THE HIDDEN LIABILITY OF AI HALLUCINATIONS: CAN BUSINESSES SUE AI VENDORS FOR INCORRECT OUTPUTS?

Vansh Sharma, LLB (Hons), OP Jindal Global University, Sonipat. Haryana

ABSTRACT

When an AI tool invents some facts and a business acts on them, who pays? This article raises a question-about whether a business can make the company that built the AI pay for these made up outputs, known as hallucinations. It argues that the answer should usually be no, but not always. The liability should depend on four questions: was the harm foreseeable, how did the contract split the risk, how much human checking stood between the output and the loss, and for what the tool was being used. A general -purpose AI vendor has not been held liable till now by the court. Courts in US and India have blamed the professional who failed to check the work. The article then asks what an Indian business can actually sue under. The Consumer Protection Act, 2019 is the best route, through deficiency of service, product liability and unfair contract terms, but it excludes pure economic loss, which is what a hallucination usually causes, and may not cover commercial buyers at all. Contract law runs into the vendor "as is" disclaimer, which Indian courts uphold between equal parties. The article closes with where a claim can still win, how to plead it, and why the evidence must be preserved.

Introduction

You have probably done it yourself. You asked an AI chatbot a question, received a confident, fluent, well-structured answer, and only later discovered that part of it was simply invented. Most of the time this costs you nothing more than a few minutes. But what happens when that confident, made-up answer finds its way into a contract you signed, a medical decision you acted on, or a brief you filed in court? The technical name for this behaviour is a hallucination. The real-world consequence is loss, and sometimes a great deal of it.

That is not a hypothetical. In 2023 a New York lawyer named Steven Schwartz used ChatGPT to research a personal injury claim against the airline Avianca. The tool produced half a dozen supporting decisions, complete with case names, quotations and citation numbers. None of them existed. When the fabrication came to light, Judge P. Kevin Castel of the Southern District of New York sanctioned the lawyers and their firm.[1] [2] What is striking is who the court held responsible. It did not blame OpenAI. It said plainly that there is nothing improper about using a reliable artificial intelligence tool, and placed the duty to check the work squarely on the human who relied on it.[3]

This raises the question that Indian law has not yet answered. When an AI system gets it wrong and your business pays the price, can you make the company that built the AI pay instead? This article argues that the answer should usually be no, but not always. AI vendors should not be held responsible liable for hallucinations. Liability should instead depend on four things: whether the harm was foreseeable, how the contract allocated the risk, how much human oversight stood between the output and the harm, and the nature of the use case. Each is examined below, with a close focus on where an Indian business realistically stands.

I. So Far, the Law Has Blamed the User, Not the Machine

The most important fact about AI hallucination litigation is that, to date, no court anywhere has held a general-purpose AI vendor liable for a hallucinated output. Every prominent decision has fallen on the human who failed to check. *Mata vs Avianca* was the first, and the courts that followed took the same view. In ***Wadsworth v. Walmart***, a federal court in Wyoming fined

[1] *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023).
[2] *Id.* at 448, 466 (imposing a $5,000 penalty on the attorneys and their firm under Fed. R. Civ. P. 11).
[3] *Id.* at 448 (observing that there is nothing inherently improper about using a reliable artificial intelligence tool for assistance).

lawyers who had filed a motion citing eight fake cases produced by their firm's own AI platform, noting that it is by now well known in the legal community that these tools generate fictitious authorities.[4] The Second Circuit has likewise referred a lawyer for discipline over an invented, AI-generated citation.[5]

The Indian courts have travelled the same road. The clearest signal came in early 2026, when the Supreme Court of India took notice of a trial court judgment that had relied on fake, AI-generated precedents in ***Gummadi Usha Rani v. Sure Mallikarjuna Rao*.** The Court did not mince words. It held that a decision based on non-existent and fake judgments is not a mere error in decision making, but would amount to misconduct, and that legal consequences would follow.[6] The Delhi High Court had earlier refused to decide a trademark dispute on the strength of ChatGPT's answers, warning of its inaccuracies and its tendency to produce fictional case law,[7] and the Punjab and Haryana High Court, in the first Indian use of the tool, was careful to say it had consulted ChatGPT only for a broad picture and not to decide the merits.[8]

There is one revealing contrast. In *Moffatt v. Air Canada*, an airline was held liable liable by the Canadian tribunal when the chatbot on its own website gave wrong advice to a passenger about bereavement fares, rejecting out of hand the airline's argument that the chatbot was a separate legal entity responsible for its own statements.[9] Liability attached to the business that deployed the tool to the public, not to the company that supplied the underlying software. That distinction, between the deployer and the vendor, runs through everything that follows.[10]

