

CAPCO

a wipro company

THE CAPCO INSTITUTE
JOURNAL
OF FINANCIAL TRANSFORMATION

REGULATION

Board decision making in the age of AI: Ownership and trust

KATJA LANGENBUCHER



GenAI

2024/2025 EDITION

THE CAPCO INSTITUTE

JOURNAL OF FINANCIAL TRANSFORMATION

RECIPIENT OF THE APEX AWARD FOR PUBLICATION EXCELLENCE

Editor

Shahin Shojai, Global Head, Capco Institute

Advisory Board

Lance Levy, Strategic Advisor

Owen Jelf, Partner, Capco

Suzanne Muir, Partner, Capco

David Oxenstierna, Partner, Capco

Editorial Board

Franklin Allen, Professor of Finance and Economics and Executive Director of the Brevan Howard Centre, Imperial College London and Professor Emeritus of Finance and Economics, the Wharton School, University of Pennsylvania

Philippe d'Arvisenet, Advisor and former Group Chief Economist, BNP Paribas

Rudi Bogni, former Chief Executive Officer, UBS Private Banking

Bruno Bonati, Former Chairman of the Non-Executive Board, Zuger Kantonalbank, and President, Landis & Gyr Foundation

Dan Breznitz, Munk Chair of Innovation Studies, University of Toronto

Urs Birchler, Professor Emeritus of Banking, University of Zurich

Elena Carletti, Professor of Finance and Dean for Research, Bocconi University, Non-Executive Director, UniCredit S.p.A.

Lara Cathcart, Associate Professor of Finance, Imperial College Business School

Géry Daeninck, former CEO, Robeco

Jean Dermine, Professor of Banking and Finance, INSEAD

Douglas W. Diamond, Merton H. Miller Distinguished Service Professor of Finance, University of Chicago

Elroy Dimson, Emeritus Professor of Finance, London Business School

Nicholas Economides, Professor of Economics, New York University

Michael Enthoven, Chairman, NL Financial Investments

José Luis Escrivá, President, The Independent Authority for Fiscal Responsibility (AIReF), Spain

George Feiger, Pro-Vice-Chancellor and Executive Dean, Aston Business School

Gregorio de Felice, Head of Research and Chief Economist, Intesa Sanpaolo

Maribel Fernandez, Professor of Computer Science, King's College London

Allen Ferrell, Greenfield Professor of Securities Law, Harvard Law School

Peter Gomber, Full Professor, Chair of e-Finance, Goethe University Frankfurt

Wilfried Hauck, Managing Director, Statera Financial Management GmbH

Pierre Hillion, The de Picciotto Professor of Alternative Investments, INSEAD

Andrei A. Kirilenko, Reader in Finance, Cambridge Judge Business School, University of Cambridge

Katja Langenbacher, Professor of Banking and Corporate Law, House of Finance, Goethe University Frankfurt

Mitchel Lenson, Former Group Chief Information Officer, Deutsche Bank

David T. Llewellyn, Professor Emeritus of Money and Banking, Loughborough University

Eva Lomnicka, Professor of Law, Dickson Poon School of Law, King's College London

Donald A. Marchand, Professor Emeritus of Strategy and Information Management, IMD

Colin Mayer, Peter Moores Professor of Management Studies, Oxford University

Francesca Medda, Professor of Applied Economics and Finance, and Director of UCL Institute of Finance & Technology, University College London

Pierpaolo Montana, Group Chief Risk Officer, Mediobanca

John Taysom, Visiting Professor of Computer Science, UCL

D. Sykes Wilford, W. Frank Hipp Distinguished Chair in Business, The Citadel

CONTENTS

TECHNOLOGY

08 Mindful use of AI: A practical approach

Magnus Westerlund, Principal Lecturer in Information Technology and Director of the Laboratory for Trustworthy AI, Arcada University of Applied Sciences, Helsinki, Finland

