



LAWFOYER INTERNATIONAL JOURNAL OF DOCTRINAL LEGAL RESEARCH

[ISSN: 2583-7753]

Volume 4 | Issue 3

2026

DOI: <https://doi.org/10.70183/lijdlr.2026.v04.306>

© 2026 LawFoyer International Journal of Doctrinal Legal Research

Follow this and additional research works at: www.lijdlr.com

Under the Platform of LawFoyer – www.lawfoyer.in

After careful consideration, the editorial board of LawFoyer International Journal of Doctrinal Legal Research has decided to publish this submission as part of the publication.

In case of any suggestions or complaints, kindly contact (info.lijdlr@gmail.com)

To submit your Manuscript for Publication in the LawFoyer International Journal of Doctrinal Legal Research, To submit your Manuscript [Click here](#)

ORCHESTRATING ACCOUNTABILITY: LIABILITY OF THE SYSTEM DEPLOYER IN AGENTIC AI UNDER INDIAN LAW

Samarth Udasin¹, Divyansh Rai² & Priyanshu Tripathi³

I. ABSTRACT

Agentic AI systems do not merely respond to queries they act, decompose goals, call APIs, enter transactions, and produce irreversible real-world consequences autonomously. That shift from tool to actor breaks every assumption on which India's intermediary liability law rests. This paper argues that existing Indian law already assigns liability to the right party: the System Deployer, the entity that selects the AI model, integrates tools, engineers data flows, and directs the agent's consequential actions. The principle is straightforward liability follows control. Section 79 of the Information Technology Act, 2000 shows that the Deployer's active orchestration role disqualifies it from safe harbour: while it qualifies as an intermediary under the statutory definition, it fails every substantive condition that Section 79(2) requires before that status translates into immunity. The Digital Personal Data Protection Act, 2023 confirms this, designating the Deployer a Data Fiduciary who will carry non-delegable duties once Section 8 commences. SEBI's algorithmic-trading framework demonstrates that Indian regulators already impose full liability on deploying entities regardless of algorithmic opacity. Tort principles of negligence and absolute liability complete the picture. Comparative analysis of the EU Digital Services Act and L'Oréal v eBay confirms global convergence. The paper concludes with a governance framework grounded in the RBI FREE-AI Report. No new legislation is needed.

II. KEYWORDS

¹ 5th year student of B.Sc.LL.B (Hons.) [Cyber Security] at National Law Institute University, Bhopal (India). Email: samarthudasin.bscllb@nliu.ac.in

² 5th year student of B.Sc.LL.B (Hons.) [Cyber Security] at National Law Institute University, Bhopal (India). Email: divyanshrai.bscllb@nliu.ac.in

³ 5th year student of B.A. LL.B (Hons.) at National Law Institute University, Bhopal (India). Email: priyanshutripathi.ballb@nliu.ac.in

Agentic AI, System Deployer, Section 79 IT Act, DPDP Act, Data Fiduciary, Intermediary Liability, SEBI Algorithmic Trading, Free-AI Report, Liability Follows Control.

I. INTRODUCTION

A. The Machine That Acts

Something fundamental changed when AI systems gained the ability to act. A language model that answers questions is a tool. An agentic AI that breaks a goal into steps, calls banking APIs, books travel, and executes trades is something else entirely. It is an autonomous actor, to which the Indian law has not caught up.

Bathae captured the technical turning point precisely: these systems are no longer executing pre-written instructions but arriving at dynamic solutions based on patterns in data that humans cannot perceive.⁴ Idowu places them in their commercial context, increasingly deployed on cloud infrastructure for its scalability and integration with big-data systems, autonomous agents have become the default interface for consequential digital action.⁵

The legal problem is not that AI is complex. The legal problem is that traditional liability doctrine assumed a human actor. It assumed someone who could form intent, foresee consequences, and be cross-examined. An autonomous AI offers none of that. And when it causes harm, everyone in the supply chain, the company that built the model, the company that deployed it, the user who gave it a goal, all points at someone else. This paper calls that diffusion of responsibility an ecosystem of blame.

The RBI's FREE-AI Committee understood the stakes. Its Sutra 5 states the governing principle without qualification: 'Accountability rests with the entities deploying AI. Accountability cannot be delegated to the model and underlying algorithm.'⁶ Sutra 1

⁴ Yavar Bathae, *The Artificial Intelligence Black Box and the Failure of Intent and Causation*, Vol. 31, HARV. J.L. & TECH., 889, 891–92 (2018).

⁵ Emmanuel Idowu, *Responsibility Attribution in Autonomous AI Agents Deployed on Cloud Infrastructure*, PREPRINTS.ORG (2025) available at <https://doi.org/10.20944/preprints202507.1337.v1>.

⁶ RESERVE BANK OF INDIA, *Report of the Committee on Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI) in the Financial Sector*, 39 (2025) ('FREE-AI Report').

adds the foundation: ‘Trust is non-negotiable and should remain uncompromised.’⁷ These are not aspirational guidelines. They are operational directives, and they already exist in the statutory text, if one reads it honestly.

B. Why the System Deployer?

Two rival candidates for liability must be addressed before the argument can proceed.