II. Foreseeability: A Known Risk Cuts Both Ways

Negligence law begins with foreseeability: a defendant is liable for harm that was a foreseeable consequence of its conduct. The difficulty for AI vendors is that hallucination is not a rare glitch but a measured, documented property of these systems. A Stanford study found that general-purpose models gave hallucinated answers to specific legal questions between 58 and

[4] *Wadsworth v. Walmart, Inc.*, 348 F.R.D. 489 (D. Wyo. 2025).
[5] *Park v. Kim*, 91 F.4th 610, 615 (2d Cir. 2024).
[6] *Gummadi Usha Rani v. Sure Mallikarjuna Rao*, 2026 SCC OnLine SC 341 (India) (order dated Feb. 27, 2026) (Narasimha & Aradhe, JJ.).
[7] *Christian Louboutin SAS v. Shoe Boutique*, CS(COMM) 583/2023 (Delhi High Court, Aug. 22, 2023) (Singh, J.).
[8] *Jaswinder Singh v. State of Punjab*, CRM-M No. 22496 of 2022 (Punjab & Haryana High Court, Mar. 27, 2023) (Chitkara, J.).
[9] *Moffatt v. Air Canada*, 2024 BCCRT 149 (Can. B.C. Civ. Resolution Trib.).
[10] *Id.* at para. 27 (rejecting the argument that the chatbot was a separate legal entity responsible for its own actions).

88 percent of the time,[11] and a later study found that even specialised legal research tools, marketed as reliable, still hallucinated between 17 and 33 percent of the time.[12] If a vendor knows its product invents facts at these rates, the argument runs, then harm to users is entirely foreseeable, and foreseeability is the foundation of a duty of care.

That is the strongest argument against the position taken here, and the answer is that foreseeability cuts both ways. The very studies that make hallucination foreseeable to the vendor also make blind reliance unreasonable for the user. Once it is common knowledge that these tools fabricate, as the court in *Wadsworth* expressly found, the immediate cause of the harm is the user's decision to rely on an unchecked output, an intervening act that breaks the chain of causation. Foreseeing that a probabilistic system will sometimes err is not the same as foreseeing the specific false statement that caused the loss, and the law has always treated a general, disclosed and unavoidable risk differently from a specific and preventable one. Foreseeability does not deliver automatic vendor liability; it explains why the user who ignores a known risk usually bears the loss.

III. The Contract: Where the Risk Is Actually Allocated

In practice, the question of who pays is decided long before any harm occurs, in the vendor's terms of service, and the market here is remarkably uniform. Every major vendor disclaims responsibility for accuracy and pushes the duty to verify onto the user.

Vendor	What the terms say about accuracy
OpenAI	Services are provided "as is." Use of output is at the user's sole risk, and output should not be relied on as a sole source of truth or a substitute for professional advice.
Anthropic	Outputs are provided "as is." Factual assertions should not be relied upon without independently checking their accuracy, as they may be false or misleading. Verification is the customer's responsibility.
Microsoft Copilot	Copilot can make mistakes. Users are told to verify the accuracy of information before relying on it, and Microsoft does not guarantee accuracy.

[11] Matthew Dahl et al., *Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models*, 16 J. Legal Analysis 64 (2024).
[12] Varun Magesh et al., *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools*, 22 J. Empirical Legal Stud. 216 (2025).

Sources: OpenAI Terms of Use; Anthropic Commercial Terms of Service; Microsoft Copilot Terms of Use.

These disclaimers are real and current.[13],[14],[15] Whether they hold up depends on the jurisdiction. In the United States, an "as is" disclaimer between commercial parties is generally enforceable, subject only to unconscionability. The United Kingdom is stricter for consumers, since the Consumer Rights Act 2015 implies quality terms that cannot be excluded, though it applies only a reasonableness test to clauses between businesses. The European Union has gone furthest: its revised Product Liability Directive of 2024 brings software and AI within strict product liability and does not permit a vendor to exclude that liability towards end users.[16]

India sits at the permissive end of this spectrum, and this is the heart of the matter for an Indian business. There is no general statutory ban on exclusion clauses between commercial parties. The only real route around a disclaimer is Section 23 of the Indian Contract Act, which voids agreements that are opposed to public policy,[17] together with the line of cases beginning with *Brojo Nath Ganguly* in which the Supreme Court struck down unconscionable terms.[18] But that doctrine has a crucial limit. It applies only where there is a gross inequality of bargaining power. In *Bharathi Knitting Co. v. DHL Worldwide Express*, the Court upheld a limitation-of-liability clause between two commercial parties,[19] and Indian courts have repeatedly held that the unconscionability doctrine does not rescue a business that struck a bad commercial bargain at arm's length.