Elisabeth Hildt, Affiliated Professor, Arcada University of Applied Sciences, Helsinki, Finland, and Professor of Philosophy and Director of the Center for the Study of Ethics in the Professions, Illinois Institute of Technology, Chicago, USA

Apostolos C. Tsolakis, Senior Project Manager, Q-PLAN International Advisors PC, Thessaloniki, Greece

Roberto V. Zicari, Affiliated Professor, Arcada University of Applied Sciences, Helsinki, Finland

14 Understanding the implications of advanced AI on financial markets

Michael P. Wellman, Lynn A. Conway Collegiate Professor of Computer Science and Engineering University of Michigan, Ann Arbor

20 Auditing GenAI systems: Ensuring responsible deployment

David S. Krause, Emeritus Associate Professor of Finance, Marquette University

Eric P. Krause, PhD Candidate – Accounting, Bentley University

28 Innovating with intelligence: Open-source Large Language Models for secure system transformation

Gerhardt Scriven, Executive Director, Capco

Tony Moenicke, Senior Consultant, Capco

Sebastian Ehrig, Senior Consultant, Capco

38 Multimodal artificial intelligence: Creating strategic value from data diversity

Cristián Bravo, Professor, Canada Research Chair in Banking and Insurance Analytics, Department of Statistical and Actuarial Sciences, Western University

46 GenAI and robotics: Reshaping the future of work and leadership

Natalie A. Pierce, Partner and Chair of the Employment and Labor Group, Gunderson Dettmer

ORGANIZATION

56 How corporate boards must approach AI governance

Arun Sundararajan, Harold Price Professor of Entrepreneurship and Director of the Fubon Center for Technology, Business, and Innovation, Stern School of Business, New York University

66 Transforming organizations through AI: Emerging strategies for navigating the future of business

Feng Li, Associate Dean for Research and Innovation and Chair of Information Management, Bayes Business School (formerly Cass), City St George's, University of London

Harvey Lewis, Partner, Ernst & Young (EY), London

74 The challenges of AI and GenAI use in the public sector

Albert Sanchez-Graells, Professor of Economic Law, University of Bristol Law School

78 AI safety and the value preservation imperative

Sean Lyons, Author of Corporate Defense and the Value Preservation Imperative: Bulletproof Your Corporate Defense Program

92 Generative AI technology blueprint: Architecting the future of AI-infused solutions

Charlotte Byrne, Managing Principal, Capco

Thomas Hill, Principal Consultant, Capco

96 Unlocking AI's potential through metacognition in decision making

Sean McMinn, Director of Center for Educational Innovation, Hong Kong University of Science and Technology

Joon Nak Choi, Advisor to the MSc in Business Analytics and Adjunct Associate Professor, Hong Kong University of Science and Technology

REGULATION

104 Mapping GenAI regulation in finance and bridging the gaps

Nydia Remolina, Assistant Professor of Law, and Fintech Track Lead, SMU Centre for AI and Data Governance, Singapore Management University

112 Board decision making in the age of AI: Ownership and trust

Katja Langenbucher, Professor of Civil Law, Commercial Law, and Banking Law, Goethe University Frankfurt

122 The transformative power of AI in the legal sector: Balancing innovation, strategy, and human skills

Eugenia Navarro, Lecturer and Director of the Legal Operations and Legal Tech Course, ESADE

129 Remuneration on the management board in financial institutions: Current developments in the framework of supervisory law, labor law, behavioral economics and practice

Julia Redenius-Hövermann, Professor of Civil Law and Corporate Law and Director of the Corporate Governance Institute (CGI) and the Frankfurt Competence Centre for German and Global Regulation (FCCR), Frankfurt School of Finance and Management

Lars Hinrichs, Partner at Deloitte Legal Rechtsanwaltsgesellschaft mbH (Deloitte Legal) and Lecturer, Frankfurt School of Finance and Management



CAPCO CEO WELCOME

DEAR READER,

Welcome to our very special 60th edition of the Capco Journal of Financial Transformation.

