

Liability for AI Generated Harms: Rethinking Indian Legal Framework.

Abhemanyu Dhakkite
Mumbai University

Abstract—The increasing integration of artificial intelligence into decision-making processes has generated new forms of harm that challenge the foundations of traditional legal liability. AI systems, characterized by autonomy, opacity, and multi-actor involvement, disrupt core doctrines of attribution, causation, and fault that underpin existing legal frameworks. This paper examines the adequacy of Indian law in addressing AI-generated harm, focusing on tort law, product liability, contract law, and relevant regulatory regimes. It argues that the current framework is structurally inadequate, not merely due to gaps in interpretation but because of a fundamental mismatch between conventional legal principles and the operational realities of AI systems. Through a detailed doctrinal analysis, the paper identifies key challenges, including the attribution vacuum, causation indeterminacy, and the limitations of fault-based liability. A comparative analysis of the European Union and the United States highlights alternative regulatory approaches, revealing India's fragmented and underdeveloped position. In response, the paper proposes a structured liability framework combining shared liability across actors, strict liability for high-risk AI systems, mandatory insurance mechanisms, and centralized regulatory oversight. The paper concludes that effective governance of artificial intelligence requires a shift from reactive legal adaptation to proactive regulatory design, ensuring that innovation is matched by accountability.

I. INTRODUCTION

Artificial intelligence is no longer confined to experimental or assistive functions; it is increasingly embedded in decision-making processes across sectors such as finance, healthcare, and governance. From algorithmic credit scoring to autonomous medical diagnostics, AI systems are now capable of generating outcomes that directly affect legal rights, economic interests, and personal well-being.

This shift has already produced tangible instances of harm. For example, algorithmic decision-making systems have been shown to replicate and amplify biases in lending and hiring, while automated trading systems have triggered rapid financial disruptions. Such cases illustrate a broader phenomenon: AI-generated harm, defined as harm caused by the autonomous or semi-autonomous operation of algorithmic systems, often without direct human intervention at the point of decision-making.

The emergence of such harm exposes a fundamental tension within existing legal frameworks. Traditional doctrines of liability particularly in tort and product liability law are premised on human agency, linear causation, and identifiable fault. These assumptions do not align with the operational realities of AI systems, which are characterized by autonomy, opacity, unpredictability, and multi-actor involvement. As a result, the application of existing legal principles to AI-generated harm produces doctrinal strain and, in some cases, outright failure.

This paper argues that the current Indian legal framework is structurally inadequate to address AI-related harm. The inadequacy arises from three interrelated challenges: the difficulty of attributing responsibility across multiple actors, the breakdown of conventional causation standards in opaque and non-linear systems, and the inability of existing legal classifications to properly categorize AI systems as either products or services. These issues are not isolated gaps, but manifestations of a deeper incompatibility between traditional legal doctrines and the nature of artificial intelligence.

To substantiate this argument, the paper proceeds in four parts. Section 2 examines the nature and types of AI-generated harm, highlighting the features that distinguish it from conventional technological risks. Section 3 analyses the existing legal framework in

India, including tort law, product liability, contract law, and relevant regulatory statutes. Section 4 identifies the doctrinal challenges that arise when these frameworks are applied to AI systems, focusing on attribution, causation, fault, and classification. Section 5 adopts a comparative perspective, evaluating regulatory approaches in the European Union and the United States to situate India's position within the global landscape. Finally, Section 6 proposes a structured liability framework combining shared liability, strict liability for high-risk systems, mandatory insurance, and institutional regulatory oversight. The central claim is straightforward: innovation in artificial intelligence has outpaced the evolution of legal accountability. Addressing this imbalance is not merely a matter of doctrinal refinement, but of systemic reform.

II. UNDERSTANDING AI GENERATED HARM.

The deep integration of Artificial Intelligence (AI) in our social and economic systems has significantly changed how we make life-altering decisions. Today, we are moving away from the conventional models where humans assumed majority control in the decision-making process, towards an autonomous and data-driven system, capable of adapting and acting on its own. As these tools become more prevalent in sectors like healthcare and finance, they don't just assist humans' decision makers but tends to replace them. This creates a significant legal challenge - a new category of algorithmic harm, as these decisions which may cause harm, are based on autonomous process without direct human control which forces us to rethink our traditional ideas of fault and liability.

2.1 Nature of AI.

AI systems, particularly those deployed in the contemporary decision making functions are machine based systems that operate with a specific objective that are either implicit or explicit using the input it receives to generate various outputs that can influence both physical or virtual environments.¹ Unlike

traditional softwares which operates based on pre-programmed instructions in a predictable and deterministic manner, these machine learning systems derive behavioural patterns for data, that produces outputs which cannot be fully determined or anticipated even by its developer at the point of its creation. As Ryan Calo puts it, such systems "accomplish tasks in ways that cannot be anticipated in advance," fundamentally distinguishing them from the passive tools that merely respond to human commands.² What makes these AI systems even more complex is that it doesn't remain static after its launched, as many AI systems are designed to keep learning and evolving based on the new data it encounter. This creates a wide spectrum of autonomy where some systems require constant human oversight while other operate entirely on their own, this gradation has direct consequences of how legal responsibility is allocated when such systems cause harm.³

2.2 Types of Harm.

The shift towards autonomous and data driven decision making is more than just a technical milestone; and has profound real-world consequences. As these tools start to dictate outcomes in sectors like healthcare and finance, the risk of harm becomes both widespread and fundamentally different from what we have dealt with before.

These injuries are unique because they do not stem from specific human mistakes but autonomous algorithmic processes. AI systems generate three legally cognisable categories of harm, financial, Physical and Reputational harm, each one challenging the existing legal frameworks in distinct but related ways -

A. Financial Harm

Financial harm perhaps has the most measurable consequences of AI-driven decision-making systems, particularly in sectors where algorithm's output directly influences economic opportunities and outcome. For instance, in the credit markets, AI

¹ OECD, Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449

² Ryan Calo, Robotics and the Lessons of Cyberlaw, 103Calif. L. Rev.513 (2015), <https://digitalcommons.law.uw.edu/faculty-articles/23>

³ Simon Chesterman, Artificial Intelligence and the Problem of Autonomy, 1 Notre Dame J. on Emerging Tech. 210 (2020)

systems often inherit biases from historical data, which can lead to discriminatory lending outcome like unjust loan approvals or skewed risk profiles. Similarly automated hiring can lead to long-term economic damage by filtering candidates through opaque and biased criteria, which directly impacts a person ability to earn and advance his career.⁴

In addition to these systemic biases, AI systems deployed in high-frequency and algorithmic trading environments, exposes risk of large-scale financial disruptions. Errors in the systems's design or unexpected interactions between different algorithms can trigger market fluctuations as seen in instances of 'Flash Crashes' wherein automated decision's leads to cascading financial losses. These risks are further exacerbated by the opacity of AI systems wherein the underlying decision-making processes remain difficult to interpret, thereby contributing to broader concerns of opacity-driven financial instability.⁵ Consequently, while financial harms arising from AI systems are often quantifiable, the mechanisms through which they occur remain complex and, at times, resistant to traditional forms of scrutiny.

