INTERNATIONAL JOURNAL FOR LEGAL RESEARCH AND ANALYSIS

Open Access, Refereed Journal Multi Disciplinary
Peer Reviewed



www.ijlra.com

ISSN: 2582-6433

DISCLAIMER

No part of this publication may be reproduced, stored, transmitted, or distributed in any form or by any means, whether electronic, mechanical, photocopying, recording, or otherwise, without prior written permission of the Managing Editor of the *International Journal for Legal Research & Analysis (IJLRA)*.

The views, opinions, interpretations, and conclusions expressed in the articles published in this journal are solely those of the respective authors. They do not necessarily reflect the views of the Editorial Board, Editors, Reviewers, Advisors, or the Publisher of IJLRA.

Although every reasonable effort has been made to ensure the accuracy, authenticity, and proper citation of the content published in this journal, neither the Editorial Board nor IJLRA shall be held liable or responsible, in any manner whatsoever, for any loss, damage, or consequence arising from the use, reliance upon, or interpretation of the information contained in this publication.

The content published herein is intended solely for academic and informational purposes and shall not be construed as legal advice or professional opinion.

Copyright © International Journal for Legal Research & Analysis. All rights reserved.

ISSN: 2582-6433

ABOUT US

The *International Journal for Legal Research & Analysis (IJLRA)* (ISSN: 2582-6433) is a peer-reviewed, academic, online journal published on a monthly basis. The journal aims to provide a comprehensive and interactive platform for the publication of original and high-quality legal research.

IJLRA publishes Short Articles, Long Articles, Research Papers, Case Comments, Book Reviews, Essays, and interdisciplinary studies in the field of law and allied disciplines. The journal seeks to promote critical analysis and informed discourse on contemporary legal, social, and policy issues.

The primary objective of IJLRA is to enhance academic engagement and scholarly dialogue among law students, researchers, academicians, legal professionals, and members of the Bar and Bench. The journal endeavours to establish itself as a credible and widely cited academic publication through the publication of original, well-researched, and analytically sound contributions.

IJLRA welcomes submissions from all branches of law, provided the work is original, unpublished, and submitted in accordance with the prescribed submission guidelines. All manuscripts are subject to a rigorous peer-review process to ensure academic quality, originality, and relevance.

Through its publications, the *International Journal for Legal Research & Analysis* aspires to contribute meaningfully to legal scholarship and the development of law as an instrument of justice and social progress.

ISSN: 2582-6433

PUBLICATION ETHICS, COPYRIGHT & AUTHOR RESPONSIBILITY STATEMENT

The *International Journal for Legal Research and Analysis (IJLRA)* is committed to upholding the highest standards of publication ethics and academic integrity. All manuscripts submitted to the journal must be original, unpublished, and free from plagiarism, data fabrication, falsification, or any form of unethical research or publication practice. Authors are solely responsible for the accuracy, originality, legality, and ethical compliance of their work and must ensure that all sources are properly cited and that necessary permissions for any third-party copyrighted material have been duly obtained prior to submission. Copyright in all published articles vests with IJLRA, unless otherwise expressly stated, and authors grant the journal the irrevocable right to publish, reproduce, distribute, and archive their work in print and electronic formats. The views and opinions expressed in the articles are those of the authors alone and do not reflect the views of the Editors, Editorial Board, Reviewers, or Publisher. IJLRA shall not be liable for any loss, damage, claim, or legal consequence arising from the use, reliance upon, or interpretation of the content published. By submitting a manuscript, the author(s) agree to fully indemnify and hold harmless the journal, its Editor-in-Chief, Editors, Editorial Board, Reviewers, Advisors, Publisher, and Management against any claims, liabilities, or legal proceedings arising out of plagiarism, copyright infringement, defamation, breach of confidentiality, or violation of third-party rights. The journal reserves the absolute right to reject, withdraw, retract, or remove any manuscript or published article in case of ethical or legal violations, without incurring any liability.