The release of this milestone edition, focused on GenAI, reinforces Capco's enduring role in leading conversations at the cutting edge of innovation, and driving the trends shaping the financial services sector.

There is no doubt that GenAI is revolutionizing industries and rapidly accelerating innovation, with the potential to fundamentally reshape how we identify and capitalize on opportunities for transformation.

At Capco, we are embracing an AI infused future today, leveraging the power of GenAI to increase efficiency, innovation and speed to market while ensuring that this technology is used in a pragmatic, secure, and responsible way.

In this edition of the Capco Journal, we are excited to share the expert insights of distinguished contributors across academia and the financial services industry, in addition to drawing on the practical experiences from Capco's industry, consulting, and technology SMEs.

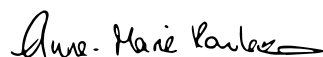
The authors in this edition offer fresh perspectives on the mindful use of GenAI and the implications of advanced GenAI on financial markets, in addition to providing practical and safe frameworks for boards and firms on how to approach GenAI governance.

The latest advancements in this rapidly evolving space demonstrate that the potential of GenAI goes beyond automating and augmenting tasks, to truly helping organizations redefine their business models, processes and workforce strategies. To unlock these benefits of GenAI, I believe that firms need a culture that encourages responsible experimentation and continuous learning across their organization, while assessing the impact of the potential benefits against a strategic approach and GenAI framework.

I am proud that Capco today remains committed to our culture of entrepreneurialism and innovation, harnessed in the foundation of our domain expertise across our global teams. I am proud that we remain committed to our mission to actively push boundaries, championing the ideas that are shaping the future of our industry, and making a genuine difference for our clients and customers – all while ensuring to lead with a strategy that puts sustained growth, integrity and security at the forefront of what we do.

I hope you'll find the articles in this edition both thought-provoking and valuable as you create your organization's GenAI strategy and future direction. As we navigate this journey together, now is the time to be bold, think big, and explore the possibilities.

My greatest thanks and appreciation to our contributors, readers, clients, and teams.



Annie Rowland, **Capco CEO**

BOARD DECISION MAKING IN THE AGE OF AI: OWNERSHIP AND TRUST

KATJA LANGENBUCHER | Professor of Civil Law, Commercial Law, and Banking Law, Goethe University Frankfurt¹

ABSTRACT

This paper describes how artificial intelligence (AI) might augment board decision making and explores legal ramifications of this development. The article begins by providing a brief overview of the use of AI as a “prediction machine”² for board decisions, and then zooms in on two core characteristics that explain what corporate law requires from board decision making: that board members fully own their decisions and that board members are trusted to form business judgments, immune from judicial second-guessing. The paper makes two contributions to the debate: it rejects the notion that black-box AI may not be used for board decision making and proposes a graphic control matrix to identify low, medium, and enhanced judicial scrutiny when boards use AI to inform their decisions.

1. INTRODUCTION

Explaining human intelligence is an intriguing topic [Langenbucher (2023b)]. For some, it represents human singularity, while others emphasize the dependence of human intelligence on mechanistic operations [Glimcher (2004), Rolffs (2023), Stiehl and Marciniak-Czochra (2021)]. 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 [Hawkins (2021), Larson (2021), Nath (2009)]. Arguably, an uncontested point of departure is that machines can sometimes surpass human performance when it comes 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?

The next section provides a brief overview on the use of AI as a “prediction machine” for board decisions. It reminds the reader that statistics has traditionally filled this role and hints at similarities and differences with using machine learning (ML). The following section 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 highlights 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 from judicial second-guessing. The expectation that boards own their decisions implies that they must not abdicate their authority.

This article explores how this principle is impacted when boards enhance their decision making with an AI. It then moves on to examine how corporate law has framed ownership of a

¹ The author is also affiliated with SciencesPo in France and a regular visiting faculty at Fordham Law School in the U.S. This paper has appeared as a working paper in the ECGI/WP series. Additional thoughts (that go deeper on decision theory-related questions) will appear at Chicago Law Review Online.

