

Responsible AI Ecosystems in Banking and Insurance

An ecosystem model and a strategy framework

Forward-Looking AI Governance (FLAIG), Work Package 2 deliverable

Matti Minkkinen, University of Turku (WP2 lead)

30 October 2026

FLAIG is a Co-Research project funded by Business Finland.

Table of Contents

Executive summary.....	3
1 Introduction.....	4
1.1 Background and purpose.....	4
1.2 Scope and boundary with Work Package 1	4
1.3 Sources.....	4
1.4 How to read this report	5
2 Why an ecosystem view?	5
2.1 AI governance as an interorganizational phenomenon.....	5
2.2 The EU’s ecosystem approach.....	6
2.3 What “ecosystem” means here.....	6
3 The responsible AI ecosystem model	8
3.1 Overview of the model.....	8
3.2 Building blocks: governance layers and institutional logics.....	10
3.3 Two patterns of governance at the environmental layer	12
3.4 The actor map for banking and insurance.....	14
3.5 Institutional translation	16
3.6 Summary of the model.....	17
4 Evidence from the Finnish and Nordic sector.....	17
4.1 Distributed responsibility and ecosystem functions: the insurance interviews	18
4.2 Customers as downstream governance actors: the retail banking interviews	19
4.3 What the evidence adds to the model	20
5 AI ecosystem strategy framework for banking and insurance.....	21
5.1 Purpose and derivation.....	21
5.2 Three positioning options	22
5.3 The seven decision areas.....	24
5.4 Using the framework.....	30
6 Piloting a horizon-scanning process	31
6.1 What was piloted.....	31
6.2 What came out	32
6.3 Running a lighter version annually	32
7 Implications and limitations.....	33
7.1 Implications for banks and insurers	33
7.2 Implications for the FLAIG project	34
7.3 Implications for research.....	34
7.4 Limitations	34
References.....	35
Appendix A. Inventory of relevant AI and IT governance standards.....	39
Appendix B. Glossary of ecosystem actors	41
Appendix C. Sources for the synthesis in the report.....	43
Appendix D. Framework self-assessment worksheet	45

Executive summary

This report delivers the two outputs of Work Package 2 of the FLAIG project (Forward-Looking AI Governance, Business Finland co-research, 2024–2026): a responsible AI ecosystem model and an AI ecosystem strategy framework for banking and insurance. It also documents a piloted horizon-scanning process that can keep both current as the ecosystem changes.

The starting point is that a bank or an insurer cannot govern its AI systems alone. Supervisors, standard-setters, auditors, technology providers, insurers, customers, and civil society all take part in deciding what responsible AI means and whether it has been achieved. The report treats this set of actors as an ecosystem and asks two questions. What does the ecosystem look like, and how should a bank or insurer position itself within it?

The **responsible AI ecosystem model** (Chapter 3, Figure 1) answers the first question. It combines three layers of governance (environmental, organizational, and AI system) with two patterns of interorganizational governance. In *polycentric governance*, independent external actors scrutinize the firm and collaborate with it. In *value chain governance*, obligations run through the firm's relationships: oversight of vendors upstream, reciprocal relations with customers, and the alleviation of AI risk through insurers. Actors enact plural, sometimes competing institutional logics, and environmental pressures do not apply themselves but are translated into practice through negotiation among internal and external actors. The model is populated with the actors of the Finnish and European banking and insurance sector.

Chapter 4 checks the model against two project studies: expert interviews in Finnish insurance and vignette-based interviews with retail banking customers. The evidence supports the model's central claim that responsibility is distributed across actors, and refines it in four respects: upstream relations involve dependence as well as oversight; customers' overseeing role is latent; governance is being built onto existing risk and compliance routines; and regulation is experienced as enabling as well as constraining.

The **AI ecosystem strategy framework** (Chapter 5) answers the second question. It turns each component of the model into a decision area: external scrutiny posture, stakeholder engagement design, vendor and upstream governance, customer-facing and downstream governance, risk transfer and insurability, translation capacity, and anticipatory capability. For each area it describes three positioning options (minimal compliance, managed participation, and active shaping) and offers four tests for choosing among them, drawn from business ecosystem theory and institutional logics.

