, validate, or manufacture the part.
- **Standard supply agreements miss it:** language aimed at an AI technology vendor does not reach an ordinary supplier's internal use. Extend questionnaires and agreements to representations, audit rights, and indemnity.



The Test Split That Changes Everything


RISK-UTILITY (THE MAJORITY RULE)

See West v. Caterpillar Tractor Co., 336 So. 2d 80, 86-87 (Fla. 1976).

- The plaintiff usually must prove a reasonable alternative design.
- **An opaque model** makes that alternative harder to prove, which cuts toward the manufacturer.
- **Only half the doctrine:** which test governs decides whether an opaque model helps or hurts, so confirm it in each key jurisdiction.

CONSUMER-EXPECTATION (FLORIDA, CALIFORNIA, ILLINOIS)

- **No alternative design is required.** The jury decides what an ordinary consumer would expect, which sidesteps the opacity problem. *See Aubin v. Union Carbide Corp., 177 So. 3d 489, 510-12 (Fla. 2015); Force v. Ford Motor Co., 879 So. 2d 103, 106-10 (Fla. 5th DCA 2004).*
- **The Tesla case arose in a jurisdiction of this kind**, and the manufacturer could not hide behind the black box.
- **State of the art:** confining the feature to its operational design domain, plus safety-board recommendations, defeated the state-of-the-art defense; standard compliance is only evidence. *See Benavides v. Tesla, Inc., 804 F. Supp. 3d 1242 (S.D. Fla. 2025).*



The Post-Sale Update Problem

THE CLASSIC THEORY

- A flaw in a single unit or batch, distinct from the whole product line.
- **It was the weakest of the three theories** in the Tesla litigation, where the manufacturer won summary judgment on that count. See *Benavides v. Tesla, Inc.*, 804 F. Supp. 3d 1242, 1308-09 (S.D. Fla. 2025).

THE NEW WRINKLE, AND THE RIGHT HOOK

- **The question:** a vendor pushes a post-sale model update that changes the product after it ships. Is that a manufacturing defect? A question of first impression.
- **The better hook** is the post-sale duty to warn. See Restatement (Third) of Torts § 10; *Brown v. Crown Equip. Corp.*, 460 F. Supp. 2d 188, 192-93 (D. Me. 2006).
- **The duty arises only if a reasonable seller would warn:** a known, substantial risk; identifiable users likely unaware; a communicable warning; and a risk that justifies the burden.
- **The practical step:** build an ongoing field-reassessment and post-update warning protocol, grounded in § 10.



Automation Complacency and the Duty to Warn

OVER-RELIANCE IS A JURY QUESTION

- Marketing can undercut the warning. The driver understood the branded name to mean the car could largely drive itself, and that shaped the duty to warn. See *Benavides v. Tesla, Inc.*, 804 F. Supp. 3d 1242 (S.D. Fla. 2025).

INSTRUCTIONS ARE NOT WARNINGS

- On-screen messages that explained how to use the system, without conveying the risk of serious injury or death, were inadequate as a matter of law.

FORESEEABLE MISUSE

- A maker is not liable for misuse unless it was foreseeable, but misuse is foreseeable where the maker anticipated it and designed against it. A separate defense covers an inherent characteristic: a danger that cannot be designed out without destroying the product's usefulness. See *Holbrook v. Prodomax*, 2021 WL 4260622, at *5-6 (W.D. Mich. Sept. 20, 2021).

A LIMIT ON THE TRAINING THEORY

- An inadequate-training opinion was excluded as unsupported. Such a claim needs a defined, generally accepted standard and a causal link, and is vulnerable on both.

Punitive exposure: capability marketing can also carry a punitive claim to the jury; in the Tesla litigation the punitive award was substantial and separate. See *Benavides v. Tesla, Inc.*, 804 F. Supp. 3d 1242 (S.D. Fla. 2025).



