

**AI AND THE BOARD:
PRACTICAL AND LEGAL CONSIDERATIONS FOR AUGMENTING
BOARD DECISION-MAKING WITH ARTIFICIAL INTELLIGENCE
AND ITS IMPLICATIONS ON CORPORATE LAW**

Jordan Cohen

In 1986, political scientist Langdon Winner wrote on the design and nature of technologies.¹ In the controversial thesis of his book *The Whale and the Reactor: A Search for Limits in an Age of High Technology*, Winner clarified the relationship between technology and politics—theorizing that technologies have politics embodying social relations, emerging from and creating social foundations.² Drawing on Lewis Mumford’s studies of the city, architecture, and history of technics, as well as the concerns voiced by other nineteenth-century critics of industrialism, Winner illustrated the political properties of technology as “the instances in which [an] invention, design, or arrangement of a specific technical device or system becomes a way of settling an issue in the affairs of a particular community.” To that end, Winner drew on studies of the arrangements of power and authority within corporations, and the compatibility of sophisticated technological systems with centralized, hierarchical managerial control.³ “The computer romantics,” as Winner explained, “are correct in noting that computerization alters relationships of social power and control . . . [t]hose who stand to benefit most obviously are large transnational business corporations[.]”⁴ A revisit of Winner’s political theory of technology is apropos three decades later in the era of intelligence where theoretical modeling and machine learning is redefining the social fabric and relationship between knowledge and decision-making.

The general framework that many decision theorists have utilized for the analysis of decision-making is the one provided by the Bradford Studies of Strategic Decision Making, a framework that integrates elements of cognitive and political theories into a comprehensive conceptual model.⁵ Under the Bradford framework, a decision-making process is characterized by the complexity of a

¹ See LANGDON WINNER, *THE WHALE AND THE REACTOR: A SEARCH FOR LIMITS IN AN AGE OF HIGH TECHNOLOGY* (1986).

² *Id.*

³ *Id.*

⁴ *Id.* at 107.

⁵ See Thomas Lawrence, *Impacts of Artificial Intelligence on Organizational Decision Making*, 4 *JOURNAL OF BEHAVIORAL DECISION MAKING* 195–214 (1991).

decision topic and the *politicality* of the decision-making process—that is, the interplay of “legitimate, expert, ideological, and political systems of power which constitute the overall pattern of influence.”⁶ To that end, corporate boards that have greater disparities of power between the directors will also have greater levels of politicality. Intrinsicly, the board-centered model of corporate governance implicates structural tensions that stem from the relationship between directors, officers, and shareholders—an interplay of power and authority amongst competing interest groups within the corporation, and society at large.⁷

Corporate governance is premised on the assumption that individuals are rational agents who seek to maximize a firm’s utility and long-term interests. Yet cognitive psychology and behavioral studies indicate that judgment, decision-making, and behavior are often divorced from logical reasoning—subject to pressures toward group conformity, heuristics, and cognitive biases.⁸ As applied to the agents of a corporation, cognitive and behavioral influences have the power to steer board members’ judgment, inferential reasoning, and behavior in a direction that defies logic and expected utility theory—violating “normative assumptions that are central to the economist’s rational model.”⁹

In *The Limits of Organization*, economist and political theorist Kenneth Arrow attributes unnecessary board error to uncertainty by virtue of information overload.¹⁰ As the complexity of an organization increases, so too does the information. Analyzing massive amounts of information becomes a daunting, if not impossible, task; some directors may have access to superior information while others simply rely on whatever information management provides to them. A recent example is Helios and Matheson Analytics board member Carl Schramm’s August 2018 resignation letter accusing management of not giving the board enough time to examine “complex documents, to review significant transactions, or to discuss how the proposed actions fit into the company’s strategic plan.”¹¹ In addition to problems related to

⁶ *Id.* at 199.

⁷ See D.N Ghosh, *Corporate Governance and Boardroom Politics*, 35 ECONOMIC AND POLITICAL WEEKLY 4010–4014 (2000).

⁸ See Oliver Marnet, *Behavior and Rationality in Corporate Governance*, 39 JOURNAL OF ECONOMIC ISSUES 613–632, 613 (2005).

⁹ *Id.*

¹⁰ See KENNETH JOSEPH ARROW, *THE LIMITS OF ORGANIZATION* 74 (2013).