B. Physical Harm -

Physical harm represents the most direct and serious risk posed by AI, largely because these systems now have the unique ability to turn data into real-world action.⁶ Unlike older software that simply gave advice to a human, modern AI is often embedded in machines that move, drive, and treat patients.⁷ This means a computer's mistake can instantly become a tangible, bodily injury. This convergence of 'thinking' and 'doing' marks a major shift in how technology can cause harm.

We see this most clearly in healthcare and autonomous systems.⁸ In a hospital, an AI tool used for diagnosis might suggest the wrong treatment because of flawed data, leading to a dangerous medical error. On our roads, self-driving cars rely on split-second data processing to navigate; any system failure here can lead to fatal accidents. A tragic example of this was the Uber autonomous vehicle incident that killed a pedestrian a clear reminder that AI can cause devastating harm entirely independent of human control.⁹

In India, these risks are becoming more foreseeable every day, yet our laws are struggling to keep up. For instance, the Motor Vehicles (Amendment) Act of 2019 is still built around the idea of a human driver. It doesn't yet account for a vehicle making its own decisions, leaving a significant 'regulatory gap' when it comes to who is liable for AI-induced physical harm.¹⁰

C. Reputational Harm -

Reputational harm is a distinct and growing threat, as AI now dictates what information is amplified or suppressed in the digital world.¹¹ Unlike physical injury, this damage often happens silently shaping public perception without the individual's knowledge. This is most dangerous in the rise of deepfakes and synthetic media, which can realistically distort a person's identity in compromising or false contexts. Because these algorithmic processes are so opaque, victims often find it impossible to track how such content is created or circulated.

While India's IT Act (Sections 66C, 66D, and 66E) and the 2021 Intermediary Guidelines offer some protection against identity theft and privacy shifts, they remain largely reactive.¹² These laws were not designed for the scale and autonomy of modern AI, leaving a significant gap in our ability to address the

⁴Solon Barocas and Andrew D Selbst, 'Big Data's Disparate Impact' (2016) 104 California Law Review 671 (2016)

⁵Add The Balck Box Ciation

⁶Ryan Calo, Robotics and the Lessons of Cyberlaw, 103Calif. L. Rev.513 (2015), <https://digitalcommons.law.uw.edu/faculty-articles/23>

⁷Ryan Calo, Robotics and the Lessons of Cyberlaw, 103Calif. L. Rev.513 (2015), <https://digitalcommons.law.uw.edu/faculty-articles/23>

⁸Abraham, Kenneth S. and Sharkey, Catherine M., Untangling AI Liability (February 23, 2026). Virginia

Public Law and Legal Theory Research Paper No. 2026-19, Virginia Law and Economics Research Paper No. 2026-05, Untangling AI Liability, 115 California Law Review

⁹Uber Technologies Inc. v. Estate of Herzberg

¹⁰Motor Vehicles (Amendment) Act 2019

¹¹Frank Pasquale, The Black Box Society: The Secret Algorithms That Control Money and Information (Harvard University Press 2015) ch2

¹²Information Technology Act 2000, ss 66C, 66D, 66E;

speed at which an algorithm can systematically dismantle a person's reputation.

Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules 2021, r 3(1)(b).

2.3 Why AI is Legally Distinct

The structural characteristics of AI-generated harm autonomy, opacity, unpredictability, and multi-actor involvement fundamentally challenge the traditional foundations of legal liability. Unlike conventional machines, AI operates along a spectrum of autonomy.¹³ This gradation complicates the allocation of responsibility, as a system's actions may not be directly attributable to a human actor at the moment of execution. Because autonomous systems can act in ways that are not fully anticipated, they undermine the legal assumptions of control and foreseeability that underpin most tort frameworks.¹⁴

This challenge is intensified by opacity, or the "black-box" nature of AI. The internal logic of these systems is often inaccessible or unintelligible to external observers, creating a significant barrier to establishing causation and fault.¹⁵ This opacity is not merely a technical hurdle but is frequently reinforced by trade secrecy, shielding the decision-making process from judicial scrutiny.¹⁶ While the U.S. case of *Loomis v. Wisconsin* saw courts struggle with algorithmic transparency, Indian jurisprudence, as seen in *State of Karnataka v. Selvi*, has historically emphasized that processes affecting individual rights must be transparent.¹⁷

Closely linked to these issues is the problem of unpredictability. AI models are probabilistic; they may produce different outputs under similar conditions,

which weakens the "foreseeability" requirement central to negligence claims.¹⁸ If an outcome cannot be reasonably anticipated by a developer or user, traditional fault-based standards often fail to provide a remedy for the victim.¹⁹

Finally, AI development involves multi-actor participation, where responsibility is fragmented across developers, data providers, and end-users.²⁰ Traditional legal models, which seek to identify a single "wrongdoer," struggle to address harm arising from this distributed network. However, Indian law has shown a unique capacity for adaptation. In *M.C. Mehta v. Union of India*, the Supreme Court established the principle of absolute liability for hazardous activities, moving away from the need to prove individual negligence in complex scenarios.²¹ While this doctrine has not yet been explicitly extended to AI, it provides a potential pathway for Indian courts to address the unique, fragmented nature of algorithmic harm.

Ultimately, the shift from human-led to machine-mediated action requires a departure from traditional liability. The law must evolve to account for systems that are no longer just tools, but active, opaque, and unpredictable participants in the legal landscape.

III. EXISTING LEGAL FRAMEWORKS IN INDIA

Currently, India lacks a unified legal regime specifically designed to handle liability for AI-generated harm. Instead, victims must rely on a patchwork of existing frameworks, primarily tort law, product liability statutes, contractual agreements, and sectoral regulations. While these pathways offer some

¹³Nathalie Nevejans, in Woodrow Barfield and Ugo Pagallo (eds), *Research Handbook on the Law of Artificial Intelligence* (Edward Elgar 2018)

¹⁴Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 *Calif. L. Rev.* 513 (2015), <https://digitalcommons.law.uw.edu/faculty-articles/23>

¹⁵Andrew D. Selbst and Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 *Fordham L. Rev.* 1085 (2018).

