



International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

The Black-Box Boardroom: Rethinking Director Accountability and Expert Reliance in Indian Corporate Law

Sumit Gurjar¹, Samriddh Sharma², Abhiman Mohapatra³

National Law Institute University, Bhopal

ABSTRACT :

Corporate boards in India increasingly rely on algorithmic tools and AI-driven valuation models to execute complex, time-sensitive restructuring and M&A transactions. This technological shift fundamentally challenges the fiduciary duty of care under Section 166(3) of the Companies Act, 2013, which strictly mandates the exercise of “independent judgment.” Existing Indian jurisprudence addresses human-to-human delegation but fails to resolve the “layered reliance” problem: whether a board can lawfully invoke the protective shield of expert reliance when the human advisor's report is generated by an opaque, “black-box” algorithm. Through a comparative analysis of Delaware’s statutory safe harbour for expert reliance (DGCL § 141(e)), this paper demonstrates that even explicit statutory protections fracture when confronted with unverified algorithmic intermediaries. To bridge this doctrinal gap, the paper proposes a novel four-factor standard of review weighing materiality, methodology comprehension, layered human verification, and documentation to distinguish legitimate reliance from fiduciary abdication. Finally, it translates this judicial test into an ex-ante corporate governance framework, proposing targeted, event-driven disclosure amendments to the SEBI (LODR) Regulations, the BRSR architecture, and the ICSI Secretarial Standards (SS-1).

Keywords: *Corporate Governance, Fiduciary Duties, Artificial Intelligence, Section 166(3), Expert Reliance, SEBI LODR, Mergers and Acquisitions.*

Introduction

Corporate transactions in India now move at a pace the drafters of the Companies Act, 2013 could scarcely have anticipated. Boards approving mergers, demergers, and acquisitions no longer rely solely on raw financial statements; they lean increasingly on algorithmic valuation models, AI-driven risk-scoring tools, and automated due-diligence platforms deployed by their financial and legal advisors.¹ These tools compress weeks of analysis into hours, and have become steadily more indispensable to time-sensitive restructuring processes.

For the purposes of this paper, algorithmic tools are divided into two categories: rule-based decision-support systems, whose logic is deterministic and can be audited, and opaque machine-learning models: “black-box” tools whose outputs cannot be fully reconstructed even by the people operating them. This distinction matters, because the duty owed by a director confronted with each is not the same, a point taken up in Section V.

The doctrinal difficulty is straightforward to state but not easy to resolve. Section 166(3) of the Companies Act requires a director to exercise “independent judgment” alongside “reasonable care, skill and diligence.”² That standard was built for a world of human-to-human delegation: a director relying on a subordinate, an auditor, or outside counsel. It says nothing about whether deferring to an algorithmically generated output satisfies this standard, or whether such deference amounts to an abdication of the very independent judgment the provision demands. Indian courts have long insisted that a director cannot function as a “mere ornamentation” on the board and must genuinely apply their mind to a transaction,³ and the Supreme Court has held that reliance on subordinates does not absolve a director of responsibility for the substance of what is ultimately placed before it.⁴ Neither principle, however, was formulated with an algorithmic intermediary in mind and it shows.

This paper argues that Section 166(3), as presently interpreted, is not equipped to answer whether, and under what conditions, a director’s reliance on an algorithmic output discharges the fiduciary duty of care. Section II surveys existing scholarship on Section 166(3) and identifies the near-total absence of Indian engagement with algorithmic reliance specifically. Section III sets out the doctrinal baseline through Indian case law, including the layered structure through which boards typically encounter these tools filtered through bankers, consultants, or law firms rather than confronted directly. Section IV turns to Delaware’s statutory reliance safe harbour under DGCL § 141(e) as a comparative lens. Section V builds a four-factor test materiality and proportionality, methodology comprehension, layered verification, and documentation and applies it to a hypothetical restructuring transaction. Section

¹ See generally Umakanth Varottil, 'Fiduciary Duties of Directors under Indian Company Law' (2015) NLS Business Law Review 1.

² Companies Act 2013, s 166(3).

³ *Official Liquidator, Supreme Bank Ltd v PA Tendolkar (Dead) by LRs* [1973] 3 SCR 364 (SC).

⁴ *N Narayanan v Adjudicating Officer, SEBI* [2013] 6 SCR 391 (SC).