Because the ecosystem is changing, Chapter 6 describes the horizon-scanning process piloted in an expert study on generative AI, links its findings to the framework, and sets out a lighter version that a bank or insurer can run annually. Chapter 7 draws implications and states limitations, and the appendices include a one-page self-assessment worksheet.

Three messages stand out. First, a bank's or insurer's responsible AI position is neither wholly in its own hands nor outside them: it depends on supervisors, providers, insurers, and customers, but the firm chooses how it is placed among them. Second, for most Finnish banks and insurers in most areas, the realistic choice today is between minimal compliance and managed participation; active shaping is realistic in a few areas, and mostly collectively, through associations, joint positions toward large providers, or shared practice with peers and research partners. Third, three areas deserve particular attention in the sector: upstream governance, where dependence on a few large model providers is greatest; customer-facing governance, where customer expectations are specific and partly coincide with what consumer credit law requires from November 2026; and risk transfer, where an emerging market for AI insurance is beginning to put a price on governance evidence.

1 Introduction

1.1 Background and purpose

FLAIG (Forward-Looking AI Governance) is a three-year Business Finland co-research project (2024–2026) carried out by the University of Turku and Aalto University with three company partners: Tietoevry as ecosystem leader, OP Financial Group, and Veritas. The project contributes to the “Building trusted digital societies” ecosystem and leading company program led by Tietoevry. Its overall aim is to secure the lawful and socially responsible development and deployment of AI in banking and insurance, sectors that are characteristically high-trust and heavily regulated.

The project has three objectives, each carried by a work package. Work Package 1 develops an AI governance framework and auditing guidelines for banking and insurance. Work Package 3 develops quantitative measures for key principles of trustworthy AI. Work Package 2, whose deliverable this report is, develops “a forward-looking view of the innovation potential of responsible AI ecosystems and the disruptions brought by generative AI in banking and insurance”. In the project plan, WP2 owes two deliverables: a *responsible AI ecosystem model*, described as “a general overview of the emerging ecosystem in Europe, laying out the types of actors involved”, and an *AI ecosystem strategy framework for banking and insurance*, which “will zoom in on the relevant issues and strategic considerations specifically for banking and insurance companies based on the general ecosystem model”. The project’s statement of goals adds that, to enable continuous foresight, the project “pilots a horizon-scanning process”.

This report delivers those items. Its purpose is to give the project partners, Business Finland, and the wider Finnish financial sector a clear account of who takes part in governing responsible AI around a bank or insurer, how those parties act on one another, and what strategic choices this opens up for a firm. It is written from an information systems (IS) research perspective, which means that it states its concepts explicitly, shows where they come from, and separates what the evidence shows from what remains to be validated. It does not attempt to speak in the voice of a practitioner’s handbook; the self-assessment worksheet in Appendix D is the part designed to be used directly.

1.2 Scope and boundary with Work Package 1

This report is about what crosses the organizational boundary: the actors around a bank or insurer, the relationships between them, and the mechanisms through which they govern AI together. What happens inside the firm (its governance framework, auditing guidelines, and regulatory obligations) is the subject of Work Package 1 and appears here only where it meets the ecosystem. This shapes how standards and regulation are treated. The ISO/IEC standards in Appendix A are presented as the artifacts through which standard-setting bodies act in the ecosystem, not as a compliance checklist, and the EU AI Act and sectoral regulation such as DORA are discussed to locate supervisors and conformity assessment bodies, not to enumerate obligations.

1.3 Sources

The model was built by synthesis. Its components come from three FLAIG studies, two of them peer-reviewed papers presented at the I3E 2026 conference: our own layered institutional perspective on AI governance, which remains an unpublished project working paper (Minkkinen et al., 2026); a conceptual synthesis of 54 academic and grey-literature

documents on interorganizational AI governance (Khan & Mäntymäki, 2027); and an interpretive study of how 23 Nordic organizations translate ethical principles and regulation into governance practice (Birkstedt et al., 2027). To these the report adds the sector's actor landscape, an inventory of standards, and empirical illustration from two Master's theses completed in the project (Heikkilä, 2026; Huopainen, 2026) and from an expert-interview study on generative AI and governance accepted for ISPIM Connects Sendai 2026 (Jylhä et al., 2026). Appendix C describes each study's method, data, and publication status, and where it is used in the report.

