

# The AAMM in Practice: Classifying AI Decision-Making Authority Across Enterprise Deployments

*Paper 2 in a Series on AI Governance and Organizational Authority*

---

Alexander Huseby  
Founder, Cognitive Liberty Institute  
Oslo, Norway – alexander@coglib.no  
2026

---

## Abstract

*This paper applies the AI Authority Maturity Model (AAMM), introduced in Huseby (2025), to four documented enterprise AI deployments across financial services, logistics, and consumer goods: BlackRock's Aladdin Copilot, JPMorgan Chase's LLM Suite and agentic systems, Maersk's NavAssist vessel routing platform and Pactum supplier negotiation system, and Unilever's Project Sky agentic supply chain. Through structured case analysis, the paper demonstrates the AAMM's utility as an analytical instrument and identifies a consistent cross-industry pattern: enterprise AI deployments are advancing from Level 1 toward Level 2–3 faster than governance frameworks are keeping pace. The paper introduces the concept of the "governance premium" – the measurable performance advantage associated with AI-savvy board oversight, quantified by MIT CISR (2025) at 10.9 percentage points in return on equity – and operationalizes the central argument of the first paper in this series: that AI governance capability is a form of strategic infrastructure that generates measurable competitive advantage. The paper also proposes two extensions to the AAMM: function-specific rather than organization-wide classification, and a governance maturity dimension that measures oversight architecture against deployment level.*

Keywords: AI Authority Maturity Model, AAMM, enterprise AI deployment, AI governance, agentic AI, bounded autonomy, governance premium, supply chain AI, financial services AI, corporate governance

---

## 1. Introduction

The first paper in this series (Huseby, 2025) introduced the AI Authority Maturity Model (AAMM) as a conceptual framework for classifying the degree of decision-making authority held by AI systems in organizational settings. It argued that the trajectory toward expanded AI authority is already underway, that this trajectory is driven by competitive economic logic rather than ideology,

and that the oversight frameworks required to make that expansion accountable are lagging significantly behind the pace of deployment. The central normative claim was that governance capability constitutes a form of strategic infrastructure.

That paper was primarily theoretical. The AAMM was introduced as a taxonomy with illustrative examples but was not systematically applied to documented cases. This paper does that work. It applies the AAMM to four well-documented enterprise AI deployments, demonstrates the model’s analytical utility, identifies patterns in how organizations navigate the transition from advisory AI to autonomous AI, and introduces empirical evidence that quantifies the financial value of the governance premium – the performance advantage associated with board-level AI oversight.

The cases selected were chosen because they represent the current leading edge of enterprise AI deployment, span multiple industries, have been publicly documented by the organizations themselves and in credible secondary sources, and collectively span AAMM Levels 0 through 2–3, allowing meaningful cross-case comparison. This paper does not claim to be a comprehensive survey of enterprise AI deployments. The cases are illustrative, not representative, and this limitation is addressed explicitly in Section 8.

## 2. The AAMM Revisited

For readers encountering this paper independently of the first in the series, a brief recapitulation of the AAMM is provided. The model classifies AI systems along six levels of organizational authority:

Level	Role	Human Oversight	Decision Authority
0	Analytics tool	Full – human acts on output	None – informational only
1	AI advisor / copilot	Human retains all decisions	Recommendatory
2	Autonomous executor (bounded domain)	Human approval gates	Executes within defined parameters
3	Strategic recommender	Human ratification required	Shapes decision space
4	Governance participant	Shared or disputed authority	Formal organizational role
5	Principal executive	Exceptional / minimal oversight	Hypothetical – not yet realized

*Table 1. The AI Authority Maturity Model (AAMM). Reproduced from Huseby (2025) with minor clarifications to the Decision Authority column based on case analysis conducted in this paper.*

Two clarifications to the original AAMM are introduced based on the case analysis that follows. First, organizations frequently operate AI systems at multiple AAMM levels simultaneously across different functions. The AAMM

should be applied at the level of specific deployments, not organizations as a whole. Second, the boundaries between levels are not always clean. Several cases straddle two levels, reflecting the transition dynamics characteristic of this moment in enterprise AI adoption. Where this occurs, the paper classifies the deployment at its effective level of autonomous authority rather than its nominal or aspirational level.