² For this term see the title of Agrawal et al. (2018), as well as Agrawal et al. (2022) and Russel and Norvig (2021).

board decision when technical support tools or human experts inform board members. Rather than analogizing AI to one of these helpers, this article introduces the second dimension: trust. It claims that the standard of judicial review has moved along these two dimensions, ownership and trust. The same logic, this article suggests, applies to board decisions that integrate AI. The paper concludes with a graphic visualization of these dimensions.

2. PREDICTION MACHINES

The term “artificial intelligence” includes various implementations, ranging from logic, ML, and neural networks to large language models (LLM) and robotics [Russell and Norvig (2021)]. These correspond with a diverse set of potential use cases in corporate life. The scenario this article explores is 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 take place. Most management decisions imply predictions of that type and employing statistics to that end is a standard tool. Statisticians work on inferences about the relationship between different variables, based on a hypothesis.³ Consider the case of the management board of a bank that decides to reduce the number of brick-and-mortar branches and 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 that: the age of a customer is a core factor in 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) [Russel and Norvig (2021)]. Its predictions about the willingness of bank customers to switch to online banking could mirror that of the statistician’s. Additionally, it might bring out unanticipated correlations. Both allow the board to react; for instance, via targeting its marketing towards specific groups.

2.1 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 [Russel and Norvig (2021)].

There are three basic forms of machine learning.⁴ In supervised learning, the AI is programmed to map input to output [Russel and Norvig (2021)]. 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). Supervised learning requires large datasets 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 dataset [Russel and Norvig (2021)]. Unsupervised learning responds to these exploratory needs. It makes the AI independently find structures and patterns. The programmer does not specify the way in which the AI performs the identification task, nor do they specify a goal or label the data. This distinguishes the technique from supervised learning, where the AI has a previously known objective. With unsupervised learning, the AI shows the user a way of sorting disordered data. This approach requires very large datasets and computers with enormous computing power. Its use for daily management will for most corporations mean buying the AI from a provider.

Learning by reinforcement occupies a space between supervised and unsupervised learning [Russel and Norvig (2021)]. The AI works without pre-labeled training data and is programmed to perform certain sequences, such as a board game [Russel and Norvig (2021)] or a robotics task [Ertel (2021)]. It receives positive or negative human feedback after completing its task. Each following round, the AI adapts its strategy to receive positive feedback [Russel and Norvig (2021)].

³ 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.

⁴ Russell and Norvig (2021), p. 671.

“

Fresh efforts must go into understanding what corporate law expects from board members who rely on AI support to augment their decision making.

”

2.2 Induction engines

To a statistician, it comes as no surprise that good data is a core ingredient for 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 [Russel and Norvig (2021)]. Predictions based on a carefully curated [Data Governance Working Group of the Global Partnership of AI (2020)], possibly even synthetic [Jordan et al. (2022)], 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 [Russel and Norvig (2021)]. Consider the example of the bank executive deciding on branch closures. If the AI is trained on a small dataset, 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 dataset, but will not 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 noise that risks producing skewed results [Bender et al. (2021)]. To bolster management decisions, a synthetic, curated, or at least “cleaned” dataset might be more useful.

Lastly, it is helpful to keep in mind that AIs are “induction engines” [Larson (2021)]. 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 [Marcus (2018)].

3. 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 [March (1994)]. 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 makes it possible 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 from decision makers at officer and employee level. Firstly, the law expects board members to fully own their decisions. By way of illustration, see Delaware General Corporate Law § 141(a) that 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. 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,⁵ rather than encouraging judicial second-guessing.

⁵ Joy v. North 692 F.2d 880, 885 (2d Cir.), cert. denied, 460 U.S. 1051 (1983); Bainbridge (2020).

4. 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 that 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, which deals with an AI enhancing board decision making. Still, three remarks are in order to hint at the 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 [Agrawal et al. (2022)]. 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 choose one AI over another. The E.U. 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.⁶

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 a reasonable cost? What type of data is needed (for instance, open source, curated, synthetic, labeled, etc.), 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 [Langenbacher (2023a)]?

