Ownership and Trust A corporate law framework for board decision-making in the age of AI

Law Working Paper N° 758/2024
March 2024

Katja Langenbucher
Goethe University Frankfurt, SciencesPo, Fordham Law School and ECGI

ECGI Working Paper Series in Law

Ownership and Trust A corporate law framework for board decision-making in the age of AI

Working Paper N° 758/2024
March 2024

Katja Langenbucher

Abstract

The paper proposes a framework for judicial review of board decisions that have been augmented by an AI. It starts from the assumption that the law treats decision-making by board members differently than decision-making by officers and employees. Against this background, the paper brings out two core characteristics. Corporate law expects board members, but not directors and employees, to fully own their decision. As a flipside of ownership, corporate law places trust in board members to form business judgments, immune against judicial second-guessing. The paper moves on to investigate how these principles play out when an AI augments board‘s decision-making. The paper makes two contributions to the debate. First, it rejects the notion that black-box AI may not be used for board decision-making. Second, it proposes a graphic control matrix to identify low, medium, and enhanced judicial scrutiny when boards use an AI to inform their decisions.

Keywords: Corporate boards, AI, business judgment rule, compliance, algorithmic discrimination

Katja Langenbucher
Professor
Goethe University Frankfurt
Theodor-W.-Adorno-Platz 3
60323 Frankfurt am Main, Germany
e-mail: langenbucher@jur.uni-frankfurt.de

Ownership and Trust

A corporate law framework for board decision-making in the age of AI

*Katja Langenbucher**

Explaining human intelligence is an intriguing topic.[1] For some, it manifests human singularity. Others emphasize the dependence of human intelligence on mechanistic operations.[2] Whether this implies a kinship between these two forms of information processing or, conversely, whether there are fundamental differences has been discussed for hundreds of years.[3] Arguably, an uncontested point of departure is that machines can sometimes surpass human performance as to speed and precision. From there, a pressing question follows for corporate decision-making. If it is advisable for doctors, lawyers, and stock exchange traders to have certain decisions augmented by machines, does this also apply to management decisions of company directors? If so, who bears the cost if things go wrong?

Part I of this paper provides a brief overview on the use of artificial intelligence (AI) as a "prediction machine"[4] for board decisions. It reminds the reader that statistics has traditionally filled this role and hints at similarities and differences when using machine learning (ML).

Part II zooms in on how corporate law frames decision-making. It starts from the assumption that the law treats decision-making by board members differently than decision-making by officers and employees. Against this background, the paper brings out two core characteristics. Corporate law expects board members, but not directors and employees, to fully own their decision. As a flipside of ownership, corporate law places trust in board members to form business judgments, immune against judicial second-guessing.

* Katja is a law professor at Goethe University, Frankfurt, Germany, She is affiliated at SciencesPo, Paris, France, and regular visiting faculty at Fordham Law School, NYC, US.

[1] Some of what follows integrates Langenbucher (2023b).

[2] Glimcher (2004); Rolffs (2023); Stiehl/Marciniak-Czochra (2021), p. 279.

[3] Overview in Hawkins (2021); Larson (2021); Nath (2009).

[4] For this term see the title of: Agrawal et al. (2018); further : Agrawal et al. (2022); Ertel (2021), pp. 201 et seq.; Russel/Norvig (2021), pp. 19-22.

The expectation that boards own their decisions implies that they must not abdicate their authority. The paper explores how this principle plays out when boards enhance their decision-making with an AI. It moves on to examine how corporate law has framed ownership of a board decision when technical support tools or human experts inform board members. Rather than analogizing AI to one of these helpers, the paper brings in the second dimension: trust. It claims that the standard of judicial review has moved along these two dimensions, ownership, and trust. The same logic, the paper suggests, applies to board decisions that integrate AI. The paper concludes with a graphic visualization of these dimensions.

The paper makes two contributions to the debate. First, it rejects the notion that black-box AI may not be used for board decision-making. Second, it proposes a graphic control matrix to identify low, medium, and enhanced judicial scrutiny when boards use an AI to inform their decisions. A detailed exploration of specific duties of care is beyond the scope of this paper.

I. Prediction machines

The term "artificial intelligence" includes various implementations, ranging from logic, ML, and neural networks to large language models (LLM) and robotics.[5] These correspond with a diverse set of potential use cases in corporate life. Logic plays a role in planning and automating processes,[6] robotics[7] can be helpful in manufacturing, and LLMs help with a vast array of knowledge work, ranging from information retrieval over summarizing studies to reviewing contracts.[8]

The scenario this paper explores implies the use of AI as a “prediction machine”. In that capacity, a board uses AI to enhance its understanding of which future events are likely to happen. Most management decisions imply predictions of that type. Employing statistics to that end is

[5] Ertel (2021), pp. 3 et seq.; Russell/Norvig (2021), chapter 1.1.
[6] Ertel (2021), pp. 29 et seq.; Russell/Norvig (2021), chapter 7.1, 11.
[7] Russell/Norvig (2021), chapter 26.1.
[8] Russell/Norvig (2021), chapter 24.

a standard tool. Statisticians work on inferences about the relationship between different variables, based on a hypothesis.[9] Consider the management board of a bank that decides on a reduction of the number of brick-and-mortar branches to move towards online banking. Studies on customer preferences, possibly also their age, occupation, or place of residence, together with mobile network coverage, and the number of bank branches can inform management. An initial hypothesis might be: The age of a customer is a core factor driving a preference for brick-and-mortar branches.

Complementing or replacing the statistician, imagine using an AI. To train it, data on customer reactions to branch closures carried out in the past is useful. The AI furnishes patterns, such as groups of bank customers with similar preferences and reactions (clustering).[10] Its predictions about the willingness of bank customers to switch to online banking could mirror the statisticians'. Additionally, it might bring out unanticipated correlations. Both allow the board to react, for instance via targeting its marketing to specific groups.

1. The machine learns

One of the intriguing features of AI is its potential to learn. Instead of being provided with an input-output pair that is specified *ex ante*, the AI is left to stroll through data, as it were. Its performance gets better after it has made observations and adjusts its reactions.[11] Self-driving cars provide a much-discussed example.[12] It would be costly (and probably impossible) to record all situations a self-driving car might encounter and program appropriate input-output pairs. Instead, it is much more efficient to program the AI to learn how human drivers have reacted in relevant situations.[13] The AI will recognize patterns in the data, for instance: cars drive slower in heavy rain. The AI will then adjust the programming of the self-driving car to match this pattern. Depending on the AI's loss function (such as: reaching the destination

[9] Exploratory data analysis precedes making inferences and producing testable hypotheses. It does not include formal statistical modeling and inference. Instead, it helps to see patterns in the data, catch mistakes, and generate potential hypotheses.
[10] On clustering see Russel/Norvig (2021), p. 671, see below I.2.
[11] Russel/Norvig (2021), p. 669.
[12] See Agrawal et al. (2018), pp. 88 et seq.
[13] Russel/Norvig (2021), p. 670.

quickly or pleasant driving experience), it will continuously adjust the programming of the self-driving car.