VI closes with regulatory recommendations for SEBI and the Ministry of Corporate Affairs.

Literature Review

Indian scholarship on Section 166(3) of the Companies Act, 2013, is not thin. Varottil has traced how the codified duties of the 2013 Act moved Indian director liability away from a purely common-law footing while retaining much of its earlier flexibility in practice.⁵ Commentators writing on the “independent judgment” limb of Section 166(3) have generally treated it as a procedural safeguard a requirement that a director’s assent to a transaction reflect genuine deliberation rather than rubber-stamping, echoing the older common-law position in *Tendolkar*.⁶ This body of work is fundamental and largely settled; it forms the doctrinal floor for Section III of this paper.

What is missing from that literature is any sustained treatment of how a director is meant to exercise independent judgment when the information placed before the board has already been produced, filtered, or scored by an algorithmic system. Following corporate governance failures like the Satyam scandal and the subsequent regulatory response to the IL&FS crisis, Indian scholars extensively debated the limits of a board’s reliance on *human* experts, such as statutory auditors and external valuers.⁷ However, existing commentary has yet to extend this skepticism to technology. It says nothing about where the line falls between legitimate reliance on a tool-assisted report and an abdication of judgment to the tool itself.

The gap is less stark internationally, though it remains recent and largely unsettled even there. Li has argued for extending fiduciary-style obligations to AI systems themselves where they function as de facto decision-makers within a firm a position more radical than anything under discussion in Indian company law, but useful for framing the outer edge of the debate.⁸ Closer to the practical concerns of this paper, Montagnani and Passador argue that the EU’s AI Act, though not directly addressed to corporate boards, is already reshaping the practical content of directors’ duties of care. By imposing expectations of traceability and oversight over the algorithmic systems a board relies on, it effectively raises the evidentiary bar for what counts as informed judgment.⁹ A smaller strand of Indonesian scholarship has approached the same problem from the opposite direction, asking whether an AI system could itself occupy a directorial role, and concluding that it presently cannot, though it may properly function as an advisory input to the board.¹⁰ None of this literature, Indian or comparative, addresses the narrower and more immediate question this paper takes up: not whether AI can be a fiduciary, but what Section 166(3) requires of a human director who relies on one. It is that gap the remainder of this paper attempts to close.

Section 166(3) Doctrinal Foundation and the “Expert” Loophole

Section 166(3) is deceptively short. It requires only that “a director of a company shall exercise his duties with due and reasonable care, skill and diligence and shall exercise independent judgment.”¹¹ The provision does not define “independent judgment,” nor does it say what counts as adequate care where a director has relied, in whole or in part, on someone else’s analysis. That gap has always existed the 2013 Act codified duties that were previously scattered across a century of common law, and it did so without resolving every question the older case law had left open.¹²

What Indian courts have supplied, in the absence of statutory detail, is a consistent refusal to let a director discharge the duty of care by pointing to someone else’s work. In *Official Liquidator, Supreme Bank Ltd v. P.A. Tendolkar*, the Supreme Court held that a director closely associated with a company’s management cannot escape responsibility merely because no specific act of dishonesty is proved against him personally; he “cannot shut his eyes” to what would be obvious on even a superficial examination of the company’s affairs.¹³ Four decades later, in *N. Narayanan v. Adjudicating Officer, SEBI*, a whole-time director argued that he bore no responsibility for fabricated financial results because the accounts had been prepared by professional staff and duly audited, and his own role was confined to human resources. The Supreme Court rejected this defence outright, holding that a director’s responsibility for the accuracy of what is placed before the board does not evaporate because the underlying work was done by someone else with greater technical competence.¹⁴ Read together, *Tendolkar* and *Narayanan* stand for a single proposition: reliance on a subordinate or an expert is not, by itself, a defence. It only becomes one where the director engaged with what was produced, rather than simply signing off on it.

This is where the doctrine runs into a structural problem that Indian scholarship has not yet confronted. Both *Tendolkar* and *Narayanan* involve a director

⁵ Umakanth Varottil, 'The Evolution of Corporate Law in Post-Colonial India: From Transplant to Autochthony' (2016) 31 American University International Law Review 253.

⁶ *Official Liquidator, Supreme Bank Ltd v PA Tendolkar (Dead) by LRs* [1973] 3 SCR 364 (SC); see also A Ramaiya, *Guide to the Companies Act* (19th edn, LexisNexis 2020).