### 3. Methodology

This paper employs structured case analysis as its primary method. Each case was selected based on three criteria: (1) the AI deployment is documented by the organization itself or in credible secondary sources; (2) sufficient detail is available to classify the deployment against AAMM criteria; and (3) the deployment is active as of late 2025 to early 2026 rather than merely announced or piloted.

Each case is analyzed against five dimensions: the technical architecture of the AI system, the decision domain in which it operates, the degree of human oversight and approval required, the consequences of AI-generated decisions, and the accountability structure in place. These dimensions map directly onto the AAMM criteria and allow systematic comparison across cases.

Sources include official company press releases and annual reports, named executive statements at investor days and in attributed media interviews, and practitioner literature on each deployment. Where sources conflict, the most conservative characterization of AI authority is adopted. This paper does not claim insider knowledge of any organization's internal systems; all classifications are based on publicly available information and are therefore subject to the limitations of corporate self-reporting, which are addressed in Section 8.

A note on inter-rater reliability: the AAMM classifications presented here represent a single researcher's judgment applied to a conceptual framework that has not yet been independently validated. Future empirical work should test whether multiple researchers, applying the AAMM independently to the same set of deployments, reach consistent classifications. This is the primary methodological requirement for moving the AAMM from conceptual framework to validated instrument, and is identified as a priority for the next stage of this research programme.

### 4. Case Analysis

#### 4.1 BlackRock Aladdin Copilot – AAMM Level 1 (Advisor)

BlackRock's Aladdin platform manages risk analytics and portfolio operations for more than \$11 trillion in assets under management. The platform's Copilot component, a generative AI layer deployed across 100+ front-end Aladdin applications, represents the most extensively documented AI advisory system operating in financial asset management.

Aladdin Copilot is built on a supervised multi-agent architecture using LangChain and LangGraph orchestration with GPT-4 function calling, coordinating specialized LLMs to handle complex queries across hundreds of domain-specific Aladdin APIs – an architecture disclosed by BlackRock's own AI engineering leads in a public presentation (Rosales & Valdez, 2025). In October 2025, BlackRock launched an "Auto Commentary" feature with Morgan Stanley as the first client, generating personalized portfolio narratives by integrating the Aladdin platform's risk analytics, the firm's CIO market outlook, and individual client portfolio data (BlackRock, 2025).

AAMM Classification: Level 1. Aladdin Copilot operates as an advisor. BlackRock has explicitly designed the system with hard governance boundaries: Copilot will not give investment advice outside Aladdin's defined parameters, employs output guardrail nodes to detect hallucinations, and maintains strict data-privacy and risk controls. The system surfaces information and recommendations; all investment decisions remain with human portfolio managers or financial advisors. Human decision-making authority is not delegated to the system.

However, the stated trajectory is toward Level 2. BlackRock and Microsoft have publicly described agentic evolution of the platform, with autonomous agents designed to "execute complex tasks" and "continuously learn from interactions" (Microsoft, 2024). Notably, BlackRock is building its governance architecture – guardrail nodes, content filters, boundary definitions – at Level 1, before the transition to autonomous execution. This sequencing represents a meaningful institutional oversight practice that distinguishes it from the other cases examined here.

#### 4.2 JPMorgan Chase LLM Suite and Agentic Systems – AAMM Levels 1-2

JPMorgan Chase is the most extensively documented example of enterprise-scale AI deployment in financial services. The bank's LLM Suite – developed entirely in-house and deployed to more than 230,000 employees globally – provides secure, scalable access to large language models across the full range of business functions: software developers use it for code review, investment bankers for presentation preparation, legal teams for contract analysis, and operations staff for workflow automation (The Digital Banker, 2026). The platform reached 200,000 users within eight months of launch and has undergone eight major upgrades.