1.4 How to read this report

Readers who want the deliverables and nothing else can read Chapter 3 (the model, with Figure 1) and Chapter 5 (the framework, with Table 5 and the worksheet in Appendix D). Chapter 2 explains why an ecosystem view is needed at all. Chapter 4 presents the sector evidence organized by what it evidences rather than by which study produced it. Chapter 6 describes the horizon-scanning pilot as a repeatable process for keeping the model and the firm's positions current. Chapter 7 closes with implications and limitations. The appendices are reference material.

Throughout, *AI governance* follows the definition used across the FLAIG project: “a system of rules, practices, processes, and technological tools that are employed to ensure an organization's use of AI technologies aligns with the organization's strategies, objectives, and values; fulfills legal requirements; and meets ethical AI principles followed by the organization” (Mäntymäki et al., 2022). *Responsible AI* is used as the broader normative aim that governance serves. *Ecosystem* is defined in Chapter 2. The *focal organization* is the bank or insurer from whose point of view the ecosystem is described. Appendix B is a glossary of actor terms.

2 Why an ecosystem view?

2.1 AI governance as an interorganizational phenomenon

The use of AI in banking and insurance has moved from analytics and machine learning to generative AI and, most recently, to AI agents that can pursue tasks across several steps with limited human intervention. Regulation has followed. The EU AI Act (Regulation (EU) 2024/1689) entered into force in August 2024 and is being phased in through 2025–2028; after the 2026 Digital Omnibus amendment (Regulation (EU) 2026/1744), the obligations for stand-alone high-risk systems apply from December 2027. The Act places several typical financial-sector uses in its high-risk category: creditworthiness assessment of natural persons, and risk assessment and pricing in life and health insurance. Ethics principles have converged around transparency, fairness, and accountability, and sectoral regulation such as DORA adds requirements for ICT risk and third-party dependencies. In Finland, the Financial Supervisory Authority's thematic review of 83 financial-sector entities found generative and general-purpose AI models to be the most frequently used AI technologies, reported by 74% of respondents; seven entities were already using AI to assess the creditworthiness of natural persons and eleven more intended to (Finanssivalvonta, 2025). Organizations that use AI systems must meet legal, societal, and ethical requirements at the same time, and how those requirements combine in sector-specific form is still poorly understood.

The response of both scholars and practitioners has been to develop *organizational AI governance*: frameworks, roles, processes, and tools inside the firm. The research group

behind FLAIG has contributed substantially to this work, including the definition of organizational AI governance cited above and a systematic review of the field (Birkstedt et al., 2023). Yet a recurring finding of that work is that the organizational frame is incomplete. Khan and Mäntymäki (2027) put the point directly: “AI systems are often developed by one organization, deployed by another, audited by third parties, and regulated by governments”, and so “AI governance is fundamentally an interorganizational phenomenon”. AI governance therefore cannot be understood purely as something that happens inside organizations and their governance systems: making it effective requires many internal and external stakeholders to act together (in the terms of this report, as an ecosystem).

This is the governance gap that gives WP2 its rationale. A bank can write an AI policy, appoint an AI governance committee, and document its models. It cannot, by itself, decide whether a foundation model it licenses has been trained responsibly, whether a supervisor will accept its explanation of a credit decision, whether an auditor exists who can certify its management system, or whether its customers will trust an automated decision. Each of these depends on other actors, and each of those actors has its own interests, competences, and ways of reasoning.

2.2 The EU’s ecosystem approach

The European Commission’s 2020 White Paper on AI framed the policy response in ecosystem terms, calling for an “ecosystem of excellence” and an “ecosystem of trust” (European Commission, 2020). FLAIG’s predecessor work analyzed how this framing was taken up. Minkinen, Zimmer, and Mäntymäki (2023) studied the emerging European ecosystem for responsible AI through EU documents and expert interviews and found that “implementing RAI at a large scale requires collaborative networks as a mediating layer between abstract AI ethics principles and organizational AI implementation”. They also found that the ecosystem has no orchestrator: “No single firm or public organization orchestrates the RAI landscape, even though large technology companies may orchestrate subsystems, and organizations (such as EU bodies) may be aspiring orchestrators” (p. 105).