¹⁶Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (Harvard University Press 2015)

¹⁷*State v. Loomis*, 881 N.W.2d 749 (Wis. 2016); *State of Karnataka v. Selvi* (AIR 2010 SC 1974)

¹⁸Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 *Calif. L. Rev.* 513 (2015), <https://digitalcommons.law.uw.edu/faculty-articles/23>

¹⁹Nathalie Nevejans, in Woodrow Barfield and Ugo Pagallo (eds), *Research Handbook on the Law of Artificial Intelligence* (Edward Elgar 2018)

²⁰Abraham and Sharkey, 'Untangling AI Liability' (2026); Solon Barocas and Andrew D. Selbst, 'Big Data's Disparate Impact' (2016) 104 *California Law Review* 671.

²¹*MC Mehta v. Union of India* AIR 1987 SC 1086

options for redress, they were fundamentally built for a pre-digital world. They struggle to account for the unique characteristics of AI, such as its autonomy, its 'black-box' opacity, and the way it breaks traditional chains of cause and effect.

This section examines how effectively these traditional frameworks can be applied to the modern reality of AI. It argues that while these laws may offer a surface-level fit, they lack the doctrinal depth and consistency needed to address the structural complexities of algorithmic harm. As we look closer, it becomes clear that trying to fit AI into these old legal categories often creates more questions than it answers.

A. Negligence:

Negligence is India's primary tool for civil liability, requiring a duty of care, a breach of that duty, and resulting damage. As articulated in *Jacob Mathew v. State of Punjab*²², liability arises when conduct falls below the "reasonable person" standard, causing foreseeable harm a principle rooted in the foundational "neighbour principle" of *Donoghue v. Stevenson*.²³

However, applying this framework to AI reveals significant doctrinal strain in three key areas:

The Foreseeability Gap: Negligence relies on the ability to anticipate harm. Because machine learning systems are adaptive and move beyond deterministic logic, they often produce unpredictable outcomes. This inherent unpredictability makes it difficult to argue that a "reasonable" developer could have foreseen a specific algorithmic error.²⁴

The Problem of "Reasonableness": The law lacks established benchmarks for AI. Without industry-wide customs or regulatory standards, courts struggle to define what "reasonable" conduct looks like in a rapidly evolving field. This challenge is worsened by the "black-box" nature of AI, as seen in *Loomis v. Wisconsin*, where opacity prevented a clear assessment of accountability.²⁵

²²*Jacob Mathew vs State of Punjab & Anr (2005) 6 SCC 1*

²³*Donoghue v Stevenson is (1932) AC 562*

²⁴Selbst, Andrew D., *Negligence and AI's Human Users* (March 11, 2019). 100 Boston University Law Review 1315 (2020), UCLA School of Law, Public Law Research Paper No. 20-01.

²⁵*Loomis v. Wisconsin*, 137 S. Ct. 2290 (2017)

²⁶Selbst, Andrew D., *Negligence and AI's Human Users* (March 11, 2019). 100 Boston University Law

Causal Complexity: Traditional law uses the "but-for" test to link a defendant's action to the victim's harm. AI breaks this chain by introducing an "inscrutable" layer between human input and machine output. With multiple actors involved such as developers, data providers, and users, isolating a single "proximate cause" becomes nearly impossible.²⁶ While courts have handled probabilistic causation in cases like *Hotson v. East Berkshire*, the scale and opacity of AI systems significantly exacerbate the issue.²⁷

Furthermore, Indian law has yet to clearly determine whether AI systems should be treated as products or services, each of which carries distinct liability standards and corresponding expectations of care.²⁸ These hurdles suggest that negligence is not merely difficult to apply; it is structurally misaligned with the realities of autonomous systems. Rather than a gap in application, there exists a deeper conflict between traditional tort principles and AI-driven harm.

B. Product Liability:

Product liability presents an alternative pathway for addressing harm caused by AI systems, primarily governed in India by the Consumer Protection Act, 2019. The Act imposes liability on manufacturers and service providers for harm caused by "defective" products or "deficiency" in services.²⁹ At first glance, this framework appears well-suited to AI-related harm, particularly where AI systems are embedded within consumer-facing technologies. However, closer examination reveals significant conceptual difficulties in applying traditional product liability principles to AI systems.

A central challenge lies in the definition of "defect." Under the Consumer Protection Act, liability is premised on the existence of a manufacturing defect, design defect, or inadequate instructions or warnings. This framework assumes that a product is static and that any defect can be identified by comparing it

Review 1315 (2020), UCLA School of Law, Public Law Research Paper No. 20-01.

²⁷*Hotson v East Berkshire Area Health Authority [1987] AC 750*

²⁸Kenneth S. Abraham & Catherine M. Sharkey, *Artificial Intelligence and the Law of Torts* (forthcoming 2026)

²⁹Consumer Protection Act, 2019 Sections 2(10), 2(11), 84 and 85

against an intended or reasonable standard of performance. AI systems, however, are inherently dynamic. Machine learning models evolve over time through continuous interaction with data, meaning that a system may function as intended at the point of deployment but produce harmful outcomes later due to changes in data inputs or environmental conditions.³⁰ This raises a fundamental question: can an outcome generated by a self-learning system be characterized as a “defect,” or is it an emergent feature of the system’s design? If the latter, traditional notions of defect become difficult to sustain. Unlike conventional products, where defects are deviations from expected performance, AI systems may produce harmful yet statistically consistent outputs based on their training data. In such cases, liability cannot easily be attributed to a flaw in the product itself.

Further complexity arises from the classification of AI systems within the product–service dichotomy. Many AI-driven applications such as recommendation engines, automated decision-making tools, and software-as-a-service platforms blur the distinction between goods and services. This classification is not merely semantic; it determines the applicable liability regime and the standard against which conduct is assessed. While product liability focuses on defects in design or manufacture, service liability hinges on deficiency and reasonableness of performance. Indian law currently provides no clear guidance on how AI systems should be categorized, leading to uncertainty in the application of liability principles.³¹

Additionally, the multi-layered nature of AI development complicates the identification of the “manufacturer” or responsible party. AI systems are often the product of collaborative inputs, including software developers, data providers, platform operators, and third-party integrators. The Consumer Protection Act, while expansive in its definition of “product manufacturer,” does not adequately address this distributed model of production.³² As a result, assigning liability within complex AI supply chains becomes both legally and practically challenging.