Additionally, the intended "workplace" for the AI might require specific features. Employees who cooperate with an AI often need special skills.⁷ 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 [Marcus (2018)] or abductive conclusions [Larson (2021)] – must be balanced against an increase in efficiency.

5. 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 into its own deliberations. 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.

Board resolutions of this type operate under the corporate law principles mentioned above [Langenbacher (2023c)]. On the one hand, the business judgment rule represents 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 says nothing against the board asking for support in its decision making [Fleischer (2023)]. 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 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.

⁶ Annex III spells out the high-risk AI systems referred to in Art. 6(2) AI Act.

⁷ For "automation bias" see Art. 14(4)(b) AI Act.

5.1 Abdicating authority: Does AI take over?

Traditionally, abdication has been understood as trading away the board's discretion.⁸ Against that background, so-called "black-box" AI has troubled some scholars [Dubovitskaya and Buchholz (2023)]. 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 [Langenbucher (2023c)].

Many scenarios are straightforward. It does not hurt to prepare a board decision by googling relevant facts. 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 from a board that trades away its authority to a human [Fleischer (2023), Möslin (2018)]. The relevant issue at stake is the same to the extent that the board does not have discretion to decide as it seems fit. From this perspective, it does not matter whether the AI is explainable or not.

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 [Langenbucher (2024), advocating for a distinction along those lines: Fleischer (2023) and Noack (2019)]. 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 [Agrawal et al. (2022)]. Arguably, building a basic understanding of technology and trying to grasp the inherent logic of algorithms provides some relief [Fleischer (2023)]. Still, few board members will become experts in AI technology.

Additionally, it does not help that humans are known to be subject to a wide variety of decision making anomalies when it comes to assessing statistical probabilities [Kahneman (2012), Kay and King (2020), Tversky and Kahneman (1983), Tversky and Kahneman (1974), Kozyreva et al. (2019)], a core element of AI. In the same way that the AI's "workplace" on any of the corporation's hierarchical levels must be carefully structured, the board's own "workplace" in cooperation with an AI also needs structure. Human cognition follows different patterns

than an AI [Burton et al. (2020)]. 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 [Burton et al. (2020)], 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 [Burton et al. (2020)]. 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 [Miller (2023)]. When the stakes are high, humans tend to demand "slow and effortful consideration of evidence", even if empirical evidence does not necessarily find that this strategy leads to better decision making [Burton et al. (2020)].

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 everyday AI, such as the Google search engine or ChatGPT, produces its results, the same is true for a pocket calculator or a GPS. The reason we do not consider the use of these tools as abdication of 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.

5.2 Informing board decision making: When to trust an AI

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 cause 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

⁸ See for Germany: Telle (2023) and Möslin (2018). For the U.S., see Bruner (2021) and Petrin (2019).

forming an informed and reasonable judgment. The prediction that an AI makes, explainable or black-box, can be just that: a known unknown.

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

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.

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.⁹ 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 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 [Langenbucher (2023a)]: 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 U.S. Equal Credit Opportunities Act or the E.U. 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.

Can the board reason as follows:

- We understand that the board must not allow credit underwriting decisions to vary along a protected variable.
- Our AI is programmed to disregard protected variables when making its prediction.
- We understand that this AI might use proxy variables, but the extent to which it does is a known unknown.
- The law trusts boards to integrate known unknowns in its decision if the board evaluates the ensuing risk.
- 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.

Assuming an affirmative duty for the board to obey the law,¹⁰ the decision regarding 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 than the installation of an input restriction for protected characteristics.¹¹ Following

⁹ 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.

¹⁰ Proponents of (some version of) the efficient breach theory do not share this assumption, for a foundation see Posner (2009); for an overview see Bigoni et al. (2014).

¹¹ See the upcoming German law on credit scoring, German Federal Government (2024).

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,¹² or prohibit the use of black-box AI altogether. What distinguishes this scenario from the AlphaFold¹³ 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.