There are three basic forms of machine learning.[14] In supervised learning, the AI is programmed to map input to output.[15] Input might be an image and output the classification as a wolf. The database that trains the AI contains labeled examples. The label tells the AI which function to find (hence the term *supervised* learning).[16] Supervised learning requires large data sets that have been processed and appropriately labeled. Using these, the AI learns to make predictions for new data.

Some situations require a more exploratory approach. The goal might be to analyze unlabeled data with a clear goal in mind. Alternatively, it might not even be clear which questions are relevant, for example when dealing with a large, unstructured data set.[17] Unsupervised learning responds to these exploratory needs. It makes the AI independently find structures and patterns.[18] The programmer does not specify the way in which the AI performs the identification task, nor does he specify a goal or label the data.[19] This distinguishes the technique from supervised learning, where the AI has a previously known objective (wolf/no wolf). With unsupervised learning, the AI shows the user a way of sorting disordered data.[20] This approach requires very large data sets and computers with enormous computing power. Its use for daily management will for most corporations mean buying the AI from a provider.

---

[14] Russel/Norvig (2021), p. 671.
[15] Russel/Norvig (2021), p. 827.
[16] Ertel (2021), p. 351.
[17] Russel/Norvig (2021), p. 827.
[18] Ertel (2021), p. 260.
[19] Depending on the use case, it can be useful to label a subset of the data set (semi-supervised learning). Consider a radiologist searching ultrasound images for anomalies. An efficient way might be to have trained radiologists label some of the data. This speeds up the search process and, at the same time, does not preclude the AI from finding other patterns from which a treating physician may be able to draw further conclusions.
[20] Like for supervised learning, image recognition is a potential use case: "when shown millions of images (...) a computer vision system can identify a large cluster of similar images which an English speaker would call 'cats'": Russel/Norvig (2021), p. 671.

Learning by reinforcement occupies a space between supervised and unsupervised learning.[21] The AI works without pre-labeled training data and is programmed to perform certain sequences, such as a board game[22] or a robotics task.[23] It receives positive or negative human feedback after completing its task. Each following round, the AI adapts its strategy to receive positive feedback.[24]

2. Induction engines[25]

To a statistician, it comes as no surprise that good data is a core ingredient of a forceful prediction. Selection biases, omitted variable biases, or the non-observance of confounding variables can be just as damaging as mathematical errors in a model. With AI, many of these issues arise in similar ways. Depending on which training data the AI receives, how that data is structured or labeled, the AI will learn to map, recognize patterns, and build models to assess future situations.[26] Predictions based on a carefully curated,[27] possibly even synthetic,[28] dataset differ significantly from the prognosis an AI makes by accessing the entire internet. If biases or past discrimination are baked into the data, the AI will suggest treating new cases in line with seasoned values. The same goes for the selection of data for the AI to learn.[29] Consider the example of the bank executive deciding on branch closures. If the AI is trained on a small data set, compiled by one bank, sampling customer reactions in one geographical area, the AI will develop a model that provides an excellent representation of this one data set, but won't necessarily generalize. The risk of error increases and the quality of the prediction decreases.

This is not to say that more data is necessarily the better solution. Take open access to the internet as an illustration. It allows for particularly precise predictions about human preferences, detecting unanticipated patterns and clusters. At the same time, much of the data is

---

[21] Russel/Norvig (2021), pp. 840 et seq.
[22] Russel/Norvig (2021), pp. 671, 840.
[23] Ertel (2021), p. 35.
[24] Russel/Norvig (2021), p. 840.
[25] For this term see: Larson (2021), pp. 115 et seq.
[26] Russel/Norvig (2021), p. 669.
[27] On curating data: Data Governance Working Group of the Global Partnership of AI (2020), pp. 19 et seq.
[28] Jordan et al. (2022).
[29] Russel/Norvig (2021), p. 19, p. 669.

noise that risks producing skewed results.[30] To bolster management decisions, a synthetic, curated, or at least "cleaned" data set might be more useful.

Lastly, it is helpful to keep in mind that AIs are "induction engines".[31] Their probabilistic estimations rely on correlations that they infer from existing data. A change in circumstances, unusual, or rare situations, technical innovations or novel human preferences arrive at an AI with a time delay.[32]

II. Decision-making and corporate law: Ownership and trust

Decision-making is one of the areas where AI has been shown to augment human capabilities. There are preformatted and rule-bound situations that provide especially fitting use cases for AI. We might be looking at robots for production, a chatbot used on a customer hotline, or automated lending decisions. Along similar lines, the AI might take over parts of rule-based decision-making. Consider a chatbot forwarding unfamiliar questions or an out-of-the ordinary credit application that human employees review further. Board decisions, by contrast, are rarely an exclusively rule-based endeavor.[33] They entail discretion, intuition and “gut”, a process of weighing and balancing different considerations, and of making value judgments. Employing AI as a prediction machine allows to build scenarios, assess their probability of materializing, and use this as a background when making an informed decision.

Corporate law adapts rights and duties to the different types of decision-makers. It treats board members differently than decision-makers at officer and employee level. Firstly, the law expects board members to fully own their decisions. By way of illustration, see section § 76 German Corporate Law Code (*Aktiengesetz*) stipulating that board members are accountable for managing the company. While they may delegate tasks, board members must not abdicate the authority the law vests in them. Similarly, Delaware General Corporate Law § 141(a) provides that a Delaware corporation is managed by or under the direction of the board of directors. In discharging their duties, they owe fiduciary duties of loyalty and care.

---

[30] Illustratively, Bender et al. (2021), p. 610.
[31] Larson (2021), p. 127.
[32] Marcus (2018), p. 9.
[33] On rule following March (1994), pp. 57 et seq.

Secondly, and as a flipside of ownership, the board allows for trust in board members. As long as they act loyally and carefully, the business judgment rule provides a generous liability regime. While board members must critically review material information, they are not required to work through any and all available information. As a second best, the law accepts what is "simply bad judgment" by board members,[34] rather than encouraging judicial second-guessing.

III. The board's role in structuring decision-making by officers and employees

Some board resolutions are purely organizational in nature. They allow for and structure decision-making by officers and employees of the corporation. Oversight duties remain with the board. Arguably, bedrock principles of corporate law are well suited to cope with these board decisions. A board must assess the value proposition of integrating AI. Gains in speed and accuracy must be balanced against the availability of an AI which is fit for the intended purpose. Relevant data and options to train personnel must be evaluated, error costs if things go wrong must be assessed.