LLM Suite itself operates at AAMM Level 1 – it is an AI capability layer through which human employees access AI tools, with humans retaining decision authority in all cases. The bank's broader AI strategy, however, is explicitly advancing toward Level 2. At JPMorgan's annual investor day, Marianne Lake, CEO of Consumer and Community Banking, stated that AI would enable the bank to reduce headcount by at least 10% in its operations and account services departments, adding that "I would take the over on this projection and bet that we will deliver more" (Entrepreneur, 2026). This projection reflects the deployment of agentic systems that handle functions previously requiring human staff in fraud detection, account management, and trade settlement.

By early 2026, the projection was already materializing: JPMorgan reported that operations staff had declined by 4% and support roles by 2%, while client-facing positions had grown by 4%, with the bank attributing part of the shift to technology-driven efficiency gains (HR Grapevine, 2026). Chief Analytics Officer Derek Waldron has publicly described the bank's ambition as building the world's first "fully AI-connected enterprise," with autonomous agents taking on "more and more responsibilities" as capabilities expand (AI News, 2025).

AAMM Classification: Level 1-2 (transitional). LLM Suite as a platform is Level 1. The bank's operational agentic systems — those executing consequential actions in fraud detection, account services, and trade settlement without per-decision human approval — are Level 2. JPMorgan is explicitly and deliberately transitioning across the Level 1-2 boundary, with named executive statements and measurable workforce data confirming that transition is underway. The governance architecture for this expansion — ensuring reliability, security, and transparency as agentic systems assume greater operational authority — is an acknowledged challenge that the bank has not yet fully resolved.

#### 4.3 Maersk NavAssist and Pactum Supplier Negotiation — AAMM Level 2

A.P. Møller-Maersk presents two distinct AI deployments that together illustrate the range of Level 2 autonomy in practice: one in operational logistics optimization, the other in commercial supplier negotiation.

The NavAssist AI-powered vessel routing platform, developed in partnership with Microsoft Azure AI, uses real-time oceanographic data, weather forecasting, and historical fuel performance to autonomously optimize routes for vessels across the Maersk fleet. The system was deployed initially on 130 container ships and is expected to expand to the full fleet. During pilot testing, vessels using NavAssist reported up to 12% reduction in fuel consumption and a 16% improvement in estimated time of arrival accuracy, according to Anders Boedker, Head of Maersk Maritime Innovation (EAN Network, 2025). Routing decisions are made and adjusted continuously by the system; human captains retain override authority but are not required to approve each routing adjustment.

The Pactum AI supplier negotiation system, in use by Maersk since 2021 for spot trucking services, deploys autonomous negotiation bots that negotiate with suppliers at scale to secure pricing and terms for routine procurement without human involvement in individual transactions. Human procurement staff set the parameters within which the system operates and review aggregate outcomes, but individual negotiation decisions are made autonomously.

AAMM Classification: Level 2 for both deployments. Both systems operate autonomously within clearly defined domains and execute consequential decisions without per-decision human approval. The Pactum case is particularly instructive: commercial negotiation has traditionally required human judgment and relationship management. Maersk has determined that for routine, parameter-bound procurement, AI autonomy is preferable to human involvement, and has maintained this deployment for more than four years. The governance structure for both deployments is domain-bounded: AI authority is

exercised within defined parameters, with human oversight concentrated at the level of parameter-setting and exception handling rather than individual decisions.

#### 4.4 Unilever Project Sky and Agentic Supply Chain – AAMM Level 2-3

Unilever's transformation of its global supply chain represents the most advanced agentic enterprise deployment examined in this paper. Working primarily with Aera Technology, Unilever has deployed agentic AI across demand planning, materials procurement, logistics coordination, and factory operations through initiatives including Project Sky and the Digital Materials Planner.