These issues suggest that while product liability offers a superficially applicable framework, it struggles to accommodate the adaptive and decentralized nature of AI systems. The difficulty in defining defect, classifying AI within existing legal categories, and identifying responsible actors underscores a broader structural limitation: product liability law, as currently formulated, is not designed to regulate systems that evolve beyond their initial design and operate across fragmented chains of control.

IV. CONTRACT LAW

Contract law provides a significant, and often underappreciated, mechanism through which liability for AI-generated harm is structured in practice. Unlike tort or statutory regimes, which impose obligations externally, contractual arrangements allow parties to pre-allocate risk through negotiated terms. In the context of AI systems particularly those delivered through software-as-a-service (SaaS) models’ liability is frequently governed by limitation of liability clauses, indemnities, and warranty disclaimers.³³

These clauses are typically designed to minimize the exposure of developers and service providers by excluding liability for indirect or consequential damages, capping monetary liability, or disclaiming warranties related to performance and accuracy.³⁴ In many AI deployment scenarios, especially involving enterprise software or platform-based services, such provisions effectively shift the risk of harm onto the end-user or client. This contractual allocation of risk often operates irrespective of fault, thereby circumventing the need to establish negligence or product defect.

However, the effectiveness of such clauses in the context of AI-generated harm raises several concerns. First, the complexity and opacity of AI systems make it difficult for users to meaningfully assess the risks they are contractually assuming. Standard-form contracts, particularly in digital services, are rarely negotiated and are often accepted on a “take-it-or-

³⁰Andrew D. Selbst, *Negligence and AI’s Human Users*, 100 B.U. L. Rev. 1315 (2020)

³¹Kenneth S. Abraham & Catherine M. Sharkey, *Artificial Intelligence and the Law of Torts* (forthcoming 2026)

³²Consumer Protection Act, 2019 Section 2(36) (definition of product manufacturer).

³³Indian Contract Act, 1872

³⁴Kenneth S. Abraham & Catherine M. Sharkey, *Artificial Intelligence and the Law of Torts* (forthcoming 2026).

leave-it” basis. This creates an imbalance of bargaining power, where users may bear disproportionate risk without a corresponding ability to influence contractual terms.³⁵

Second, the enforceability of limitation and exclusion clauses is not absolute under Indian law. Courts have, in certain circumstances, refused to uphold clauses that are unconscionable, contrary to public policy, or that seek to exclude liability for fundamental obligations. In *Central Inland Water Transport Corporation v. Brojo Nath Ganguly*, the Supreme Court recognized that unfair and unreasonable contract terms in standard-form agreements may be invalidated where there is inequality of bargaining power.³⁶ Similarly, in *Bharathi Knitting Company v. DHL Worldwide Express Courier Division*, the Court upheld limitation clauses but emphasized that their enforceability depends on the specific terms and circumstances of the contract.³⁷ However, the application of these principles to AI-related harm remains uncertain, particularly where harm arises from autonomous system behavior rather than identifiable human conduct.

Third, contractual frameworks do not resolve the broader issue of third-party harm. While contracts can allocate risk between the immediate parties, they do not address situations where AI systems cause harm to individuals who are not privy to the contractual arrangement. In such cases, reliance on contract law provides no direct remedy, necessitating recourse to tort or statutory frameworks.

Finally, the widespread use of contractual risk allocation contributes to a fragmentation of liability. Rather than establishing consistent standards of responsibility, liability becomes contingent on the specific terms of individual agreements. This undermines predictability and may allow developers and deployers to systematically externalize risk without corresponding accountability.

In sum, while contract law offers a flexible tool for managing risk in AI transactions, it does so in a manner that prioritizes private ordering over public

accountability. The heavy reliance on limitation clauses and standard-form agreements highlights a structural limitation: contractual frameworks are ill-suited to address the broader societal implications of AI-generated harm, particularly where issues of fairness, transparency, and third-party impact are involved.

V. REGULATORY FRAMEWORK

Beyond private law mechanisms, elements of liability for AI-generated harm in India are indirectly addressed through a range of regulatory statutes. These include the Information Technology Act, 2000, the Digital Personal Data Protection Act, 2023, and the *Bharatiya Nyaya Sanhita*, 2023.

While none of these frameworks are specifically designed to govern AI liability, they collectively shape the legal landscape within which AI systems operate. The Information Technology Act, 2000 provides a foundational framework for intermediary liability and digital governance. Under Section 79, intermediaries are granted conditional immunity from liability for third-party content, subject to due diligence requirements.^{38,38} While this provision has been central to regulating online platforms, its applicability to AI systems is limited. AI-driven decision-making systems do not fit neatly within the traditional concept of intermediaries, particularly where harm arises not from third-party content but from autonomous system outputs. As a result, the safe harbour regime offers only partial and uncertain coverage in the context of AI-generated harm.

The Digital Personal Data Protection Act, 2023 introduces obligations relating to the processing of personal data, including requirements of consent, purpose limitation, and data security.³⁹ To the extent that AI systems rely on personal data, this framework provides an important layer of accountability. However, its focus remains on data protection rather than harm caused by AI decision-making. The Act does not directly address issues of liability arising

³⁵Friedrich Kessler, *Contracts of Adhesion: Some Thoughts About Freedom of Contract*, 43 Colum. L. Rev. 629 (1943).

³⁶*Central Inland Water Transport Corporation v. Brojo Nath Ganguly*, (1986) 3 SCC 156.

³⁷*Bharathi Knitting Company v. DHL Worldwide Express Courier Division*, (1996) 4 SCC 704.

³⁸Information Technology Act, 2000 Section 79

³⁹Digital Personal Data Protection Act, 2023

from erroneous, biased, or harmful outputs generated by AI systems, thereby limiting its relevance to broader questions of civil liability.

The Bharatiya Nyaya Sanhita, 2023, which replaces the Indian Penal Code, introduces provisions that may be invoked in cases involving AI-related misconduct, particularly in areas such as fraud, misinformation, or cyber-enabled offences.⁴⁰ However, as a criminal statute, its focus is on punishment rather than compensation. Moreover, its application presupposes identifiable human intent or culpability, which may be difficult to establish in cases involving autonomous or semi-autonomous systems.