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 from situations where humans support board decision making, especially if the human help is not an employee of the corporation. The law expects 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 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.¹⁴

5.3 Visualizing ownership and control

To illustrate how courts review board decision making, I have provided a four-square control matrix in Figure 1. 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-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, which 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, and 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.

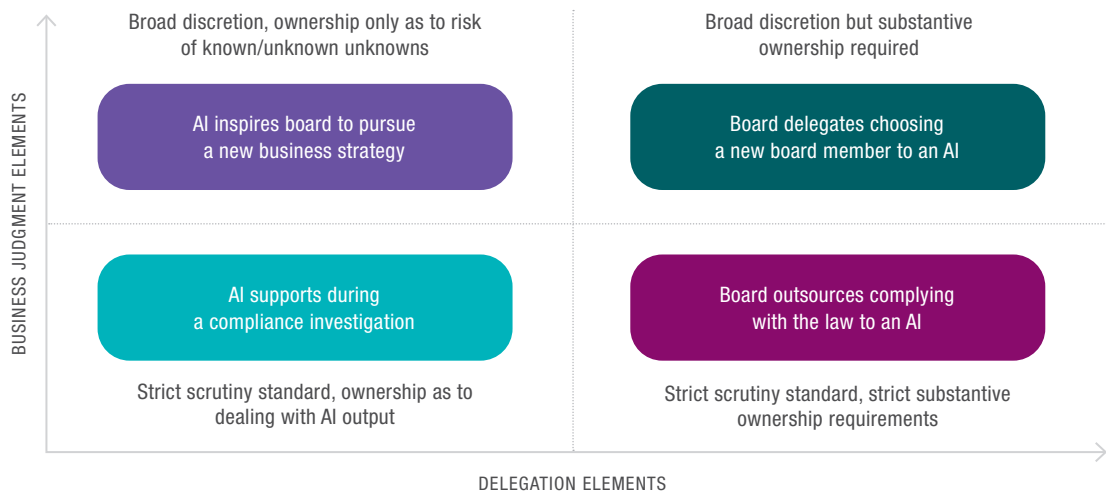
The fourth square, which is located in the lower left-hand corner, shows decisions that were reached with much outside help, hence, score low on ownership. Additionally, these decisions score low on discretion, because they include few or no business judgment elements. This square symbolizes the highest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

¹² See Art. 18(3) Consumer Credit Directive (EU) 2023/2225 on data gathered from social media.

¹³ AlphaFold is an artificial intelligence program developed by DeepMind, a subsidiary of Alphabet, which performs predictions of protein structure. The program is designed as a deep learning system (<https://tinyurl.com/wmvkjfha>).

¹⁴ *Brehm v. Eisner*, 746 A.2d 244, 262 (Del. 2000), <https://tinyurl.com/ywnswc9c>

Figure 1: How courts review board decision making

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 us 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, 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 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 the 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 own sets 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 between 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 law allows the board to trust outside experts but tightens the requirements for doing so, by stressing the careful selection of the expert, including their field of expertise.

A board that uses an AI prediction as its stepping stone 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 the 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. By way of illustration: Delaware judges will double-check the board’s selection of an 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 unflinchingly analogize an AI to a human 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 everyday AI search engine resembles the purely technical support tool that corporate law has not deemed to be 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 that inspires the board to move ahead with a novel product. Its decision concerns a business judgment that the law entrusts to the board. 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. Visualizing it in the control matrix, we look at the upper right-hand corner.

Using an AI in a credit underwriting scenario is a counter example. Assume an 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. 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 makes it possible to assess individual credit underwriting decisions. The board might not be able to converse with the AI like it would with a human peer, but it has access to an explanation regarding why the AI preferred one loan over another.