The first is the Model Developer, which is the company that built the large language model. The Developer is a plausible defendant where the harm flows from the model’s inherent design. Bathaee’s analysis is clear: when autonomous AI causes harm in a mission-critical setting, the ‘risks of deploying a black-box AI autonomously in such settings should fall on the AI’s user or creator.’⁸ But the Model Developer creates a general-purpose capability. It does not choose which APIs to integrate, which workflows to build, or which commercial context to operate in. Those decisions belong to the Deployer. Holding the Developer liable for a Deployer’s operational choices would be like suing the car manufacturer every time a taxi driver causes an accident.

The second candidate is the User, which is the individual who gave the agent its goal. Users can of course be liable where they misuse the technology.⁹ But the User issues only a natural-language objective: ‘book the cheapest flight.’ Every subsequent decision, such as which API to call, how to structure the request, which third-party services to authorize, is made entirely by the Deployer’s orchestration layer. Placing primary liability on the User perpetuates what this paper terms the moral crumple zone: harm that the architecture routes to the party least able to prevent it.

The System Deployer, by contrast, controls all the decisions that produce harm. It selects the AI model. It integrates the APIs. It engineers the data pipelines. It architects the

⁷ FREE-AI Report, *supra* note 3, Sutra 1, 38.

⁸ Bathaee, *supra* note 1, 934–36.

⁹ Shreya Goyal, *AI and Legal Liability: Who Is Responsible When Machines Go Wrong?*, Vol. 7, Indian J.L. & Legal Rsch., 9404, 9407 (2025).

workflow that translates a human goal into a real-world action with legal consequences. Liability must follow that control.

C. The Policy Stakes

The CCI's market study found that 37% of startup respondents identified AI-facilitated collusion as a significant concern, and 22% identified new entry barriers.¹⁰ A liability vacuum amplifies both problems. Without a clear accountability framework, dominant incumbents can deploy powerful agents whose failures are absorbed by users and third parties, while their market positions are reinforced. The CCI's mandate requires that India resolve this question and resolve it in a way that assigns costs to the party best placed to prevent harm.

D. Research Objectives

1. To identify the party that properly bears liability for harms caused by agentic AI systems under existing Indian laws.
2. To test the System Deployer's claim to safe harbour under Section 79 of the Information Technology Act, 2000.
3. To evaluate whether corroborating liability frameworks, the DPDP Act, 2023, SEBI's algorithmic-trading regime, and tort law, converge with the IT Act's structural argument.
4. To propose a governance framework for System Deployers grounded in existing Indian regulatory guidance.

E. RESEARCH QUESTIONS

1. Does the System Deployer qualify as an 'intermediary' under Section 2(1)(w) of the Information Technology Act, 2000?

¹⁰ Competition Commission of India, Market Study on Artificial Intelligence and Competition, conducted in partnership with Management Development Institute Society, Gurugram, 52, tbl. 4.1 (2025) ('CCI Market Study').

2. Can the Deployer satisfy the cumulative conditions of Section 79(2)(b) of the IT Act so as to claim safe harbour?
3. Do the DPDP Act, 2023, SEBI's algorithmic-trading framework, and tort law converge with the structural argument advanced under the IT Act?
4. What governance framework should apply to System Deployers of agentic AI under Indian law?

F. RESEARCH METHODOLOGY

This paper follows a doctrinal, black-letter methodology. It proceeds by close textual analysis of the primary statutory sources governing intermediary liability, data protection, and algorithmic trading in India, principally the Information Technology Act, 2000 and its 2021 and 2026 Amendment Rules, the Digital Personal Data Protection Act, 2023 and the Digital Personal Data Protection Rules, 2025, and SEBI's circulars on algorithmic trading, read alongside the judicial decisions that interpret them, from *Shreya Singhal v. Union of India* through the Delhi High Court's active-participant jurisprudence. Regulatory instruments that do not carry the force of statute, including the RBI's FREE-AI Committee Report and MeitY's India AI Governance Guidelines, are treated as persuasive policy material illustrating how Indian regulators already apply a deployer-centric liability logic in practice. The comparative discussion of the European Union's Digital Services Act and AI Act is undertaken on a functional basis: these instruments are examined not as binding authority but as evidence of convergent reasoning in a mature comparable jurisdiction, tested against the doctrinal claims developed from Indian sources. Throughout, the analysis is normative as well as descriptive: it does not merely report the state of the law but argues, from within existing doctrine, for a specific allocation of liability.

III. THE TECHNICAL PREDICATE: WHAT A DEPLOYER ACTUALLY DOES

A. The Four-Stage Cycle

An agentic system operates through a repeating four-stage cycle that this paper calls Goal-Plan-Act-Reflect, or GPAR. Understanding this cycle is essential to understanding why the Deployer, and not the Developer or the User holds the decisive control.

1. **Goal:** The User states a high-level objective in natural language: ‘Rebalance my portfolio to reduce equity exposure.’ The User’s contribution ends here.
2. **Plan:** The AI model, acting within the Deployer’s orchestration layer, breaks the goal into a structured sequence of steps. This is internal to the Deployer’s system. The User sees none of it.
3. **Act:** The agent executes each step by calling external tools via APIs. This is the legally decisive stage. Here, digital instruction becomes real-world action: a financial transaction, a contractual commitment, a data access. The Deployer has already selected these APIs, granted the agent its credentials, and configured the scope of permitted action.
4. **Reflect:** The agent processes feedback, learns from the result, and refines its approach, increasing its operational independence over time.

The RBI’s FREE-AI Committee identified this autonomous cycle as a distinct governance challenge, noting the move toward systems ‘acting independently in financial decision-making’ and the complexity that Agent-to-Agent communication adds.¹¹ The legally critical point is that the Act stage, where abstract goal becomes concrete consequence, is a stage entirely engineered by the Deployer.