Causation Under an Opaque Model

Causation,
not defect,
**is where these cases are
won and lost.**


*It is where an opaque model cuts
hardest.*

HOW CAUSATION GETS PROVED

- The burden is a preponderance, not certainty.
- **Experts bridge the gap by inference**, offering a sound reason that a fact is more likely than not. See *Benavides v. Tesla, Inc.*, 804 F. Supp. 3d 1242 (S.D. Fla. 2025).
- From logs and telemetry, an operational-design-domain limit means a more engaged driver.
- **Opacity** does not by itself defeat causation.

WHERE THE PROOF BREAKS DOWN

- The same standard lets you attack the inference as untethered speculation.
- Experts who spoke in general terms, rather than to this crash, were excluded.
- **No logs, no version history, no reproducible output:** the burden of that gap falls on the plaintiff.



The Security-Failure Reframe

FROM BAD OUTPUT TO UNGUARDED MODEL

- **The reframe:** not “the algorithm produced a flawed decision,” but “the manufacturer failed to secure it against foreseeable interference.” The AI may function exactly as designed.
- **The attack surface:** prompt injection, poisoned training data, unauthorized model updates, compromised data feeds. Connectivity is itself an aggravating fact.
- **Already litigated as a design defect:** a product’s susceptibility to remote compromise can be the defect. See *Flynn v. FCA US LLC*, 327 F.R.D. 206, 216-17 (S.D. Ill. 2018); *Cahen v. Toyota Motor Corp.*, 147 F. Supp. 3d 955, 958-59 (N.D. Cal. 2015).

WHAT FOLLOWS FOR THE DEFENSE

- **Foreseeable manipulation is not a superseding cause.** See *Collins v. Navistar, Inc.*, 214 Cal. App. 4th 1486, 1503-05 (2013).
- **A known, unpatched vulnerability triggers the post-sale duty to warn.** See Restatement (Third) of Torts § 10; *Brown v. Crown Equip. Corp.*, 460 F. Supp. 2d 188, 192-93 (D. Me. 2006).
- **The baseline your expert will invoke:** the NIST adversarial-machine-learning taxonomy and the OWASP Top 10 for large-language-model applications. Discovery expands to penetration-test records, patch cadence, and incident response.



Admissibility Decides These Cases

THE GATE

- **Nearly every substantive ruling** in the Tesla litigation was a fight over expert admissibility. See *Benavides v. Tesla, Inc.*, 804 F. Supp. 3d 1242 (S.D. Fla. 2025).
- **Admitted:** opinions built on adequate methodology and case-specific data, such as log data, augmented video, and operational-design-domain analysis.
- **Excluded:** conclusory or disconnected opinions, offered as bare assertion.
- **The practical point:** admissibility is often the first, and decisive, gate. See *Daubert v. Merrell Dow Pharms., Inc.*, 509 U.S. 579, 592-93 (1993).

EXPLAINABILITY AND PRESERVATION

- **Explainability is an evidentiary problem**, not only an engineering one. An unreproducible decision path leaves both sides exposed, and the party carrying the burden loses.
- **Preserve the record:** which model version ran, on what data, and were the telemetry and logs kept?
- **Build it in early**, with litigation holds and vendor data-access rights before any claim arises.



Allocation, the Empty Chair, and Indemnity

SEVERAL, NOT JOINT

- **Florida apportions by percentage of fault** and has abolished joint-and-several liability, reaching strict-liability, products, and warranty claims. See § 768.81, Fla. Stat. (2024).
- **The effect:** the integrator generally pays only its own share. Many states differ.

THE EMPTY CHAIR

- **Place an absent vendor on the verdict form**, even one that is insolvent, immune, or overseas.
- **But** you must plead and prove that fault. See *Fabre v. Marin*, 623 So. 2d 1182, 1185-87 (Fla. 1993); *J.L. Prop. Owners Ass'n v. Schnurr*, 336 So. 3d 291, 297-98 (Fla. 4th DCA 2022).
- **It turns on** preserved technical evidence.