Taken together, these regulatory frameworks reflect a fragmented and indirect approach to AI governance in India. Each statute addresses a specific aspect of digital activity intermediary conduct, data protection, or criminal wrongdoing but none provide a coherent framework for attributing liability for harm caused by AI systems. The absence of a unified regulatory approach results in gaps, overlaps, and uncertainty, particularly in cases where AI-generated harm does not fit neatly within existing legal categories.

This fragmentation underscores a broader structural issue: current regulatory regimes are reactive and sector-specific, rather than designed to address the cross-cutting challenges posed by AI. As a result, while elements of accountability can be derived from existing laws, they fail to collectively provide a consistent or adequate model for addressing AI-generated harm.

VI. DOCTRINAL CHALLENGES IN APPLYING EXISTING LAW

The application of existing legal frameworks to AI-generated harm reveals a series of deep doctrinal tensions. While the preceding section demonstrates that Indian law contains multiple liability mechanisms, these frameworks are built on assumptions that do not hold in the context of autonomous and data-driven systems. The challenge, therefore, is not merely one of

interpretation, but of structural incompatibility.

This section examines four core doctrinal problems that of attribution, causation, fault, and classification, along with the broader regulatory vacuum in which they operate. Together, these challenges illustrate that traditional liability models struggle to accommodate the distributed, opaque, and autonomous nature of artificial intelligence, thereby undermining their effectiveness in addressing AI-related harm.

A. Attribution Problem

The attribution of liability in cases of AI-generated harm exposes a fundamental limitation in traditional tort doctrine. Classical liability frameworks presuppose the existence of a human actor whose conduct can be evaluated against a legal standard of care. However, as AI systems increasingly operate with a degree of autonomy, this assumption begins to collapse. David C. Vladeck characterises such systems as “machines without principals,” where no identifiable human actor exercises direct control over the decision that ultimately causes harm.⁴¹

This problem is further compounded by the distributed nature of AI systems. As Kenneth S. Abraham and Catherine M. Sharkey observe, modern AI operates through a multi-actor ecosystem involving developers, deployers, data providers, and end-users, each contributing to the system’s functioning without exercising complete control.⁴² This fragmentation makes it difficult to assign responsibility to any single actor under existing legal standards. In the Indian context, this attribution challenge is particularly pronounced. Meghna Bal and N. S. Nappinai highlight that current legal frameworks do not provide clear mechanisms for allocating liability across these actors, resulting in uncertainty and potential gaps in accountability.⁴³ Accordingly, the issue is not merely one of evidentiary difficulty, but of doctrinal inadequacy. Where no clear “principal” can be identified, the foundational basis of fault-based liability itself becomes unstable, giving rise to what may be described as an attribution vacuum.

⁴⁰Bharatiya Nyaya Sanhita, 2023

⁴¹David C. Vladeck, Essay, *Machines Without Principals: Liability Rules and Artificial Intelligence*, 89 Wash.L. Rev. 117 (2014).

⁴²Kenneth S. Abraham & Catherine M. Sharkey, *Untangling AI Liability*, SSRN (forthcoming in 115 Calif. L. Rev., 2026)

⁴³Meghna Bal & N. S. Nappinai, *Crafting a Liability Regime for AI Systems in India* (Esysa Centre & Cyber Saathi Foundation, 2024), at 15–22.

B. Causation Problem

Beyond attribution, AI-generated harm presents a profound challenge to the doctrine of causation. Traditional tort law relies on relatively linear and explainable causal chains, most commonly operationalised through the “but-for” test. However, the internal functioning of AI systems, particularly those based on machine learning, disrupts this model. As Andrew D. Selbst argues, AI systems rely on probabilistic and statistical correlations rather than deterministic reasoning, making it difficult to isolate a single causal factor responsible for a given outcome.⁴⁴ This difficulty is compounded by the structural opacity of many AI systems. Kenneth S. Abraham and Catherine M. Sharkey note that AI decision-making often involves complex interactions between training data, algorithmic design, and deployment context, resulting in non-linear and multi-factorial causal pathways.⁴⁵ In such environments, harm cannot be easily traced back to a specific act or omission, thereby weakening both factual and legal causation.

From a law-and-economics perspective, Miriam Buiten, Alexandre de Stree, and Martin argue that this opacity creates systematic evidentiary barriers for claimants. Plaintiffs may be unable to demonstrate causation even where harm is clearly linked to the operation of an AI system, leading to under-compensation and reduced deterrence.⁴⁶

Accordingly, the challenge is not merely that causation is difficult to prove, but that existing doctrinal tools are ill-equipped to address systems in which causal relationships are inherently diffuse and partially inscrutable. In such cases, the requirement of proving causation risks becoming a barrier to liability rather than a principled filter.

C. Fault vs. Strict Liability

The combined breakdown of attribution and causation in AI systems exposes a deeper doctrinal tension: the

inadequacy of fault-based liability in addressing AI-generated harm. Negligence, as the dominant framework, presupposes that harm results from a failure to exercise reasonable care.

However, in the context of AI, harmful outcomes may arise even where all actors such as developers, deployers, and users, have complied with prevailing standards. As Andrew D. Selbst argues, this reflects a structural mismatch between individual fault and system-level behaviour, where harm is produced not by carelessness but by the inherent complexity and probabilistic nature of algorithmic decision-making.⁴⁷ This challenge is amplified in scenarios where no identifiable human “principal” directs the system’s actions. David C. Vladeck contends that in such cases, the normative foundation of fault-based liability collapses, as the doctrine depends on linking harm to human agency.⁴⁸ Similarly, Kenneth S. Abraham and Catherine M. Sharkey observe that in multi-actor AI ecosystems, attempts to assign fault often become artificial, distributing responsibility in ways that do not accurately reflect control or risk creation.⁴⁹

From an economic standpoint, Miriam Buiten, Alexandre de Stree, and Martin Peitz argue that fault-based regimes may lead to inefficient outcomes, particularly under conditions of uncertainty where actors cannot anticipate or internalize the risks generated by AI systems.⁵⁰ In such cases, negligence fails to achieve its core objectives of deterrence and loss allocation.

In response, a growing body of scholarship supports the adoption of strict liability for certain categories of AI-related harm. Anat Lior advances this position on the basis of non-reciprocal risk, arguing that where individuals are exposed to harm without participating in or benefiting from the underlying activity, strict

⁴⁴Andrew D. Selbst, *Negligence and AI’s Human Users*, 100 B.U. L. Rev. 1315, 1325–1335 (2020).

