Capability Configurations and Delegation Governance for Agentic AI

Domenico Pochiero, Hans Mulder and Steven De Haes
Antwerp Management School, University of Antwerp, Belgium

domenico.pochiero@student.ams.ac.be
hans.mulder@ams.ac.be
steven.dehaes@uantwerpen.be

Abstract: Agentic AI (AAI) systems are a step beyond decision-support AI, where AI acts as an agent in an autonomous way for multiple steps. There are two related organizational problems: the conditions for organizational readiness for adoption, and the governance of delegated decision authority once autonomous AI agents are deployed. Based on a mixed methods Design Science Research (DSR) study, this paper reports on semi-structured interviews with four senior practitioners in the field of AI from regulated European industries as well as a quantitative survey of 302 practitioners. The findings from the empirical analysis show four organizational capability configurations related to AAI adoption: an organizational governance configuration, a graduated delegation logic, a structural governance capacity pattern, and a financial-experimentation tension. Based on the above findings, the study proposes a prescriptive artifact for readiness assessment based on Agentic AI - Agentic AI Readiness and Integration Framework (AAIRF), which consists of five dimensions: organizational capabilities, technological readiness, human factors, institutional context, and agentic dimensions. It then proposes the Agentic AI Delegation Governance Model (AADGM) to operationalize the governance of the authority shift from AI-assisted to AI-autonomous decision making, by defining task characteristics for each delegation threshold and governance mechanism. The answer to the question of governance that readiness frameworks leave unanswered is how much and on what basis to give autonomy to organizations.

Keywords: Agentic AI, Capability configurations, Delegation governance, Authority threshold, AI governance, Design science research, EU AI Act, Organizational readiness

1. Introduction

The use of Generative AI is becoming more and more agentic: systems that plan, orchestrate, and execute sequences of actions with little or no human involvement (Acharya et al. 2025, Hughes et al. 2025). AI decision support systems, which recommend decisions for human decision makers to execute, are shifting to systems that perform decisions, which initiate and complete decision action in organizational workflows (Dwivedi et al. 2021, Feuerriegel et al. 2024). This is possible using technologies like large language models (LLMs), tool using agents, and multi-agent orchestration frameworks (Baird and Maruping 2021, Li et al. 2024, Wang et al. 2024). This uncertainty is in two aspects. The first is about readiness – does the organization have the leadership, data infrastructure, culture and governance knowledge needed to incorporate autonomous AI systems into their decision making processes? The second is about delegation governance: When an organization is ready enough, how much autonomy should be given to the AI agents and what should be put in place to protect that autonomy? While there are adopted frameworks in the dominant IS literature that focus on the first dimension, the second dimension (Raisch and Krakowski 2021, Berente et al. 2021) is largely overlooked. This gap is exacerbated by the EU Artificial Intelligence Act (EU AI Act), which contains requirements for risk-based compliance specific to the use of autonomous systems in Europe (Veale and Borgesius 2021, Hacker et al. 2023). For organizations that have not taken the time to think through the question of delegation, it is possible that they are not compliant because their governance system is not expressed. Both dimensions are addressed in this paper using a Design Science Research (DSR) paradigm (Hevner et al. 2004, Gregor and Hevner 2013), which includes a qualitative interview approach and a quantitative survey approach. The Main Research Question (MRQ) is how are organizational capability configurations related to readiness to embrace Agentic AI and how should organizational decision authority be delegated as the AI system moves from assistive to autonomous use? The Agentic AI Readiness and Integration Framework (AAIRF) (see Section 4.3) and the AADGM (see Section 4.4) are the two design artifacts presented to address the two questions above: readiness and governance.

2. Literature Review

The Technology-Organization-Environment (TOE) theory (Tornatzky et al. 1990) is still the most commonly used perspective for technology adoption studies, which postulates that technology adoption is a blending of technological attributes, organizational capabilities and environmental influences. Data management infrastructure, technical skills and learning mechanisms in the organization have been identified as enablers in the AI context (Jöhnk 2021). However, the frameworks assume something that restricts their use in AAI: they envision using AI to support a human decision maker, not to replace a set category of decisions. They ask, but do not answer, the question of whether an organization is able to embed AI systems in its operations, but not