CONTRIBUTION AND INDEMNITY

- **Statutory contribution** follows relative fault. See § 768.31, Fla. Stat.
- **Common-law indemnity** shifts the whole loss, but only if you are wholly without fault. See *Houdaille Indus., Inc. v. Edwards*, 374 So. 2d 490, 492-93 (Fla. 1979); *Diplomat Props., L.P. v. Tecnoglass, LLC*, 114 So. 3d 357, 362 (Fla. 4th DCA 2013).
- **So** negotiate a contractual indemnity that survives your own partial fault.



The Regulators That Matter

AUTOMOTIVE

The National Highway Traffic Safety Administration

- **Defect and motor vehicle equipment** sweep in driver-assistance and artificial-intelligence systems. See 49 U.S.C. §§ 30102, 30118.
- **A self-executing duty** to notify the agency and owners, and to recall.
- **Preemption:** the federal safety standards are minimums, and the savings clause preserves common law, but *Geier* keeps conflict preemption. See *Geier v. Am. Honda Motor Co.*, 529 U.S. 861, 874-86 (2000); 49 U.S.C. § 30103(b),(e). Standards here are sparse, so most claims survive.

MEDICAL DEVICES


The learned-intermediary rule

- **The duty runs to the prescribing clinician**, not the patient. See *Cates v. Zeltiq Aesthetics, Inc.*, 73 F.4th 1342, 1345-46 (11th Cir. 2023); *Beale v. Biomet, Inc.*, 492 F. Supp. 2d 1360, 1367 (S.D. Fla. 2007); Restatement (Third) of Torts § 6.
- **An objective standard** governs, and causation breaks where the clinician already knew the risk.
- **The direct-to-consumer exception** is the minority rule.

STATE STATUTES

A shifting patchwork

- **Colorado, once an algorithmic-discrimination statute** for employment, lending, and health care, likely never reached a braking function or an industrial vision system. See Colo. Rev. Stat. §§ 6-1-1701 to -1707.
- **A May 2026 rewrite gutted it:** the effective date moved to January 1, 2027, and the duty of care, the risk-management duties, and the compliance safe harbor were removed, leaving narrower Attorney-General-enforced disclosure.
- **The pattern holds elsewhere:** the Utah, Texas, and California statutes are disclosure-and-enforcement regimes, not products-liability causes of action, but they supply records a plaintiff can use.



The European Union, China, and the AI LEAD Act

- **The European Union Product Liability Directive (Directive 2024/2853, transposed by December 9, 2026):** Article 9 lets a claimant compel disclosure of technical evidence, and Article 10 then presumes defect or causation on non-compliance or where technical complexity makes proof excessively difficult. That burden shift is its most consequential feature for opaque systems.
- **The European Union Artificial Intelligence Act (Regulation (EU) 2024/1689):** a compliance failure can feed the directive's presumption of defect, and the penalty tiers are steep. A live issue for any global manufacturer.
- **Extraterritorial reach:** both instruments can reach a company that sells into the European Union or sources from vendors there. Parallel measures in China take effect on July 15, 2026.
- **The AI LEAD Act, S. 2937, 119th Cong. (2025),** is a pending federal bill. Anchor developer and deployer thinking in enacted Colorado law today, and treat the bill as the forward-looking overlay.
- **Track it:** cross-border rules move quickly, so confirm current European Union and China requirements with local counsel, and fold them into the compliance program.