B. The Black Box and Its Legal Consequence

Agentic AI systems are black boxes. Bathaee explains why: deep neural networks distribute information across thousands of artificial neurons in a way that ‘may be as impenetrable as that of the human brain.’¹² Even their creators cannot trace the path from

¹¹ FREE-AI Report, *supra* note 3, 22–24.

¹² Bathaee, *supra* note 1, 892.

input to output. The system's reasoning is opaque, not because it was designed to deceive, but because complexity at this scale is inherently opaque.

This opacity has a direct legal consequence. It defeats the traditional route to liability: showing that a specific human intent caused a specific outcome.¹³ Goyal identifies this as one of the three core weaknesses of existing frameworks: AI systems 'behave as black boxes that cannot provide explanations for their decisions.'¹⁴ When the traditional route is blocked, the law must take the structural route, identifying the party that controlled the system's design and deployment, not the party whose internal reasoning the harm traceable to. That party is the Deployer.

India's regulators have already accepted this logic. SEBI classifies black-box algorithms, those 'where the logic is not known to the user and is not replicable', as a permitted category.¹⁵ It does not exempt them from liability. It does the opposite: it fixes full, non-delegable liability on the broker that chose to deploy them. Opacity is a managed risk, not a defense.

C. The API as a Grant of Authority

When a Deployer integrates a banking API into an agentic workflow, it does not merely enable a technical connection. It grants the AI agent authority to act on its behalf: to authenticate, to instruct, to commit. Idowu's responsibility attribution framework captures the significance: API integration is a deliberate architectural decision that confers operational capability.¹⁶ This paper terms it a digital power of attorney, and it is a power that only the Deployer can grant or revoke.

IV. THE DOCTRINAL CORE: SECTION 79 OF THE IT ACT

The System Deployer's best defense under Indian statute is Section 79 of the IT Act.¹⁷ Section 79 gives intermediaries a safe harbor: conditional immunity from liability for

¹³ Bathaee, *supra* note 1, 893.

¹⁴ Goyal, *supra* note 6, 9409.

¹⁵ SEBI Algo Circular, *infra* note 37, ¶2(b)(ii).

¹⁶ Idowu, *supra* note 2, 2.

¹⁷ The Information Technology Act, 2000, §79 [hereinafter IT Act].

unlawful third-party content they host or transmit. This Chapter shows that the Deployer qualifies as an intermediary under the statutory definition, but that it still cannot claim the safe harbor, because it fails every substantive condition the provision requires.

A. How Section 79 Works: The Threshold Question

The analysis proceeds in two steps. The first question is whether the System Deployer qualifies as an ‘intermediary’ at all. The second, and more consequential, question is whether a Deployer that does qualify can still claim the safe harbor that Section 79 provides. Both questions must be answered. As this section shows, the Deployer clears the first hurdle but fails decisively at the second.

Section 2(1)(w) defines an intermediary as any person who, on behalf of another, ‘receives, stores or transmits’ an electronic record or provides any service with respect to that record.¹⁸ That definition is deliberately broad. A System Deployer receives the user’s natural-language goal, processes and stores intermediate outputs across the agentic workflow, and transmits structured API calls to external services at each step occurring ‘on behalf of’ the user who issued the initiating instruction. Gupta and Srinivasan confirm that the amended Section 2(1)(w) was designed to extend beyond network service providers to encompass the full range of entities providing services ‘with respect to’ electronic records, however enhanced their functionality.¹⁹ Devadasan similarly acknowledges that the definition is capacious enough to cover entities well beyond passive conduits.²⁰ The System Deployer, on this analysis, falls within the statutory definition of an intermediary.

But meeting the definition is only the first step. Section 79 does not protect every intermediary. It protects only those that satisfy the three cumulative conditions in Section 79(2): the intermediary must not have initiated the transmission, must not have selected

¹⁸ *Id.*, §2(1)(w).

¹⁹ Indranath Gupta & Lakshmi Srinivasan, *Evolving Scope of Intermediary Liability in India*, Vol. 37, INT’L REV. L. COMPUTERS & TECH., 294, 296–300 (2023).

²⁰ Vasudev Devadasan, *Conceptualizing India’s Safe Harbour in the Era of Platform Governance*, Vol. 19, INDIAN J.L. & TECH., 8, 19–21 (2024).

the receiver, and must not have selected or modified the information. These conditions operationalize a single underlying concept, passive neutrality. An intermediary that fails any one of them cannot claim the safe harbor, regardless of whether it qualifies as an intermediary under Section 2(1)(w). The System Deployer satisfies the definition but violates all three conditions. It is therefore an intermediary without immunity.

B. The Shreya Singhal Standard and Its Limits

Shreya Singhal remains the governing authority on safe-harbor conditions.²¹ The Supreme Court held that forcing platforms to proactively police content would chill expression.²² Actual knowledge of unlawful content, which is the trigger for liability, arises only from a court order or government notification, not from private complaints or the platform's own assessment.²³ This standard was affirmed in Google India²⁴ and operationalized in the IT Rules 2021.²⁵

Shreya Singhal governs the conditions under which a qualifying intermediary forfeits protection. But qualifying as an intermediary is merely the starting point: Section 79 then imposes three further conditions in subsection (2)(b) that must all be met before safe harbor is available. A court must examine these conditions independently of whether the defendant meets the definition in Section 2(1)(w). For the System Deployer, as the following sections demonstrate, this is where the safe harbor collapses.