¹¹ See Nathan McAlone, *MOVIEPASS PARENT COMPANY BOARD MEMBER RESIGNS AND ACCUSES MANAGEMENT OF WITHHOLDING MATERIAL INFORMATION FOR MONTHS* BUSINESS INSIDER (2018), <https://www.businessinsider.com/board-member-resigns-moviepass-parent->

accessing and understanding complex information, under the cognitive bias analysis, directors tend to filter information in accordance with their preconceptions—it is easier for someone to understand and accept information that aligns with prior experiences and previous beliefs than it is to overcome cognitive dissonance.¹²

Programmed as a knowledge-intensive problem solver, AI can be engineered as an expert system to overcome and correct the risks inherent in human directors. Directors are presumably selected by virtue of their knowledge and expertise in a corporation’s line of business, but as modern corporate governance practices reveal, an individual director’s contribution often has little to do with expertise. Consider JPMorgan Chase’s 2012 “London Whale” that led to \$6 billion in trading losses. There, the financial institution had no directors with risk expertise on the board’s risk committee. Such failures motivate activists and corporate governance rating agencies to sense that board members often fail to reflect the skill sets companies need. The premise of AI is that, as an expert system, it employs “knowledge specific to a problem domain to provide ‘expert quality’ performance in that application area.”¹³ An important feature of AI is in the heuristic problem-solving techniques that it uses to solve a wide range of problems including:

Interpretation—forming high-level conclusions from collections of raw data.

Prediction—projecting probable consequences of given situations.

Diagnosis—determining the cause of malfunctions in complex situations based on observable symptoms.

Design—finding a configuration of system components that meets performance goals while satisfying a set of design constraints.

Planning—devising a sequence of actions that will achieve a set of goals given certain starting conditions and run-time constraints.

Monitoring—comparing a system’s observed behavior to its expected behavior.

company-helios-and-matheson-analytics-2018-8?r=UK&IR=T (last visited Nov 26, 2018).

¹²ARROW, *supra* note 10, at 75.

¹³ See GEORGE F. LUGER, ARTIFICIAL INTELLIGENCE: STRUCTURES AND STRATEGIES FOR COMPLEX PROBLEM SOLVING 277 (6 ed. 2009).

Instruction—assisting in the education process in technical domains.

Control—governing the behavior of a complex environment.¹⁴

In the context of corporate governance, an AI system would be capable of perceiving the environment and market in which a business operates by processing raw data that includes market information, financial and industrial data, competitor data, news articles, and the like, all within seconds.

AI can be tasked with resolving information gaps that many independent directors suffer from. Specifically, AI can draw sound inferences and useful conclusions from poorly formed and uncertain data. In *Netflix Approach to Governance: Genuine Transparency with the Board*, David F. Larker and Brian Tayan explain that because of an independent director’s limited exposure to a firm’s day-to-day activities—and independence from the business itself—they often have a “less-complete understanding of the company and the market than executives.”¹⁵ Data-heavy and analytically-week, the information provided to independent directors impair the ability to make good decisions.¹⁶ The problem is further exacerbated by boardroom dynamics that impede information flow “particularly in settings where the CEO maintains strict control over the content presented, when presentations are carefully scripted, when follow-up beyond one or two questions is discouraged due to time, and when presentations are made by only a limited number of executives—such as the CEO, CFO, general counsel, and not others.”¹⁷

There too exists the ability of AI systems to be embedded in a corporation’s database—continuously observing all data and information on the company’s internal shared systems, as well as continuously monitoring competitor information, news, and the market at large. Consider Cisco’s Kinetic platform which helps to extract, analyze, and distribute data from disparate devices and applications within an organization. Applying the Kinetic platform to capture the data produced by an organization’s various devices and applications would provide directors streamlined access to high-

¹⁴ *Id.* at 278.

¹⁵ See David F. Larker & Brian Tayan, *Netflix Approach to Governance: Genuine Transparency with the Board*, STANFORD CLOSER LOOK SERIES CORPORATE GOVERNANCE RESEARCH INITIATIVE (2018).