Decisions of that type are not the focus of this paper that deals with an AI enhancing board decision-making. Still, three remarks are in order to hint at relevant duties of care. The availability of an AI model that is fit for the intended purpose is an obvious first consideration. Some departments, such as trading, compliance, or risk management might be especially prone to using AI in the form of ML. Marketing and customer services might profit enormously from LLMs. In other cases, integrating AI might require a rewiring of the entire workflow.[35] Balancing the potential gain against the probable costs is a business judgment for which the law grants boards considerable discretion. This includes the suitability of the selected product, extends to its ongoing control, and follow-up product monitoring. In most cases, corporations will purchase the AI from a third party. Selecting an appropriate provider and making sure the offer can be tuned to data that is relevant for the corporation is relevant for the board's choice of an AI. Over time, standard practices will develop, shaping the business judgment on why to

[34] Joy v. North 692 F.2d 880, 885 (2d Cir.), cert. denied, 460 U.S. 1051 (1983); Bainbridge (2020), p. 125.
[35] Agrawal et al. (2022), pp. 85 et seq.

choose one AI over another. The EU AI Act encourages certifications and provides guidelines. In what it terms "high-risk applications", it includes mandatory requirements that will shape board choices for an appropriate AI.[36]

If the choice of the AI model is the first step, the availability and relevance of data comes next. Business judgments concern questions such as: Does the corporation have proprietary data or can it obtain third party data at reasonable cost? What type of data is needed (for instance: open source, curated, synthetic, labeled), how high is the probability of flawed data, and how high are estimated costs when proceeding with it? Will the AI be helpful as a "cognitive fix" for standard flaws of human decision-making? How high is the risk of the AI learning from biased decisions?[37]

Additionally, the intended "workplace" for the AI might require specific features.[38] Employees who cooperate with an AI often need special skills.[39] This involves basic training, as required for every new machine or technology, to be able to correctly classify its mode of operation and risks. The risk of known shortcomings of AI, for example problems with the coding of known knowledge[40] or abductive conclusions,[41] must be balanced against an increase in efficiency.

IV. The board structuring its own decision-making

Using AI as a prediction machine when preparing a board decision is different from the board adopting AI as a tool to enhance the corporation's workflow. Rather than programming (partly or fully) automated decisions, the board integrates the AI in its own deliberation. It hopes to enhance the quality of its decision-making by gaining a good understanding of, for instance, how markets, customer preferences, capital allocation, or investor appetite will evolve.

[36] Annex III spells out the high-risk AI systems referred to in Art. 6(2) AI Act.
[37] Langenbucher (2023a).
[38] Hacker (2020), p. 2145.
[39] For „automation bias" see Art. 14(4)(b) AI Act.
[40] Marcus (2018), pp. 11 et seq.
[41] Larson (2021), pp. 99, 162 et seq.

Board resolutions of this type operate under the corporate law principles mentioned above.[42] On the one hand, the business judgment rule manifests the law trusting board members with decision-making and keeping judicial second-guessing to a minimum. On the other hand, the law expects the board to own its decisions, ruling out an abdication of authority or an overreliance on experts. The requirement to own a decision leaves no room for the board to have an AI decide in its place. At the same time, the law has nothing against the board asking for support in its decision-making.[43] An emerging discussion has revolved around how to draw the line between an AI merely supporting and entirely taking over decision-making. I argue below that (where we stand today) it is unlikely to see a board so comprehensively integrate an AI in its decisions that we would be looking at an abdication of board authority. Instead, I suggest that fresh efforts must go into understanding what corporate law expects from board members who rely on support to augment their decision-making. I use Delaware and German corporate law to illustrate legal rules for human experts who assist the board and suggest adapting these to the challenges brought about by integrating AI into board decision-making.

1. Abdicating authority: Does the AI take over?

Traditionally, abdication has been understood as trading away the board's discretion.[44] Against that background, so-called "black-box" AI has troubled some scholars.[45] They view integrating a black-box AI as an abdication of authority to an "AI-oracle", as it were. The problem with AI as a tool augmenting decision-making, they claim, is especially prominent if its predictions and recommendations cannot be explained. This view is rejected here as focusing overly on one element of a decision, losing sight of the broader board judgment.[46]

a) Two straightforward scenarios

[42]See above II; Langenbucher (2023c), pp. 728 et seq.
[43] Fleischer (2023) § 76 AktG margin no. 20, 74.
[44] See for Germany: Telle (2023), pp. 32 et seq.; Uwer (2019), p. 154, who points to corporate law principles on outsourcing; see further: Möslein (2018a), pp. 649, 658 et seq.; on the U.S.: Bruner (2021-23), p. 9; Petrin (2019), pp. 965, 1014 et seq.
[45] Dubovitskaya/Buchholz (2023), p. 63.
[46] Langenbucher (2023c) pp. 725 et seq.

Many scenarios are straightforward. It does not hurt to prepare a board decision by googling relevant facts. Google is a familiar AI tool to support decision-making without taking agency away from a board member. Arguably, the same rule applies if, instead of googling, a board member asks a LLM to answer well-defined questions. One *caveat* is in order: ChatGPT or Bard are not yet as familiar a tool as google. Board members should therefore have a basic understanding of what an AI can (and cannot) deliver. However, requirements of that type have to do with the board's duties of care. They do not implicate an abdication of authority.

Another clear case is the (more theoretical) scenario of a board that formally or effectively commits to follow an AI's recommendation. Arguably, the law will not treat this situation any different than a board that trades away its authority to a human.[47] The relevant issue at stake is the same to the extent that the board does not have discretion to decide as it seems fit. Under this angle, it does not matter whether the AI is explainable or not.

b) The hard case scenario

The hard cases are situated between these two scenarios. With AI developing into a standard tool, board judgments will look and feel differently than today. AI outperforms humans in many tasks and continues to evolve, taking over ever more areas. A clear distinction between the AI preparing the decision and the board making the decision will often look artificial.[48] The more closely a decision follows the AI's recommendation, the more the board's role might seem reduced to implementing what the AI has proposed.[49] Arguably, building a basic understanding of technology and trying to grasp the inherent logic of algorithms provides some relief.[50] Still, few board members will become experts in AI technology.

Additionally, it doesn't help that humans are known to be subject to a wide variety of decision-making anomalies when it comes to assessing statistical probabilities,[51] a core element of AI.

[47] Fleischer (2023) § 76 AktG marginal no. 78; Möslein (2018), p. 209
[48] See Langenbucher (2024); however, advocating for a distinction along those lines: Fleischer (2023) § 76 AktG marginal no. 74, 77; Grigoleit (2020) § 76 AktG marginal no. 87; Möslein (2018b) p. 208; Noack (2019) p. 119.
[49] On the "foundational nature" of predictions for decisions: Agrawal et al. (2022), p. 24.
[50] Fleischer (2023) § 76 AktG marginal no. 78; Möslein (2018b), p. 208.
[51] Kahneman (2012), p. 411, with references to de-biasing in the corporate context on p. 417; Kay/King (2020), pp. 53 et seq., 154 et seq.; Tversky/Kahneman (1983), p. 299; Tversky/Kahneman (1974), p. 1124; historical overview at Kozyreva et al. (2019), pp. 343 et seq.