Two things were left open by that work and by the policy discussion, and they are what this report addresses. First, the *shape* of the ecosystem for a specific sector remained unclear: which actors matter for a bank or insurer, and through which mechanisms they act. Second, the *strategic position* of a firm within that ecosystem was untreated: whether a bank or insurer should simply comply with what the ecosystem demands or take an active part in shaping it. Chapter 3 addresses the first, and Chapter 5 the second.

2.3 What “ecosystem” means here

Sections 2.1 and 2.2 have used the word “ecosystem” in the loose sense common in policy discussion. The AI governance literature often uses it in the same way. Wirtz et al. (2022), for example, build an ecosystem framework of AI governance from the policymaker’s standpoint and distinguish two readings of the term: a proprietary platform environment, governed privately by its owner, and a macro-level network of actors and systems, which calls for public governance. Neither reading describes the position of a bank or insurer, which owns no platform and makes no public policy, and which has to govern its AI through relationships it only partly controls. Stahl (2022, 2023) uses the term in a third way, asking what would make an AI ecosystem as a whole responsible. His question is normative and system-wide, whereas the question here is how a single firm is governed by, and can act within, the ecosystem around it. This report gives the term a more specific meaning, drawn from the business ecosystem literature (Adner, 2017; Granstrand & Holgersson, 2020;

Tsujimoto et al., 2018) and combined with institutional logics and multi-actor AI governance research. The ecosystem for responsible AI in banking and insurance is the structure of actors, relationships, and institutions that must align for AI in the sector to be responsibly governed and audited. Three features of this definition carry the argument of the report: the ecosystem is a structure rather than a community, it contains institutions as well as actors, and its actors reason from different logics.

The ecosystem is a structure, not a community. Adner (2017) distinguishes ecosystem-as-affiliation (a community of actors loosely connected by a platform or a network) from ecosystem-as-structure, “the alignment structure of the multilateral set of partners that need to interact in order for a focal value proposition to materialize” (p. 42). The second view is the one taken here. The value proposition is responsibly governed and audited AI in banking and insurance, and the ecosystem consists of the actors whose contributions are needed for that proposition to hold, whether or not the firm deals with them directly. This extends Adner’s construct, which was developed for commercial offerings. The extension is justified because responsibly governed AI has the property that makes the construct useful: it is something the firm must be able to demonstrate to supervisors, auditors, customers, and insurers but cannot produce alone, so that whether it materializes depends on the alignment of others’ contributions. Unlike a commercial value proposition, however, it is defined partly by actors who are not the firm’s partners at all, such as legislators and supervisors, which is why the second feature below matters. This is where the governance gap of Section 2.1 becomes an ecosystem question: a bank depends on model providers, supervisors, auditors, and customers alike, and the structural view “explicitly extends the strategic view to include activities and actors over which the focal organization may have no control, and with whom they have no direct contact” (Adner, 2017, p. 44). It is also why the model in Chapter 3 is organized around mechanisms of governance rather than around a list of organizations, and why it includes actors, such as supervisors and standard-setters, with which a firm may have no direct commercial relationship.

The ecosystem contains institutions as well as actors. Granstrand and Holgersson (2020) define an innovation ecosystem as “the evolving set of actors, activities, and artifacts, and the institutions and relations, including complementary and substitute relations, that are important for the innovative performance of an actor or a population of actors”. On this understanding, regulation, standards, and supervisory expectations are part of the ecosystem, not an external constraint on it, which is also how the EU’s “ecosystem of trust” was conceived (Section 2.2). In the model of Chapter 3 they sit in the environmental layer; that layer is the part of the ecosystem lying outside the focal organization, not something outside the ecosystem itself. Empirical ecosystem research bears this out. Compared with supply-chain and alliance research, it “includes the non-physical, non-business, informal, and invisible relationships”, and “bureaucrats are among the most important actors in the ecosystem analysis; they affect the other actors through the use of regulations and non-official behavior” (Tsujimoto et al., 2018, p. 56). For a sector as heavily regulated as banking and insurance, a definition that left institutions out would miss much of what shapes AI governance.

