

HOW LEGAL FRAMEWORKS SHAPE AI DECISION-MAKING IN CORPORATE MANAGEMENT

SPULBAR LUCIAN

PHD STUDENT, UNIVERSITY OF CRAIOVA, DOCTORAL SCHOOL OF ECONOMICS SCIENCES „EUGENIU CARADA”, CRAIOVA, ROMANIA

e-mail: lucian.spulbar01@gmail.com

Abstract

The relationship between legal systems and artificial intelligence decision-making in managing business organizations is becoming increasingly crucial in light of fast-paced technological progress and growing legal complexities. The topic for this research is the role of binding legal norms, exemplified by the EU's AI Act and data governance laws, in both directly and indirectly conditioning managerial decisions through artificial intelligence systems. The study will be conducted through qualitative comparative research on large global organizations operating in high-risk business environments and will isolate instances of patterns and trends in managing legal risks through artificial intelligence systems and processes. The research will draw on semi-structured, qualitative research on views obtained from compliance professionals, artificial intelligence software engineers, and CXO personnel, and on content and text analysis of corporate governance reports to evaluate the effectiveness of legal compliance in altering organizational behavior patterns. The research will show how legal systems are no longer external influences on business organizations, as they are increasingly being constitutive of organizational behavior through various internal mechanisms, including accountability algorithms, legal risk software, and cross-functional ethics committees, thereby increasingly subjecting artificial intelligence decision processes to a set of legally informed constraints and determinations. This will show how this synthesis of artificial intelligence decision systems and legal systems is leading to a paradigm shift in management decisions and inference, thereby leading to greater resilience and greater trust-building in corporate organizations by increasingly positioning the law as a resource, and thereby leading to greater confidence and trust and trust-building in artificial corporate management systems.

Keywords: Artificial intelligence governance, corporate management, legal compliance, algorithmic decision-making, regulatory frameworks.

Clasificare JEL : L21, M48, K23.

1. Introduction

The relationship between law and artificial intelligence has, in recent years, moved beyond the technical or legal *background* and has come to directly affect the way organizations are run. As companies increasingly use AI systems to support managerial decisions, from risk assessment and investment prioritization to supplier selection, customer scoring, or human resource management, it is becoming increasingly clear that legal rules can no longer be treated as simple external conditions, checked off at the end of a project. On the contrary, with the acceleration of technological progress and the densification of legal obligations, the legal framework begins to *sit* inside organizational processes and shape the very way in which decisions are constructed, justified, and validated.

In this paper, we start from the idea that legal systems, through binding norms such as the upcoming (and already influential) AI Act of the European Union, together with legislation on data governance, not only limit the use of artificial intelligence, but also organize it. More precisely, I am interested in how these norms condition management both directly, through explicit requirements (risk assessments, transparency, auditability, clearly distributed responsibilities), and indirectly, through internal mechanisms that emerge in response to the pressure of compliance: accountability algorithms, legal risk management software, ethics committees or cross-functional structures that translate law into operational rules and technical parameters. In other words, we try to capture how law ends up being *encapsulated* in the decision-making architecture of AI and, through this, influences the behavior of the organization.

The question that guides me is simple in formulation, but difficult in consequences: how do legal frameworks shape the way AI makes decisions in corporate management, especially in large organizations operating in high-risk environments? Behind this question lies a practical stake: if AI becomes a real governance tool, then the way we *anchor* it legally will determine not only how safe it is, but also how credible and accepted it becomes inside and outside the company. And the theoretical stake is equally important: I want to show that law no longer functions only as a constraint coming from the outside, but increasingly as a constitutive element of the organization, that is, a factor that enters the internal mechanisms and changes the way decisions are made.

To explore this dynamic, I draw on comparative qualitative research that follows global organizations in contexts where risk is structural, either by the nature of the industry or by exposure to regulation and liability. My method combines semi-structured interviews with compliance professionals, AI software engineers, and executives (CXOs), with content analysis and textual analysis of corporate governance documents. What I look for is not a *snapshot* of compliance, but patterns: how certain solutions are repeated, what tensions arise between algorithmic efficiency and legal requirements, how accountability processes are built, and what kind of organizational change they generate.

Finally, I propose to argue that this synthesis between AI systems and legal systems produces a paradigm shift in management. The decision is no longer just an optimized inference, but becomes an inference made under embedded legal constraints, with audit trails, distributed responsibilities and internal control mechanisms. And this transformation, although it may seem at first glance to be a brake, often functions as a form of resilience: it creates the conditions for trust, both inside the organization (between teams and decision levels) and outside (regulators, partners, customers). In the logic I follow here, law is not just a limit, but becomes a managerial resource that can stabilize and legitimize the decision assisted by artificial intelligence.