6. CONCLUSION

This article has explored the legal ramifications of board members employing 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 everyday 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 or internal and external 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 and a non-business judgment that the board takes with little help from 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, A., J. Gans, and A. Goldfarb, 2018, *Prediction machines: the simple economics of artificial intelligence*, Harvard Business Review Press
- Agrawal, A., J. Gans, and A. Goldfarb, 2022, *Power and prediction: the disruptive economics of artificial intelligence*, Harvard Business Review Press, Boston
- Bainbridge, S. M., 2020, *Corporate law*, 4th revised edition, Foundation Press
- Bender, E., T. Gebru, A. McMillan-Major, and S. Shmitchell, 2021, "On the danger of stochastic parrots: can language models be too big?" *Proceedings of FAccT '21*, 610–623
- Bigoni, M., S. Bortolotti, F. Parisi, and A. Porat, 2014, "Unbundling efficient breach," *University of Chicago Coase-Sandor Institute for Law & Economics research paper no. 695*, <https://tinyurl.com/s9mmb4m9>
- Bruner, C. M., 2021, "Artificially intelligent boards and the future of Delaware Corporate Law," *University of Georgia School of Law, Legal Studies research paper no. 2021-23*, <https://tinyurl.com/ys7p5m3d>
- Burton, J. W. M. K. Stein, and T. B. Jensen, 2020, "A systematic review of algorithm aversion in augmented decision making," *Journal of Behavioral Decision Making* 33, 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*, <https://tinyurl.com/273rpdht>
- Dubovitskaya, E., and A. Buchholz, 2023, "Die Geschäftsleitung und der Rat des Algorithmus," *ZIP* 2023, 63-73
- Ertel, W., 2021, *Grundkurs Künstliche Intelligenz: Eine praxisorientierte Einführung*, Springer
- Fleischer, H., 2020, *Beck-online Großkommentar German Federal Government*, 2024, *Entwurf eines Ersten Gesetzes zur Änderung des Bundesdatenschutzgesetzes*, <https://tinyurl.com/48v25m42>
- Glimcher, P., 2004, *Decisions, uncertainty, and the brain: the science of neuroeconomics*, The MIT Press
- Hawkins, J., 2021, *A thousand brains. A new theory of intelligence*, Basic Books
- Jordan, J., F. Houssiau, G. Cherubin, S. N. Cohen, L. Szpruch, M. Bottarelli, C. Maple, and A. Weller, 2022, "Synthetic data – what, why and how?" *The Royal Society*, <https://tinyurl.com/52nuxsby>
- Kahneman, D., 2012, *Thinking, fast and slow*, Penguin Books
- Kay, J., and M. R. King, 2020, *Uncertainty, decision-making for an unknowable future*, The Bridge Street Press
- Koch, J., 2023, in *Aktiengesetz*, 17th edition, Beck
- Kozyreva, A., T. Pleskac, T. Pachur, and R. Hertwig, 2019, "Interpreting uncertainty: a brief history of not knowing," in Hertwig, R., T. Pleskac, and T. Pachur (eds.), *Taming uncertainty*, MIT Press
- Langenbacher, K., 2023a, "Consumer credit in the age of AI – beyond anti-discrimination law," *ECGI working paper series in Law no. 663/2022*, <https://tinyurl.com/ye8mh3nd>
- Langenbacher, K., 2023b, "KI in der Leitungsentscheidung des Vorstands der Aktiengesellschaft," *SAFE white paper no. 96*, <https://tinyurl.com/4s3yd3rf>
- Langenbacher, K., 2023c, "Künstliche Intelligenz in der Leitungsentscheidung des Vorstands," *ZHR* 187, 723-738
- Langenbacher, K., 2024, "AI judgment rule(s)," *forthcoming, Chicago Law Review online*
- Larson, E. J., 2021, *The myth of artificial intelligence. Why computers can't think the way we do*, Harvard University Press
- March, J. G., 1994, *A primer on decision making. How decisions happen*, The Free Press
- Marcus, G., 2018, "Deep learning: a critical appraisal," <https://tinyurl.com/z99rczsn>
- Miller, T., 2023, "Explainable AI is dead, long live explainable AI!. Hypothesis-driven decision support," <https://tinyurl.com/mr275aut>
- Möslein, F., 2018, "Robots in the boardroom: artificial intelligence and corporate law," in Barfield, W., and U. Pagallo (eds.), *Research handbook on the law of artificial intelligence*, Edward Elgar Publishing
- Nath, R., 2009, *Philosophy of artificial intelligence: a critique of the mechanistic theory of mind*, Universal Publishers
- Noack, U., 2019, "Organisationspflichten und -strukturen kraft Digitalisierung," *ZHR* 183, 105-144
- Petrin, M., 2019, "Corporate management in the age of AI," *Columbia Business Law Review* 3, 965-1030
- Posner, R., 2009, "Let us never blame a contract breaker," *Michigan Law Review* 107, 1349-1364
- Rolffs, M., 2023, "Causality and mental causation, a defense of nonreductive physicalism," *Stanford Encyclopedia of Philosophy*, <https://tinyurl.com/4ca94f65>
- Russell, S., and P. Norvig, 2021, *Artificial intelligence: a modern approach*, 4th edition, Pearson
- Stiehl, T., and A. Marciniak-Czochra, 2021, "Intelligente Algorithmen und Gleichungen? – Eine Annäherung an die Intelligenz mathematischer Konzepte," in Holm Hadulla, R. M., J. Funke, and M. Wink, (ed.), *Intelligence: theoretical foundations and practical applications*, Heidelbergberger Jahrbücher vol. 6
- Spindler, G., 2023, *Münchener Kommentar zum Aktiengesetz*, 6th edition, Beck
- Telle, M., 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, A., and D. Kahneman, 1974, "Judgment under uncertainty: heuristics and biases," *Science, New Series* 185:4157, 1124-1131
- Tversky, A., and D. Kahneman, 1983, "Extensional versus intuitive reasoning: the conjunction fallacy in probability judgement," *Psychological Review* Volume 90:4, 293-315