C. The Active-Participant Doctrine: MySpace and Amazon

The System Deployer is an intermediary under Section 2(1)(w), but it cannot satisfy the active-participation conditions that Section 79(2) requires. Indian courts have consistently held that exercising control over content or transactions defeats any claim to safe harbor.

In *MySpace v Super Cassettes*, the Delhi High Court's Division Bench ultimately strengthened intermediary safe harbour in the copyright context by holding that liability

²¹ *Shreya Singhal v. Union of India*, (2015) 5 SCC 1.

²² *Id.*, ¶¶116–18.

²³ *Id.*, ¶119.

²⁴ *Google India Pvt. Ltd. v. Visaka Industries Ltd.*, (2020) 4 SCC 162 (Supreme Court of India), ¶53.

²⁵ The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021, R. 3(1)(d).

turns on specific or actual knowledge of infringing material, and not on general awareness of infringement or the mere existence of platform safeguards.²⁶ The Court also examined the statutory conditions in Section 79(2)(b), including whether the intermediary initiates the transmission, selects the receiver, or modifies the information, and held on the facts that MySpace satisfied those conditions.²⁷ MySpace therefore should not be overstated as a freestanding authority for the proposition that “active participation” alone defeats immunity. Its relevance here is narrower: it confirms that Section 79(2)(b) requires a functional inquiry into initiation, receiver selection, and modification. The Deployer’s position is materially different from MySpace because an agentic AI orchestration layer does not merely host user-generated content; it converts a natural-language instruction into fresh, structured API calls and directs consequential external action.

The e-commerce jurisprudence complicates, but does not defeat, the analysis. The Single Judge in *Amway India Enterprises Pvt. Ltd. v. IMG Technologies Pvt. Ltd. & Anr.* found that warehousing, packaging, and payment-collection services constituted active participation sufficient to defeat passive-intermediary status.²⁸ On appeal, however, the Division Bench in *Amazon Seller Services Pvt. Ltd. v. Amway India Enterprises Pvt. Ltd.* set aside that reasoning, holding that Section 79 draws no textual distinction between ‘active’ and ‘passive’ intermediaries, and that logistical or qualitative involvement, without more, does not by itself forfeit safe harbor.²⁹ This paper acknowledges that the Division Bench’s holding cuts against, rather than supports, a judicially-created active-participant test. But the Deployer’s claim to safe harbor does not rest on that judicial doctrine. It rests on Section 79(2)(b)’s own conjunctive statutory conditions, examined independently in Section III.D below, conditions the Division Bench did not disturb, and

²⁶ *MySpace Inc. v. Super Cassettes Industries Ltd.*, 236 (2017) DLT 478, ¶74.

²⁷ *Id.*, ¶¶68–72, 74.

²⁸ *Amway India Enterprises Pvt. Ltd. v. IMG Technologies Pvt. Ltd. & Anr.*, (2019) 260 DLT 690, ¶¶38–41.

²⁹ *Amazon Seller Services Pvt. Ltd. v. Amway India Enterprises Pvt. Ltd.*, 2020 SCC OnLine Del 454, ¶¶44–47.

which the Deployer fails on their own terms regardless of whether ‘active participation’ survives as a freestanding judicial test.

D. Failure of the Statutory Conditions

Even treating a Deployer as an intermediary, it fails all three conditions for safe harbor under Section 79(2)(b).³⁰ These conditions are conjunctive: all three must be satisfied. The Deployer satisfies none.

1. **On initiation:** the Deployer’s system initiates new transmissions that would not exist but for its orchestration layer. When a user says ‘find the cheapest flight,’ the system does not relay that request, it generates new, structured API queries to airline systems.
2. **On receiver selection:** the Deployer’s system actively chooses between available APIs and external services based on its internal logic. This is deliberate selection, not neutral routing.
3. **On modification:** the Deployer’s system transforms a natural-language goal into a structured, machine-executable plan of API calls. This is not reformatting. It is a substantive reconstitution of the user’s intent into a specific legal act.

E. A Framework Built for a Different World

Section 79’s notice-and-takedown architecture, a court order triggers a 36-hour removal window, was designed for static content on a hosting platform. It cannot govern a system whose harms materialize instantly and irreversibly. More fundamentally, content removal is not the remedy for a structural flaw in an autonomous actor. The harms that agentic AI causes flow from its deployment architecture, not from any specific piece of third-party content. Consistent application of the existing statute produces the right result: the entity that controls the architecture bears the liability.

F. The 2026 Amendment Rules: A New Layer, Not a New Answer

³⁰ IT Act, §§79(2)(b)(i)–(iii).

On 10 February 2026, the Ministry of Electronics and Information Technology notified the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026, effective from 20 February 2026, which amend the 2021 Rules to regulate ‘synthetically generated information’ (SGI), audio, visual, or audio-visual content artificially or algorithmically created or altered to appear real and indistinguishable from a natural person or event.³¹ The 2026 Rules impose new due-diligence obligations under Rule 3(3): intermediaries that enable or facilitate the creation or dissemination of SGI must deploy technical measures to prevent unlawful SGI and prohibit specified high-risk categories, and Significant Social Media Intermediaries must verify and label such content, on pain of losing safe harbor for non-compliance. The Rules also compress the takedown window from the 36 hours contemplated under the 2021 framework to three hours for court- or government-directed removal, and two hours for non-consensual intimate imagery and deepfakes.