In the same way as the AI's "workplace" on any of the corporation's hierarchical levels must be carefully structured,[52] the board's own "workplace" in cooperation with an AI needs structure. Human cognition follows different patterns than an AI. [53] This entails thinking about the appropriate cognitive cooperation with the AI. Sometimes, the AI can be very helpful if it acts as a "cognitive fix" for human behavioral anomalies.[54] However, the more behavioral anomalies have been baked into the data the AI was trained on, the more these are amplified at scale, rather than reduced. Additionally, scholars have highlighted human preferences for social interaction instead of receiving algorithmic advice.[55] If offered the choice, humans seem to go for a discursive back and forth, rather than receiving a blunt prognosis without the option to engage in arguments and counter-arguments.[56] Especially if the stakes are high, humans tend to demand "slow and effortful consideration of evidence", even if empirical evidence does not necessarily show that this strategy leads to better decision-making.[57]

Against this background, the tough question corporate law must answer is what it expects as a minimum from board members in terms of owning decisions that rest on predictions by an AI. Arguably, the prohibition to abdicate board authority is too coarse a tool to provide a meaningful answer. While few board members have a precise understanding of how every-day AI such as a google search engine or ChatGPT produces its results, the same is true for a pocket calculator or a GPS. The reason why we do not understand these tools as abdicating board authority to a machine is that they contribute but one element to a decision that the board fully owns. It follows from there that the relevant question is not *if* but *how* board members integrate AI in their overall judgment. Short of a situation where the board commits to following the AI's recommendation "no matter what", most cases are not about abdicating authority. Instead, they have to do with delegating (increasingly large) parts of the decision-making process.

c) Known unknowns – part I

[52] See above III.
[53] On the question on which decision theory AI is based, see Burton et al. (2020), p. 226.
[54] Burton et al. (2020), p. 226.
[55] Burton et al. (2020), p. 224 for social structures that encourage this ("existing research points out that organizational and social structures favor the expert intuiter over a cold algorithmic decision maker").
[56] Miller (2023) suggests the adoption of such a "back and forth" approach to AI advice, on the calibration of trust see already below IV.2.
[57] Burton et al. (2020), p. 224.

Most board decisions rest on a large variety of assumptions and predictions. Many of these are known unknowns: How will the market react to the bank closing brick-and-mortar branches? Will the self-driving car produce terrible accidents? Which percentage of my debtors will perform on their loan? When will customer preferences for my product change? How will a geopolitical crisis affect my firm? In these scenarios, the board owning its decision translates as: Understanding the risk of working with a known unknown, evaluating it, and forming an informed and reasonable judgment. The prediction that an AI makes, explainable or black-box, can be just that: a known unknown.

Take the board of a pharmaceutical company deciding on an investment in further research as an illustration. Assume the board members have been discussing the viability of protein folding structures suggested by the AI Alphafold.[58] The board won't be overly interested in an explanation of how the AI went about detecting these structures. From the board's perspective, it might not even make a difference whether an AI or a human researcher did the work. Critical to the board's decision are the reliability and testability of the results, combined with the need and costs for double-checking. If explainability is at all relevant for this part of the decision, it still does not necessarily follow that the board must not work with a black-box AI. Instead, the board must balance the potential profit to be made if the AI got it right against the costs if it did not. It is a known unknown, just like many business judgment rule scenarios.[59] The law expects the board to own its decision. At the same time, it trusts the board to handle a known unknown situation. In no part of this decision is the board abdicating its authority to an AI such as Alphafold.

2. Informing board decision-making: When to trust an AI

So far, we have seen that the law allows boards to delegate individual parts of a decision-making process. This includes a decision in the face of known unknowns. With these, the law trusts board members to come to a reasoned business judgment. Nonetheless, boards do not get a *carte blanche*. Generally, a board must evaluate and double-check information it receives.

[58] https://alphafold.com.
[59] Langenbucher (2023c), p. 727.

On closer inspection, the law distinguishes among decisions (business judgments and others) and among support tools (technical help, humans integrated in the corporation, outside experts).

a) Business judgments

Board members' duties of care vary depending on the decision at hand. For business judgments, the law largely trusts the board, lowering its standard of judicial review. As to doctrinal detail, jurisdictions follow different approaches. Under Delaware law, it is for the plaintiff to prove that the board did not collect appropriate information before making a business judgment. By contrast, under German law various test prongs apply before admitting the board to the safe harbor of the business judgment rule, § 93 para. 1 s. 2 Aktiengesetz. One of these test prongs is the obligation to collect appropriate information with reasonable care. When establishing that this was done, board members bear the burden of proof. Courts have been strict and, at times, required the board to collect any and all available information. Many scholars disagree.[60] Similarly to Delaware courts, they stress that rational shareholders put trust in boards and accept that boards under time pressure gather material information only.

b) Known unknowns – part II

Outside of business judgments, courts apply an enhanced scrutiny standard. Compliance and risk management are paradigm examples. Courts assess the board's decision-making process, including the information the board collected and evaluated.[61] Sometimes, this can restrict the use of AI, especially of the black-box variant.

Consider a board that wishes to cut down on costs. It is impressed by an AI that performs better at predicting credit default risk of borrowers or suitability of potential new hires. It decides to restructure its human resources or its credit underwriting department. Many elements of this plan qualify as a business judgment – the need to cut down costs, the choice

[60] Koch (2023), § 93 AktG marginal no. 42; Fleischer (2023), § 93 AktG marginal no. 94; Spindler (2023), § 93 AktG marginal no. 57.

[61] Delaware law adds major decisions such as change of control transactions, the sale of the company, or the implementation of defenses in a takeover situation. Under German (and European) law, these are business judgments but require shareholder consent.

between different AI models, the decision to remodel the entire department or start with small steps. However, some elements of the board's decision do not qualify as a business judgment with its ensuing broad discretion. The lively debate on algorithmic discrimination provides ample examples for such elements:[62] The decision to restructure human resources must not lead to hiring decisions that systematically discriminate between applicants. Automating credit underwriting must not allow for discriminatory lending practices. Assume, as an integral part of restructuring credit underwriting, the board installs a black-box AI to help with assessing credit default risk. Anti-discrimination laws such as the US Equal Credit Opportunities Act or the EU Consumer Credit Protection Directive prohibit a denial of credit based on protected characteristics. Assume further that the AI collects publicly available data on retail consumers, develops personalized credit default risk assessments, recommends underwriting decisions, or even extends an automated contractual offer. To respect anti-discrimination law, the AI is programmed to disregard all protected attributes. However, given the big data it draws on, the AI is still likely to use proxy variables. Proxy variables stand in for protected characteristics. First names may double as gender or ethnicity, social media friends can be a proxy for age, and activities on a Saturday a proxy for religious faith. The use of proxy variables (first name) can lead to a disparate outcome between minority and majority groups (women and men), even if no protected characteristic (gender) was used. Neither the board nor the corporation's credit officers or even data scientists and coders of the AI will necessarily be able to identify the variables that the black-box AI used.