The counterargument is that AI terms are standard-form contracts of adhesion, accepted with a single click and open to no negotiation, and that they should therefore be treated as unconscionable. There is force in this where the user is an individual consumer. But most

[13]OpenAI, *Terms of Use*, https://openai.com/policies/row-terms-of-use/ (last visited June 21, 2026) (providing the services on an "as is" basis and stating that use of output is at the user's sole risk).
[14]Anthropic, *Commercial Terms of Service*, https://www.anthropic.com/legal/commercial-terms (last visited June 21, 2026) (disclaiming that outputs are accurate, complete or error-free and placing the duty to verify on the customer).
[15]Microsoft, *Copilot Terms of Use*, https://www.microsoft.com/en-us/microsoft-copilot/for-individuals/termsofuse (last visited June 21, 2026) (advising users to verify the accuracy of information before relying on it).
[16]Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on Liability for Defective Products, 2024 O.J. (L 2024/2853) 1 (bringing software and AI within strict product liability).
[17]Indian Contract Act, No. 9 of 1872, § 23 (India) (rendering void agreements whose object the court regards as opposed to public policy).
[18]*Cent. Inland Water Transp. Corp. v. Brojo Nath Ganguly*, (1986) 3 SCC 156 (India).
[19]*Bharathi Knitting Co. v. DHL Worldwide Express*, (1996) 4 SCC 704 (India) (upholding a limitation-of-liability clause between commercial parties).

businesses that deploy AI in consequential workflows, whether a law firm, a hospital or a bank, have the procurement power to negotiate accuracy warranties, indemnities and audit rights. They usually do not, because the tools are cheap and convenient. A business that could have bargained for protection and chose not to will struggle to persuade an Indian court that the bargain was unconscionable.

IV. Human Oversight: The Decisive Variable

If one factor decides these cases, it is human oversight. Going through again to the decisions. *Mata*, *Wadsworth* and *Gummadi Usha Rani* are, at bottom, all the same case. A tool produced a draft, a professional adopted it without checking, and the harm arose in the gap between the two. Where oversight is real, the user catches the error and there is no loss. Where oversight is not there, it is the user who has been careless, and not the vendor.

This is why the law increasingly leans on the idea of keeping a human in the loop. Scholarship on AI negligence argues that where the system is a tool that supplies a human with information, the person operating the tool is the one ultimately responsible for the result.[20] The European Union has written this principle into law, requiring that high-risk AI systems be designed so that a human can effectively oversee them.[21] The same logic underlies the learned intermediary doctrine in medicine, where the manufacturer's duty runs to the doctor who exercises independent judgment, not to the patient.

The honest counterargument is that human oversight is sometimes a fiction. A reviewer who rubber-stamps whatever the machine produces has not really supervised anything, and when a tool is deployed to thousands of employees, genuine review of every output is impossible. Automation bias, the human tendency to trust a confident machine, is real and well documented. But this does not argue for shifting liability to the vendor. It argues for matching the level of oversight to the stakes of the task. The reviewer who defers entirely to the machine on a high-stakes decision has not discharged the duty, and that failing belongs to the user, not the supplier.

[20] Andrew D. Selbst, *Negligence and AI's Human Users*, 100 B.U. L. Rev. 1315 (2020).
[21] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence, art. 14, 2024 O.J. (L 2024/1689) 1 (AI Act) (requiring effective human oversight of high-risk AI systems).

V. The Use Case: Not All Hallucinations Are Equal

The final factor is the nature of the use. A hallucination in a brainstorming session is of no harm; the same hallucination in a cancer diagnosis, a credit decision or a court filing can be very problematic. Modern AI regulation is built around exactly this idea of graduated, risk-based treatment. The European Union's AI Act sorts systems into risk tiers, and the GDPR gives individuals a right not to be subject to purely automated decisions that significantly affect them,[22] a right that has no equivalent in India's Digital Personal Data Protection Act of 2023.

Regulators, too, have aimed at high-stakes and deceptive conduct rather than at error as such. The United States Federal Trade Commission's "Operation AI Comply" targeted vendors who misrepresented what their products could do, not vendors whose products simply made mistakes.[23] Vendor liability, then, is not off the table. It attaches where a vendor markets a tool for a high-risk use, hides a known accuracy problem, or makes a false claim about performance. But in those cases the wrong is misrepresentation or a defective product, not the hallucination itself.