These changes reinforce, rather than disturb, this Chapter’s argument. Where an agentic Deployer’s workflow generates or manipulates SGI, the 2026 Rules layer additional, content-specific due-diligence obligations onto the Section 79(2)(b) analysis developed above, obligations a Deployer that already fails the conjunctive statutory conditions cannot escape by pointing to compliance with the narrower SGI regime. And the drastic compression of the takedown window, itself a legislative acknowledgment that synthetic and AI-mediated harms move faster than the 2021 framework assumed, underscores Section III.E’s point above: a notice-and-takedown architecture, however accelerated, remains poorly suited to harms that an agentic system’s own architecture makes instantaneous and irreversible.

V. CORROBORATING FRAMEWORKS

The IT Act analysis does not stand alone. Three independent legal frameworks converge on the same conclusion: liability follows control, and control rests with the Deployer.

³¹ The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026, R. 2(1)(wa), R. 3(3).

A. The DPDP Act: The Deployer as Data Fiduciary

Agentic systems process personal data constantly. Every goal they pursue, every API they call, every action they take generates, transmits, or transforms personal data. The Digital Personal Data Protection Act 2023 therefore applies, and it places the System Deployer squarely in the role of Data Fiduciary.

Section 2(i) of the DPDP Act defines a Data Fiduciary as the entity that ‘determines the purpose and means of processing of personal data.’³² Malhotra and Misra confirm that this concept is the Act’s primary accountability mechanism.³³ The Deployer satisfies the definition. The User determines the purpose in the broadest sense, ‘book a flight.’ But the Deployer determines the means in the technical and operational sense that Section 2(i) intends: the AI model selected, the APIs integrated, the data pipelines configured, the memory architecture designed. The difference is between abstract intent and engineered execution. The Deployer supplies the latter.

This status will be non-delegable once the relevant provisions commence. Section 8(1) makes the Fiduciary responsible for all processing undertaken by it or on its behalf.³⁴ The resulting obligations are concrete: accuracy (Section 8(3));³⁵ security safeguards (Section 8(5));³⁶ breach notification (Section 8(6)).³⁷ Goyal notes that the DPDP Act does not yet explicitly address autonomous AI liability.³⁸ But purposive interpretation fills that gap: the entity that engineers the means of processing is the Data Fiduciary, regardless of whether it also wrote the underlying model.

A timing caveat is necessary. On 13 November 2025, MeitY notified the Digital Personal Data Protection Rules, 2025 alongside the phased commencement of the DPDP Act

³² The Digital Personal Data Protection Act, 2023, §2(i) [hereinafter DPDP Act].

³³ Palak Malhotra & Amita Misra, *Accountability and Responsibility of Artificial Intelligence Decision-Making Models in Indian Policy Landscape* in Proceedings of the AISafety Workshop at the International Joint Conference on Artificial Intelligence (2022).

³⁴ DPDP Act, §8(1).

³⁵ DPDP Act, §8(3).

³⁶ DPDP Act, §8(5).

³⁷ DPDP Act, §8(6).

³⁸ Goyal, *supra* note 6, 9407–08.

itself.³⁹ That notification brought only the Data Protection Board's constitutive provisions (Sections 18–26) and certain ancillary provisions into immediate effect; the substantive obligations in Sections 3–17, including Section 8 in its entirety, commence eighteen months later, around mid-May 2027. Section 2(i)'s definition of 'Data Fiduciary,' on which this Section's analysis rests, is already in force, so the identification of the Deployer as Data Fiduciary holds today. But the non-delegable duties described above, accuracy, security safeguards, and breach notification under Section 8, are not yet directly enforceable. The DPDP Act's present contribution to the liability case is therefore its settled definitional architecture and its clearly signalled direction of travel; its contribution of concrete, enforceable non-delegable duties becomes operative only once Section 8 commences.

B. SEBI's Algorithmic-Trading Framework: The Clearest Precedent

India already has a regulatory regime that assigns liability to the deploying entity regardless of algorithmic opacity. SEBI's Circular on Algorithmic Trading is the clearest precedent available.⁴⁰

The Circular's principal holding is unambiguous: 'For all API-based trading, the broker shall be the principal while any algo provider, strategy provider or vendor providing the trading algorithm shall act as its agent.'⁴¹ The deploying broker bears full liability. It cannot transfer that liability to the algorithm vendor, even where the algorithm is a black box whose logic it cannot inspect. SEBI then adds operational safeguards such as kill switches and simulation testing, that the Circular requires before deployment.⁴² And it closes with a direct statement of the point: 'Brokers shall be solely responsible for handling investor grievances related to algo trading.'⁴³

³⁹ Ministry of Electronics and Information Technology, Notification bringing into force the Digital Personal Data Protection Act, 2023 and notifying the Digital Personal Data Protection Rules, 2025 (Nov. 13, 2025), read with DPDP Act, §1(2).

⁴⁰ Securities & Exchange Board of India, *Safer Participation of Retail Investors in Algorithmic Trading*, Circular No. SEBI/HO/MIRSD/MIRSD-PoD/P/CIR/2025/0000013 (Feb. 4, 2025) ('SEBI Algo Circular').

⁴¹ SEBI Algo Circular, ¶4(a).

⁴² SEBI Algo Circular, ¶¶6(c), 7(b).

⁴³ SEBI Algo Circular, ¶4(c).