Is this, like the Alphafold scenario explored above,[63] a question of known unknowns? Can the board reason as follows: (i) We understand that the board must not allow credit underwriting decisions to vary along a protected variable. (ii) Our AI is programmed to disregard protected variables when making its prediction. (iii) We understand that this AI might use proxy variables, but (iv) the extent to which it does is a known unknown. (v) The law trusts boards to integrate known unknowns in its decision if the board evaluates the ensuing risk. (vi) As long as we assess the profit to be made with the credit-underwriting AI and balance it against the risk of potential litigation, we are fine to use the black-box AI.

[62] On what follows: Langenbucher (2023a) with further sources.
[63] See above IV.1.c).

Assuming an affirmative duty for the board to obey the law,[64] the decision whether we face a known-unknown-scenario depends on an interpretation of the anti-discrimination rule. Courts might decide that compliance with that rule requires nothing more but the installation of an input restriction for protected characteristics.[65] Following this interpretation, the black-box AI could be used, as long as the input restriction was in place. Courts that prefer a tougher reading of the anti-discrimination rule might introduce further restrictions on permitted data[66] or prohibit the use of black-box AI altogether. What distinguishes this scenario from the Alphafold example is the degree of trust accorded to the board. The decision to restructure credit underwriting as such is up to the board. However, the decision to install a black-box AI to hand out loan contracts is not entirely discretionary. As far as protected groups are concerned, the law requires some degree of scrutiny as to the known unknown element. This stands in contrast with the Alphafold scenario. The board was able to treat Alphafold and its findings on protein folding structures as a known unknown, qualifying as a classic business judgment.

c) Technical support tools, inside and outside experts

When informing the board, technical support tools, ranging from a pocket calculator to high-powered computers, have been a standard feature. There are no rules stipulating distinct duties for boards that employ a machine to assist decision-making. This is different for humans who support board decision-making, especially if the human help is not an employee of the corporation.

For illustration, I once again use German and Delaware law. Both jurisdictions expect some level of engagement from a board that has humans inform its decision-making. DGCL § 141(e) distinguishes between "information, opinions, reports, or statements presented by any of the corporation's officers, employees, or committees" and input "by any other person". A board may draw on sources from inside the company as long as this is done in good faith. For outside experts, the rule adds extra test prongs. The input must stem from "any other person as to matters the member reasonably believes are within such other person's professional or expert

[64] Proponents of (some version of) the efficient breach theory do not share this assumption, for a foundation see Posner (2009), p. 1349; for an overview see Bigoni et al. (2014).
[65] See the upcoming German law on credit scoring German Federal Government (2024).
[66] See Art. 18(3) Consumer Credit Directive (EU) 2023/2225 on data gathered from social media.

competence". Additionally, such person must have "been selected with reasonable care by or on behalf of the corporation". Hence, for outside experts the court will explore two issues: the reasonable belief that the expert is competent to deliver the relevant input, and the reasonableness of the director's selection of the expert. For both issues, the standard of review is strict, and the business judgment rule is not available.[67]

German law has no explicit statutory rule on integrating external input into board decision-making. However, when evaluating information by an outside expert, a landmark German precedent has added an extra test prong to general duty of care obligations.[68] The case had an executive board decide on a capital increase. One of the supervisory board members had suggested a specific strategy that the courts later declared illegal. This supervisory board member was the partner of a law firm, mandated to work on structuring the capital increase. The members of the executive board claimed they had in good faith relied on the law firm partner on their supervisory board, along with the work done by his firm. The German court did not accept this defense. It stressed that the board's duty of care included figuring out the legal situation. The risk of misunderstanding the law, so the court held, was to be borne by the board, even if they were not legal experts. The court highlighted an obligation for individual board members to make sure the expert opinion was "plausible". Board members were to double-check if what the expert had proposed was in line with their own market knowledge, experience, and, possibly, intuition.

3. Visualizing ownership and control

As an illustration of how courts review board decision-making, I have visualized decisions in a four-square control matrix. The y-axis represents the level of allowance for board discretion according to the decision's subject matter (trust). Boards enjoy broad discretion for those elements of a resolution that qualify as a business judgment. Little discretion is accorded to parts of a decision that have to do with compliance, risk management, and similar, non-business-

---

[67] Brehm v. Eisner, 746 A.2d 244, 262 (Del. 2000); https://www.professorbainbridge.com/professorbainbridgecom/2023/09/does-the-bDelawareiness-judgment-rule-apply-to-a-boards-choice-of-expert.html.

[68] Bundesgerichtshof 20.September 2011, II ZR 234/09, "Ision"; Langenbucher (2023b), pp. 14 et seq.; see Smith v. Van Gorkom 488 A.2d 858, 875 (Del. 1985) for a Delaware decision denying admittance to the safe harbor provided by § 141(e).

judgment issues. The x-axis looks at the intensity of information support (ownership). Boards have been free to use technical support tools, ranging from pocket calculators to high-powered computer networks. Human helpers have attracted more scrutiny. This is true for input by officers, committees, or employees of the corporation. Even more scrutiny concerns outside experts.

Following this representation, four squares emerge. In the upper right-hand corner, we find the first square. It symbolizes business judgments that score high on discretion and do not rely on human support. They face the lowest degree of judicial review.

The second square is situated in the upper left-hand corner. It symbolizes decisions that still score high on discretion but have drawn on considerable help, including from outside the corporation. For those, the standard of review is higher than for the first square, given the dominant role of support tools.

A third square is situated in the lower right-hand corner. It symbolizes decisions that are about non-business judgment issues but do not rely much on human support. Its standard of review resembles the one just described. It is higher than for the first square, given its low score on trust.

The fourth square is located in the lower left-hand corner. It shows decisions that were reached with much outside help, hence, score low on ownership. Additionally, these decisions score low on discretion, because they include few or no business judgment elements. This square symbolizes the highest intensity of judicial review.