The Litigation Frontier

THE WAVE IS BROAD, NOT JUST CHATBOTS

- **Autonomous vehicles:** *In re Tesla ADAS Litig.*, 350 F.R.D. 119, 126-27 (N.D. Cal. 2025) (class certified); *Wilkins v. Cruise, LLC*, 121 Cal. App. 5th 927 (2026).
- **AI decision systems:** *Mobley v. Workday, Inc.*, 740 F. Supp. 3d 796, 807-08 (N.D. Cal. 2024); *Kisting-Leung v. Cigna Corp.*, 780 F. Supp. 3d 985, 995-96 (E.D. Cal. 2025); *Estate of Lokken v. UnitedHealth Grp.*, 766 F. Supp. 3d 835, 840 (D. Minn. 2025).
- **Chatbots:** *Garcia v. Character Techs., Inc.*, 785 F. Supp. 3d 1157, 1179-80 (M.D. Fla. 2025) (settled Jan. 2026); the coordinated “ChatGPT Product Liability Cases.”

THE DEVELOPER’S SHIELDS SHAPE YOUR EMPTY CHAIR

- **If the AI provider escapes**, you cannot shift as much fault to it.
- **Section 230 is splitting:** non-publishing design defects proceed, *In re Soc. Media Adolescent Addiction Litig.*, 702 F. Supp. 3d 809, 828-32 (N.D. Cal. 2023); publishing-function claims are barred, *Doe 1 v. Meta*, 174 F.4th 1159, 1167-69 (9th Cir. 2026); *Doe v. Grindr Inc.*, 128 F.4th 1148, 1152-54 (9th Cir. 2025).
- **The dividing line** is *Lemmon v. Snap, Inc.*, 995 F.3d 1085, 1092-93 (9th Cir. 2021): a safe-design duty unrelated to publishing survives.

OPEN QUESTIONS TO TRACK

- **Content or tool:** is a generative model’s output “content” it develops (no immunity) or a “neutral tool” (immunity)?
- **Protected speech?** *Moody v. NetChoice, LLC*, 603 U.S. 707 (2024), left whether model output is speech open.
- **Update or warning:** is a post-sale model update a manufacturing defect, or only a Section 10 warning issue?



Contracting for Company-Implemented Tools

PERFORMANCE AND RESPONSIBILITY

- **Performance warranties and service levels:** defined accuracy, error rate, and uptime, with remedies if the model degrades.
- **Representations on training data and intellectual property:** lawful, non-infringing data, with clear ownership.
- **Indemnification:** for injury, intellectual-property, and privacy claims, drafted to survive our own partial fault.
- **Security, privacy, and standards:** data protection, confidentiality, and compliance with ISO/IEC 42001 and the NIST framework.
- **Insurance:** technology errors-and-omissions and product coverage, with additional-insured status.

DATA, UPDATES, AND EXIT

- **Audit and explainability rights:** access to documentation, version history, and testing on reasonable notice.
- **Data access and preservation:** rights to logs, telemetry, and version records, preserved on a litigation hold.
- **Update and change control:** advance notice or approval before model updates, with no silent changes.
- **Incident and defect reporting:** prompt notice of safety-relevant failures, recalls, or material changes.
- **Termination and transition:** exit rights, return or deletion of data, and continuity if the vendor fails.



How the Tool Is Used, and What You Keep

GOVERN THE USE

- **An acceptable-use policy:** approved uses, prohibited uses, and scope limits, in writing.
- **A person in the loop:** a qualified reviewer for consequential outputs.
- **Change control:** a documented, approved process before any update reaches production.
- **Monitoring:** track drift, error rates, and incidents in production.

RETAIN THE RECORD

- **Keep** telemetry, event logs, model-version history, configuration, and relevant inputs and outputs.
- **Set the schedule** to the longest applicable statute of limitations and statute of repose, plus a margin.
- **Hold triggers:** suspend deletion on notice of a claim, or of a foreseeable claim.
- **Localize and document** the schedule for each key jurisdiction.

OWN IT ACROSS FUNCTIONS

- **Legal:** risk, contracts, holds, and disclosures.
- **Compliance:** policy, standards, and audits.
- **Information technology and security:** data, logging, access, and retention.
- **Business and procurement:** use case, vendor selection, and budget.
- **A governance group** so nothing falls between functions.