how an organization should govern AI systems that start to take action consequentially without the need for human intervention (Shrestha 2019). AI output trust has been shown to predict AI adoption, and it is influenced by perceived transparency, explainability, and reliability (Glikson and Woolley 2020, Rai 2020). But for AAI, the trust issue is turned upside down from confidence in an output to confidence in an autonomous process; can the organization ensure that their AI agents operate safely within acceptable boundaries, know that they will fail if they don't, and who will be held accountable if they do (Hughes et al. 2025, Kasirzadeh and Gabriel 2025)? This is a governance architecture question, and it's only a beginning that existing adoption research has begun to address (Papagiannidis et al. 2025). Governance design is becoming increasingly influenced by regulatory frameworks. The EU AI Act introduces risk-based classifications for AI systems and introduces audit documentation and human oversight requirements for high-risk autonomous applications (Hacker et al. 2023, Veale and Borgesius 2021). New research on governance teaches that regulatory requirements can influence responsible delegation decisions by offering thresholds that depend on compliance (Papagiannidis et al. 2025, Ransbotham et al. 2021). Unlike the previous generations of AI, AAI systems are capable of self-planning, coordinating, and executing the interdependent tasks without continuous human intervention (Kasirzadeh and Gabriel 2025, Islam et al. 2026). Multi-agent architectures break down a complex task into specialized agents that are managed by orchestration mechanisms (Baird and Maruping 2021, Li et al. 2024). These capabilities are further enhanced by LLM-based agents with natural language reasoning, tool usage, and adaptive replanning, allowing them to maintain goal-directed behavior in more complex environments (Bommasani et al. 2021, Wang et al. 2024, Yao et al. 2022, Park et al. 2023). The impact on the organization's governance is significant. (Shrestha 2019) state that it is important that an organization re-organizes its decision making processes in a systematic manner within explicit authority arrangements that define those decisions which can be completely delegated, those that need to be overseen, and those that need to be approved. (Jarrahi 2018) suggests that explicit boundary agreements are not optional but are necessary conditions for successful human-AI collaboration. (Berente et al. 2021) show that scaling up AI governance across organizations will need a governance framework that considers the performance of the AI models as well as delegation, monitoring, and institutional control. (Tarafdar et al. 2023) demonstrates that the interaction between the human agent and AI agent should be formally defined as work roles. Multi-agent systems lacking explicit delegation governance carry risks of systematic accountability failures, as warned by (Chaffer et al. 2024).

This research is motivated by three related gaps. First, although TOE and dynamic capabilities describe organizational readiness for AI adoption, they are insufficient to identify the governance mechanisms needed once the autonomous AI systems are given decision authority. Secondly, the empirical relationship between the organizational capability configurations and the delegation governance decisions is still not well studied in IS research (Shrestha 2019, Berente et al. 2021). Third, the EU concept of a governance calculus is unique to EU-based organizations, given the EU AI Act's risk tiers and requirement for compliance documentation. All these gaps drive an empirical investigation of capability configurations and the design of a governance model that effectively implements the delegation decision. The AAIRF is to bridge the readiness gap, and the AADGM is to bridge the delegation governance gap.

3. Methodology

This study follows the DSR, which sees empirical investigation and artifact design as a single process. The development of the DSR cycle started with problem identification (Section 2), followed by a mixed-method empirical design, and then followed by building the artifact (Section 4.3 and 4.4) and demonstrating it via grounded derivation. The mixed-method approach (Venkatesh et al. 2013) stems from the nature of the research questions, which are ones that can be answered by both qualitative and quantitative methods: how organizations actually govern delegation (qualitative); and how widely observed capability patterns are observed across a sample of a greater size of European organizations (quantitative). Agentic AI was defined as AI systems that plan, co-ordinate, and take actions in multiple steps processes that involve some or conditional human participation, as opposed to advisory AI characterized by the following three properties: the capacity to execute rather than advise on action; the capacity to coordinate across interdependent processes; and delegated decision making authority within boundaries.

3.1 Qualitative Phase

A semi-structured interview protocol was used to interview 4 senior practitioners. Sectors were chosen for their complexity of regulation and operational risks: aviation and insurance have very strict compliance requirements, and autonomous decisions have safety implications and financial consequences; technology and consulting are contexts with a high level of autonomy and cautious delegation of authority, respectively. This intentional

change does not seek confirmatory consistency; rather, it provides a range of possibilities in delegation governance. The participants were selected using criterion-based purposive sampling (Eisenhardt 1989), where all were directly involved in the governance of AI deployment decisions and had at least five years of experience with AI in the relevant field. This study focuses on the analytically relevant characteristic of their governance context (analyzed in Table 1).