2. Theoretical and conceptual framework

Before I get into the analysis, I feel the need to establish a *common language* about what AI-assisted decision-making in management actually means. In organizations, AI does not make decisions in a vacuum: it structures options, prioritizes risks, recommends actions, and produces justifications in formats that end up weighing in front of management. This is where I start with the idea of human-AI *symmetry*: management is not replaced, but reconfigured, because humans and algorithmic systems end up collaborating under conditions of uncertainty and complexity (Jarrahi, 2018). And this collaboration is not neutral: once the organization gets used to scores, classifications, and predictions, the algorithm becomes an organizational actor that can redraw boundaries of responsibility, modes of control, and even internal realities of work (Faraj et al., 2018).

In my point of view, this is where law comes in as a shaping force, but not just as a limit. The literature on *accountable algorithms* insists that simple transparency (*open source*) does not solve the problems of fairness and compliance, but rather a combination of verifiable techniques and procedures that can demonstrate that a system complies with relevant standards is more useful (Kroll, 2015). This is important for my topic, because it moves the discussion from the abstract (*AI must be ethical*) to the operational area: what needs to exist in the company for the decision assisted by AI to be controllable, demonstrable and, if necessary, contestable.

When I bring up the European framework, I do not do it as a decoration, but as a concrete mechanism for *translating* norms into internal requirements. The logic of the EU AI Act is risk-based: certain uses are prohibited, others are treated as *high-risk* and attract governance, documentation and control obligations (European Union, 2024; Veale & Zuiderveen Borgesius, 2021). For me, this is the key: regulation does not just constrain behavior, but forces organizations to build internal infrastructure (processes, roles, records, auditability) that ends up sticking to the way AI is *allowed* to function in management.

At the same time, data governance remains, in practice, the ground on which any AI decision rests. The GDPR sets rules for accountability and limitation of processing, and for automated decisions there is a tension between what is legally required and what is technically feasible when it comes to *explanations* (European Union, 2016). It seems useful here to me not to promise more than the law can deliver: the discussion of a *right to explanation* in the GDPR is much more nuanced than is popularly believed, and the literature shows why triumphalist interpretations can be misleading (Wachter et al., 2017). Along these lines, interpretability becomes less of a fad and more of a governance condition: you need explanations not to *philosophically understand* the model, but to be able to manage risk, justify the decision and control it procedurally (Doshi-Velez & Kim, 2017).

To describe how the law reaches *inside* the organization, I am greatly helped by the literature on audit, verification and the institutionalization of risk management. The idea of *audit society* captures precisely the tendency of organizations to transform compliance into rituals and internal verification systems, which are no longer just one-off reactions, but become a stable part of how the institution works (Power, 1997). AI amplifies this tendency: if the system produces decisions at scale, then the organization also needs *systems at scale* to prove responsibility.

This brings us to the second conceptual block: concrete accountability mechanisms. In the literature on algorithm auditing, the idea that organizations need to build end-to-end internal audit, with documents, checkpoints, and traceability throughout the system’s lifecycle, not just at the end, is a recurring theme (Raji et al., 2020). In parallel, proposals for *algorithmic impact assessments* show how impact assessments can become a governance technology, but only if they are treated as institutions (who is responsible, who validates, who can stop the system), not as simple formulas (Selbst, 2021; Metcalf et al., 2021). This fits perfectly with my hypothesis about the constitutive character of law: when impact assessments, auditing, and documentation become mandatory or standardized, they begin to constrain the very design and use of AI in management.

Finally, I cannot ignore the normative-ethical layer, because in practice, many companies justify their AI governance through ethical language (principles, values, *trustworthy AI*), and the literature shows both convergences and limits. There is a global convergence around some principles (transparency, accountability, fairness, confidentiality), but their implementation remains uneven (Jobin et al., 2019). In addition, there is an important caveat: if ethics is used as a substitute for obligations, it can degenerate into *ethics washing*, i.e. into an appearance without a mechanism (Bietti, 2020). That is why, for me, the most interesting part is the intersection area: how the principles are translated into verifiable procedures and legally compatible controls (Floridi et al., 2018).