This framework makes the argument by analogy exact. The SEBI Circular is not a perfect template for all agentic AI, it addresses financial trading specifically. But its underlying principle could not be broader: the entity that chooses to deploy a complex autonomous system for commercial gain bears the legal consequences, whatever the system's internal logic.

C. Tort Law: Three Routes to the Same Conclusion

Independent of statute, the Deployer faces liability under three overlapping tort principles.

1. Negligence

The neighbor principle, established in *Donoghue v Stevenson*⁴⁴ and affirmed by the Indian Supreme Court in *Chairman, Railway Board v Chandrima Das*,⁴⁵ requires that one take reasonable care to avoid acts or omissions that one can foresee would likely injure one's neighbor. The System Deployer is a commercial entity that deploys powerful software capable of causing irreversible real-world harm. Its users and third parties affected by its agent's actions are plainly its neighbors. The standard of care is calibrated to the known risks, and those risks are substantial.

A breach occurs when the Deployer fails to implement robust pre-deployment testing, build meaningful human-oversight checkpoints, or restrict the agent's scope of action to what the risk profile justifies. The unpredictability of the AI is not a defense to this duty. It is precisely the source of the duty.

2. Absolute Liability

Where an agentic deployment is capable of causing catastrophic, irreversible harm to third parties, the MC Mehta absolute-liability doctrine may apply.⁴⁶ Under MC Mehta,

⁴⁴ *Donoghue v. Stevenson*, [1932] AC 562, 580.

⁴⁵ *Chairman, Railway Board v. Chandrima Das*, (2000) 2 SCC 465, ¶¶16–18.

⁴⁶ *M.C. Mehta v. Union of India (Oleum Gas Leak)*, (1987) 1 SCC 395, ¶¶31–33.

an enterprise engaged in a hazardous activity bears a non-delegable duty to ensure no harm results, without the exceptions available under *Rylands v Fletcher*.

Bathae argues that strict liability is a poor fit for AI where harm flows from unpredictable outputs, because one cannot take the targeted precautions that strict liability is designed to incentivize.⁴⁷ This paper accepts that argument, however, it applies narrowly. MC Mehta's rule attaches to the decision to deploy a hazardous system, not to each individual unpredictable output. The deployment decision is fully within the Deployer's foresight and control.

3. The Moral Crumple Zone

Perhaps the most important tort principle here is structural: courts must not allow liability to rest on the party least able to prevent the harm. The User who issues a natural-language goal has no visibility into the agent's architecture, no access to its decision process, and no contractual power to alter its design.

Bathae's sliding-scale framework makes this point in doctrinal terms: liability should reflect 'the degree of the AI's transparency, the constraints its creators or users placed on it, and the vigilance used to monitor its conduct.'⁴⁸ Idowu's attribution model reaches the same conclusion from an architectural perspective: the deploying entity is 'the central node of accountability.'⁴⁹ Both analyses lead to the Deployer. Tort law must follow.

VI. COMPARATIVE AND COMPETITION PERSPECTIVES

A. The EU Digital Services Act: The Same Principle, Codified

The European Union's Digital Services Act reached the same legal conclusion through the same reasoning, and codified it.

⁴⁷ Bathae, *supra* note 1, 931–32.

⁴⁸ Bathae, *supra* note 1, 933–36.

⁴⁹ Idowu, *supra* note 2, 3.

Article 6 of the DSA provides a hosting exemption for intermediary services.⁵⁰ Recital 18, which is the interpretive preamble, explains when the exemption does not apply: where the provider ‘plays an active role of such a kind as to give it knowledge of, or control over’ the relevant information.⁵¹ The operative legal test, however, comes from the CJEU’s ruling in *L’Oréal v eBay*. The Court held that a marketplace operator providing assistance that optimizes the presentation of sellers’ offers has played an ‘active role’ and cannot claim the exemption.⁵² The test is functional control and not mere knowledge.

Devadasan’s analysis draws an explicit parallel to India’s evolving platform-governance approach, arguing that the forthcoming Digital India Bill will need to distinguish between pre-conditions to safe harbor and direct statutory duties.⁵³ The System Deployer’s orchestration function of decomposing goals, selecting APIs, directing actions, is the paradigm case of the ‘active role’ that both the DSA framework and the CJEU’s jurisprudence treat as incompatible with immunity.

B. The EU AI Act: Deployer-Centric Liability as Global Standard

The EU AI Act goes further, establishing a risk-based framework that places significant operational compliance burdens on the deploying entity.⁵⁴ Article 26 requires deployers of high-risk AI systems to implement appropriate technical and organisational measures, ensure human oversight, monitor use, and maintain logs where such logs are under their control.⁵⁵ A timing qualification is necessary: Regulation (EU) 2026/1744, the Digital Omnibus on AI, has deferred the application of Chapter III, Sections 1, 2 and 3, including the high-risk obligations applicable to Article 26 deployers, to 2 December 2027 for Annex

⁵⁰ Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022, *Digital Services Act*, O.J. E.U. L277/1, Art. 6 (‘DSA’).

⁵¹ DSA, *supra* note 47, Recital 18.

⁵² *L’Oréal SA v. eBay International AG*, Case C-324/09, July 12, 2011 (Ct. J. E.U.) ¶¶116–20.

⁵³ Devadasan, *supra* note 17, 40–46.

⁵⁴ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024, *Artificial Intelligence Act*, O.J. L, 2024/1689, 12.7.2024, Art. 26 (‘EU AI Act’).