⁴⁵Kenneth S. Abraham & Catherine M. Sharkey, *Untangling AI Liability*, SSRN (forthcoming in 115 Calif. L. Rev., 2026), at 18–28.

⁴⁶Miriam Buiten, Alexandre de Stree & Martin Peitz, *The Law and Economics of AI Liability*, 48 *Comput. L. & Sec. Rev.* 105733, 6–12 (2023).

⁴⁷Andrew D. Selbst, *Negligence and AI’s Human Users*, 100 B.U. L. Rev. 1315, 1335–1345 (2020)

⁴⁸David C. Vladeck, *Machines Without Principals: Liability Rules and Artificial Intelligence*, 89 *Wash. L. Rev.* 117, 130–138 (2014)

⁴⁹Kenneth S. Abraham & Catherine M. Sharkey, *Untangling AI Liability*, SSRN (forthcoming in 115 Calif. L. Rev., 2026), at 28–38.

⁵⁰Miriam Buiten, Alexandre de Stree & Martin Peitz, *The Law and Economics of AI Liability*, 48 *Comput. L. & Sec. Rev.* 105733, 12–18 (2023).

liability ensures fairer risk distribution.⁵¹ This approach also addresses concerns that victims may otherwise be left without remedies due to doctrinal barriers in proving fault or causation.

Accordingly, the persistence of fault-based liability as the default framework in AI contexts appears increasingly untenable. While negligence may retain relevance in limited circumstances, the structural characteristics of AI systems necessitate a calibrated shift toward strict liability, particularly in high-risk domains where the potential for harm is significant and difficult to predict.

D. Classification Problem

A further doctrinal difficulty arises in the classification of AI systems within existing legal categories. Liability frameworks, particularly in product liability law, often depend on whether the subject matter is characterised as a “product” or a “service.” However, AI systems resist this binary distinction. As Karni A. Chagal-Feferkorn observes, algorithmic decision-makers exhibit characteristics of both categories: they are frequently delivered as services (such as software-as-a-service platforms), yet function autonomously in a manner analogous to products.⁵²

This hybrid nature is further complicated by the dynamic and evolving character of AI systems. Unlike traditional products, which are static at the point of sale, AI systems may continuously update, learn, and adapt post-deployment. This undermines core assumptions of product liability law, which are premised on identifiable defects at a fixed point in time¹.

Chagal-Feferkorn therefore argues for a purposive approach to classification one that looks beyond formal labels and instead considers the underlying rationale of products liability, particularly the allocation of risk to those best positioned to prevent harm¹. However, the need for such reinterpretation underscores a broader doctrinal issue: existing

classification frameworks are not designed to accommodate autonomous, evolving systems.

As a result, the classification of AI systems becomes not merely a technical question, but a strategic one, with significant implications for liability, available remedies, and the allocation of risk. This ambiguity further contributes to the instability of existing legal doctrines when applied to AI-generated harm.

E. Regulatory Vacuum

The doctrinal challenges outlined above are compounded by the absence of a coherent, AI-specific liability framework in India. While existing legal regimes such as tort law, consumer protection, and data protection may incidentally apply to AI-related harm, they operate in isolation and were not designed to address the unique risks posed by autonomous systems. As Meghna Bal and N. S. Nappinai observe, this results in a fragmented legal landscape lacking clear rules on liability allocation across developers, deployers, and other stakeholders.⁵³

From a comparative and economic perspective, Miriam Buiten, Alexandre de Streel, and Martin Peitz highlight that the absence of structured liability rules leads to inefficient outcomes, including under-deterrence and uncertainty in risk allocation.⁵⁴ Without clear ex ante standards, actors may lack incentives to internalize the risks created by AI systems, while victims may face significant barriers in obtaining compensation.

Further, Anat Lior argues that the lack of a defined liability regime can result in systemic under-compensation, particularly in contexts involving non-reciprocal risks, where those harmed neither control nor benefit from the underlying technology.⁵⁵ This concern is particularly acute in high-risk AI applications, where the scale and impact of harm may be substantial.

In contrast to jurisdictions that are moving toward structured, risk-based regulatory models, India

⁵¹Anat Lior, *AI Strict Liability Vis-à-Vis AI Monopolization*, 22 Colum. Sci. & Tech. L. Rev. 90, 110–120 (2020).

⁵²Karni A. Chagal-Feferkorn, *Am I an Algorithm or a Product? When Products Liability Should Apply to Algorithmic Decision-Makers*, 30 Stan. L. & Pol’y Rev. 61, 70–90 (2019).

⁵³Meghna Bal & N. S. Nappinai, *Crafting a Liability Regime for AI Systems in India* (Esysa Centre & Cyber Saathi Foundation, 2024), at 10–18, 30–35.

⁵⁴Miriam Buiten, Alexandre de Streel & Martin Peitz, *The Law and Economics of AI Liability*, 48 Comput. L. & Sec. Rev. 105733, 18–25 (2023).

⁵⁵Anat Lior, *AI Strict Liability Vis-à-Vis AI Monopolization*, 22 Colum. Sci. & Tech. L. Rev. 90, 120–130 (2020).

continues to rely on piecemeal adaptation of existing doctrines. The result is not merely a gap in regulation, but a broader regulatory vacuum one that generates uncertainty for claimants, weakens accountability for developers and deployers, and places undue interpretive burdens on courts.

VII. COMPARATIVE PERSPECTIVE

A. European Union

The European Union has adopted a comprehensive and forward-looking approach to artificial intelligence regulation through the Regulation (EU) 2024/1689 (Artificial Intelligence Act). This framework establishes a risk-based classification system, categorizing AI systems into prohibited, high-risk, limited-risk, and minimal-risk tiers, with regulatory obligations calibrated accordingly¹.

Under this model, high-risk AI systems such as those used in healthcare, critical infrastructure, and law enforcement are subject to extensive *ex ante* requirements, including conformity assessments, risk management systems, technical documentation, and post-market monitoring obligations.⁵⁶ These requirements are imposed primarily on providers and deployers, thereby clarifying responsibility across the AI lifecycle.

As Miriam Buiten, Alexandre de Streel, and Martin Peitz observe, this risk-tiered framework effectively redistributes liability by aligning regulatory obligations with the capacity of actors to control and mitigate risk.⁵⁷ In doing so, the EU moves beyond traditional *ex post* liability models and embeds accountability directly into the design, development, and deployment stages of AI systems.