ISSN: 2582-6433

FROM BOARDROOM TO BOT: HOW ARTIFICIAL INTELLIGENCE IS CHANGING CORPORATE GOVERNANCE

AUTHORED BY - TRISHA CHAKRABORTY & HRISITA SAHA

Introduction

Artificial intelligence is changing how corporate boards gather, test, and act on information, and corporate law has barely begun to ask who answers when a machine does part of the thinking. This paper traces that shift from the traditional boardroom, where Delaware's fiduciary duties of care, loyalty, and oversight assumed that every participant was a person who could be questioned or sued, to the arrival of machine learning and generative AI in business decisions. It examines how boards use AI today, how Delaware's reliance doctrine applies to AI output, and how continuous monitoring can strengthen oversight under *Caremark* and *Marchand*. It then analyzes the risks of bias, opacity, unreliable outputs, and misleading disclosure, and surveys emerging regulation and standards, including the EU AI Act, the NIST framework, and ISO/IEC 42001. The paper argues that AI changes how boards obtain and check information, but the ultimate duty of oversight remains with human directors, and it closes with practical steps for boards to document that oversight.

For most of the history of the modern corporation, running a board meant relying on what people knew and how they judged things. Directors read the materials management sent them, questioned officers and advisers, and then decided. That is changing. Software can now condense a board package of hundreds of pages into a few, flag payments that look odd, and test how a strategic plan would perform under different economic conditions. The change leaves the law with a question it has barely begun to answer: if a machine does part of the thinking, who is responsible for the result?[1]

This paper follows the story step by step. Part I describes how boards and fiduciary law worked before artificial intelligence (AI). Part II traces how AI entered business decision-making. Part III examines what AI is used for in boardrooms today and how Delaware's reliance doctrine applies to it. Part IV considers AI's promise for oversight, and Part V its dangers: bias, opacity, and misleading disclosure. Part VI surveys how regulators and standard-setters are responding and what boards should do about it. My main point is simple. AI changes how boards get and

ISSN: 2582-6433

check information. The final duty of oversight still sits with the directors, who are human.

I. The Traditional Boardroom: How Corporate Governance Worked Before AI

A. The Board's Place in the Company

Governance is simply the way a company is directed and kept accountable, and the board of directors sits at the center of it. Boards exist because of a problem that Adolf Berle and Gardiner Means described in the 1930s: when a company's owners are spread across thousands of shareholders, none of them can watch the managers closely, so someone has to do it for them.[2] In Delaware, where many of the largest American companies are incorporated, the statute places the management of the company's business and affairs under the direction of its board.[3]

B. The Duty of Care

With that power come fiduciary duties, and the two principal ones are care and loyalty. Courts ordinarily defer to board decisions under the business judgment rule, a presumption that directors acted on an informed basis, in good faith, and in the honest belief that the action served the company.[4] The presumption depends on the word "informed." The duty of care requires directors to inform themselves of the material information reasonably available before they decide anything. In *Smith v. Van Gorkom*, the directors approved a cash-out merger after a brief meeting, without reading the merger agreement and without any valuation study, and the Delaware Supreme Court held that they had been grossly negligent and had failed this duty.[5] The case remains the classic warning about decisions made on thin information. It matters here because AI tools make information easier to gather and also make it easier to accept a conclusion on trust.

C. The Duty of Oversight

A third strand, the duty of oversight, comes from *In re Caremark International Inc. Derivative Litigation*. It asks directors to make a real, good-faith effort to put in place a system that brings problems to the board's attention.[6] The Delaware Supreme Court later treated this duty as part of the duty of loyalty, so liability requires bad faith: either a sustained failure to implement any reporting system, or a conscious decision to ignore red flags the system produced.[7] In *Marchand v. Barnhill*, the board of an ice cream maker had no committee, process, or reporting

ISSN: 2582-6433

line for food safety, which was the company's central compliance risk, and the Delaware Supreme Court reversed the dismissal of claims after a listeria outbreak.[8] Bad faith in this setting means an intentional dereliction of duty or a conscious disregard of known responsibilities, not mere error.[9] Later decisions of the Court of Chancery have applied the same logic to other "mission critical" safety risks and have extended oversight duties to corporate officers.[10] [11]

D. Federal Law

Federal law added further layers. The Sarbanes-Oxley Act of 2002 required public companies to have independent audit committees, required chief executive and financial officers to certify the accuracy of financial reports, and required management to assess internal controls over financial reporting.[12]

Behind all these rules sits an idea that nobody bothered to say out loud. Directors got their information from managers, auditors, and advisers, and then used their own judgment on it. Everyone in that chain was a person who could be questioned, replaced, or sued. That assumption was never written into a statute because it never needed to be. AI puts it under strain.