Table 1: Delegation governance practices by interviewed organization

ID	Sector	Delegation Threshold Criteria	Human Oversight Level and Governance Mechanism
P1	Aviation	Operational risk level; EASA safety classification; task reversibility	Human-on-the-Loop for scheduling and routing optimization; Human-in-the-Loop for final route approval and pilot decisions. No autonomous execution of safety-critical choices.
P2	Insurance	Claim complexity; coverage ambiguity; financial exposure; automated storm/weather threshold triggers	"AI-first, Human-in-the-Loop" board program. Human-on-the-Loop for standard claims; Human-in-the-Loop for high-value, ambiguous, or fraud-flagged cases. EU AI Act compliance questionnaires gate deployment.
P3	Technology / Cybersecurity	Known vs. unknown behavioral anomaly classification; threat scoring confidence level	Full Autonomy for known-pattern threats (agent shuts down traffic, closes ports, isolates devices). Human-in-the-Loop only for unknown-classification events. 20-30M events/month processed autonomously.
P4	Academic / Consulting	All client-facing decisions: any task with reputational or professional risk remains with human	Human-in-the-Loop for all outputs. AI provides intelligence augmentation only. Autonomous client deliverables explicitly rejected following documented accountability failures in the sector.

It involved governance and accountability frameworks, how decisions are made on delegation, the design of human override and escalation procedures, audit infrastructure, employee responses to AI-initiated actions, and the design for the EU AI Act compliance within governance. The governance themes were identified through open coding, axial coding, and selective coding in NVivo software and are reported in Section 4.1 as five interrelated themes.

3.2 Quantitative Phase

The survey on the impact of AI in the workplace was sent via professional networks, LinkedIn AI and digital transformation groups, and IS research panels in Belgium, the Netherlands, Germany, France, Italy, Spain, and the United Kingdom to 302 AI-related professionals in European companies. For this purpose, seven constructs were measured using multi items Likert scales (1-5) such as Leadership AI Readiness (3 items, (Jöhnk 2021)); Innovation Culture (4 items, (Vial 2021)); Data Readiness (3 items); AI Trust (3 items, (Glikson and Woolley 2020)); Regulatory Pressure (2 items); ROI Orientation (3 items); and Agentic AI Adoption (4 items developed for this study). Table 2 provides descriptive statistics for survey items that were specific to governance and directly informed the design of the AADGM.

Table 2: Governance-specific survey findings (N = 302)

Governance Dimension	Key Finding	Governance Implication for AADGM
Trust in AI outputs in operational processes (Q5.6).	43.4% moderate; 17.8% high; 6.6% none; 32.2% low or very low	Most organizations are not ready for Full Autonomy; monitored autonomy or Human-on-the-Loop is the appropriate threshold for the majority.
Importance of autonomy in AI tools (Q4.2).	35.6% moderate; 20.9% very important; 12.3% not important; 31.2% low importance.	Autonomy is valued but organizations are divided; no single threshold applies universally; per-task threshold setting is essential.
AI agents valuable for risk assessment decisions (Q5.2).	60.9% agree; 21.7% neutral; 14.3% disagree.	Practitioner support exists for AI in risk work; this maps to monitored autonomy depending on consequence severity

Governance Dimension	Key Finding	Governance Implication for AADGM
Regulatory compliance perceived as barrier to AI adoption (Q6.1).	47.8% agree; 27.3% neutral; 20.1% disagree.	Near-majority view compliance as constraining; AADGM Human-in-the-Loop level directly maps to EU AI Act high-risk mandatory oversight requirements
Compliance procedures formally in place (Q6.2)	39.2% yes; 45.3% partial; 15.5% no	Only 39% have complete governance infrastructure; majority cannot yet enforce monitored autonomy or Human-on-the-Loop mechanisms reliably
Speed prioritized over accuracy in AI-driven decisions (Q5.8)	46.5% agree; 28.7% neutral; 21.7% disagree	Speed-first logic may lead organizations to select lower-oversight thresholds than risk profiles justify; financial-experimentation tension applies here