3. Research methodology

Given the emerging and complex nature of the intersection between legal frameworks and artificial intelligence systems, this paper adopts a comparative qualitative research strategy. The central methodological objective was not to statistically measure the degree of compliance, but to understand the causal mechanisms and internal processes through which the legal norm ends up modifying the decision-making architecture. The choice of the qualitative method is justified by the need to explore phenomena that usually take place behind the corporate curtain, where tensions between algorithmic efficiency and legal rigor are informally negotiated before being codified in standardized procedures.

The study is based on a purposive sampling logic, targeting global organizations operating in high-risk environments, where regulatory pressure, generated by the upcoming AI Act or GDPR norms, is maximum. Entities from the financial-banking, healthcare and technology sectors were selected for analysis, areas where automated decisions have a direct impact on fundamental rights or economic stability. This sectoral diversity allowed the isolation of compliance patterns that

transcend the specifics of a single industry, thus validating the hypothesis that law acts as a constitutive factor of the organization, regardless of the field of activity.

To avoid the bias inherent in a single perspective and to ensure the robustness of the conclusions, the research used a data triangulation method, combining semi-structured interviews with documentary analysis. The interviews targeted key actors located at different points in the decision-making chain, including software engineers and data scientists, compliance professionals and members of the strategic leadership (CXO). This approach aimed to capture the tensions between the different internal languages: the technical one, oriented towards statistical performance, and the legal one, oriented towards minimizing risks. In addition, a corpus of corporate governance documents, such as sustainability reports and AI codes of ethics, was analyzed to identify discrepancies between public commitments and the operational reality of control mechanisms.

An important aspect of this methodology is the ethical management of data and the inherent limitations of access to strategic information. Given the highly sensitive nature of the intellectual property on algorithms and risk management strategies, access to subjects was subject to strict confidentiality clauses. For this reason, the data were completely anonymized and the identity of the participating organizations is not disclosed, the research focusing on the typology of processes and not on nominal case studies. Although this constraint limits the possibility of exact replication of the study, it allowed obtaining honest information about the real implementation difficulties, information that would otherwise have been censored in a public context. The collected data were subjected to rigorous thematic analysis, aiming to identify recurring patterns in the way organizations *translate* abstract legal requirements into technical parameters and audit procedures.

4. Results and discussions

The analysis of the collected data reveals that the impact of legal frameworks, such as the AI Act or GDPR, on corporate management is not a linear one, limited to the simple fulfillment of bureaucratic requirements. On the contrary, the results indicate a profound transformation in the way in which organizations structure their decision-making processes. It should be noted once again that, given the sensitivity of information regarding risk strategies and intellectual property on algorithms, the organizations included in the study requested strict anonymity. This deliberate opacity indirectly confirms the huge stake that legal compliance has in protecting corporate reputation. Beyond these access limitations, the aggregated data outlines three major trends that describe the new organizational reality.

A first distinct result, observed cross-sectionally in the analyzed organizations, is the emergence of hybrid governance structures and so-called “friction zones”. While legal and technical departments used to operate in separate spheres, the pressure of new regulations has forced an operational merger that is not without tensions. Engineers tend to define risk in terms of statistical error, while compliance officers view it in terms of legal liability. The study shows that, to mediate this conflict of language and vision, companies have created new “translation” roles, such as AI ethicists or algorithm auditors. These actors, while not always holding the final decision-making power, exercise a procedural veto, blocking systems whose documentation is incomplete. Thus, the managerial decision becomes the result of a continuous internal negotiation between technical performance and legal certainty.

Perhaps the most interesting finding of the research concerns the way in which law acts as a silent architect of technical systems, modifying the code before it is put into operation. In high-risk sectors, such as banking or human resources, a clear trend of technological self-limitation has been identified, which we can call the “performance paradox”. Several respondents indicated that their organizations have deliberately chosen to avoid “black box” models, such as deep unstructured neural networks, in favor of simpler models, such as decision trees. The motivation behind this choice is not technical, but strictly governance: simple models are easier to explain to regulators.

This finding validates the hypothesis that the legal framework functions as a selection filter, with management preferring an explainable and defensible decision to the detriment of an optimal but opaque decision.

Last but not least, the analysis of internal documents revealed an intense concern for traceability, which turns auditability into a managerial defense mechanism. In practice, AI systems are surrounded by dense layers of records that document the entire history of a decision, from model training to the final result. This bureaucratization of AI suggests a shift toward a defensive management style, in which human managers use compliance reports to deflect responsibility for errors. When a system produces a controversial result, the organization’s defense is not based on the infallibility of the algorithm, but on demonstrating compliance with procedural steps. Overall, these findings show that organizations are beginning to internalize the law as a standard of quality and stability, even if this entails operational costs and a slowdown in innovation processes.