The distinguishing characteristic of Unilever's deployment is cross-functional scope. Where Maersk's agentic systems operate in single-function domains, Unilever's systems coordinate across multiple functions simultaneously. Aera Technology's documentation indicates that when Unilever's supply chain encounters a disruption, the agentic system plans and executes a response across planning, procurement, and logistics without requiring human initiation of each component action (Aera Technology, 2025). In February 2026, Unilever announced a five-year strategic partnership with Google Cloud explicitly designed to create "agentic workflows" in which AI agents "autonomously orchestrate campaigns in real time while keeping employees in the loop for oversight" (AI Magazine, 2026).

AAMM Classification: Level 2-3 (transitional, with Level 3 interpretive). The supply chain AI operates autonomously across multiple functions, which is consistent with Level 2. The classification edges toward Level 3 – strategic recommender – based on the cross-functional coordination that shapes the decision options available to human managers rather than simply executing within a pre-defined space. It is important to note that the Level 3 component of this classification is interpretive: Unilever has not publicly described its system as shaping strategic options in those terms, and the evidence for this characterization comes from the system's documented architecture rather than from company statements. A more conservative classification would place the deployment at Level 2 with emerging Level 3 characteristics. As the Google Cloud integration matures, the system's effective authority is likely to advance further regardless of classification.

## 5. Cross-Case Analysis: Patterns and Observations

### 5.1 The Governance-Deployment Gap Is Consistent Across Industries

Each case demonstrates a version of the governance-deployment gap identified in the first paper: AI decision-making authority is expanding faster than the accountability structures designed to govern it. In each case, the technical capability to expand AI autonomy is ahead of the organizational oversight architecture required to make that autonomy accountable. BlackRock is building institutional oversight infrastructure for a Level 2 transition it has not yet made. JPMorgan is making the Level 1-2 transition with acknowledged challenges it has not fully resolved. Maersk operates Level 2 systems with domain-bounded governance that has not been tested against genuinely novel disruptions.

Unilever is advancing toward Level 3 with oversight mechanisms that are not yet fully specified in the public record.

This pattern is not specific to any industry or organizational type. It reflects the structural dynamic identified in the first paper: the competitive pressure to expand AI authority is immediate and tangible, while the governance costs are deferred and diffuse. Individual actors rationally prioritize deployment over governance, producing a collective outcome – a systematic accountability deficit – that is suboptimal for the system as a whole.

## 5.2 Level Transitions Are Incremental but Cumulative

None of the organizations examined here moved directly from Level 0 to Level 2 or 3. Each deployment represents an incremental extension of AI authority from a prior baseline: BlackRock's Copilot built on years of Aladdin analytics infrastructure; Maersk's NavAssist built on earlier voyage optimization and predictive maintenance pilots; JPMorgan's agentic systems built on LLM Suite adoption by 230,000 employees; Unilever's Project Sky built on prior supply chain digitization.

This incrementalism is individually rational and collectively significant. Each step builds organizational familiarity with AI decision-making, creates dependencies that make reversal costly, and normalizes a degree of AI authority that serves as the baseline for the next step. Accountability interventions at earlier AAMM levels have substantially more leverage than interventions at later levels, because the organizational dependencies and cultural normalizations that make later governance difficult have not yet accumulated.

## 5.3 Governance Architecture Varies by Level

The oversight architectures in place across the cases are broadly appropriate to their AAMM level but reveal characteristic gaps at each transition point. Level 1 deployments rely primarily on output guardrails, content filters, and usage policies – necessary but insufficient for Level 2, where they address output quality but not accountability for AI-initiated action.

Level 2 deployments rely primarily on domain bounding and parameter-setting. This is more appropriate to autonomous execution, but creates a gap at the boundary of the defined domain: what happens when the AI system encounters a situation outside its defined parameters? In each case, the answer is human escalation – but the escalation protocols are not systematically documented in the public record for any of the deployments examined here.