The ecosystem’s actors reason differently. Tsujimoto et al. (2018) describe an ecosystem as a “multilayer social network” of “actors that have different attributes, decision principles, and beliefs” (p. 55). Where no single actor orchestrates the ecosystem, as Minkkinen, Zimmer, and Mäntymäki (2023) found for responsible AI in Europe, these differences are not smoothed out from above and need to be understood in their own right. In IS research, the established way of analyzing them is through institutional logics, defined as “the socially

constructed, historical patterns of material practices, assumptions, values, beliefs, and rules by which individuals produce and reproduce their material subsistence, organize time and space, and provide meaning to their social reality” (Thornton & Ocasio, 1999, p. 804). Institutional logics have been used to study ecosystems in IS research before; Qiu, Gopal, and Hann (2017), for example, show app developers on a platform managing professional and market logics through both conflict and synthesis. A supervisor, a data scientist, a chief financial officer, and a consumer advocate do not merely have different interests in an AI system; they hold different views of what counts as a good reason. Chapter 3 uses four such logics to explain why the same external pressure produces different responses in different firms, and Chapter 5 uses them to explain why some positioning choices are harder to sustain than others.

These three features (structure, institutions, and logics) give the ecosystem view its content, and they set the task for the next chapter: to specify, for banking and insurance, which actors and institutions make up the structure, through which mechanisms they act on the firm, and why their different logics matter for how the firm responds.

3 The responsible AI ecosystem model

3.1 Overview of the model

The responsible AI ecosystem model can be stated in a paragraph, and the rest of this chapter unfolds it.

Responsible AI in banking and insurance is governed across three layers (the environmental layer, the organizational layer, and the AI system layer), each populated by different actors enacting plural, sometimes competing institutional logics. At the environmental layer, governance operates through two patterns: polycentric governance, in which authority is distributed among independent external actors who scrutinize and collaborate with the firm, and value chain governance, in which the firm is itself a link and obligations are embedded in its upstream relations with vendors, its reciprocal relations with customers, and the alleviation of risk through insurers. Pressures generated at this layer do not apply themselves: they are translated into organizational practice through ongoing negotiation among internal and external actors. A bank’s or insurer’s responsible AI position therefore depends as much on how it is placed within, and acts upon, this surrounding structure as on its internal governance framework.

Figure 1 draws the model. The three layers are stacked from top to bottom. The environmental layer holds the two governance patterns, each with its components and the banking and insurance actors that populate them, and the figure draws them so that their different directions of flow are visible. Polycentric governance acts on the focal bank or insurer from outside: its arrow down to the organizational layer is labeled *institutional translation* rather than *implementation*, for reasons that Section 3.5 explains, and the return arrow carries the firm’s evidence and accountability back to those who scrutinize it.

Value chain governance runs through the focal organization, which is itself a link in the chain. The figure therefore draws the firm once, in the organizational layer, and connects its vendors, insurers, and customers to it across the layer boundary: the firm oversees its vendors upstream, exchanges preparation and feedback with its customers downstream, and receives assessment and pricing from insurers, whose cover alleviates risk for the firm and its

customers. A dashed line links insurers to collaborative stakeholder engagement, because they also assess AI risk as external stakeholders. The AI systems the firm operates sit at the system layer.

Colored tags mark the institutional logics that each group of actors characteristically enacts (Table 2). The tags appear at every layer, because the logics cut across layers rather than belonging to any one of them, and they show that the two patterns lean toward different logics: polycentric actors mainly toward regulatory oversight and principlism, value chain actors mainly toward managerial rationalism and IT professionalism, with insurers and customers as hybrids. Inside the focal organization, the logics carried in from outside meet the logics of the firm's own functions, and the figure names the outcomes that Sections 3.2.2 and 3.5 describe: loose coupling, compromise, polarization, or synthesis. At the system layer, a note records that the properties of AI systems enable and constrain which logics can be enacted (cf. Faik et al., 2020).