¹⁶ *Id.*

¹⁷ *Id.*

value data.¹⁸ The Netflix approach entails allowing board members to observe monthly and quarterly senior management meetings, and receive board communications that are structured “as approximately 30-page online memos in narrative form that not only include links to supporting analysis but also allow open access to all data and information on the company’s internal shared systems, including the ability to ask clarifying questions of the subject authors.”¹⁹ An AI system like Kinetic that is fully integrated into a corporation’s internal network would provide independent directors with direct exposure to the corporation’s activities, and a deeper understanding of high-level data.

Companies like Salesforce and Deep Knowledge Ventures have recognized that AI—free from management’s influence and internal boardroom dynamics—is capable of providing their boards with the unbiased opinion of an independent director without the knowledge deficits that independent directors suffer from. In March 2017, Salesforce and IBM announced a deal that would integrate the CRM developer’s forecasting and modeling technology, Einstein, with Watson. Created as a question answering computing system, Watson applies “advanced natural language processing, information retrieval, knowledge representation, automated reasoning, and machine learning technologies to the field of open domain question answering.”²⁰ Thus, Watson is capable of analyzing precise natural language questions by identifying sources such as encyclopedias, dictionaries, news articles, literary works, and other databases, finding and generating hypotheses based on the data, merging and ranking such hypotheses, and returning a precise answer.²¹

As applied to Salesforce, the Einstein-Watson joint AI partnership combines customer insights from Salesforce’s Einstein with Watson’s structured and unstructured data—information on weather, financial services, retail markets, and the like—to provide Salesforce’s Chairman and CEO with recommendations.²² By

¹⁸ See Cisco, TRANSFORMING BUSINESS WITH ARTIFICIAL INTELLIGENCE CISCO (2018) <https://www.cisco.com/c/dam/en/us/solutions/collateral/digital-transformation/ai-whitepaper.pdf> (last visited Nov 27, 2018).

¹⁹ Larcker & Tayan, *supra* note 15.

²⁰ See IBM, THE DEEPQA RESEARCH TEAM IBM RESEARCH (2016), https://researcher.watson.ibm.com/researcher/view_group.php?id=2099 (last visited Nov 27, 2018).

²¹ See David Ferrucci et al., *Building Watson: An Overview of the DeepQA Project*, 59 AI MAGAZINE, 2010.

²² See IBM and Salesforce Announce Landmark Global Strategic Partnership, IBM (2017), <https://www-03.ibm.com/press/us/en/pressrelease/51707.wss> (last visited Nov 26, 2018).

integrating Watson APIs (application program interface) into Salesforce, decision-making is accelerated.²³

Deep Knowledge Ventures, a venture capital firm based in Hong Kong, has not only incorporated AI into their corporate governance regime but has given its AI system, Vital, a seat in the boardroom.²⁴ Vital, which the company dubbed as an acronym for Validating Investment Tool for Advancing Life Sciences, helps the biotechnology fund's board of directors make smarter decisions by analyzing large data sets to reveal patterns critical for assessing risk factors.²⁵ Because of the biotech sector's high failure rate, Vital prevented Deep Knowledge Ventures from potentially fatal investments on "overhyped projects."²⁶ Functionally, Vital integrates data from "scientific literature, grants, patent applications, clinical trials and even the biographies of individual team members of companies in which Deep Knowledge Ventures is interested."²⁷ Legally, Vital is treated as a member of Deep Knowledge Ventures' board of directors with observer status; the firm's directors agreed that absent corroboration by Vital, no positive investment decisions would be made.²⁸

Sitting at the apex of a corporation's governing structure, the board of directors operate under state laws which impose on them the duties of obedience, loyalty, and due care—a triad of fiduciary duties to act in good faith, with reasonable care, and in the best interest of the corporation and its shareholders.²⁹ Acting as a safe harbor to these three broad duties of corporate directors is the business judgment rule, a legal doctrine that generally provides directors with broad discretion to make good faith business decisions "absent fraud, gross negligence or other misconduct."³⁰ In this sense, the business judgment rule pervades every aspect of corporate law and sits at the intersection of authority and accountability.³¹ Delaware plays a pervasive role in business judgment rule jurisprudence because Delaware chancellors "sit at

²³ *Id.*

²⁴ See Nicky Burrige, ARTIFICIAL INTELLIGENCE GETS A SEAT IN THE BOARDROOM NIKKEI ASIAN REVIEW (2017), <https://asia.nikkei.com/Business/Companies/Artificial-intelligence-gets-a-seat-in-the-boardroom> (last visited Nov 26, 2018).