This observation suggests a refinement to the governance framework proposed in the first paper: accountability mechanisms should be specified not only by AAMM level but by transition point. The critical requirement is not the steady-state operation of an AI system at a given level, but the boundary conditions – what happens when the system reaches the edge of its defined domain, and who bears accountability for the consequences.

## 5.4 Advanced Deployments Concentrate in Operations, Not Strategy

Across all cases, the highest AAMM levels are reached in operational and logistical functions – supply chain execution, vessel routing, procurement negotiation – rather than in strategic decision-making. No case examined here involves AI systems at Level 3 or above in strategic planning, capital allocation, M&A, or organizational design.

This pattern is consistent with the theoretical limits identified in the first paper. Strategic decisions involve value conflicts and genuine uncertainty that current AI systems cannot navigate reliably. Operational decisions within defined parameters are more tractable for AI autonomy precisely because parameters can be specified in advance. The trajectory toward Levels 3 and above will require AI systems to operate in domains where parameter specification is not possible – and the cases examined here do not yet test that boundary.

## 6. The Governance Premium: Empirical Evidence

The first paper in this series argued that treating AI oversight as a competitive capability rather than a compliance cost would differentiate organizational performance. This section presents empirical evidence that this argument is not merely theoretical.

Research published by the MIT Center for Information Systems Research (Weill, Woerner, & Banner, 2025) examined the relationship between board-level AI expertise and organizational financial performance across large US organizations. Updating their 2019 digital savviness research, the team found that only 26% of large company boards qualify as "digitally and AI savvy" under criteria that include experience with generative AI, AI agents, and robotics. Organizations with AI-savvy boards outperformed their industry peers by 10.9 percentage points in return on equity; organizations with non-AI-savvy boards underperformed their industry average by 3.8 percentage points. The same research found that AI-savvy board companies averaged \$15.5 billion higher market capitalization than their industry average, while non-savvy companies averaged \$5.4 billion below (MIT Sloan Management Review, 2025).

This finding – a total performance spread of 14.7 percentage points in ROE between AI-savvy and non-AI-savvy boards – quantifies what this paper terms the *governance premium*: the measurable financial advantage associated with board-level competence in AI oversight. The mechanism is consistent with the theoretical argument: boards with AI expertise make better strategic decisions about AI deployment, reduce costly oversight failures, and build the institutional trust that enables responsible expansion of AI authority.

It is important to note that this finding is correlational rather than causal. Organizations that perform better may be more likely to invest in board AI expertise, rather than the reverse. Establishing causal direction would require longitudinal research design that the available data does not support. The governance premium, as presented here, is a performance association that is consistent with the theoretical argument, not a proof of it.

The McKinsey (2025) governance statistics provide essential context: fewer than 25% of companies have board-approved AI policies; 66% of directors report

limited or no knowledge of AI; and nearly one in three say AI does not appear on their board agendas. These figures represent the baseline from which the governance premium must be understood: most organizations are not only failing to capture the governance premium, they are accumulating accountability deficits as they deploy AI at AAMM Levels 1 and 2 without the oversight architecture to manage those deployments responsibly.

## 7. Implications for the AAMM

The case analysis and the governance premium evidence together suggest several refinements and extensions to the AAMM as originally presented.

AAMM classifications are function-specific, not organization-wide. The case analysis consistently shows that large organizations operate AI systems at multiple AAMM levels simultaneously. JPMorgan's LLM Suite is Level 1; its operational agentic systems are Level 2. Unilever's supply chain AI approaches Level 3 while its marketing tools remain at Level 1. Future applications of the AAMM should specify the function and domain being classified, not the organization as a whole.

Transition governance is as important as steady-state governance. The accountability mechanisms most relevant to each AAMM level differ from those most relevant at the transition between levels. Transition from Level 1 to Level 2 requires explicit decisions about which decisions will be delegated to AI autonomy, what the accountability structure for those decisions is, and what the escalation protocol is when the AI system encounters boundary conditions. These transition requirements are currently absent from both the AAMM framework and the organizations examined here.