⁷ Jayati Sarkar and Subrata Sarkar, *Corporate Governance in India* (SAGE Publications 2012) ch 7.

⁸ Zhaoyi Li, 'Artificial Fiduciaries' (2024) University of Pittsburgh School of Law Working Paper.

⁹ Maria Lilla Montagnani and Maria Lucia Passador, 'Fiduciary Duties and the Business Judgment Rule 2.0 in the AI Act Age' (Stanford-TTLF Working Paper No 135, 2026).

¹⁰ Yolanda Simbolon, 'Considering the Use of Artificial Intelligence in the Board of Directors: A Corporate Law Review' (2025) Proceedings of ICRTLAW.

¹¹ Companies Act 2013, s 166(3).

¹² Varottil (n 1).

¹³ *Official Liquidator, Supreme Bank Ltd v PA Tendolkar* (n 2).

¹⁴ *N Narayanan v Adjudicating Officer, SEBI* (2013) 12 SCC 152.

relying on identifiable human intermediaries an auditor, an accounts department, a managing director whose work could, at least in principle, be interrogated, questioned, or second-guessed by a sufficiently attentive board member. An algorithmic tool sits differently in this chain. A financial advisor's valuation model, or a due-diligence platform's risk score, is rarely presented to the board as a raw output; it arrives pre-filtered through the advisor's own report, already narrated and contextualised in prose the board is used to reading. The board's engagement, in other words, is almost always with the advisor's account of what the tool produced, not with the tool itself.

Call this the layered reliance problem. If a director's duty under Section 166(3) is satisfied by genuinely engaging with what is placed before the board as *Tendolkar* and *Narayanan* require it becomes genuinely unclear where that engagement is meant to stop when an algorithm sits one layer behind the human expert the director is entitled to question. Is it enough for the board to interrogate the advisor's conclusions, or does independent judgment now require some engagement with the tool that generated them? This ambiguity is particularly consequential for independent directors, whose liability under Section 149(12) is already conditioned on whether they "acted diligently" a standard the provision leaves largely undefined, and one that becomes harder to satisfy in the layered-reliance scenario this paper describes. If diligence requires meaningful engagement with what is placed before the board, it is unclear what engaging diligently with an algorithmically-derived report actually demands.¹⁵

The following section turns to Delaware's DGCL § 141(e) to examine how a jurisdiction with a highly developed statutory safe harbour for expert reliance resolves this tension.

Comparative Analysis: DGCL § 141(e)

To understand how Indian corporate law might resolve the layered reliance problem, it helps to look at a jurisdiction that has actually legislated around it. Delaware, as the primary forum for most large American corporate litigation, offers the most extensively tested comparator for India's own maturing transactional practice. Where the Companies Act leaves the mechanics of expert reliance almost entirely to judicial interpretation, the Delaware General Corporation Law (DGCL) addresses it directly, by statute.

Section 141(e) of the DGCL provides that directors are "fully protected" in relying in good faith on the records of the corporation, or on information, opinions, reports, or statements presented by corporate officers, employees, or outside experts, provided the director reasonably believes the matter falls within that person's professional competence and the person was selected with reasonable care.¹⁶

Point of Comparison	India: Section 166(3)	Delaware: DGCL § 141(e)
Statutory Standard	Demands "independent judgment" and "due and reasonable care," leaving the boundaries of reliance undefined by statute	Explicitly provides a safe harbour for good-faith reliance on corporate officers and outside experts
Protection Level	Open-ended; relies on judicial interpretation of whether a director "shut his eyes"	Facially protective, but narrower in application than its text suggests
The Algorithmic Ambiguity	Unclear whether an algorithmic output compromises "independent judgment"	Unclear whether an algorithm qualifies as a "person" under the provision, or whether reliance on a black-box tool is genuine reliance at all

On its face, Section 141(e) offers a cleaner template than anything Section 166(3) currently provides. But two qualifications, both drawn from how Delaware courts have actually applied the provision, complicate the comparison — and matter directly to the algorithmic question this paper is asking. The first is textual: Section 141(e) protects reliance on a "person" — an officer, employee, or outside expert. An algorithm is neither. A board cannot invoke the safe harbour directly to protect reliance on a standalone algorithmic tool, however sophisticated; the statute was written for a delegation chain that ends in a human being.¹⁷