²⁵ *Id.*

²⁶ *Id.*

²⁷ *Id.*

²⁸ *Id.*

²⁹ See S. REP. NO. 107-70, at 5 (2002).

³⁰ *Id.*

³¹ See Stephen M. Bainbridge, *The Business Judgment Rule as Abstention Doctrine*, UCLA L. LAW & ECONOMICS RESEARCH PAPER SERIES (2003).

the center of the corporate law universe.”³² For this reason, this paper analyzes the business judgment rule through the lens of Delaware’s general corporation law.

Under Delaware General Corporation Law (DGCL), the business judgment rule exists to “protect and promote the full and free exercise of the managerial power granted to Delaware directors.”³³ Just as the basic corporate law principle of limited liability insulates shareholders from being held personally liable for the debts and torts of a corporation, the business judgment rule exists to encourage director risk-taking and innovation without the fear of being held personally liable for errors of judgment. The rule achieves this by establishing a presumption that the directors of a corporation “acted on an informed basis, in good faith and in the honest belief that the action taken was in the best interests of the company.”³⁴

If a board implements the use of AI to augment their decision-making, directors may effectively lose the protective shield of the business judgment rule in situations where their poor business decision significantly departed from the recommendations of the AI system. Although at present the use of AI in boardrooms is rare, as the practice continues to spread into more and more corporate boardrooms its use may very well set a new standard of director behavior and serve as evidence of a board’s failure to exercise due care. In cases where the business judgment rule is treated as a standard of liability, departures from the recommendations of an AI system that have a negative effect on the corporation will likely constitute a breach of the duty of care.

Alternatively, if a board implements the use of AI to augment their decision-making and the decision is detrimental to the company, directors could possibly rely on the additional safe harbor found in DGCL § 141(e) to protect themselves from liability. Directors often rely on the expert advice of others, and this reliance is encouraged by § 141(e)’s protection over a director’s good faith reliance “upon such information, opinions, reports or statements presented to the corporation[.]”³⁵ In order to invoke the protective shield of § 141(e) the information, opinions, reports or statements must come from the corporations “officers, or employees, or committees of the board of directors, or by any other person as to matters the member reasonably believes are within such other person’s professional or expert competence and who has been

³² See D. Gordon Smith, *Chancellor Allen and the Fundamental Question*, 21 SEATTLE U. L. REV. 577, 578 (1998).

³³ See *Smith v. Van Gorkom*, 488 A.2d 858 (Del. 1985).

³⁴ See *Aronson v. Lewis*, 473 A.2d 805, 812 (Del. 1984).

³⁵ See DEL. CODE ANN. tit. 8, § 141(e) (2016).

selected with reasonable care by or on behalf of the corporation.”³⁶ Here, the question turns on the long continuing debate of legal personhood, and whether a machine or nonhuman entity should be considered a legal person with a set of legal rights and duties.³⁷

Corporations and government entities are the most familiar example of nonhuman legal persons, yet under current law nonhuman autonomous systems are not legal persons.³⁸ Nevertheless, the flexibility that modern business entities like the limited liability company provide demonstrates a means to providing legal personhood to autonomous systems without the need for significant legal reform.³⁹ The extreme flexibility of limited liability company law allows for any arrangement that has an operating agreement to be given legal entity status.⁴⁰ Thus, an operating agreement can adopt “as the acts of a legal entity, the state or actions of arbitrary physical systems.”⁴¹

DGCL § 141(b) allows for only natural persons to act as directors of a corporation, and until Delaware follows the approach taken by Swiss corporate law—allowing legal entities to have an autonomous system take management decisions on behalf of the entity—an AI system like Deep Knowledge Venture’s VITAL (located in Hong Kong) cannot serve as a director. Yet, the modern Model Business Corporation Act (MBCA) provides a solution that comes close to giving AI a voice in the boardroom. Under § 7.32(a) of the MBCA:

An agreement among the shareholders of a corporation that complies with this section is effective among the shareholders and the corporation even though it is inconsistent with one or more other provisions of this Act in that it:

³⁶ *Id.*

³⁷ See Lawrence B. Solum, *Legal Personhood for Artificial Intelligences*, 70 N.C. L. REV. 1231 (1992).

