Liability for AI Outputs in Multi-Layered AI Supply Chains

Authors: Tuhin Batra, Partner
Niharika Singh, Senior Associate

I. Context

The commercial deployment of artificial intelligence systems increasingly relies on layered supply chains involving foundation model providers, API intermediaries, system integrators, and enterprise deployers. This multi-tiered structure complicates a fundamental legal question, liability for outputs generated by AI systems. As regulatory frameworks evolve, the allocation of risk through contract is becoming the primary mechanism through which parties attempt to manage uncertainty.

Unlike traditional software, where outputs are largely deterministic and attributable to coded logic, AI systems introduce probabilistic outputs that may be inaccurate, biased, or infringing. In the absence of a comprehensive statutory liability regime specific to AI in India, contractual arrangements between vendors and customers play a decisive role in determining how risk is distributed across the supply chain.

Structure of AI Supply Chains and Points of Legal Exposure

A typical AI deployment involves multiple contractual layers. A foundation model provider licenses access to a model, often through an API. A downstream vendor integrates this capability into a product or service. An enterprise customer then deploys the system in a commercial setting. Each layer introduces distinct legal exposures.

At the model layer, risks include copyright infringement arising from training data, and output generation that reproduces protected content. At the integration layer, risks relate to misuse, insufficient safeguards, or failure to implement appropriate content filters. At the deployment layer, liability may arise from reliance on AI outputs in decision-making processes affecting consumers or third parties.

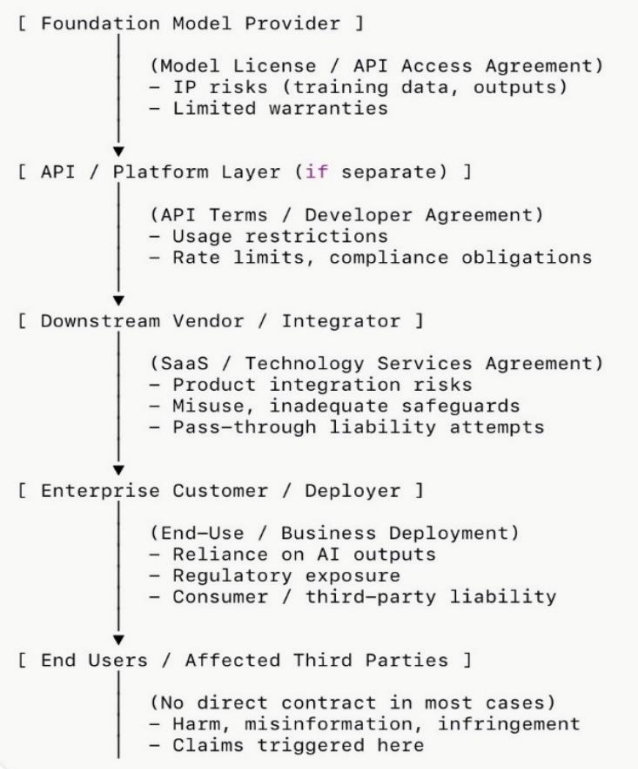

Typical AI Supply Chain Flow

This fragmentation creates a misalignment between control and liability. The entity closest to the end user may face regulatory or litigation exposure, even where the underlying risk originates upstream.

Recent Policy Trends

Recent regulatory and policy narratives[1] in India suggest that the policy makers want to "Govern" the applications, not the "technology" itself. The recently issued

[1] "*India AI Governance Guidelines: Enabling Safe and Trusted AI Innovation*", a document published by MeitY on 05th Nov 2025.

guidelines by MeitY, "India AI Governance Guidelines: Enabling Safe and Trusted AI Innovation" ("MeitY AI Guidelines"), suggests that AI should be regulated by empowering the sectoral regulators, rather than enacting a single overarching AI law governing the underlying technology. The official narratives also suggest "Techno-legal" approach as a governance tool and "Voluntary Industry frameworks" as the first layer of accountability.