4. Preliminary Results

4.1 Qualitative Findings

Based on data analysis from the four interviews by inductive deductive coding with NVivo software, five interrelated themes of governance emerged. The extent of autonomy is set by leadership. In all four organizations, senior executives offered not just strategic support for the use of AI, but also explicitly stated and enforced the limits of acceptable freeplay. The C-suite management team was actively engaged in an ongoing governance issue in aviation, which involved permissible limits of autonomous decision making (Fountaine et al. 2019). In the insurance industry, a board-driven program has been specifically developed to define the scope for decisions to be made independently, and the scope for decisions to be made on a case-by-case basis. The strongest adoption predictor in the quantitative data is leadership readiness, which works partially via this authority channel. Governance mechanisms will only work if they are supported by the infrastructure. Having reliable data pipelines is necessary for audit trails, having compatible system architectures is necessary for real-time monitoring, and integration of escalation procedures is necessary in many legacy environments. Aviation and insurance practitioners recounted cases in which the technical substrate was lacking, so the governance designs were not feasible. Therefore, data readiness is not only a prerequisite to adopting data, but it's also a governance requirement. While the preconditions for delegation did not prevent the process from going further, override mechanisms and escalation paths consistently made it possible for the process to move forward. Organizations were able to expand the scope of autonomous authority to more categories of tasks when employees were aware that AI-driven actions could be easily corrected. The technological organization had exclusive control over 90-99% of events, since the exception channel for events of unknown classification was clearly designed and enforced. This convergence aligns with the argument of (Makarius et al. 2020) regarding the need for sociotechnical co-design in order to achieve effective integration of human and AI elements and the multilevel governance framing by (Benbya et al. 2020) that autonomous delegation can only be sustained when the individual, organizational, and institutional dimensions are aligned. After a high-profile case in the sector showed accountability gaps where AI was not autonomous, the firm decided to not use autonomous AI at all for client needs. All four of the interviews mentioned the EU AI Act as a governance structuring force. In insurance, pre-deployment questionnaires for compliance are the first step toward the introduction of AI projects. In technology, an EU AI Act scoring tool offers automated compliance testing of models used. As of the writing of this article, GDPR consent restrictions in aviation already restricted the scope of how personalization works for AI recommendations. Reducing governance ambiguity was said to be regulatory clarity, not uncertainty, as the teams had a structured basis for delegating decisions on the risk tier that would otherwise be made inconsistently across the business units. Delegation boundaries were defined by risk and reversibility, not AI ability, in all four organizations. Tasks were suitable for autonomous execution if they were low-risk and reversible tasks that could be easily corrected if there were an error. Tasks that had safety, monetary, or regulatory implications continued to be considered mandatory human tasks. This confirms (Shrestha 2019) authority configuration argument: that delegation governance is more about managing risk than it is an assessment of technology.

> *"We don't deploy these systems in one step. We start with small use cases, test them, and only if everything aligns across data, trust, and internal processes do we scale." (P2, Insurance Expert)*

The most consistent across the four organizations was that delegation was not subject to any organizational policy, but rather assessed for each task based on risk and reversibility. A limited amount of supervision was applied to tasks that had bounded consequences, those that had easily correctable errors, but not to tasks that had cascading effects, regulatory exposure, or irreversibility. In aviation, no matter how mundane the task of scheduling is, a human-on-the-loop checkpoint was necessary. No client-facing task was delegated despite AI's abilities, as the issue of an irreversible impact on the reputation of AI-generated professional deliverables was raised in consulting. This pattern directly addresses (Shrestha 2019) permission structure argument in governance, which states that governance structures must be created and based on explicit permission boundaries and not informal use. We don't drop a system in one go, said the insurance practitioner. We typically scale up from small use cases and only scale up if we've got alignment on data, trust, and internal processes.

4.2 Quantitative Findings

A total of 302 respondents from across Europe were found to have similarly confirmed, through Spearman correlation analysis (Conover 1999), that the capability constructs identified qualitatively as being relevant to governance co-vary systematically with AAI adoption. Leadership readiness was the strongest predictor ($\rho = .327$, $p < .001$), followed by data readiness ($\rho = .227$, $p < .001$), innovation culture ($\rho = .219$, $p < .001$), regulatory pressure ($\rho = .213$, $p < .001$), and AI trust ($\rho = .202$, $p < .001$). Adoption was nominally correlated with company size, which was .218 ($p < .001$), and there was no significant correlation with turnover, suggesting a capacity of governance formalization over company size for distinguishing between those ready for adoption. The findings presented here are part of a larger companion readiness study (Pochiero 2026), and the analysis below is based on the governance implications.