5. Conclusions

In this research, I started from an intuition that, over time, seemed increasingly difficult to ignore: with the entry of AI into the areas of managerial decision-making, law no longer remains a *fence* placed outside the organization, but begins to be an element of its internal infrastructure. We are not just talking about rules that prohibit or allow, but about rules that oblige the organization to build mechanisms, roles and procedures through which AI becomes governable. The logic of risk-based regulation, specific to the AI Act, pushes companies towards a type of management in which documentation, traceability and control over the entire life cycle of systems are no longer optional, but operating conditions. In this sense, I came to see the legal framework not as a simple constraint, but as a force that shapes the way in which AI-assisted decision-making is designed and used.

What emerged most clearly, conceptually, is that AI-assisted managerial decision-making becomes a decision *under embedded constraints*. Where previously management could assess a model’s performance almost exclusively through efficiency and accuracy, now a second, equally important axis emerges: the demonstrability of compliance and accountability. In practice, this means that organizations end up developing a kind of internal grammar of accountability: accountability algorithms, legal risk management tools, internal audit, cross-functional committees, and escalation procedures when a system produces sensitive results. The literature on accountable algorithms and internal audit helped me understand that this shift is not just *good practice* but a realistic condition for organizational control when AI operates at scale. At the same time, I also recognized the risk that these mechanisms could become rituals, that is, forms without substance, especially in organizations that are tempted to confuse compliance with document production.

On the topic of transparency and explanation, my conclusion is more nuanced than I would have initially expected. In the public sphere, the idea that *the law requires explanation* is sometimes taken for granted, but the literature shows that things are more complicated, especially when automated decisions and opaque models come into play. From a corporate management perspective, explanation is not just a philosophical issue about interpretability, but a governance condition: you need explanations that are good enough to justify the decision, to verify it, and to challenge it internally when necessary. This also changes the way technical solutions are chosen, because sometimes the organization prefers more *controllable* models to the detriment of marginal performance gains.

Perhaps the most important thing I take with me is the idea of a paradigm shift in management: instead of treating the law as an obstacle, organizations start using it as a resource for stabilization and legitimization. In high-risk conditions, compliance is no longer just about avoiding sanctions, but about building trust, internally, between technical teams, compliance, and management, and externally, towards regulators, partners, and the public. In organizational terms, this produces a form of human-AI symbiosis in which roles are redistributed: humans do not disappear from decision-making, but are pushed into tasks of oversight, justification, and

accountability, while AI becomes a tool for inference under rules. From an ethical perspective, I also found it useful to keep the warning that the language of *trusted AI* can be used superficially, without verifiable mechanisms, ethics risks becoming just a reputational layer. In this sense, it is precisely the meeting of ethics and law, when translated into procedures and audits, that seems to me to be the point where organizations show their real maturity.

Of course, this approach has limits that I consider part of the honesty of research. A comparative qualitative research gives you depth, but it does not promise statistical generalization. In addition, access to internal data about real patterns, incidents, audits, or tensions can be filtered by confidentiality and the organizational interest in presenting itself in a good light. That is why we rely on triangulation and comparison between cases as a form of strengthening credibility. For future research, it seems natural to me to continue in two directions: a longitudinal one, which would track how AI governance changes as the AI Act is effectively implemented in practice, and a mixed one, which would combine qualitative depth with quantitative indicators (e.g., types of controls implemented, frequency of incidents, response times and compliance costs).

I conclude with an idea that I consider central: as AI becomes a decision infrastructure, law becomes an organizational infrastructure. And this encounter produces not only limitations, but also a reorganization of trust, a shift from *we rely on the system because it performs* to *we rely on the system because it performs under controlled, verifiable, and assumed conditions*. If there is a practical meaning to this work, it is precisely this: in contemporary corporate management, competence no longer means just implementing AI, but governing it legally and responsibly, as an integrated part of the organization.