II. Enter the Machine: How AI First Found Its Way into Business Decisions

A. From Records to Rules to Learning

Computers helped with business decisions long before the current excitement about AI. They began as tools for storing records and printing reports, and over time they grew into decision-support software and rule-based programs that handled routine choices automatically. The bigger step arrived with machine learning, in which a system finds patterns in large amounts of data without a programmer specifying each task. During the 2010s companies began using it for credit scoring, fraud detection, algorithmic trading, forecasting, and hiring.

Bank regulators noticed early. In 2011 the Federal Reserve issued supervisory guidance on model risk management, which told banks to validate their models, document their limits, and give senior management and boards responsibility for the risk that models create.[13] That guidance is an early statement of a principle that still holds: using a model does not move accountability away from the institution that uses it.

ISSN: 2582-6433

B. Scholars Take Notice

At that stage AI mostly stayed out of the boardroom. Managers used it inside operations, and directors saw its results only secondhand, in dashboards and risk reports. Scholars nevertheless saw what was coming. Florian Möslein argued that as machines become more capable, corporate law will have to deal with algorithms performing tasks that once belonged to directors.[14] Mark Fenwick and Erik Vermeulen made the related point that digital technology was already changing how companies are structured and how they make decisions.[15]

C. The Generative Turn

The pace changed in late 2022, when OpenAI released ChatGPT and generative AI spread through workplaces within months.[16] Suddenly executives and directors had writing and analysis tools at their fingertips, sometimes before their companies had adopted any rules about using them. These tools also brought a new kind of failure. Generative models can produce fluent text that is simply false. In one widely reported case, lawyers were sanctioned for filing a brief that cited cases a chatbot had invented.[17] A director who relies on an AI-generated summary faces the same hazard, in a setting where an error may be even harder to spot.

III. When the Algorithm Takes a Seat at the Table

A. A Seat Without a Vote

A very well-known example dates to 2014. Deep Knowledge Ventures, a venture capital firm in Hong Kong, was reported to have given an algorithm called VITAL a seat on its board to help judge investments.[18] Part of that was publicity, since company law generally assumes that a director is a person. Delaware, for instance, requires every director to be a natural person.[19] Scholars have debated whether the entity laws of other jurisdictions could be used to build organizations that software effectively controls.[20] But the episode showed where the conversation was heading.

B. What Boards Actually Do with AI

Today what happens in practice is more modest and far more common. Directors and staff use AI to shorten board packs, compare the company with its competitors, test plans under different scenarios, and draft minutes. Researchers describe a range that starts with tools that simply help directors, moves through systems that advise them, and ends with systems that would decide on their own.[21] Right now most real use sits near the helping end. A director who asks a chatbot to summarize a 300-page diligence report is in a very different legal position from a

ISSN: 2582-6433

board that lets software approve a transaction, and the law should treat them differently.

C. The Reliance Doctrine

Delaware law fits this picture. The statute protects directors who rely in good faith on the records of the corporation and on reports of officers, employees, and outside experts chosen with reasonable care.[22] Courts have insisted that reliance be informed: a director cannot shelter behind a report on a matter outside the reporter's competence, or behind a report that the director did not take the trouble to understand.[23] No one has yet settled whether an AI system amounts to an "expert" of this kind, and the statutory language, which speaks of persons, does not obviously cover software. The idea behind reliance suggests a workable answer. Directors can treat what an AI produces as one input among several, tested against other sources, without handing over their responsibility for the final decision. The algorithm may have a seat at the table. It has no vote, and it cannot be held liable.

D. Practical Frictions

Several everyday problems accompany this use. Board materials are often confidential, privileged, or market-sensitive, and feeding them into a third-party tool may expose them. Prompts and AI-generated drafts may themselves become records that a litigant could seek in discovery. And a board that cannot say which tools its directors and officers use cannot say what it is relying on. None of these problems requires new legal doctrine, but each calls for a policy, and in most companies the policy is still missing.