The AADGM design is closely related to four structural patterns. There is a fairly high level of interdependency among the four leadership readiness, data readiness, innovation culture, and regulatory pressure factors (pairwise $\rho = .354$-$.543$, all $p < .001$), indicating that these four factors form a coherent governance system, not independent predictors. Overall consistency of weak yet significant correlations across all constructs leads to a gradual delegation logic, following the qualitative finding that an increasing level of autonomous authority is granted as a result of a cumulative increase of governance confidence over a number of different dimensions, not in a single decision. Size-governance correlation without turnover-governance correlation confirms a governance formalization capacity pattern: large organizations are more developed not because they are simply bigger, but because they have a more developed governance. The only negative inter-construct relationship in the data set is between ROI orientation and innovation culture ($\rho = -.147$, $p = .020$), indicating that there can be a tension between financial justification logic and exploratory governance approaches that early delegation entails (Benner and Tushman 2015).

> *"The real challenge is not funding the technology; it is having the governance structure and organizational architecture to control and scale it responsibly." (P4, Consulting)*

The descriptive results from survey items not included in the capability analysis are reported separately for each governance sector in Table 2 (Section 3.2). Interestingly, 39.2% said they have compliance procedures in place, 47.8% believe regulatory compliance is a hurdle to the adoption of AI, and 43.4% expressed moderate trust in the output of AI in their operations. This suggests that the majority of European organizations are currently in a middle-ground state of governance, recognizing the importance of granting AI autonomy but not ready to go beyond a certain threshold. The AADGM has been created for this intermediate stage.

4.3 Agentic AI Readiness and Integration Framework (AAIRF)

The AAIRF is the main outcome product of this DSR program, grounded in the following of the iterative process of the empirical findings. It groups the five organizational dimensions that have been identified as important to adoption: the two most important are Organizational Capabilities (leadership preparedness and innovation culture), while Technological Readiness (data infrastructure and integration maturity), Human Factors (AI trust and workforce acceptance), and Institutional Context (regulatory pressure and formalization of compliance) are other dimensions. The final is the AAIRF's most unique: as opposed to previous AI readiness models which assumed a technology classification of AI, the AAIRF specifically incorporates agentic properties as cross-cutting dimensions that change how readiness is understood across all other dimensions. The framework is based on the idea that adoption is the result of capability configurations: in an organization with strong leadership but weak data infrastructure, it is difficult to implement governance rules and procedures that require and rely on audit and accountability, which would otherwise enable effective delegation, regardless of the strategy. The complete description and assessment of the AAIRF is provided in a companion study (Pochiero 2026). In the

present paper, readiness of the AAIRF is considered as the organizational groundwork on which delegation governance is built.

4.4 Agentic AI Delegation Governance Model (AADGM)

The AADGM is offered as an extension of the AAIRF which is governance focused. While the AAIRF study is concerned with whether conditions are suitable for the use of Agentic AI, the AADGM study is concerned with decision authority distribution once conditions are suitable for the use of Agentic AI. The AADGM is composed of three layers of integrated AADGM (Figure 1) made up of five qualitative themes and four quantitative patterns. Layer 1, Delegation Assessment, is the qualitative result of the finding that the two key factors that govern the delegation of a task are the level of risk and its potential reversibility, and the translation of these two factors into four operational dimensions: task complexity and routineness; level of risk and reversibility; strategic significance and accountability requirements; data sensitivity and regulatory classification. These dimensions are based on the actual considerations practitioners reported applying.