Figure 1. The responsible AI ecosystem model for banking and insurance.

The model is a synthesis: it integrates four strands of work we have conducted or commissioned over the course of the project. The three layers and the four institutional logics are taken from our layered institutional perspective, set out in an unpublished project working

paper (Minkkinen et al., 2026). The two governance patterns and their five components build on Khan and Mäntymäki’s (2027) conceptual synthesis of interorganizational AI governance. The translation mechanism is adapted from our actor-network study of Nordic organizations (Birkstedt et al., 2027). The actor map is this report’s own contribution, assembled from the FLAIG standards inventory, the sector’s supervisory and vendor landscape, and the two Master’s theses commissioned for the project. The sections below take each element in turn.

3.2 Building blocks: governance layers and institutional logics

3.2.1 Governance layers

AI governance takes place at several levels at once, and layered models have been a recurring device in the literature for making sense of this. Gasser and Almeida (2017) distinguished technical, ethical, and social/legal layers; Shneiderman (2020) distinguished team, organization, and industry. Our layered institutional perspective builds on these and identifies three interconnected layers, each with its own governance mechanisms and stakeholders. Table 1 summarizes them.

Table 1. The three governance layers.

Layer	Governance mechanisms	Typical stakeholders
Environmental	Institutional–regulatory (polycentric): formal regulation (EU AI Act, GDPR, sectoral law); collaborative governance (co-regulation, standards, multi-stakeholder initiatives). Value chain: procurement and contractual control upstream, disclosure and feedback downstream, risk alleviation through insurance (see Section 3.3)	Institutional–regulatory: regulators and supervisors, international bodies, standards organizations, external auditors, academia, civil society, investors, affected individuals and groups. Value chain: cloud and foundation-model providers, application vendors, distribution partners, customers, insurers
Organizational	Corporate governance, IT governance, data governance; AI governance as the part that goes beyond these	Managers and decision-makers, AI deployers, developers, AI system providers, internal auditors, review boards, and compliance functions
AI system	System-level governance (model documentation, testing, monitoring); foundation model governance	Developers, AI system providers including foundation model providers, auditors, suppliers, end users

Two features of the layered model matter for the ecosystem model. First, governance pressures flow in both directions. Regulation and expectations flow down from the environmental layer to the organization and to its systems; but the properties of AI systems, such as their opacity and their dependence on a foundation model trained elsewhere, flow upward and create regulatory and reputational risk at the environmental layer, which in turn increases internal governance requirements. Second, the organizational layer is a *mediating* layer. It is where requirements from the environment are turned into practices, and where the constraints of specific systems are turned into policy. This is why the strategy framework in Chapter 5 is addressed to the organization: it is the layer at which choices are made.

The layered model also fixes the meaning of “ecosystem” in this report in relation to the layers. The ecosystem is not one of the layers but the whole structure that the model describes, including the focal organization and its AI systems. The environmental layer is the part of the ecosystem that lies outside the focal organization, seen from the organization’s standpoint, and it is where the actors of Chapter 2 sit. “Environmental” is meant in the organization-theory sense of an organization’s environment, not the natural environment. The

term replaces the label “societal/ecosystem” used by Minkkinen et al. (2026), so that “ecosystem” is reserved for the structure as a whole. The model’s two governance patterns describe how those actors and the focal organization govern one another: polycentric actors act on the firm from outside, whereas in the value chain the firm is itself a link that governs as well as being governed.

It also explains why the environmental layer is not divided into a societal layer and a value chain layer. The layers are ordered by scope (a system sits within an organization, which sits within its environment), whereas the difference between a supervisor and a vendor is not a difference in scope but in the kind of relation through which the actor is connected to the focal organization: authority in the one case, transaction in the other. Several actors occupy both positions at once. An insurer is a market counterparty and, through pricing, a governance authority; an external auditor is both a contracted service provider and an independent scrutineer; a corporate customer both transacts with the organization and assesses it. Splitting the layer would force such actors into one box and would restate, less precisely, the distinction that Section 3.3 already draws as the difference between polycentric and value chain governance. The layer is therefore kept whole and differentiated internally: Table 1 marks its two halves, and Section 3.3 sets out how each of them operates.