This approach directly addresses the doctrinal failures identified earlier. By imposing traceability, documentation, and compliance obligations, it mitigates attribution ambiguity and evidentiary barriers in causation, while functionally incorporating strict liability principles in high-risk contexts.

Accordingly, the EU model represents a proactive regulatory paradigm, where the focus shifts from post-harm adjudication to pre-harm risk governance.

B. United States

In contrast to the European Union's harmonised regulatory framework, the United States has adopted a fragmented and largely reactive approach to AI liability. Rather than enacting a comprehensive federal statute, the U.S. relies on the incremental adaptation of existing legal doctrines, primarily tort and product liability. Through judicial decisions and state-level developments. As Miriam Buiten, Alexandre de Streel, and Martin Peitz note, U.S. courts have increasingly attempted to fit AI-related harms within traditional product liability frameworks, treating algorithmic systems as products where possible and extending doctrines such as design defect and failure to warn.⁵⁸ However, this approach remains doctrinally strained, as it depends on analogies that do not fully capture the autonomous and evolving nature of AI systems. Further, Kenneth S. Abraham and Catherine M. Sharkey highlight that AI liability in the United States is developing through a patchwork of state-level litigation, resulting in inconsistent standards and uncertain outcomes.⁵⁹ Courts continue to grapple with familiar issues such as attribution, causation, and fault without a unified framework to guide their resolution. While this case-law driven model offers a degree of flexibility, allowing doctrines to evolve alongside technological developments, it also reproduces many of the structural challenges identified in Section 4. Attribution remains contested in multi-actor systems, causation is often difficult to establish due to technical opacity, and fault-based liability continues to dominate despite its limitations. Accordingly, the U.S. approach may be characterised as reactive and evolutionary, relying on *ex post* adjudication rather than *ex ante* regulatory design. While it allows for doctrinal experimentation, it lacks the coherence and predictability of a unified liability regime.

⁵⁶Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 (Artificial Intelligence Act), arts. 5–29, 61–72 (OJ L, 12 July 2024)

⁵⁷Miriam Buiten, Alexandre de Streel & Martin Peitz, EU Liability Rules for the Age of Artificial Intelligence (CERRE, 2021), at 10–22.

⁵⁸Miriam Buiten, Alexandre de Streel & Martin Peitz, EU Liability Rules for the Age of Artificial Intelligence (CERRE, 2021), at 22–30.

⁵⁹Kenneth S. Abraham & Catherine M. Sharkey, Untangling AI Liability, SSRN (forthcoming in 115 Calif. L. Rev., 2026), at 40–55.

C. Insight India in Comparative Perspective

A comparative analysis of the European Union and the United States reveals two distinct regulatory trajectories in addressing AI liability. The EU, through the Regulation (EU) 2024/1689 (Artificial Intelligence Act), adopts a proactive, risk-based model that embeds accountability into the design and deployment stages of AI systems through *ex ante* obligations.⁶⁰ In contrast, the United States relies on a reactive, litigation-driven approach, where liability evolves incrementally through judicial interpretation and the adaptation of existing doctrines.⁶¹

As Miriam Buiten, Alexandre de Streel, and Martin Peitz observe, these approaches reflect fundamentally different strategies for risk allocation: the EU emphasizes regulatory coordination and preventive compliance, while the U.S. prioritizes doctrinal flexibility and *ex post* adjudication.⁶² Meanwhile, Kenneth S. Abraham and Catherine M. Sharkey highlight that the U.S. model, despite its adaptability, continues to struggle with fragmentation and doctrinal inconsistency in AI-related litigation.

Against this backdrop, India's position is markedly distinct and significantly weaker. Unlike the EU, it lacks a unified, risk-based regulatory framework governing AI systems. Unlike the United States, it also lacks a sufficiently developed body of jurisprudence capable of incrementally adapting existing doctrines to technological change. Instead, India remains dependent on the piecemeal application of legacy legal frameworks, resulting in fragmentation, uncertainty, and doctrinal strain.

The comparison therefore yields a clear conclusion:

- The EU represents structured, *ex ante* governance
- The U.S. represents adaptive, *ex post* evolution
- India represents doctrinal reliance without systemic reform

This positioning is not merely a matter of regulatory lag, but of structural absence. India is not choosing between competing models it is currently operating without a coherent liability framework altogether.

⁶⁰Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 (Artificial Intelligence Act), arts. 5–29.

⁶¹Kenneth S. Abraham & Catherine M. Sharkey, *Untangling AI Liability*, SSRN (forthcoming in 115 Calif. L. Rev., 2026), at 40–55.

VIII. PROPOSED FRAMEWORK

The preceding analysis demonstrates that the challenges posed by artificial intelligence are not merely gaps within existing legal doctrines, but structural incompatibilities that render traditional frameworks inadequate. Attribution failures, causation breakdowns, and the limitations of fault-based liability collectively indicate the need for a reconstructed liability architecture, rather than incremental doctrinal adjustment. Accordingly, this section proposes a calibrated framework for AI liability in India, grounded in three core objectives: effective risk allocation, meaningful victim compensation, and regulatory clarity. The proposed model combines shared liability across actors, strict liability for high-risk systems, mandatory insurance mechanisms, and institutional regulatory oversight.

IX. SHARED LIABILITY MODEL

The attribution vacuum identified earlier necessitates a shift from singular liability models to a distributed framework of responsibility that reflects the multi-actor nature of AI systems. AI-generated harm does not arise from a single point of failure, but from the cumulative interaction of developers, deployers, and users across the system lifecycle.

A doctrinal foundation for such an approach is provided by Anat Lior, who draws on the principle of respondent superior to conceptualise AI systems as functional agents within a broader chain of human actors.⁶³ Under this analogy, liability may be extended across those who design, control, and benefit from the operation of AI systems, even in the absence of direct involvement at the moment of harm.

Operationally, this translates into a model of proportionate liability:

- Developers bear responsibility for design architecture, training data, and system-level risks

⁶²Miriam Buiten, Alexandre de Streel & Martin Peitz, *EU Liability Rules for the Age of Artificial Intelligence* (CERRE, 2021), at 30–40.

⁶³Anat Lior, *AI Entities as AI Agents: Artificial Intelligence Liability and the AI Respondeat Superior Analogy*, 46 Mitchell Hamline L. Rev. 1043, 1065–1080 (2020).