© 2024 The Capital Markets Company (UK) Limited. All rights reserved.

This document was produced for information purposes only and is for the exclusive use of the recipient.

This publication has been prepared for general guidance purposes, and is indicative and subject to change. It does not constitute professional advice. You should not act upon the information contained in this publication without obtaining specific professional advice. No representation or warranty (whether express or implied) is given as to the accuracy or completeness of the information contained in this publication and The Capital Markets Company BVBA and its affiliated companies globally (collectively "Capco") does not, to the extent permissible by law, assume any liability or duty of care for any consequences of the acts or omissions of those relying on information contained in this publication, or for any decision taken based upon it.

ABOUT CAPCO

Capco, a Wipro company, is a global management and technology consultancy specializing in driving transformation in the energy and financial services industries. Capco operates at the intersection of business and technology by combining innovative thinking with unrivalled industry knowledge to fast-track digital initiatives for banking and payments, capital markets, wealth and asset management, insurance, and the energy sector. Capco's cutting-edge ingenuity is brought to life through its award-winning Be Yourself At Work culture and diverse talent.

To learn more, visit www.capco.com or follow us on LinkedIn, Instagram, Facebook, and YouTube.

WORLDWIDE OFFICES

APAC

Bengaluru – Electronic City
Bengaluru – Sarjapur Road
Bangkok
Chennai
Gurugram
Hong Kong
Hyderabad
Kuala Lumpur
Mumbai
Pune
Singapore

MIDDLE EAST

Dubai

EUROPE

Berlin
Bratislava
Brussels
Dusseldorf
Edinburgh
Frankfurt
Geneva
Glasgow
London
Milan
Paris
Vienna
Warsaw
Zurich

NORTH AMERICA

Charlotte
Chicago
Dallas
Houston
New York
Orlando
Toronto

SOUTH AMERICA

São Paulo

THIS UNIQUE IMAGE WAS GENERATED USING MID-JOURNEY, STABLE DIFFUSION AND ADOBE FIREFLY

WWW.CAPCO.COM



CAPCO
a wipro company