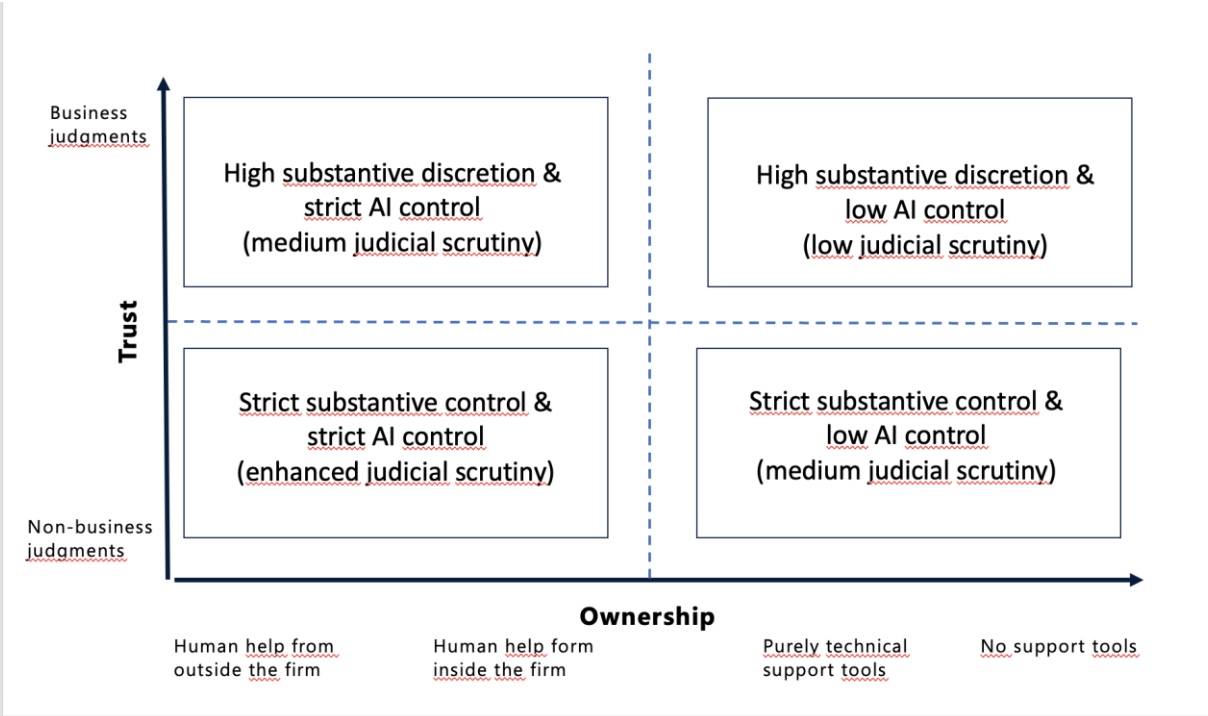

The graphical representation is helpful given that board decisions rarely fall into one neat category. The above example on restructuring credit underwriting showed how board decisions combine different elements.[69] Some of these are about developing and deciding on a novel business strategy, involving market knowledge, experience, intuition, and gut. All these are characteristics of a low-judicial-scrutiny decision. However, other parts of the decision might depend on the professional evaluation of a particular market niche or of a new product that only outside experts can deliver. Legal issues might be decisive for the success of the new strategy because a new product requires regulatory approval. These legal issues could be small and resolvable in-house or complex, calling for outside counsel. Visualizing the matrix and „moving" the decision, as it were, allows to understand the degree of judicial review that a comprehensive board resolution, with its various sub-parts, will attract.

The legal logic underlying the matrix reflects the tension between boards owning their decisions and the law trusting boards without holding them accountable for "simply bad judgment". As explained above,[70] the law expects that board members are accountable and will own their decisions. It follows from there, that the law does not allow the board to abdicate

[69] See above IV.2.b).
[70] See above II.

its authority and hide behind an alternative decision-maker, as it were. It does not matter whether an alternative decision-maker might be more capable than the board: it is not the one the shareholders voted for. Along similar lines, the board may not delegate core parts of its decision-making to non-board members. The more a board decision looks like nodding to what someone else has proposed, the less it conforms with the law's expectation of the board owning its decision.

At the same time, a board cannot sensibly own a decision unless it fully understands its pros and cons. If the board lacks relevant knowledge or if it would take too much time to gain comprehensive insight, it makes sense to bring in help. However, human helpers come with their set of thoughts, approaches, and incentives that are not necessarily transparent to the board. Additionally, the board members might lack the expert knowledge to evaluate their input. Delaware law is mindful of that, distinguishing the type of human helpers a board brings in. If these are officers or employees of the corporation, the trust the law places on boards by and large extends to these helping hands. With outside experts, it is less clear that their incentives are aligned with the corporation in the way officers and corporate committees are. Against this background, Delaware and German law allow the board to trust outside experts but tighten the requirements for doing so. DGCL § 141(e) stresses the careful selection of the expert, including its field of expertise. The German plausibility check mentioned above highlights the expectation that a board owns its decisions. Its members must critically quiz the expert and explore to what extent his suggestions match the board members' experience and intuition. In this way, the board still owns a decision that was informed by an outsider.

4. Judicial review of AI-augmented board decisions

There are countless ways for an AI to enhance the quality of a board decision. It might inform board members on market developments or consumer preferences, offer an exploration of data to understand a take-over target or the corporation's own risk exposure, produce an executive summary of a report or allow for a "conversation", as it were, with a LLM. This paper has been concerned with AI as a prediction machine, helping the board to assess and evaluate future scenarios. Judicial review comes into play if things go wrong. The foundation model, bought by the corporation, would have needed trimming to fulfil the board's expectations. The model's loss function might not have captured adequately what the board wanted to do, leading to a prediction that did not reflect the real world.[71] The training data might have been biased, triggering a skewed result.

A board that uses an AI prediction as its steppingstone is likely to face liability if an overreliance on the flawed AI-prediction led to a bad business decision. Following the control matrix' visualization, a first line of defense shows. Corporate law trusts board members to exercise discretion whenever a business judgment is at stake. Substantive control of what the board considered the best business strategy is low because the law is reluctant to make judges second-guess managerial decisions.

However, the trust placed on board members comes with the expectation that they own their decisions. This points towards the second line of defense. A board that painstakingly double-checks information it receives fully owns its decision. By contrast, the more a board outsources important parts of decision-making to inside or outside help, the stricter the judicial review, the more intense the relevant duties of care for selecting help. Delaware judges will, for instance, double-check the board's selection of an expert.[72] German judges will insist on board members to independently evaluate how plausible an expert's prediction was.

These finely-tuned rules have been developed against the background of human cooperation. They assume incentives for human behavior, potential for communication, the chance to build interpersonal trust or, alternatively, the need for skepticism and critical inquiry. An AI, by contrast, does not offer opinions or engages with board members for a critical discussion among peers (even if a LLM can be programmed to seem like one). Instead, it produces a data-driven statistical prediction. How does this fit into the control matrix' expectation of ownership? Is an AI like a technical support tool, a pocket calculator on steroids, as it were? Alternatively, should we treat an AI like a corporate officer or even like an outside expert?

The visualization of the control matrix shows how it is neither necessary to comprehensively define any AI as a purely technical support tool, nor to unfailingly analogize an AI to a human

[71] See Langenbucher (2023a), p. 40 for discussion of an example that involves an algorithm incorrectly predicting which patients in a hospital needed special care.
[72] See above IV.2.c).

expert, be it inside or outside the corporation. Instead, the matrix allows to move the needle, as it were, along the x-axis, ranging from low to high ownership.