- Deployers are accountable for implementation, monitoring, and contextual use
- Users may incur liability where misuse or unreasonable reliance contributes to harm

Such a framework aligns liability with control, capability, and risk creation, rather than attempting to artificially isolate a single culpable actor. It also ensures that gaps in attribution do not translate into gaps in accountability.

By distributing responsibility across the AI ecosystem, the shared liability model preserves the corrective function of law while adapting it to the realities of autonomous and distributed decision-making systems.

A. Strict Liability for High-Risk AI

While a shared liability model addresses the problem of distributed responsibility, it does not fully resolve situations in which harm arises despite the exercise of reasonable care by all actors. In such cases, the continued reliance on fault-based liability risks leaving victims uncompensated and undermining the deterrent function of the law. This is particularly evident in the context of high-risk

AI systems, where the scale, opacity, and unpredictability of harm exceed the capacity of traditional negligence frameworks.

A compelling case for strict liability in such contexts is advanced by Anat Lior, who argues that AI systems often generate non-reciprocal risks that is, risks imposed on individuals who neither control nor benefit from the underlying activity.⁶⁴ In such scenarios, fairness and efficiency considerations justify shifting the burden of harm onto those who introduce and profit from the technology, regardless of fault.

Strict liability also addresses the evidentiary barriers identified earlier. By removing the requirement to prove negligence or causation in complex, opaque systems, it ensures that victims are not denied remedies due to structural limitations in proving fault. At the same time, it incentivizes developers and deployers to internalize the risks associated with AI systems, encouraging higher standards of safety, testing, and oversight.

Importantly, concerns that strict liability may stifle innovation are not decisive. As Lior demonstrates,

well-calibrated strict liability regimes can coexist with technological development by promoting responsible innovation and efficient risk allocation¹. Rather than deterring innovation, such frameworks may enhance trust in AI systems, thereby facilitating broader adoption. Accordingly, strict liability should be imposed for clearly defined categories of high-risk AI applications, including but not limited to healthcare systems, autonomous vehicles, and critical infrastructure technologies. In these domains, the potential magnitude of harm and the limitations of fault-based doctrines justify a shift toward liability regimes that prioritize victim protection and systemic accountability.

X. MANDATORY INSURANCE

Even a well-calibrated liability framework combining shared and strict liability may fail to ensure effective compensation if responsible actors lack the financial capacity to satisfy claims. This risk is particularly acute in the context of AI, where harm may be large-scale, diffuse, and difficult to predict. Accordingly, a mandatory insurance regime is a necessary complement to liability rules, ensuring that victims are compensated irrespective of the solvency of individual actors. A comprehensive framework for such an approach is developed by Anat Lior, who argues that insurance should function as a central pillar of AI regulation rather than a peripheral mechanism¹. Under this model, developers and deployers of AI systems, particularly those operating in high-risk domains would be required to maintain liability insurance covering potential harms arising from system operation. Mandatory insurance serves multiple functions. First, it guarantees accessible and timely compensation for victims, reducing reliance on prolonged litigation. Second, it enables risk pooling and distribution, allowing insurers to absorb and spread the financial impact of AI-related harm.

Third, it introduces a layer of private regulation, as insurers are incentivized to assess, price, and monitor risk, thereby indirectly enforcing safety standards across the industry.⁶⁵

Moreover, insurance mechanisms can be designed to

⁶⁴Anat Lior, *AI Strict Liability Vis-à-Vis AI Monopolization*, 22 *Colum. Sci. & Tech. L. Rev.* 90, 100–120 (2020).

⁶⁵Anat Lior, *Insuring AI: The Role of Insurance in Artificial Intelligence Regulation*, 35 *Harv. J.L. & Tech.* 467, 480–510 (2022).

operate alongside strict liability regimes, creating a no-fault compensation structure in appropriate contexts. This is particularly valuable where causation is difficult to establish, as it allows victims to recover without navigating complex evidentiary barriers. In effect, mandatory insurance transforms liability from a purely corrective mechanism into a forward-looking risk management tool, aligning financial incentives with safer AI development and deployment practices.

XI. REGULATORY AUTHORITY

The effectiveness of the proposed liability framework ultimately depends on the existence of an institutional mechanism capable of coordinating, enforcing, and continuously adapting regulatory standards. In the absence of such an authority, even well-designed liability rules risk fragmentation and inconsistent application. Accordingly, the establishment of a dedicated AI regulatory authority is essential to operationalize shared liability, strict liability, and insurance-based mechanisms within a coherent system.

Policy support for such an approach can be found in the NITI Aayog's framework on responsible AI, which emphasizes the need for structured governance, algorithmic accountability, and multi-stakeholder oversight.⁶⁶ The report highlights mechanisms such as algorithmic audits, impact assessments, and continuous monitoring as critical tools for ensuring responsible deployment of AI systems.

Building on this, a specialized regulatory body in India could perform several key functions:

- Certification and classification of AI systems based on risk levels
- Mandatory audit and compliance frameworks for high-risk applications
- Oversight of insurance and liability mechanisms
- Standard-setting for transparency, documentation, and explainability

Such an authority would also play a crucial role in reducing doctrinal uncertainty by providing ex ante guidance, thereby easing the burden on courts to

resolve complex technical questions through ex post adjudication.

Importantly, this model does not require the creation of an entirely new regulatory philosophy, but rather the institutional consolidation of existing policy directions within a specialized framework tailored to AI governance. By integrating regulatory oversight with liability rules, India can transition from a fragmented, reactive system to a more structured and predictable regime.

XII. CONCLUSION

The rapid integration of artificial intelligence into decision-making processes has exposed fundamental limitations in existing legal frameworks. As this paper has demonstrated, traditional doctrines of attribution, causation, and fault-based liability are ill-equipped to address the distributed, opaque, and autonomous nature of AI systems.

A coherent response requires more than incremental adaptation. The framework proposed in this section, combining shared liability, strict liability for high-risk systems, mandatory insurance, and institutional regulatory oversight offers a structured approach to reconciling innovation with accountability.

The urgency of reform lies not only in addressing present harms, but in anticipating future risks as AI systems become increasingly embedded in critical sectors. Without a clear and comprehensive liability regime, the law risks falling behind technological development, leaving both victims and market participants in a state of uncertainty.

Ultimately, the objective is not to constrain innovation, but to ensure that it operates within a framework of responsibility. In the context of artificial intelligence, accountability is not a constraint on progress. It is a precondition for it.

⁶⁶NITI Aayog, Approach Document for India: Part 2 Operationalising Principles for Responsible AI 18–35 (2021).