At Level 2, Authority Threshold, the results of the assessments are aligned to four points of governance. For low-ticket reversible tasks, where errors are easily identified and corrected, full autonomy will be used, and for the 90-99% of known-pattern threats, autonomously operating cybersecurity will be used. The similarity with the insurance claims determination is that monitored autonomy is also about automation, and a “rolling” audit trail, as well as anomaly alerts. In contrast to aviation, the override function is always ON in Human-on-the-Loop and allows no prior approval of each action for scheduling and routing optimization. Human-in-the-Loop allows humans to intervene in the AI's actions and is needed for AI systems with potentially serious consequences for safety and/or the law, as per the EU AI Act when dealing with high-risk AI systems. The four levels are aligned with the EU AI Act's respective levels of human oversight, which are minimal-risk, limited-risk, human-on-the-loop, and human-in-the-loop.

Layer 3, Governance Mechanisms, details the organizational safeguards that will be needed at each threshold: override control and escalation paths, audit trails and accountability mechanisms, trust calibration, human AI interface design, and failure detection and recovery mechanisms. The four mechanisms are set out as ways of implementing Themes 3, 4 and 5 from the qualitative analysis. There's a constant feedback loop from Layer 3 to Layer 1, with threshold recalibrations as the governance confidence of organizations rises. It builds on organizational readiness (represented at the bottom of the diagram) and relies on the AAIRF dimensions: governance mechanisms in Layer 3 are dependent on the presence of organizational readiness with leadership readiness, data maturity, innovation culture, and regulatory awareness. An obvious example of this dependency is the 60.8% of the organizations surveyed that reported they did not consistently have full or partial compliance procedures in place with a fully prepared readiness foundation. AADGM is thus a governance design instrument for organizations that are ready to implement it and a diagnostic for those that are not ready yet and lack readiness to inhibit governance options.

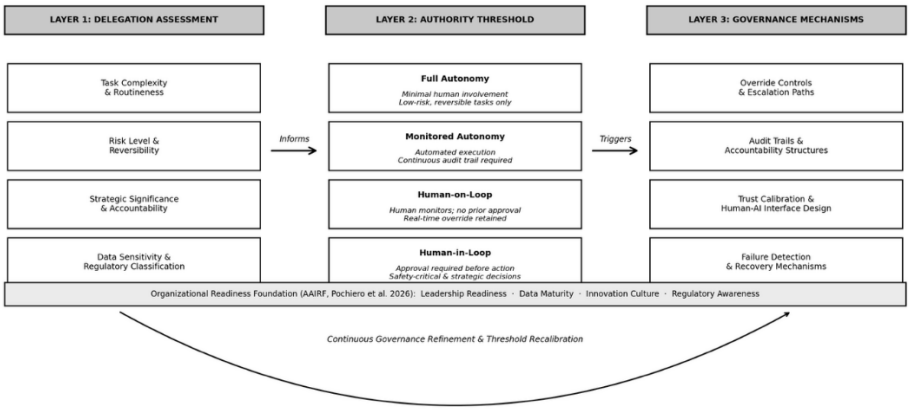

Figure 1: Agentic AI Delegation Governance Model (AADGM). Note: delegation threshold is determined per task context, not as a fixed organizational policy. The AADGM operationalizes the governance question that arises once AAIRF readiness assessment is complete

5. Conclusion

In this paper, both the readiness and governance aspects of Agentic AI adoption in European organizations have been discussed. Empirically, it shows that there are four different types of organizational capabilities: one

governance formalization capacity pattern indicating that organizational maturity, as measured by the presence of structures in administrative processes, is the key factor instead of scale, and a financial-experimentation tension indicating that the logic applied to govern ROI processes may structurally influence the exploratory approaches of delegating processes that are necessary to adopt early. It introduces a governance characteristic to the governance construct separate from the question of readiness to adopt, the AAIRF being about whether capability configurations allow adoption, the AADGM being about how much autonomy to grant, and under which organizational safety nets. It adds a governance characteristic to the governance construct, separate from the question of readiness to adopt (where the AAIRF asks the question whether a capability configuration is ready for adoption, the AADGM asks how much autonomy to grant and under what organizational safety nets). This distinction is based on (Shrestha 2019) authority configuration argument, (Berente et al. 2021) AI management framework, and (Jarrahi 2018) human-AI symbiosis model, building on their work by adding a design instrument tailored for the EU AI Act compliance context. In practice, both artifacts provide an organization with a linear governance process: Determine readiness using the AAIRF, then plan delegation governance using the AADGM, whose threshold levels are directly aligned to the EU AI Act risk classification, thereby providing a useful compliance planning tool as well as a governance tool. In practice, the discovery that override mechanisms are not a barrier to delegation, but rather a facilitator, has practical ramifications for the investment in AI governance: in a world where an increasing number of organizations are investing in building governance structures for human control (even if not all), they are paradoxically investing in more delegation than those that try to go fully autonomous without governance scaffolding. This aligns with (Steyvers and Kumar 2024) three challenges for AI assisted decision making and the need to have calibrated trust and a capability to be overwritten by humans in order to be able to deploy AI effectively at scale. This paper is a component of a doctoral research program which will explore how regulatory environments and organizational culture shape the governance dynamics identified in this paper, focusing on cross-sector variance between aviation, financial services and technology sectors.