The second qualification is more instructive, and comes from how Delaware courts have twice narrowed the protection even where a genuine human expert sits at the end of that chain. In *Valeant Pharmaceuticals International v Jerney*, a director sought to invoke Section 141(e) after relying on a compensation consultant's report to justify a self-interested bonus decision. The Court of Chancery declined to treat the reliance as protective, in part because the consultant had been retained only after the bonus figures were effectively already decided, and was tasked with justifying an outcome rather than independently arriving at one.¹⁸ The reliance, in other words, was reliance in form only the report supplied a post-hoc narrative for a decision already made. That is precisely the failure mode this paper is concerned with when a black-box model sits behind a human advisor's report: a narrated justification is not an independent analysis, whether the narrator is a consultant or an algorithm the consultant is fronting for.

Separately, in *In re Rural Metro Corp Stockholders Litigation*, Section 141(e) was not itself in dispute, but the Court of Chancery observed that a director

¹⁵ Companies Act 2013, s 149(12).

¹⁶ Del Code Ann tit 8, § 141(e).

¹⁷ See generally Li (n 4).

¹⁸ *Valeant Pharmaceuticals International Inc v Jerney* 921 A 2d 732 (Del Ch 2007).

could follow expert advice, be found to have breached the duty of care in the underlying decision, and still be “fully protected” from *monetary liability* for that breach.¹⁹ Commentators have since read this as narrowing what “full protection” means in practice: it may shield a director’s pocket without ever validating the reasonableness of the underlying decision.²⁰ For a comparative framework meant to inform Indian reform, this matters even Delaware’s codified safe harbour does not resolve the reliance question so much as decide who pays once reliance turns out to have been misplaced.

Read together, *Valeant* and *Rural Metro* suggest that Delaware’s safe harbour, robust as it looks on paper, already strains under ordinary human-expert reliance once the expert’s independence or methodology is genuinely in question. India’s judicially constructed “independent judgment” standard offers no comparable protection to begin with. Layering an algorithmic intermediary onto either regime does not create a new problem so much as sharpen an old one. The following section builds a four-factor test that attempts to give this question a workable shape.

Applying the Framework — The Restructuring Hypothetical

To bridge the gap between Section 166(3)’s demand for “independent judgment” and the reality of algorithmic intermediaries, courts and regulators require a functional standard of review. This section proposes a four-factor test to determine when a board’s reliance on an algorithmically-derived expert report crosses the line from legitimate delegation into an abdication of fiduciary duty.

Consider the following hypothetical. The board of a listed Indian entity, TargetCo, proposes a complex demerger of its high-growth data analytics division under Section 230 of the Companies Act, 2013.²¹ To determine the share entitlement ratio, the board retains a premier investment bank. The bank presents a valuation report that heavily utilizes a proprietary, black-box machine learning model (“Valu-ML”), trained on a decade of global tech spin-offs, to project post-demerger synergies and risk weightings. The board reviews the bank’s final prose report and approves the scheme. Six months post-demerger, it emerges that Valu-ML systematically misweighted Indian macroeconomic regulatory friction because its training data was heavily skewed toward US transactions. Massive shareholder value is destroyed, and derivative litigation ensues alleging a breach of Section 166(3).

Under the proposed framework, a court or regulator would evaluate the board’s liability by interrogating four distinct factors.

- i. **Materiality and Proportionality.** The level of scrutiny a director must apply to an algorithmic output should scale with the materiality of the transaction. A Section 230 demerger fundamentally alters corporate capital structure and minority shareholder rights. In such existential transactions, the requirement for active, independent judgment is at its zenith. If the board were merely using an algorithmic tool to screen routine vendor contracts or low-value operational compliance, a higher degree of passive deference would be permissible. However, when an algorithm dictates a share exchange ratio or asset valuation, the board cannot rely on the tool’s output without intense scrutiny.
- ii. **Methodology Comprehension (The “Red Flag” Inquiry).** Directors are not expected to understand the mathematical architecture or source code of a neural network. However, to satisfy Section 166(3), they must comprehend the fundamental parameters, assumptions, and limitations of the tool deployed by the expert. In the TargetCo hypothetical, a diligent director would not ask to inspect Valu-ML’s proprietary code; rather, they would interrogate the investment banker on what data the model was trained on, whether it accounted for Indian regulatory timelines (such as National Company Law Tribunal approvals), and what its margin of error is in volatile emerging markets. A failure to ask these basic scoping questions renders the reliance legally unreasonable.²²
- iii. **Layered Verification (The Human Stress-Test).** This factor directly resolves the *Valeant* problem identified in Section IV. A board must ascertain whether the human expert actually verified the algorithm’s output, or merely provided a post-hoc narrative justification for a computer-generated number. The board’s fiduciary duty requires probing the human filter: did the analysts stress-test these algorithmic assumptions against conventional discounted cash flow or asset-based valuation methods? If the investment bank effectively acted as a printing press for the Valu-ML output, the human expert has exercised no independent judgment, and the board cannot safely rely on the expert’s report. The chain of reliance breaks at the expert level.²³
- iv. Historically, Indian board minutes in M&A transactions are notoriously sterile, frequently recording only that the board “perused the valuation report and approved the scheme.” Under this proposed framework, satisfying the duty of care requires the minutes to explicitly document the board’s interrogation of the algorithmic reliance — that the board actively questioned the investment bankers about the tool’s limitations, its training data biases, and the human verification processes applied to it.