Amendments Proposed in existing Laws

The 'MeitY AI Guidelines' also suggests amendments being introduced in existing laws as a short-term measure. The suggested amendments are:

(i) Amendments in Copyright Act, 1957.
(ii) Amendments in laws for graded liability system based on the function performed, level of risk, and whether due diligence was observed.
(iii) Amendments in current definitions under IT Act, 2000: Definitions of 'intermediary', 'publisher', 'computer system', etc. to bring clearly in the roles of various actors in the AI value chain (i.e., a 'developer', 'deployer', 'users', etc.) and how they will be governed.

On the sectoral side, Indian sectoral regulators like SEBI[2], Reserve Bank of India (RBI)[3], Insurance Regulatory and Development Authority of India (IRDAI)[4], and recently the Indian Council of Medical Research (ICMR) and Central Drugs Standard Control Organisation (CDSCO) have already issues guidelines, frameworks, and consultation papers for AI related issues concerning their respective sectors.

The issue of Copyright infringement in generative AI

It is interesting to note that soon after a report published by INDIAAI in January, 2025, titled "Report on AI governance guidelines development", the Department for Promotion of Industry and Internal Trade (DPIIT) established a committee in April, 2025 to deliberate on the issue of Copyright in relation to generative AI systems and implications of using copyrighted materials in the training and development of AI models. As per Section 52 of the Indian Copyright Act, 1957, limited 'fair dealing' exceptions apply for private or personal use, including research. The DPIIT committee is currently examining the legality of using copyrighted material in training AI systems and its implications. It is also evaluating the copyrightability of works produced by generative AI systems, and reviewing international practice to propose a copyright framework for India.

Absence of Clear Statutory Allocation of Liability

Indian law does not currently provide a dedicated framework for allocating liability for AI-generated outputs. Existing doctrines under contract law, tort law, and intellectual property law apply, but do not directly address the distributed nature of AI systems.

Judicial developments in the digital context provide partial guidance but do not resolve AI-specific questions. In *Shreya Singhal v. Union of India*[5], the Supreme Court clarified the contours of intermediary liability and held that platforms are required to act upon receiving actual knowledge through lawful orders. While this framework is relevant to platform liability,

[2] (1) The SEBI algorithmic trading framework (updated Feb 2025, effective April 1, 2026), accessible at here; (2) Reporting for Artificial Intelligence (AI) and Machine Learning (ML) applications and systems offered and used by Market Infrastructure Institutions (MIIs), accessible here.

[3] (FREE-AI) Committee Report: Framework for Responsible and Ethical Enablement of Artificial Intelligence. Accessible here.

[4] (1) Emerging Technologies in Insurance - Adoption Strategies by Various Stakeholders, 09th Dec 2020, accessible here; and (2) "Report on InsurTech - Working Group Findings & Recommendations", 31st July 2018. Accessible here.

[5] AIR 2015 Supreme Court 1523.

its application to AI-generated outputs remains uncertain, particularly where content is generated autonomously rather than hosted or transmitted.

Similarly, in *Google India Pvt. Ltd. v. Visaka Industries*[6], the Court examined intermediary protections in the context of defamation and recognised that liability may arise where an intermediary has knowledge and fails to act. However, AI systems complicate this analysis, as the concept of "knowledge" becomes difficult to apply to probabilistic and dynamically generated outputs.

In the intellectual property context, courts have recognised the importance of originality and authorship in copyright law, as seen in a 2007 Supreme Court judgement in *Eastern Book Company v. D.B. Modak*[7], The judgment emphasises human authorship and minimal creativity, raising unresolved questions about the ownership and liability implications of AI-generated works that lack direct human authorship.

These decisions illustrate that while existing legal principles provide a foundation, they are not designed to address the distributed and autonomous nature of AI systems. This creates an environment where contractual allocation of risk becomes central.

Ongoing Litigations

In an ongoing litigation before the Delhi High Court (*ANI Media Pvt. Ltd. v. OPEN AI Opco LLC*)[8], a court in India for the first time is facing the challenge of deciding on the difficult issues that emerge from the use of copyrighted materials by a leading frontier AI model.[9] The Court has already heard final arguments in the matter and the judgement is reserved. It would be interesting to see the court's stand in the final judgement particularly in light of the fact that a DPIIT committee is already examining the legality of using copyrighted material in training AI systems and its implications.

II. Contractual Mechanisms for AI Risks Allocation

Parties are increasingly relying on detailed contractual provisions to manage liability exposure arising from AI outputs. These provisions typically operate across four key dimensions.