⁵⁵ EU AI Act, *supra* note 51, Art. 26.

III high-risk systems and 2 August 2028 for Annex I high-risk systems.⁵⁶ This deferral does not undermine the comparative point. It clarifies that deployer-centric responsibility remains the EU's legislative choice, although the enforceability timeline for the relevant high-risk obligations has been extended.

C. The Competition Argument

Assigning liability to the Deployer is also a pro-competitive choice.

The CCI, in its 2025 Market Study on Artificial Intelligence and Competition, found that dominance in the AI value chain, built on proprietary cloud infrastructure, foundation models, and datasets, creates structural entry barriers.⁵⁷ It also identified hub-and-spoke coordination risks: a common AI platform serving multiple competing firms could facilitate tacit price-signaling or output coordination.⁵⁸ An agentic Deployer whose platform serves multiple downstream competitors could fall within Section 3 of the Competition Act 2002.⁵⁹

A liability vacuum amplifies these harms. Without accountability, dominant incumbents can deploy powerful agents whose failures are absorbed by users and third parties, while their structural advantages accumulate. Fixing liability at the Deployer level internalizes those costs and creates a direct incentive to design safer, more transparent systems. The threshold objection, that Deployer liability might itself become a barrier to entry, is answered by graduated regulatory accommodation for smaller entities, as FREE-AI Recommendation 7 already proposes. The liability principle itself must hold.

VII. GOVERNANCE FRAMEWORK AND CONCLUSION

A. Five Arguments, One Conclusion

Five independent lines of analysis all reach the same destination.

⁵⁶ Regulation (EU) 2026/1744 of the European Parliament and of the Council on the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI), O.J. L, 2026/1744, 27.7.2026, Art. 1.

⁵⁷ CCI Market Study, *supra* note 7, 61–63.

⁵⁸ CCI Market Study, *supra* note 7, 54–56.

⁵⁹ The Competition Act, 2002, §3.

1. **Technically:** the Deployer controls everything that matters. It selects the model, integrates the APIs, engineers the pipelines, and architects the workflow that turns human intent into consequential action.
2. **Doctrinally:** the IT Act's safe harbor is unavailable. The Deployer qualifies as an intermediary under Section 2(1)(w) but it fails all three of the substantive conditions in Section 79(2)(b) that are required before that status attracts immunity.
3. **Statutorily:** the DPDP Act makes the Deployer a Data Fiduciary, who is the entity that determines the means of processing, and will carry non-delegable obligations for security, accuracy, and breach notification once Section 8 commences.
4. **Regulatorily:** SEBI's algorithmic-trading framework has already proved the model. The deploying broker is the principal. Opacity is a managed risk, not a defense.
5. **In tort:** negligence, the absolute-liability doctrine of MC Mehta, and the need to prevent the moral crumple zone all point to the Deployer as the party best positioned to prevent harm, and therefore the party that must bear it.

The comparative analysis from the EU and the competition analysis from the CCI confirm the same conclusion. No new statute is needed. The existing law, read honestly, is enough.

B. Suggestions and Recommendations

This Section sets out the paper's suggestions and recommendations: the sector-specific governance measures a System Deployer in a regulated market should adopt, and the compliance architecture every Deployer should build regardless of sector.

On 5 November 2025, MeitY released the India AI Governance Guidelines, a general, cross-sectoral framework issued under the IndiaAI Mission and built on seven guiding

principles, or ‘Sutras,’ including Accountability and human-centred oversight.⁶⁰ The Guidelines recommend a path this paper has already traced: adapting existing law, including the IT Act and the DPDP Act, to clarify the liability of developers, deployers, and users, rather than enacting standalone AI legislation, alongside a domestic risk-classification framework and voluntary, risk-proportionate compliance measures. Where the RBI’s FREE-AI Report supplies sector-specific rigour for financial-sector Deployers, MeitY’s Guidelines supply the general, cross-sectoral counterpart; the two are complementary rather than competing, and the framework below draws on both. The FREE-AI Report translates the accountability principle into concrete operational requirements. Four of its recommendations form a minimum governance framework for System Deployers.

1. Board-Level Policy and Risk Classification (Recommendation 14)

Every Deployer must establish a Board-approved AI policy that classifies all AI use cases as low, medium, or high risk.⁶¹ This is not a paper exercise. It places the deployment decision at the level of institutional seniority commensurate with its risk. A board that approves high-risk agentic deployment cannot later disclaim knowledge of the consequences.

2. System Governance Across the AI Lifecycle (Recommendation 16)

The governance framework must cover design, development, deployment, and decommissioning.⁶² For medium and high-risk systems, this means mandatory human oversight, kill switches, and immutable audit logs, the same technical safeguards that SEBI requires for black-box algorithms.

3. Independent Audit for High-Risk Systems (Recommendation 24)

⁶⁰ Ministry of Electronics and Information Technology, *India AI Governance Guidelines* (Nov. 5, 2025).

⁶¹ FREE-AI Report, *supra* note 3, Recommendation 14.

⁶² FREE-AI Report, *supra* note 3, Recommendation 16.

High-risk deployments require independent third-party audit covering data inputs, model and algorithm, and decision outputs.⁶³ The EU AI Act's conformance-testing requirements provide a useful benchmark for what such audits should examine.