IV. Reading the Room with Data: Smarter Oversight and Risk Spotting

A. Continuous Watching

The strongest reason for using AI in governance is oversight. No director can personally check every transaction, contract, or email, and the usual reports arrive now and then, after management has already shaped them. An AI system can watch the data all the time. It can flag unusual entries in the accounts, read through contracts and internal messages looking for compliance problems, and keep an eye on outside developments such as supplier trouble, cyber threats, or new regulations. Audit committees can use it to test whole populations of transactions rather than samples, and deal teams can use it to review thousands of contracts during due diligence. Problems may therefore come to light sooner, and the board depends a little less on management's own account of events.

ISSN: 2582-6433

B. Why It Matters Legally

This matters in law as well. *Caremark* and *Marchand* expect boards to keep systems that bring risks to their attention, and AI monitoring can be part of such a system.[24] A board that ignores tools it could easily use to track a risk may arguably have a harder time defending itself later. Courts have not said so, and *Caremark* liability still requires bad faith rather than a failure to adopt the best available technology. The more realistic effect is gradual. As AI monitoring becomes standard in an industry, plaintiffs will point to it as evidence of what a reasonable reporting system looks like.

C. The Limits of Detection

The advantages rely on good data and a well-built model. A system trained on past records cannot spot what it has never seen, and too many false alarms can bury the warnings that genuinely count. There is a human problem too. Research on automation bias shows that people tend to over-trust automated recommendations and to search less for evidence that contradicts them.[25] A detection tool helps only if directors keep the habit of asking what the AI might have left out, and a dashboard that shows no problems is not proof that none exists. The board's role is to keep that skepticism alive.

V. Prejudice, Black Boxes, and the Transparency Problem

A. Bias

AI has the features that make it powerful, and the same features make it risky. One risk is bias. Models learn from data, and if that data carries old prejudices, a model can repeat them. Amazon reportedly gave up an experimental hiring tool after finding that it scored résumés linked to women lower.[26] Regulators are responding. New York City, for example, requires bias audits of automated employment decision tools.[27] For a board, a biased system is more than an ethical problem. It is a source of legal exposure and reputational harm that the oversight system should be built to catch.

B. Opacity

Another risk is opacity, meaning that nobody can see how an AI reached its answer. In *State v. Loomis*, the Wisconsin Supreme Court let a sentencing judge consider a proprietary risk score, but it warned that the score must be used carefully, partly because the company that built it kept its method secret.[28] That case is about sentencing, but its lesson reaches the boardroom: relying on a tool whose reasoning cannot be examined is hard to justify. Computer scientists

ISSN: 2582-6433

have gone further and argued that for high-stakes decisions organizations should prefer models that are interpretable by design over opaque models explained after the fact.[29] Boards need not adopt that view in full, but they can ask management why a given use case justifies a black box.

C. Accountability Stays with People

These risks clash with the duty of care. A director who cannot explain why an AI recommended something has a weak claim to having been properly informed, and *Van Gorkom* shows how courts view decisions made on little understanding.[30] So accountability stays with people. An algorithm is not a person and owes no fiduciary duties, which leaves responsibility with the company, its directors, and its officers. Delaware's direction cuts the same way. Officers owe oversight duties, and a 2022 amendment lets corporations exculpate officers from monetary liability for breach of the duty of care in direct claims, though not in derivative suits brought on the company's behalf.[31] More individuals may therefore have to explain how AI systems were chosen and supervised.

D. Honest Disclosure

Honesty in disclosure matters as well. In March 2024 the Securities and Exchange Commission charged two investment advisers with making false and misleading statements about their use of AI, a practice that has come to be called "AI washing."[32] Public companies face the same risk under the general antifraud rules.[33] Tribunals also tend to treat what a company's automated tools say as the company's own words; a Canadian tribunal held an airline responsible for incorrect information given by its website chatbot.[34] The board needs to be sure that what the company says about its AI matches what it really does, and that disclosure controls reach AI-related claims.