The every-day AI-search engine resembles the purely technical support tool that corporate law has not deemed in need of special judicial scrutiny. This is true for both, business judgments and non-business judgments.

The same can be true for a very sophisticated AI such as Alphafold. The board that decided to invest heavily after having learned what Alphafold can do, scored very high on trust. Its decision was about a business judgment that the law entrusts to the board. It required the determination to move forward with further research and development, fully understanding that Alphafold's findings might not be as useful as initially hoped for. Deciding in the face of a known unknown along those lines is anything but unusual for a corporate board. Putting a probability on different outcomes and deciding which risk to take when faced with uncertainty is what the law trusts the board to do. Furthermore, the level of decision support by the AI in this scenario is low, hence, the board's ownership is high. The input delivered by the Alphafold AI was merely a trigger for a strategic business decision that the board fully owned. The intuition to go ahead with it, despite some known unknowns, controversies on whether spending the money on more research would pay off, and similar debates are classic issues of board discretion. Visualizing it in the control matrix, we look at the upper right-hand corner.

The AI credit underwriting scenario is a counterexample. The AI furnishes an assessment of credit default risk. One element of the decision to restructure the credit underwriting department concerns pricing loans, a standard business judgment that qualifies for a high level of trust towards the board. However, a major part of credit underwriting has to do with compliance with anti-discrimination laws. For those parts, there is low discretion accorded to the board. The board is not faced with a known-unknowns-situation, as was the case with Alphafold. It is not the board's task to put a probability on its credit model breaking the law and then move forward, in line with its risk appetite. Instead, we face a scenario where strict substantive control is in order. For a board to fully own a decision about complying with the law, it must make sure it has gathered enough information to not break the law. Visualizing the y-axis of the matrix helps to identify the level of judicial scrutiny. A black-box model that produces automated underwriting decisions achieves a very low score of ownership and, in turn, makes a case for intense judicial scrutiny. By contrast, an explainable model, working exclusively with a limited list of known data points, scores high on ownership. It allows to assess individual credit underwriting decisions. The board might not be able to converse with the AI like it would with a human peer. However, it has access to an explanation why the AI preferred one loan over another. Judicial scrutiny is still higher than in the Alphafold case. The reason for this is the trust dimension. Both, Alphafold and the explainable credit underwriting case concern a fully owned decision. However, Alphafold, in addition, scored high on trust, given that a pure business judgment was at stake. This distinguishes it from the explainable credit underwriting scenario, where parts of the decision concerned compliance with the law.

V. Summa

The paper has explored legal ramifications when board members employ AI to augment their decision-making. It focuses on AI as "prediction machines" that offer a glance into the future. I submit that predictions, with or without AI, are an every-day element of board decision-making. They imply an assessment and a risk evaluation of known unknowns, a paradigmatic example for a business judgment. Corporate law is well aware of the necessity to trust boards with making such decisions. Still, the law requires board members to eventually own their decisions, rather than diffuse responsibility among the various helpers that inform boards.

Two dimensions, ownership, and trust, provide the framework for understanding how corporate law shapes board decision-making. I introduce a "control matrix" to graphically illustrate these dimensions. If the law accords high levels of trust to the board, we look at business judgments that offer considerable discretion. Low levels of trust are characteristic for rule-bound decisions such as compliance. High levels of ownership characterize decisions that the board takes, by and large, without external support. The more elements of a decision a board outsources to officers, committees, or outside experts, the lower its ownership of the final board decision.

Augmenting decision-making via an AI, I claim, does not necessarily amount to a loss of ownership. Importantly, it does not involve a novel form of abdicating board authority. This applies to both, explainable and black-box AI. Rather, using an AI to inform boards can be understood in the broader context of boards drawing on support in the form of technical tools, inside or outside experts.

To fully understand the relevant standard of judicial review, the dimension of ownership must be complemented by its twin dimension of trust. I introduce a graphic representation to allow for situating a board decision along these two dimensions. Business judgments score high on trust. This makes for a flexible standard of judicial review. By contrast, non-business judgments fall under an enhanced standard of judicial review.

A board that comprehensively builds a non-business judgment on an AI prediction scores low on both dimensions, ownership, and trust. It faces intense judicial review. By contrast, a board that uses AI merely to inspire a classic business judgment scores high on both dimensions, entailing low judicial review. Two scenarios sit in between. A business judgment that relies predominantly on an AI prediction scores high on trust but low on ownership. A non-business judgment that the board takes with little help of an AI scores low on trust but high on ownership. Building on this framework, future research endeavors will have to spell out the details of relevant duties of care.

References

Agrawal, Ajay/Gans, Joshua/Goldfarb, Avi (2018): Prediction Machines: The Simple Economics of Artificial Intelligence, Harvard Business Review Press, Boston, Massachusetts.

Agrawal, Ajay/Gans, Joshua/Goldfarb, Avi (2022): Power and Prediction: The Disruptive Economics of Artificial Intelligence, Harvard Business Review Press, Boston, Massachusetts.

Bainbridge, Stephen M. (2020): Corporate Law, 4th revised edition, St. Paul, MN : Foundation Press.

Bender, Emily/Gebru, Timnit/McMillan-Major, Angelina/Shmitchell, Shmargaret (2021): On the Danger of Stochastic Parrots: Can Language Models Be Too Big?, FAccT'21, p. 610-623.

Bigoni, Maria/Bortolotti, Stefania/Parisi, Francesco/Porat, Ariel (2014): Unbundling Efficient Breach, University of Chicago Coase-Sandor Institute for Law & Economics Research Paper No. 695, available at: https://chicagounbound.uchicago.edu/cgi/viewcontent.cgi?article=2384&context=law_and_economics (last access 28 February 2024).

Bruner, Christophe M (2021-23): Artificially Intelligent Boards and the Future of Delaware Corporate Law, University of Georgia School of Law, Legal Studies Research Paper No. 2021-23; available at: htps://papers.ssrn.com/sol3/papers.cfm?abstract_id=3928237 (last access 28 February 2024).

Burton, Jason W/Stein, Mari Klara /Jensen, Tina Blegind (2020): A Systematic Review of Algorithm Aversion in Augmented Decision Making, 33 Journal of Behavioral Decision Making, pp. 220-239.

Data Governance Working Group of the Global Partnership of AI (2020): The Role of Data in AI. Report for the Data Governance Working Group of the Global Partnership of AI, Edinburgh, November 2020, available at: https://gpai.ai/projects/data-governance/role-of-data-in-ai.pdf (last access 28 February 2024).

Dubovitskaya, Elena/Buchholz, Annika (2023): Die Geschäftsleitung und der Rat des Algorithmus, ZIP 2023, pp. 63-73.