5.1 Limitations and Future Research

The qualitative sample (n = 4) provides governance-focused depth through purposive selection but cannot support sector-level generalization. The quantitative data are cross-sectional; co-variation patterns reflect organizational states at a point in time rather than governance dynamics over time. AAIRF and AADGM are first-cycle DSR artifacts requiring systematic expert validation through structured workshops, Delphi-style refinement of threshold specifications, and controlled deployment in partner organizations to assess practical utility. Sector-level comparative analysis, particularly contrasting aviation and financial services with technology and consulting, will test whether governance patterns differ significantly across regulatory environments. The financial-experimentation tension also opens a research agenda on ambidexterity in AI governance investment: if ROI-oriented governance logic is structurally incompatible with the exploratory delegation approaches early adoption requires, organizations may need to develop separate evaluation frameworks for governance investment that do not apply the same metrics as operational AI spending (Benner and Tushman 2015).

Ethics Declaration: Ethical clearance for this research was obtained from the University of Antwerp. All interview participants provided informed consent prior to participation, and all data were anonymized to protect respondent and organizational identity.

AI Declaration: This paper was prepared using an AI language model (Gemini) for condensation to meet the conference's length and formatting requirements. Any analysis, interpretation of the findings, and theory contributions are the authors' own. All AI generated text has been reviewed and checked for accuracy by the authors before being submitted.

References

Acharya, D. B., Kuppan, K. & Divya, B. 2025. Agentic AI: Autonomous intelligence for complex goals—A comprehensive survey. *IEEe Access,* 13**,** 18912-18936.

Baird, A. & Maruping, L. M. 2021. The next generation of research on IS use: A theoretical framework of delegation to and from agentic IS artifacts. *MIS quarterly,* 45**,** 315-341.

Benbya, H., Davenport, T. H. & Pachidi, S. 2020. Artificial Intelligence in Organizations: Current State and Future Opportunities. *SSRN Electronic Journal*.

Benner, M. J. & Tushman, M. L. 2015. Exploitation, exploration, and process management: The productivity dilemma revisited. *Academy of management review,* 40**,** 497-514.

Berente, N., Gu, B., Recker, J. & Santhanam, R. 2021. Managing artificial intelligence. *MIS Quarterly,* 45**,** 1433-1450.

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A. & Brunskill, E. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.

Chaffer, T. J., Goins II, C. v., Okusanya, B., Cotlage, D. & Goldston, J. 2024. Decentralized governance of autonomous AI agents. *arXiv preprint arXiv:2412.17114*.

Conover, W. J. 1999. *Practical nonparametric statistics*, john wiley & sons.

Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J. & Eirug, A. 2021. Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management,* 57**,** 101994.

Eisenhardt, K. M. 1989. Building theories from case study research. *Academy of management review,* 14**,** 532-550.

Feuerriegel, S., Hartmann, J., Janiesch, C. & Zschech, P. 2024. Generative AI. *Business & Information Systems Engineering,* 66**,** 111-126.

Fountaine, T., McCarthy, B. & Saleh, T. 2019. Building the AI-powered organization. *Harvard business review,* 97**,** 62-73.

Glikson, E. & Woolley, A. W. 2020. Human trust in artificial intelligence: Review of empirical research. *Academy of management annals,* 14**,** 627-660.

Gregor, S. & Hevner, A. R. 2013. Positioning and presenting design science research for maximum impact. *MIS quarterly,* 37**,** 337-355.

Hacker, P., Engel, A. & Mauer, M. 2023. Regulating ChatGPT and other large generative AI models. *Proceedings of the 2023 ACM conference on fairness, accountability, and transparency***,** 1112-1123.