E. Reliability and Dependence

Finally, there are quieter risks. Models can degrade as the world changes and the data they were trained on becomes stale. Most companies buy their AI from a small number of vendors, so a vendor's outage, security breach, or change in terms can become the company's own crisis. A board that asks only whether the system works today, and not what happens when it stops working or when the vendor changes it, has asked half the question.

ISSN: 2582-6433

VI. Rewriting the Rulebook: Regulation, Accountability, and the Future-Ready Director

A. The Regulatory Landscape

Lawmakers and standard-setters are now building rules for these risks. The European Union's AI Act is risk-based: the strictest requirements fall on high-risk systems and cover risk management, data governance, documentation, and human oversight.[35] [36] The requirements take effect in stages over several years. European data-protection law adds a further limit by restricting significant decisions based solely on automated processing.[37] The United States has relied on sector-specific rules and voluntary guidance, such as the framework from the National Institute of Standards and Technology, which organizes good practice into governing, mapping, measuring, and managing AI risk.[38] There is also an international standard, ISO/IEC 42001, which sets out what an AI management system should include.[39]

B. What Boards Should Do

I believe that inside companies the wise course is to manage the tools themselves. In practice that means:

1. Giving AI oversight to the full board or to a committee, such as audit or risk, and recording that allocation in the committee's charter.
2. Keeping an inventory of the AI systems in use, including those embedded in vendor products.
3. Asking management to report regularly on testing, incidents, and vendors.
4. Setting rules for when high-impact decisions are escalated to a human, especially decisions that materially affect customers, employees, or financial reporting.
5. Reviewing public statements about AI before they are released.

These steps turn principles into a documented process, and that is exactly what courts look for when they judge oversight. The absence of any such process is what sank the board in *Marchand*.

Oversight should also be proportionate. A tool that drafts routine correspondence deserves less scrutiny than one that screens job applicants, prices credit, or monitors product safety. Sorting uses by risk lets the board spend its limited attention where an error would do the most damage, which is also how the leading frameworks are organized.

ISSN: 2582-6433

C. The Future-Ready Director

I also see that the role of a director is changing. Directors do not need to write code. They need enough technical understanding to question where the data comes from, what the model cannot do, and how it might fail. They need the judgment to weigh efficiency against fairness and privacy, and they need the humility to keep learning. Regular briefings, outside experts who report directly to the board, and a willingness to ask basic questions without embarrassment are inexpensive ways to build that capacity.

D. Should Machines Decide?

Some scholars go further and ask whether autonomous governance might eventually outperform human boards, which are limited by time, bias, and conflicts of interest.[40] The question is a serious one, but present law leaves no room for it. Fiduciary duties presuppose a human who can be held to account, and the legal devices that might let software hold control remain untested in this setting.[41] Until that changes, the more productive project is to design human oversight that works.

Conclusion

To sum up, AI is changing governance by changing how information is collected, analyzed, and tested. Boards that use it well will see problems sooner, test their assumptions more thoroughly, and rely less on management's own account of events. Boards that use it carelessly will inherit new risks of bias, opacity, and misstatement without gaining the benefit of better information.

AI is not replacing the directors who answer for the outcome. The law still assumes that humans make the decisions and that humans are responsible for them. The boards that will be best prepared are the ones that use machines to see more while still insisting on understanding what the machines show them.

Notes

[1] *See generally* John Armour & Horst Eidenmüller, *Self-Driving Corporations?*, 10 Harv. Bus. L. Rev. 87 (2020).

[2] Adolf A. Berle, Jr. & Gardiner C. Means, *The Modern Corporation and Private Property* (1932).

[3] Del. Code Ann. tit. 8, § 141(a) (2025).

ISSN: 2582-6433

[4] Aronson v. Lewis, 473 A.2d 805, 812 (Del. 1984).

[5] Smith v. Van Gorkom, 488 A.2d 858, 872–74 (Del. 1985).

[6] *In re* Caremark Int'l Inc. Derivative Litig., 698 A.2d 959, 967–70 (Del. Ch. 1996).

[7] Stone v. Ritter, 911 A.2d 362, 369–70 (Del. 2006).

[8] Marchand v. Barnhill, 212 A.3d 805, 809, 821–24 (Del. 2019).

[9] *In re* Walt Disney Co. Derivative Litig., 906 A.2d 27, 67 (Del. 2006).