A governance maturity dimension should be added to the AAMM. The current model classifies AI systems by their level of decision-making authority. A more complete framework would add a second dimension: the maturity of the oversight architecture in place for that level of authority. An organization operating a Level 2 deployment with a Level 2 oversight architecture is in a fundamentally different risk position from one operating a Level 2 deployment with only Level 1 oversight architecture. The gap between deployment level and oversight maturity is a more precise measure of organizational risk than either dimension alone. Developing and validating this two-dimensional version of the AAMM is identified as the primary objective of the next paper in this series.

The governance premium appears level-sensitive. The MIT CISR finding measures board AI expertise against overall organizational performance. A more granular hypothesis — which future research should test — is that the governance premium is largest for organizations operating at AAMM Levels 2 and 3, where the consequences of oversight failure are largest and the complexity of accountability requirements is greatest.

## 8. Limitations

Case selection is not representative. The four cases were selected for documentation quality and AAMM relevance, not as a random or stratified sample. They represent large, well-resourced, technologically sophisticated organizations. The AAMM dynamics observed here may not generalize to smaller organizations, less-resourced sectors, or different regulatory contexts.

All information is from public sources. This paper has no insider access to any organization's internal systems, governance documents, or decision protocols. AAMM classifications are based on what organizations have chosen to disclose publicly, which may systematically underrepresent governance failures and overrepresent governance strengths.

The AAMM lacks inter-rater validation. The framework introduced in Huseby (2025) and applied here remains a conceptual instrument. The classifications presented are a single researcher's judgment. Future research should test inter-rater reliability by recruiting multiple evaluators to independently classify a larger set of deployments, measuring agreement, and refining classification criteria where divergence occurs. This is the primary methodological step required to move the AAMM toward a validated instrument.

The governance premium finding is correlational. As noted in Section 6, the MIT CISR finding establishes association rather than causation. The direction of the relationship between board AI expertise and organizational performance cannot be established from the available cross-sectional data.

Deployments are evolving rapidly. Enterprise AI deployments are advancing quickly. Classifications accurate as of late 2025 to early 2026 may be outdated within 12–24 months.

## 9. Conclusion

This paper has applied the AI Authority Maturity Model to four documented enterprise AI deployments, demonstrating the framework's utility as an analytical instrument and identifying consistent patterns in how organizations navigate the expansion of AI decision-making authority.

The cases confirm the central argument of the first paper: the trajectory toward expanded AI authority is economically rational, incrementally normalized, and consistently ahead of the accountability architecture required to make it responsible. BlackRock, JPMorgan, Maersk, and Unilever — among the most sophisticated and well-resourced organizations in their respective industries — all exhibit versions of the governance-deployment gap. If the gap is this consistent among leading organizations, it is likely to be more pronounced among organizations with fewer resources and less institutional sophistication.

The governance premium evidence — a 14.7 percentage point spread in ROE between AI-savvy and non-AI-savvy boards — operationalizes the normative argument that oversight capability is strategic infrastructure. The financial case for investing in AI governance capacity is now empirically supported, not merely theoretically argued, even acknowledging the correlational rather than causal nature of the evidence.

The two extensions proposed to the AAMM – function-specific classification and a governance maturity dimension – are offered as inputs to future empirical research. The third paper in this series will examine the technical architecture of AI oversight mechanisms: how the design of AI systems themselves can embed or undermine the accountability properties required at each AAMM level, and what software engineering practices are most effective at maintaining human oversight as AI authority expands.