Ertel, Wolfgang (2021): Grundkurs Künstliche Intelligenz. Eine praxisorientierte Einführung, 5th edition, Springer, Wiesbaden.

Fleischer, Holger (2020): Beck-online Großkommentar, 1 October 2023, Beck, Munich.

German Federal Government (2024): Entwurf eines Ersten Gesetzes zur Änderung des Bundesdatenschutzgesetzes, available at: https://www.bmi.bund.de/SharedDocs/gesetzgebungsverfahren/DE/Downloads/kabinettsfassung/VII4/aendg-bdsg.pdf;jsessionid=F13370FB1AA19FA96EA8DD532C186914.live882?__blob=publicationFile&v=3 (last access 28 February 2024).

Glimcher, Paul (2004): Decisions, Uncertainty, and the Brain: The Science of Neuroeconomics, The MIT Press, Cambridge.

Grigoleit, Hans Christoph (2020): Aktiengesetz, 2nd edition, Beck, Munich.

Hacker, Philipp (2020): Europäische und nationale Regulierung von Künstlicher Intelligenz, in: NJW 2020, pp. 2142-2147.

Hawkins, Jeff (2021): A Thousand Brains. A New Theory of Intelligence, Basic Books, New York.

Jordan, James/Houssiau, Florimond/Cherubin, Giovanni/Cohen, Samuel N./Szpruch, Lukasz/Bottarelli, Mirko/Maple, Carsten/Weller, Adrian (2022): Synthetic data – what, why and how?, available at https://royalsociety.org/-/media/policy/projects/privacy-enhancing-technologies/Synthetic_Data_Survey-24.pdf (last access 28 February 2024).

Kahneman, Daniel (2012): Thinking, fast and slow, Penguin Books, New York.

Kay, John/King, Mervyn Radical (2020): Uncertainty, Decision-making for an unknowable future, The Bridge Street Press, London.

Koch, Jens (2023): in Aktiengesetz, 17th edition, Beck, Munich.

Kozyreva, Anastasia/Pleskac, Timothy/Pachur, Thorsten/Hertwig, Ralph (2019): Interpreting Uncertainty: A Brief History of Not Knowing, in: Hertwig, Ralph/Pleskac, Timothy/Pachur, Thorsten: Taming Uncertainty, pp. 343-362.

Langenbucher, Katja (2023a): Consumer Credit in The Age of AI – Beyond Anti-Discrimination Law, ECGI Working Paper Series in Law, Working Paper N° 663/2022, February 2023, available at: https://www.ecgi.global/working-paper/consumer-credit-age-ai-—-beyond-anti-discrimination-law (last access 28 February 2024).

Langenbucher, Katja (2023b): KI in der Leitungsentscheidung des Vorstands der Aktiengesellschaft, SAFE White Paper No. 96, available at: https://safe-frankfurt.de/policy-center/policy-publications/policy-publ-detailsview/publicationname/ki-in-der-leitungsentscheidung-des-vorstands-der-aktiengesellschaft.html (last access 28 February 2024).

Langenbucher, Katja (2023c): Künstliche Intelligenz in der Leitungsentscheidung des Vorstands, in: ZHR 187 (2023) pp. 723-738.

Langenbucher, Katja (2024): AI judgment rule(s), forthcoming in: Chicago Law Review online.

Larson, Eric J. (2021): The Myth of Artificial Intelligence. Why Computers Can´t Think the Way We Do, Harvard University Press, Cambridge.

March, James G. (1994): A Primer on Decision Making. How decisions happen, The Free Press, New York.

Marcus, Gary (2018): Deep Learning: A Critical Appraisal, 2 January 2018, available at: https://arxiv.org/abs/1801.00631 (last access 28 February 2024).

Miller, Tim (2023): Explainable AI is Dead, Long Live Explainable AI!. Hypothesis-driven decision support, 14 March 2023, preprint available at: https://arxiv.org/abs/2302.12389 (last access 28 February 2024).

Möslein, Florian (2018a): Robots in the Boardroom: Artificial Intelligence and Corporate Law, in: Barfield, Woodrow/Pagallo, Ugo: Research Handbook on the Law of Artificial Intelligence, Northhampton, pp. 649-669.

Möslein, Florian (2018b): Digitalisierung im Gesellschaftsrecht: Unternehmensleitung durch Algorithmen und künstliche Intelligenz?, in: ZIP 2018, pp. 204-212.

Nath, Ratjakishore (2009): Philosophy of Artificial Intelligence: A Critique of the Mechanistic Theory of Mind, Universal Publishers, Florida.

Noack, Ulrich (2019): Organisationspflichten und -strukturen kraft Digitalisierung, in: ZHR 183 (2019), pp. 105-144.

Petrin, Martin (2019): Corporate Management in the Age of AI, in: Columbia Business Law Review 3 (2019), p. 965-1030.

Posner, Richard (2009): Let Us Never Blame a Contract Breaker, in: 107 Michigan Law Review (2009) pp. 1349-1364.

Rolffs, Mathias (2023): Causality and Mental Causation, A Defense of Nonreductive Physicalism, available at: https://plato.stanford.edu/entries/physicalism/ (last access 28 February 2024).

Russell, Stuart/Norvig, Peter (2021): Artificial Intelligence: A Modern Approach, 4th edition, Pearson, Harlow.

Stiehl, Thomas/Marciniak-Czochra, Anna (2021): Intelligente Algorithmen und Gleichungen? – Eine Annäherung an die Intelligenz mathematischer Konzepte, in: Holm Hadulla, Rainer M./Funke, Joachim/Wink, Michael: Intelligence: Theoretical Foundations and Practical Applications, Heidelberger Jahrbücher Volume 6.

Spindler, Gerald (2023): Münchener Kommentar zum Aktiengesetz, 6th edition, Beck Munich.

Telle, Marc (2023): Einsatz Künstlicher Intelligenz zur vorbereitenden Unterstützung von Leitungsentscheidungen des Vorstands einer AG, in: Abhandlungen zum Deutschen und Europäischen Gesellschafts- und Kapitalmarktrecht (AGK), Volume 209.

Tversky, Amos/Kahneman, Daniel (1974): Judgment under Uncertainty: Heuristics and Biases, in: Science, New Series, Vol. 185, No. 4157, pp. 1124-1131.

Tversky, Amos/Kahneman, Daniel (1983): Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgement, in Psychological Review Volume 90 Number 4, pp. 293-315.

Uwer, Dirk (2019): Outsourcing digitaler Funktionen, in: ZHR 183 (2019) pp. 154-168.

about ECGI

The European Corporate Governance Institute has been established to improve *corporate governance through fostering independent scientific research and related activities.*

The ECGI will produce and disseminate high quality research while remaining close to the concerns and interests of corporate, financial and public policy makers. It will draw on the expertise of scholars from numerous countries and bring together a critical mass of expertise and interest to bear on this important subject.

The views expressed in this working paper are those of the authors, not those of the ECGI or its members.