³⁸ See Shawn Bayern, *The Implications of Modern Business–Entity Law for the Regulation of Autonomous Systems*, 7 EUROPEAN JOURNAL OF RISK REGULATION 297–309 (2016).

³⁹ See Shawn Bayern, Thomas Burri, Thomas D. Grant, Daniel M. Häusermann, Florian Möslein, and Richard Williams, *Company Law and Autonomous Systems: A Blueprint for Lawyers, Entrepreneurs, and Regulators*, 9 HASTINGS INT’L & COMP.L. REV. 135, 136 (2017).

⁴⁰ *Id.* at 139.

⁴¹ *Id.* at 136.

(1) eliminates the board of directors or restricts the discretion or powers of the board of directors[.]⁴²

Thus, a corporation may specify in its shareholder agreement legal obligations and conditions of the human directors' decision-making discretion by specifying that the directors shall take any legal actions determined by the AI system.⁴³ Similarly, a shareholder agreement may specify that the directors shall not make decisions that materially depart from the guidance of AI. Consider, then, the following illustration of how such a provision might be drafted in a limited liability company's operating agreement:

Sample Operating Agreement Provision

Consultation with AI Advisory System.

Directors shall consult with and rely upon the Artificial Intelligence Advisory System on all matters connected with the discharge of his or her responsibilities. Notwithstanding anything to the contrary in this Agreement, the Board of Directors shall not take, consent to, authorize, approve, ratify or effect any action on behalf of or with respect to the Company or its business or affairs that materially depart from the recommendations or guidance of the Artificial Intelligence Advisory System.

The sample provision above treats a board's duty to consult with an AI system in the same way that some operating agreements require directors to confer with the advisory board on certain matters.⁴⁴ By treating an AI system as an advisory board (an informal committee of experts— with no fiduciary duties or voting rights—that exists to assist the board of directors in the decision-making process), AI systems can theoretically exist in harmony with American corporate law. As of the date of this paper, no examples of such a provision could be found and the above sample provision drafted by the author is highly conceptual. Although the analysis above demonstrates how existing law can enable autonomous systems without the need for radical legal reform, still, under American corporate law, an AI system cannot serve as a director. Nevertheless, as this area of technology and corporate governance continues to develop, it will

⁴² See MODEL BUS. CORP. ACT § 8.03(a) (2002).

⁴³ Bayern, *supra* note 38, at 303.

⁴⁴ See LIMITED LIABILITY COMPANY OPERATING AGREEMENT OF B.N ENTERTAINMENT US LLC (2015), <https://www.nybusinessdivorce.com/wp-content/uploads/sites/94/2018/04/Gildea-OA.pdf> (last visited Nov 26, 2018).

be interesting to observe how drafters of operating agreements will approach the issue.

Ultimately, under this theoretical model of corporate governance, shareholders will need to balance the benefits of employing AI in the corporation's decision-making process with the drawbacks such a system would create. To wit, as an AI system's authority increases in the boardroom, directors will likely invoke the system's guidance as justification for corporate disasters. Courts will likely find that those directors who relied on the guidance of AI—as required by the director's legal obligations and conditions in the shareholder agreement—should be indemnified from any liability. In such a scenario, any derivative suit brought against directors whose decision-making was governed by an AI system would likely be dismissed for failure to state a claim.

In the context of corporate governance, AI-augmented decision-making has clearly defined advantages. Humans, by their very nature, are limited in their heuristic problem-solving techniques because of information asymmetries, competing interests, and cognitive biases. These limitations, combined with the agency costs associated with the separation of ownership and control, have accounted for corporate disasters that have attracted an extraordinary degree of public scrutiny and proposals from institutional investors, exchanges, and government regulators alike. While following the Salesforce or Deep Knowledge Ventures approach to AI-augmented decision-making is not the be-all and end-all to corporate governance, it is a step forward in resolving some of the most pressing challenges facing corporate law today. Corporations will need to balance the benefits of rational, data-driven decisions with questions of accountability, liability, and fiduciary rights of directors that have relied on an AI system's guidance. These questions will likely be answered as Congress and individual states develop a corporate law framework that accounts for the rapid development of AI and its implications on corporate governance.