[10] *In re* Boeing Co. Derivative Litig., No. 2019-0907-MTZ, 2021 WL 4059934 (Del. Ch. Sept. 7, 2021).

[11] *In re* McDonald's Corp. Stockholder Derivative Litig., 289 A.3d 343 (Del. Ch. 2023).

[12] Sarbanes-Oxley Act of 2002, Pub. L. No. 107-204, 116 Stat. 745 (codified as amended at 15 U.S.C. §§ 78j-1(m), 7241, 7262).

[13] Bd. of Governors of the Fed. Rsrv. Sys., SR 11-7: Guidance on Model Risk Management (Apr. 4, 2011).

[14] Florian Möslein, *Robots in the Boardroom: Artificial Intelligence and Corporate Law*, *in* Research Handbook on the Law of Artificial Intelligence 649 (Woodrow Barfield & Ugo Pagallo eds., 2018).

[15] Mark Fenwick & Erik P.M. Vermeulen, *Technology and Corporate Governance: Blockchain, Crypto, and Artificial Intelligence*, 48 Tex. J. Bus. L. 1 (2019).

[16] *Introducing ChatGPT*, OpenAI (Nov. 30, 2022), https://openai.com/index/chatgpt/.

[17] Mata v. Avianca, Inc., 678 F. Supp. 3d 443 (S.D.N.Y. 2023).

[18] Rob Wile, *A Venture Capital Firm Just Named an Algorithm to Its Board of Directors—Here's What It Actually Does*, Bus. Insider (May 13, 2014).

[19] Del. Code Ann. tit. 8, § 141(b) (2025).

[20] *See* Shawn Bayern, *The Implications of Modern Business-Entity Law for the Regulation of Autonomous Systems*, 19 Stan. Tech. L. Rev. 93 (2015).

[21] *See generally* Martin Petrin, *Corporate Management in the Age of AI*, 2019 Colum. Bus. L. Rev. 965.

[22] Del. Code Ann. tit. 8, § 141(e) (2025).

[23] Brehm v. Eisner, 746 A.2d 244, 261–62 (Del. 2000); *see also Van Gorkom*, 488 A.2d at 875.

[24] *Marchand*, 212 A.3d at 824; *Caremark*, 698 A.2d at 970; *see also Stone*, 911 A.2d at 370.

[25] Kate Goddard et al., *Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators*, 19 J. Am. Med. Informatics Ass'n 121 (2012).

[26] Jeffrey Dastin, *Insight—Amazon Scraps Secret AI Recruiting Tool That Showed Bias*

ISSN: 2582-6433

Against Women, Reuters (Oct. 10, 2018).

[27] N.Y.C. Admin. Code §§ 20-870 to -874.

[28] State v. Loomis, 881 N.W.2d 749, 760–64 (Wis. 2016).

[29] Cynthia Rudin, *Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead*, 1 Nature Mach. Intel. 206 (2019).

[30] *See Van Gorkom*, 488 A.2d at 872–74.

[31] Del. Code Ann. tit. 8, § 102(b)(7) (2025); *see also McDonald's*, *supra* note 11.

[32] Press Release, SEC, SEC Charges Two Investment Advisers with Making False and Misleading Statements About Their Use of Artificial Intelligence (Mar. 18, 2024).

[33] *See* 17 C.F.R. § 240.10b-5 (2025).

[34] Moffatt v. Air Canada, 2024 BCCRT 149 (Can. B.C. Civ. Resol. Trib.).

[35] Regulation 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence, 2024 O.J. (L 2024/1689) 1.

[36] *See id.* arts. 4, 9–11, 14.

[37] Regulation 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data, art. 22, 2016 O.J. (L 119) 1.

[38] Nat'l Inst. of Standards & Tech., U.S. Dep't of Com., NIST AI 100-1, Artificial Intelligence Risk Management Framework (AI RMF 1.0) (2023).

[39] Int'l Org. for Standardization, ISO/IEC 42001:2023, Information Technology—Artificial Intelligence—Management System (2023).

[40] *See* Armour & Eidenmüller, *supra* note 1.

[41] *See* Bayern, *supra* note 20.