---

## References

- Aera Technology. (2025). *How Unilever is envisioning the autonomous supply chain with agentic AI*. AeraHUB 2025 London. <https://meet.aeratechnology.com/how-unilever-is-envisioning-the-autonomous-supply-chain-with-agentic-ai>
- AI Magazine. (2026, February 22). Why Unilever is placing a five-year AI bet on Google Cloud. <https://aimagazine.com/news/unilever-signs-five-year-google-cloud-deal-for-ai-backbone>
- AI News. (2025, December 16). JPMorgan Chase AI strategy: US\$18B bet paying off. <https://www.artificialintelligence-news.com/news/jpmorgan-chase-ai-strategy-2025/>
- BlackRock. (2025, October 2). Aladdin Wealth launches AI-enabled commentary tool at Morgan Stanley. <https://www.blackrock.com/aladdin/discover/aladdin-wealth-launches-ai-enabled-commentary-tool-at-morgan-stanley>
- Rosales, B., & Valdez, P. V. (2025). *How BlackRock orchestrates \$11T in assets with production AI agents*. Public engineering presentation by BlackRock AI Engineering Leads. Documented by ZenML LLMOps Database. <https://www.zenml.io/llmops-database/agentic-ai-architecture-for-investment-management-platform>
- Deloitte. (2026). *State of AI in the Enterprise 2026*. Deloitte Global. <https://www.deloitte.com/global/en/issues/generative-ai/state-of-ai-in-enterprise.html>
- EAN Network. (2025, April 18). Maersk launches AI-powered vessel routing platform to cut emissions and improve efficiency. <https://ean-network.com/maersk-launches-ai-powered-vessel-routing-platform-to-cut-emissions-and-improve-efficiency/>
- Entrepreneur. (2026, January 30). JPMorgan Chase says AI could cut headcount by 10% in some divisions. <https://www.entrepreneur.com/business-news/jpmorgan-to-cut-headcount-in-some-divisions-due-to-ai/491864>
- HR Grapevine. (2026, February 25). JPMorgan CEO Jamie Dimon says society must consider AI job impact. <https://www.hrgrapevine.com/us/content/article/2026-02-25-jpmorgan-ceo-jamie-dimon-says-society-must-consider-ai-job-impact>
- Huseby, A. (2025). Governance in the age of AI leadership: From advisory systems to organizational authority. Zenodo. <https://doi.org/10.5281/zenodo.20330936>
- Jensen, M. C., & Meckling, W. H. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4), 305–360.
- McKinsey & Company. (2025, December 4). The AI reckoning: How boards can evolve. <https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/the-ai-reckoning-how-boards-can-evolve>

- McKinsey & Company. (2025, November 5). The state of AI in 2025: Agents, innovation, and transformation. McKinsey Global Survey.
- Microsoft. (2024, September 30). Elevating investment management tech: AI-powered leadership from BlackRock and Microsoft. <https://www.microsoft.com/en-us/microsoft-cloud/blog/financial-services/2024/09/30/elevating-investment-management-tech-ai-powered-leadership-from-blackrock-and-microsoft/>
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.
- Russell, S. J. (2019). *Human Compatible: Artificial Intelligence and the Problem of Control*. Viking.
- Suchman, M. C. (1995). Managing legitimacy: Strategic and institutional approaches. *Academy of Management Review*, 20(3), 571–610.
- The Digital Banker. (2026, March 6). JPMorgan Chase's LLM Suite drives AI transformation across the enterprise. <https://thedigitalbanker.com/jpmorgan-chases-llm-suite-drives-ai-transformation-across-the-enterprise/>
- Weill, P., Woerner, S. L., & Banner, J. (2025, December 8). AI-savvy boards drive superior performance. *MIT Sloan Management Review*. Based on MIT Center for Information Systems Research Briefing released March 20, 2025. <https://sloanreview.mit.edu/article/ai-savvy-boards-drive-superior-performance/>

---

Published under Creative Commons Attribution 4.0 International License (CC BY 4.0). Citation:  
Huseby, A. (2026). *The AAMM in Practice: Classifying AI Decision-Making Authority Across Enterprise Deployments*. Zenodo.