VI. The Indian Toolkit: What a Business Can Actually Sue On

The four factors describe how a court should give their reasoning on this. But an Indian business that has lost its money to a hallucination wants to know something more concrete: which law do I sue under, in which forum, and what can I recover? The honest answer is that the routes are very narrow, and they run in a clear order of preference.

The most promising route, and the right place to start, is the Consumer Protection Act of 2019. It offers two different ways. The first is a complaint of deficiency in service, which the Act defines as any fault, imperfection or inadequacy in the quality or manner of performance of a service.[24] An AI tool sold on a subscription is, a service, and a tool that returns fabricated information can be argued to be deficient. The second way is the product liability action introduced in Chapter VI of the Act, which allows a complainant claim compensation where a

[22]Regulation (EU) 2016/679, art. 22, 2016 O.J. (L 119) 1 (GDPR) (right not to be subject to a solely automated decision producing legal or similarly significant effects).
[23]Press Release, Fed. Trade Comm'n, FTC Announces Crackdown on Deceptive AI Claims and Schemes (Sept. 25, 2024), https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes ("Operation AI Comply").
[24]Consumer Protection Act, No. 35 of 2019, § 2(11) (India) (defining "deficiency" in service).

defective product or a deficient service causes harm.[25] For the first time in Indian law, a product service provider can be made to pay for a service that was faulty or inadequate, and a complainant need not prove a contract at all.

This is where the analysis has to be honest about the limits, because they are the core of the problem. The product liability route carries a built-in trap for exactly this kind of claim. The Act confines recoverable harm to death, personal injury and damage to property, and it pointedly excludes pure commercial or economic loss and loss of profit.[26] Yet economic loss is precisely what a hallucination usually causes. A business that acts on a fabricated market figure, a fabricated clause or a fabricated citation loses money, not life or limb. So the very law that looks most promising fails at the most common fact pattern, and succeeds mainly where the hallucination caused physical harm, such as a defective AI medical device or a faulty diagnostic tool. There is also the unsettled question of whether software counts as a "product" at all, since the Act's definition was written with goods in mind.

The Act has one more tool that is easy to overlook. It empowers the consumer forums to strike down an unfair contract term, including a term that limits liability in a way that defeats the purpose of the contract.[27] This is the most honest answer to the vendor's disclaimer, but it is subject to the same limitation that runs through this discussion. These protections belong to a "consumer," a person who buys for personal use and not for resale or commercial purpose. An individual who depended on a hallucinating consumer app fits comfortably. A company that deployed the tool across its business may fall outside the definition altogether, and with it outside the Act's protection.

When the Consumer Protection Act does not fit, the business is left with ordinary contract law, and here the ground is harder. A claim for breach runs straight into the vendor's disclaimer, and the only way past it is to argue that the exclusion is unconscionable and void under Section 23 of the Indian Contract Act. As shown earlier, that argument works for a weak consumer but rarely for a commercial party bargaining at arm's length. Even if a business gets past the disclaimer, Section 73 caps its damages at losses that naturally flowed from the breach or that

[25]Consumer Protection Act, No. 35 of 2019, § 2(34) (India) (defining "product liability"); *id.* §§ 82–87 (the product-liability action).
[26]Consumer Protection Act, No. 35 of 2019, § 2(34), 2(22) (India) (confining "harm" in product-liability claims to death, personal injury and property damage, and excluding loss to the product itself and pure commercial or economic loss).
[27]Consumer Protection Act, No. 35 of 2019, § 2(46)–(47) & § 49(2) (India) (defining "unfair contract" and empowering the State Commission to declare contractual terms null and void).

both parties knew were likely, and the disclaimer itself is powerful evidence that the vendor never accepted that risk.[28] Contract law, in other words, gives the business a theory but takes away most of the recovery.

Two further instruments shape the battlefield without giving the business a weapon. The Information Technology Act of 2000 offers vendors a safe harbour for intermediaries, but it is a shield, not a sword, and it is far from certain it even applies. A generative-AI provider does not only host another person's content; it generates the output itself, which may place it outside the definition of an intermediary entirely.[29] That uncertainty cuts against the vendor as much as for it. Finally, the Digital Personal Data Protection Act of 2023 is notable for what it leaves out. India made much of its data law on the European framework, but it declined to copy the European right not to be subject to a purely automated decision. A person harmed by an automated, hallucinated decision in India therefore has no data-protection remedy that is given to a European, which is a real and intentional gap.