Hevner, A. R., March, S. T., Park, J. & Ram, S. 2004. Design science in information systems research. *MIS quarterly***,** 75-105.

Hughes, L., Dwivedi, Y. K., Malik, T., Shawosh, M., Albashrawi, M. A., Jeon, I., Dutot, V., Appanderanda, M., Crick, T. & De', R. 2025. AI agents and agentic systems: A multi-expert analysis. *Journal of Computer Information Systems,* 65**,** 489-517.

Islam, M. A., Somu, S. & Aldaihani, F. M. F. 2026. The rise of agentic AI: Synthesis of current knowledge and future research agenda. *Global Business and Organizational Excellence,* 45**,** 402-416.

Jarrahi, M. H. 2018. Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons,* 61**,** 577-586.

Jöhnk, J., Weißert, M. and Wyrtki, K. 2021. Ready or not, AI comes: an interview study of organizational AI readiness factors. *Business and Information Systems Engineering,* 63**,** 5-20.

Kasirzadeh, A. & Gabriel, I. 2025. Characterizing AI agents for alignment and governance. *arXiv preprint arXiv:2504.21848*.

Li, X., Wang, S., Zeng, S., Wu, Y. & Yang, Y. 2024. A survey on LLM-based multi-agent systems: workflow, infrastructure, and challenges. *Vicinagearth,* 1**,** 9.

Makarius, E. E., Mukherjee, D., Fox, J. D. & Fox, A. K. 2020. Rising with the machines: A sociotechnical framework for bringing artificial intelligence into the organization. *Journal of Business Research,* 120**,** 262-273.

Papagiannidis, E., Mikalef, P. & Conboy, K. 2025. Responsible artificial intelligence governance: A review and research framework. *The Journal of Strategic Information Systems,* 34**,** 101885.

Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P. & Bernstein, M. S. 2023. Generative agents: Interactive simulacra of human behavior. *Proceedings of the 36th annual acm symposium on user interface software and technology***,** 1-22.

Pochiero, D., Mulder, H. and De Haes, S. A mixed-method study of Agentic AI adoption capabilities and decision transformation in European organizations. 18th International Conference on Knowledge Management and Information Systems (KMIS 2026), 2026 Porto, Portugal. SCITEPRESS.

Rai, A. 2020. Explainable AI: From black box to glass box. *Journal of the academy of marketing science,* 48**,** 137-141.

Raisch, S. & Krakowski, S. 2021. Artificial intelligence and management: The automation–augmentation paradox. *Academy of management review,* 46**,** 192-210.

Ransbotham, S., Candelon, F., Kiron, D., LaFountain, B. & Khodabandeh, S. 2021. The cultural benefits of artificial intelligence in the enterprise. *MIT Sloan Management Review and Boston Consulting Group,* 1.

Shrestha, Y. R., Ben-Menahem, S.M. and Von Krogh, G, 2019. Organizational Decision-Making Structures in the Age of Artificial Intelligence. *California Management Review*.

Steyvers, M. & Kumar, A. 2024. Three challenges for AI-assisted decision-making. *Perspectives on Psychological Science,* 19**,** 722-734.

Tarafdar, M., Page, X. & Marabelli, M. 2023. Algorithms as co-workers: Human algorithm role interactions in algorithmic work. *Information Systems Journal,* 33**,** 232-267.

Tornatzky, L. G., Fleischer, M. & Chakrabarti, A. K. 1990. *The processes of technological innovation*, Lexington Books.

Veale, M. & Borgesius, F. Z. 2021. Demystifying the Draft EU Artificial Intelligence Act — Analysing the good, the bad, and the unclear elements of the proposed approach. *Computer Law Review International,* 22**,** 97-112.

Venkatesh, V., Brown, S. A. & Bala, H. 2013. Bridging the qualitative–quantitative divide: Guidelines for conducting mixed methods research in information systems. *MIS quarterly,* 37**,** 21-54.

Vial, G. 2021. Understanding digital transformation: A review and a research agenda. *Managing digital transformation***,** 13-66.

Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X. & Lin, Y. 2024. A survey on large language model based autonomous agents. *Frontiers of Computer Science,* 18**,** 186345.

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K. R. & Cao, Y. 2022. React: Synergizing reasoning and acting in language models. *The eleventh international conference on learning representations*.