Applying the framework to TargetCo’s board, the picture is not favourable. The transaction was unambiguously material (factor one), yet nothing in the facts suggests the board asked the investment bank about Valu-ML’s training data or its geographic scope (factor two). There is no indication the bank’s human analysts cross-checked Valu-ML’s output against a conventional DCF model (factor three), and the board minutes recorded only that the report

¹⁹ *In re Rural Metro Corp Stockholders Litigation* 88 A 3d 54, n 13 (Del Ch 2014).

²⁰ Peter A Atkins, ‘Reliance by Directors: What’s a Conscientious Director to Do?’ (Skadden, Arps, Slate, Meagher & Flom LLP, 2014).

²¹ Companies Act 2013, s 230.

²² Cf *Brehm v Eisner* 746 A 2d 244, 262 (Del 2000)

²³ Cf *Valeant Pharmaceuticals International Inc v Jerney* (n 14) 752.

was “perused and approved” (factor four). On these facts, the board’s conduct would fail three of the four factors precisely the kind of unexamined algorithmic deference Section 166(3), properly read, should not permit. The board’s error was not retaining Valu-ML; it was treating the bank’s prose summary as a substitute for the scrutiny a transaction of this magnitude demanded.

. A foreseeable objection to this framework is that subjecting algorithmic reliance to heightened fiduciary scrutiny will chill technological adoption in Indian boardrooms, forcing companies to rely entirely on slower, conventional analytics. The first factor is proportionality directly answers this concern. The framework does not penalise the use of AI; it penalises the unexamined use of AI in highly material transactions. By requiring boards to simply ask human experts to explain and justify their algorithmic tools, the framework preserves the efficiency gains of data-driven technology while maintaining the human accountability that Section 166(3) fundamentally demands.

Regulatory Recommendations

A purely judicial resolution to the layered reliance problem is inadequate. Courts intervene only post facto, usually after shareholder value has already been destroyed. A proactive corporate governance solution requires ex-ante regulatory intervention. Because the deployment of bespoke, black-box algorithmic tools is currently capital-intensive and heavily concentrated among large, publicly traded companies, the regulatory response should begin with the Securities and Exchange Board of India (SEBI) rather than a blunt, market-wide amendment to the Companies Act.

SEBI’s Business Responsibility and Sustainability Reporting (BRSR) framework, mandatory for the top 1,000 listed entities by market capitalisation, offers a ready-made architectural template.²⁴ The BRSR already forces boards to quantify, standardise, and disclose complex, non-financial operational risks. SEBI should introduce an “Algorithmic Governance” parameter within this reporting framework, tied to a two-tier disclosure mechanism under the SEBI (Listing Obligations and Disclosure Requirements) Regulations, 2015.

At the periodic level, the annual BRSR filing should require companies to disclose their overarching board oversight policies regarding technology and algorithmic decision-support systems.

At the transactional level, where algorithmic valuation is most consequential specifically, a Scheme of Arrangement under Regulation 37 or a material Related Party Transaction under Regulation 23 disclosure must be event-driven.²⁵ Whenever a board approves a material restructuring in substantial reliance on an algorithmically-derived valuation or risk model, the explanatory statement and the valuation report submitted to the stock exchanges must disclose:

1. **Tool Typology:** Whether the advisory input relied upon a deterministic rule-based system or an opaque machine-learning model;
2. **Human Verification:** A declaration from the human expert (e.g., the registered valuer or merchant banker) confirming they independently stress-tested the algorithmic output against conventional methodologies; and
3. **Board Engagement:** Confirmation that the board was apprised of the tool’s primary data dependencies, training biases, and material limitations.