VII. Where a Business Is Most Likely to Win, and How to Sue

Putting the four factors and the Indian toolkit together produces a practical map. An Indian business suing an AI vendor for a hallucinated output is unlikely to succeed in most of the cases , but its chances increases in four situations

1. **The vendor made a specific accuracy claim.** Where a vendor advertised that its tool was reliable for a particular professional use but it was not, the business can claim it as misrepresentation or as a deficiency in service, which avoids the disclaimer because a party cannot disclaim its way out of its own false promise.

2. **The buyer was a consumer, not a business.** The protection given by the Consumer Protection Act against unfair terms and its product liability route are far stronger for an individual consumer than for a commercial buyer, because the equal-bargaining-power exception that saves the disclaimer does not apply.

3. **The vendor also deployed the tool.** As *Air Canada* shows, a company that puts an AI in

[28]See Indian Contract Act, No. 9 of 1872, § 73 (India) (limiting damages to loss that naturally arose from the breach or that the parties knew was likely).
[29]*Tata Sons Ltd. v. Greenpeace Int'l*, (2011) 178 DLT 705 (India) (on the function of an intermediary). The status of a generative-AI provider under § 79 remains untested.

front of the public and presents its output as its own answer cannot hide behind the tool. The deployer is responsible for what it tells its customers.

4. **The harm was not purely economic but physical.** Because the product liability route excludes pure commercial loss, a hallucination that causes physical injury, through a faulty medical or diagnostic tool, fits the regime far better than one that merely costs money.

The strategy follows from the map. A claimant should plead in the alternative, leading with deficiency of service and unfair-term relief under the Consumer Protection Act, and pleading misrepresentation rather than mere breach so that the disclaimer becomes a shield the vendor cannot raise against its own false claim. The forum choice matters too. The consumer forums, tiered by the value of the claim from the District Commission upward, are faster and cheaper in comparison to a civil suit and do not require the claimant to prove a contract, which is why they are the natural first stop for a smaller loss; a large commercial loss with no consumer hook may have nowhere to go but a civil court or arbitration under the vendor's own terms.

Whatever the forum, the case is won or lost on evidence, and the evidence is fragile. The only important thing a business can do is preserve the record at the moment the harm occurs: the exact prompt, the exact output, the version of the tool, and the date. A hallucination is ephemeral. The same prompt run a week later may produce a right answer, and the vendor will say the tool works. A claimant who cannot show the court the false output as it appeared has, in practice, no case. Causation is the next problem, because the business must show that the hallucination is the reason behind the loss and not its own failure to check, and that failure is exactly what the vendor will point to. This is why the realistic remedy in most cases is compensation for a proven deficiency in a consumer forum, not the large damages award a business imagines, and why the disclaimer, the economic-loss exclusion and the verification failure together explain why most claims settle small or never begin at all.

Outside the four favourable situations, the business is on weak ground. Its claim runs into the vendor's disclaimer, the equal-bargaining-power exception that protects that disclaimer, the difficulty of treating software as a "product," and the exclusion of pure economic loss from the product liability regime. The intermediary safe harbour in the Information Technology Act, whatever its precise reach for generative AI, is in any event a shield for the vendor and not a sword for the business.

VIII. Conclusion: A Calibrated Rule, Not a Blanket One

For now,India has chosen not to write a special law for artificial intelligence. The November 2025 governance guidelines say as much, preferring to rely on existing law wherever possible.[30] That choice is sound, because the existing law, applied with care, already points to the right answer. A rule of automatic vendor liability would be easy to state and wrong in practice. It would punish vendors for a known and disclosed limitation of their technology, ignore the contracts that businesses freely signed, and excuse the careless professional who filed the brief without reading it.

A calibrated rule works better. Was the harm foreseeable, and to whom? How did the contract divided the risk, and whether the parties bargained as equals. How much human judgment stood between the machine and the harm, and what the tool was being used for. Answered honestly, these four questions will usually place the loss on the business that relied on an unchecked output, sometimes on the company that deployed the tool to the public, and occasionally on a vendor that lied about what its product could do. That is not a gap in the law but the law doing its work, sorting a new problem into categories it has long understood. For the businesses now using these tools into their daily tasks, the message is simpler still. The cheapest insurance against an AI hallucination is not a lawsuit after the loss. It is a human being who reads the output before anyone relies on it.

[30] Ministry of Electronics & Info. Tech., Gov't of India, *India AI Governance Guidelines* (Nov. 2025).