4. Graded Supervisory Response (Recommendation 8)

A rigid, punitive liability framework would discourage responsible experimentation. The FREE-AI Report proposes a graduated response: where a Deployer has followed prescribed safeguards and takes prompt corrective action, a first failure should attract an accommodative supervisory approach.⁶⁴ This accommodation disappears for repeated violations or gross negligence. Bathaee's sliding scale, calibrated to the level of human supervision and AI transparency, provides the theoretical foundation.⁶⁵

5. A Compliance Toolkit for Deployers

In practice, every System Deployer needs four things to manage its liability exposure.

- **Contractual clarity:** User agreements that are transparent about the scope and limits of autonomous agent actions. Contracts with model developers and API providers that include clear indemnification clauses and audit rights.
- **Technical safeguards:** Human-in-the-loop checkpoints for high-risk actions. Kill switches the Deployer can activate. Audit trails that capture every agent action in a form that courts and regulators can inspect.
- **Corporate governance:** An AI Ethics Review Board with real authority to pause deployments. A tested AI Incident Response Plan. AI liability insurance sized to the deployment's risk profile.

⁶³ FREE-AI Report, *supra* note 3, Recommendation 24.

⁶⁴ FREE-AI Report, *supra* note 3, Recommendation 8 and ¶4.4.30.

⁶⁵ Bathaee, *supra* note 1, 936. FREE-AI Recommendation 8 operationalizes that scale for Indian financial-sector regulation.

- **Public accountability:** Annual AI governance disclosures consistent with FREE-AI Recommendation 25. Voluntary participation in the sector-wide AI Compliance Toolkit under Recommendation 26.

VIII. CONCLUSION

When an autonomous agent acts in the world, doing tasks such as booking the flights, executing trades, or entering contracts, someone must answer for what it does. The law cannot allow that question to dissolve into an ecosystem of blame. The answer is already in the existing statute. The System Deployer, as the Orchestrator, controls every decision that translates a human goal into a real-world consequence. Section 79 of the IT Act classifies the Deployer as an intermediary – but then withholds the safe harbor, because the Deployer’s active orchestration role violates every condition that Section 79(2) requires for immunity. The DPDP Act imposes non-delegable fiduciary duties on the entity that determines the means of processing. SEBI’s framework proves the model in financial markets. Tort law points to the same party as best placed to prevent harm. The global consensus, from the CJEU to the EU AI Act, confirms the direction. The RBI FREE-AI Committee has already stated the principle: accountability cannot be delegated to the model and the algorithm. The law is ready. The question is whether those who apply it will be.

IX. REFERENCES / BIBLIOGRAPHY

A. Primary Legislation

1. Competition Act, 2002.
2. Digital Personal Data Protection Act, 2023.
3. Information Technology Act, 2000.
4. Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021.

5. Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026.
6. Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022, Digital Services Act.
7. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024, Artificial Intelligence Act.
8. Regulation (EU) 2026/1744 of the European Parliament and of the Council on the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI).

B. Case Law

1. Amazon Seller Services Pvt. Ltd. v. Amway India Enterprises Pvt. Ltd., 2020 SCC OnLine Del 454.
2. Amway India Enterprises Pvt. Ltd. v. 1MG Technologies Pvt. Ltd. & Anr., (2019) 260 DLT 690.
3. Chairman, Railway Board v. Chandrima Das, (2000) 2 SCC 465.
4. Donoghue v. Stevenson, [1932] AC 562.
5. Google India Pvt. Ltd. v. Visaka Industries Ltd., (2020) 4 SCC 162.
6. L'Oréal SA v. eBay International AG, Case C-324/09, Court of Justice of the European Union, 12 July 2011.
7. M.C. Mehta v. Union of India (Oleum Gas Leak), (1987) 1 SCC 395.
8. MySpace Inc. v. Super Cassettes Industries Ltd., 236 (2017) DLT 478.
9. Shreya Singhal v. Union of India, (2015) 5 SCC 1.

C. Regulatory Instruments and Reports

1. Competition Commission of India, Market Study on Artificial Intelligence and Competition, conducted in partnership with Management Development Institute Society, Gurugram, 2025.
2. Ministry of Electronics and Information Technology, India AI Governance Guidelines, 2025.
3. Ministry of Electronics and Information Technology, Notification bringing into force provisions of the Digital Personal Data Protection Act, 2023 and notifying the Digital Personal Data Protection Rules, 2025.
4. Reserve Bank of India, Report of the Committee on Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI) in the Financial Sector, 2025.
5. Securities and Exchange Board of India, Safer Participation of Retail Investors in Algorithmic Trading, Circular No. SEBI/HO/MIRSD/MIRSD-PoD/P/CIR/2025/0000013, 4 February 2025.

D. Secondary Sources

1. Yavar Bathaee, 'The Artificial Intelligence Black Box and the Failure of Intent and Causation' (2018) 31 *Harvard Journal of Law & Technology* 889.
2. Vasudev Devadasan, 'Conceptualizing India's Safe Harbour in the Era of Platform Governance' (2024) 19 *Indian Journal of Law and Technology* 8.
3. Shreya Goyal, 'AI and Legal Liability: Who Is Responsible When Machines Go Wrong?' (2025) 7 *Indian Journal of Law and Legal Research* 9404.
4. Indranath Gupta and Lakshmi Srinivasan, 'Evolving Scope of Intermediary Liability in India' (2023) 37 *International Review of Law, Computers & Technology* 294.
5. Emmanuel Idowu, 'Responsibility Attribution in Autonomous AI Agents Deployed on Cloud Infrastructure' (2025) *Preprints.org*.