This approach adapts the risk-tiering logic of the EU AI Act—under which high-risk deployments trigger specific transparency obligations—without importing its heavier, cross-sectoral compliance machinery.²⁶ By nesting these requirements into the existing filing workflows of Regulation 37 and Regulation 23, SEBI can enforce algorithmic transparency through channels boards and their advisors already navigate, transforming algorithmic accountability into a routine compliance check rather than an esoteric technology audit.

Amending the Secretarial Standards (SS-1) To operationalise the fourth factor of the proposed test Documentation of Deliberation the Institute of Company Secretaries of India (ICSI), with the approval of the Ministry of Corporate Affairs, should amend Secretarial Standard-1 (SS-1) on Meetings of the Board of Directors.²⁷

Crucially, compliance with SS-1 is not a voluntary best practice; it is a mandatory statutory obligation under Section 118(10) of the Companies Act, 2013. An amendment should mandate that whenever an expert report relying on algorithmic models is tabled for a material transaction, the minutes must explicitly record the board’s inquiries regarding the tool’s methodology and the expert’s verification process. This imposes a modest documentary burden, but it forces the “red flag” inquiry into the open. If a board knows the Company Secretary is statutorily bound to minute its algorithmic due diligence, directors will be compelled to ask the human expert the necessary scoping questions in the first place—neutralising the risk of passive rubber-stamping before a transaction is ever executed.

²⁴ Securities and Exchange Board of India, 'Business Responsibility and Sustainability Reporting by Listed Entities' (Circular SEBI/HO/CFD/CMD-2/P/CIR/2021/562, 10 May 2021); Securities and Exchange Board of India (Listing Obligations and Disclosure Requirements) Regulations 2015, reg 34(2)(f).

²⁵ Securities and Exchange Board of India (Listing Obligations and Disclosure Requirements) Regulations 2015, regs 23, 37.

²⁶ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L2024/1689, art 6, Annex III.

²⁷ Institute of Company Secretaries of India, *Secretarial Standard on Meetings of the Board of Directors* (SS-1); Companies Act 2013, s 118(10).

VII. Conclusion

The introduction of algorithmic tools into the Indian boardroom represents a fundamental stress test for corporate fiduciary duties. Section 166(3) of the Companies Act, 2013, was designed to police human-to-human delegation, demanding an exercise of “independent judgment” that courts have consistently interpreted as an active, interrogative duty. However, as this paper has demonstrated, this doctrinal framework fractures when confronted with the “layered reliance” problem—where a board relies on a human expert who, in turn, relies on a black-box algorithmic model. Neither Indian jurisprudence nor the comparative haven of Delaware’s statutory safe harbour currently provides a coherent standard for determining when reliance on an algorithmically-derived report constitutes informed judgment versus an abdication of duty.

The four-factor test proposed in this paper—weighing materiality, methodology comprehension, layered human verification, and explicit documentation—offers a functional standard of review. Crucially, by translating these factors into targeted regulatory amendments within the SEBI (LODR) Regulations, the BRSR framework, and the ICSI Secretarial Standards, this framework ensures that algorithmic accountability is operationalized as a routine governance mechanic rather than left to the unpredictable ex-post scrutiny of derivative litigation.

This analysis is not without limitations. The proposed standard remains untested before the National Company Law Tribunal or the Indian Supreme Court, and its application to highly technical valuation disputes relies on a judiciary equipped to parse the nuances of algorithmic deployment. Furthermore, the framework presented here is strictly bounded to the present reality, wherein algorithms function as *advisory* inputs.

As corporate technology inevitably matures from advisory machine-learning models to autonomous AI agents capable of executing transactions independently, the legal inquiry will shift. The question will no longer be how a director must *rely* on an algorithmic tool, but whether a director can lawfully *delegate* their fiduciary authority to one entirely. Until that paradigm shift occurs, Indian corporate law must address the tools currently in the boardroom. Regulators cannot freeze technological adoption, but by insisting on rigorous, documented human verification, they can ensure that the ultimate responsibility for a company’s trajectory remains firmly in human hands.