First, representations and warranties are used to address the integrity of training data, compliance with applicable law, and the absence of known infringement risks. Upstream vendors may limit such representations, particularly in relation to outputs, given the probabilistic nature of AI systems.

Second, indemnity clauses are used to allocate specific risks, including third-party intellectual property claims, regulatory penalties, and damages arising from harmful outputs. The scope of indemnity often reflects bargaining power, with model providers seeking to narrow coverage and downstream vendors attempting to pass through liability.

Third, limitation of liability clauses are heavily negotiated in AI contracts. Vendors typically seek to exclude liability for indirect damages and cap overall exposure. However, customers may resist broad exclusions where AI outputs are integrated into critical business functions.

Indian courts have generally upheld contractual limitation clauses subject to reasonableness and public policy constraints. In Kailash Nath Associates v. Delhi Development Authority[10], the Supreme Court emphasised that damages must reflect actual loss, reinforcing the importance of carefully structuring liability caps and damage

[6] AIR 2020 Supreme Court 350.
[7] AIR 2008 Supreme Court 809.
[8] CS(COMM) 1028/2024
[9] First Order capturing the facts and issues framed by the Court.
[10] 2015 (4) SCC 136

frameworks in contracts involving uncertain AI risks.

Fourth, use restrictions and acceptable use policies are deployed to shift responsibility to the customer. These provisions may prohibit certain use cases, such as high-risk decision-making or regulated activities, thereby limiting vendor exposure.

III. Challenges in Multi-Layered AI Risk Transfer

While contractual mechanisms provide a framework for allocating risk, their effectiveness is constrained by structural challenges inherent in AI supply chains.

One key issue is the enforceability of back-to-back indemnities. A downstream vendor may assume liability towards its customer but face limitations in recovering equivalent amounts from upstream providers due to narrower indemnity coverage or liability caps.

Another challenge is information asymmetry. Upstream providers have greater visibility into model training and limitations, while downstream vendors bear responsibility for implementation and deployment. This imbalance affects the ability of parties to negotiate informed contractual protections.

There is also a risk of regulatory override. Contractual allocation of liability does not bind regulators or courts, which may impose liability based on statutory obligations or public policy considerations. Indian jurisprudence has consistently recognised that contractual terms cannot override statutory mandates or public policy concerns, particularly in contexts involving consumer harm.

IV. Emerging Global Approaches and Their Contractual Implications

International regulatory developments indicate a trend towards imposing obligations on multiple actors within the AI value chain. Frameworks such as the European Union's AI regulatory regime distinguish between providers, deployers, and importers, assigning responsibilities accordingly. While these frameworks are jurisdiction-specific and subject to phased implementation, they signal a move towards more structured allocation of responsibility.

For Indian entities participating in global AI supply chains, these developments have contractual implications. Vendors may be required to provide compliance assurances aligned with foreign regulatory standards, and customers may demand contractual commitments reflecting these obligations.

V. Practical Contract Drafting Recommendations

From a transactional perspective, contracts involving AI systems require a departure from standard software templates. Parties should consider clearly defining the scope of AI functionality, identifying high-risk use cases, and aligning contractual obligations with the technical architecture of the system.

Indemnity provisions should be carefully scoped to reflect realistic risk allocation, taking into account upstream dependencies. Liability caps should be evaluated in the context of potential exposure, particularly where AI outputs influence business-critical decisions.

In addition, audit rights, transparency obligations, and incident response mechanisms may be incorporated to address ongoing compliance risks. These provisions can help bridge the gap between contractual allocation and operational realities.

VI. Conclusion

The question of liability for AI outputs is unlikely to be resolved solely through statutory intervention in the near term. In the interim, contractual arrangements will continue to serve as the primary mechanism for allocating risk across multi-layered AI supply chains. However, the effectiveness of such arrangements depends on their ability to

reflect the technical and operational complexities of AI systems.

Indian case law provides important principles on intermediary liability, contractual enforcement, and intellectual property, but does not directly resolve the challenges posed by AI systems. As AI adoption accelerates, the precision of contractual risk allocation, informed by both domestic jurisprudence and emerging global regulatory trends, will increasingly determine the legal and commercial viability of AI-enabled services.