Bibliography:

- [1] **Bietti, E.** From ethics washing to ethics bashing: a view on tech ethics from within moral philosophy. In Proceedings of the 2020 conference on fairness, accountability, and transparency (pp. 210-219), 2020. Available at: <https://dl.acm.org/doi/abs/10.1145/3351095.3372860>
- [2] **Doshi-Velez, F., & Kim, B.** Towards a Rigorous Science of Interpretable Machine Learning. arXiv:1702.08608, 2017. Available at: <https://arxiv.org/abs/1702.08608>
- [3] **European Union.** Regulation (EU) 2016/679 (General Data Protection Regulation). Official Journal of the European Union, 2016. Available at: <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>
- [4] **European Union.** Regulation (EU) 2024/1689 (Artificial Intelligence Act). Official Journal of the European Union, 2024. Available at: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- [5] **Faraj, S., Pachidi, S., & Sayegh, K.** Working and organizing in the age of the learning algorithm. Information and Organization, 28(1), 62–70, 2018. Available at: https://www.sciencedirect.com/science/article/pii/S1471772718300277?casa_token=fDFKr pCbgScAAAAA:Py5dddSkglyz0wWpchgKdbfIKMg0CpHcNzMYt-K11RWGc5ns4la1D3gDp7Q6ing4hCkIxN5HIQ
- [6] **Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... & Vayena, E.** AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. Minds and machines, 28(4), 689-707, 2018. Available at: <https://link.springer.com/article/10.1007/S11023-018-9482-5>
- [7] **Jarrahi, M. H.** Artificial intelligence and the future of work: Human–AI symbiosis in organizational decision making. Business Horizons, 61(4), 577–586, 2018. Available at: <https://doi.org/10.1016/j.bushor.2018.03.007>
- [8] **Jobin, A., Ienca, M., & Vayena, E.** The global landscape of AI ethics guidelines. Nature Machine Intelligence, 1, 389–399, 2019. Available at: <https://www.nature.com/articles/s42256-019-0088-2>

- [9] **Kroll, J. A.** Accountable algorithms (Doctoral dissertation, Princeton University), 2015. Available at: <https://www.proquest.com/openview/a29166818f9cf2ffad47c9778da8354d/1?pq-origsite=gscholar&cbl=18750>
- [10] **Metcalf, J., Moss, E., Watkins, E. A., Singh, R., & Elish, M. C.** Algorithmic impact assessments and accountability: The co-construction of impacts. Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT), 2021. Available at: <https://doi.org/10.1145/3442188.3445935>
- [11] **Power, M.** The Audit Society: Rituals of Verification. Oxford University Press, 1997. Available at: [https://books.google.ro/books?hl=ro&lr=&id=q4U3AwAAQBAJ&oi=fnd&pg=PP1&dq=Power,+M.+\(1997\).+The+Audit+Society:+Rituals+of+Verification.+Oxford+University+Press.&ots=DKhPolDUMg&sig=PbHIVtH1aSIGP5jEZkuAQX9876s&redir_esc=y#v=onepage&q=Power%2C%20M.%20\(1997\).%20The%20Audit%20Society%3A%20Rituals%20of%20Verification.%20Oxford%20University%20Press.&f=false](https://books.google.ro/books?hl=ro&lr=&id=q4U3AwAAQBAJ&oi=fnd&pg=PP1&dq=Power,+M.+(1997).+The+Audit+Society:+Rituals+of+Verification.+Oxford+University+Press.&ots=DKhPolDUMg&sig=PbHIVtH1aSIGP5jEZkuAQX9876s&redir_esc=y#v=onepage&q=Power%2C%20M.%20(1997).%20The%20Audit%20Society%3A%20Rituals%20of%20Verification.%20Oxford%20University%20Press.&f=false)
- [12] **Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P.** Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT*), 33–44, 2020. Available at: <https://doi.org/10.1145/3351095.3372873>
- [13] **Selbst, A. D.** An institutional view of algorithmic impact assessments. Harvard Journal of Law & Technology, 35(1), 2021. Available at: <https://jolt.law.harvard.edu/assets/articlePDFs/v35/Selbst-An-Institutional-View-of-Algorithmic-Impact-Assessments.pdf>
- [14] **Veale, M., & Zuiderveen Borgesius, F.** Demystifying the Draft EU Artificial Intelligence Act. Computer Law Review International, 22(4), 97–112, 2021. Available at: <https://www.degruyterbrill.com/document/doi/10.9785/cr-2021-220402/html>
- [15] **Wachter, S., Mittelstadt, B., & Floridi, L.** Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. International Data Privacy Law, 7(2), 76–99, 2017. Available at: <https://doi.org/10.1093/idpl/ix005>

Acknowledgement:

This work was supported by a grant from the Romanian Ministry of Research, Innovation and Digitalization, the project with the title „Economics and Policy Options for Climate Change Risk and Global Environmental Governance” (CF 193/28.11.2022, Funding Contract no. 760078/23.05.2023), within Romania's National Recovery and Resilience Plan (PNRR) - Pillar III, Component C9, Investment I8 (PNRR/2022/C9/MCID/I8) - Development of a program to attract highly specialised human resources from abroad in research, development and innovation activities.