3.2.2 Institutional logics

Actors at every layer reason from institutional logics. Our layered institutional perspective, drawing on earlier published work (Minkkinen & Mäntymäki, 2023), identifies four logics that are relevant to AI governance and shows which stakeholder groups characteristically enact each. Table 2 summarizes them.

Table 2. Four institutional logics in AI governance (after Minkkinen & Mäntymäki, 2023).

Logic	Emphasis	Characteristically enacted by
AI ethics principlism	Adherence to principles such as fairness, transparency, and non-maleficence; embedding principles in practice	Civil society, academia, end users, affected groups
Managerial rationalism	Efficiency, return on investment, cost control, risk management	Investors and markets, managers, AI deployer organizations
IT professionalism	Technical effectiveness, data management, bias audits, safety testing, sound development methodology	Standards organizations, AI developers, AI system providers
Regulatory oversight	Compliance with binding requirements; enforcement of rules; standardization	Regulators and supervisors, international regulatory bodies, external auditors, internal review boards

The four logics are not a taxonomy of good and bad actors. Each is a coherent way of judging what counts as responsible AI, with its own sources of legitimacy: ethical expertise and guidelines for principlism, measurable efficiency and profitability for managerial rationalism, technical competence and being in control for IT professionalism, and the coercive force of law for regulatory oversight. The point of naming them is that organizations face *institutional pluralism* (pressures from several logics that may complement or conflict with one another) and that the response to pluralism is rarely straightforward compliance.

IS research on plural logics shows why, and its findings recur in the FLAIG evidence. Berente and Yoo (2012) showed that when a system embodies managerial rationalism and its users reason from a professional logic, the result is *loose coupling*: practices that satisfy both

logics on the surface while diverging underneath. Boonstra, Eseryel, and van Offenbeek (2018) found that stakeholders enacting competing logics in IT governance produce *polarization*, *compromise* or, more rarely, *synthesis*, and that an organization can reach synthesis on one issue while polarizing on another. Bernardi and Exworthy (2020) found that people with *hybrid identities* (professionals who have also become managers) can broker between logics, but that their brokering capacity degrades as conflict intensifies. Hansen and Baroody (2020) add that the same two logics can be complementary around one activity and in conflict around another, so that conflict is a property of the activity, not of the logics as such.

For the ecosystem model these findings carry a simple message. When a supervisor, a standard-setter, or a customer group presses a requirement on a bank, the requirement arrives carrying a logic, and it meets other logics inside the bank. What happens next depends on who enacts which logic, how well they perform their roles, and whether the organization has people able to translate between them. Section 3.5 returns to this.

The four logics were identified for AI governance across sectors, and Minkkinen and Mäntymäki (2023) deliberately bracketed out sector-specific logics, recommending that as few logics as possible be added for the sake of parsimony. Banking and insurance nonetheless have a strong sectoral logic of their own: a prudential and actuarial logic of risk, in which risk is measured, capitalized, priced, and supervised. This report does not add it as a fifth logic. In the sector it is carried by functions that already combine two of the four: risk management, actuarial, and internal audit functions answer both to supervisory requirements (regulatory oversight) and to the firm's risk appetite and capital efficiency (managerial rationalism). The insurer's pricing of AI risk in the risk alleviation component combines the same two. The model therefore treats the sectoral risk logic as a hybrid of regulatory oversight and managerial rationalism, shown in Figure 1 by double tags on the firm's risk functions and on insurers. Whether it behaves as a distinct logic, with its own sources of legitimacy and professional identity, is an empirical question that the FLAIG evidence can only begin to answer, and one for further research with the sector's risk, actuarial, and audit professionals.

3.3 Two patterns of governance at the environmental layer

The layers say *where* governance happens. They do not say *how* governance runs between the focal organization and the actors at the environmental layer. That is the contribution of Khan and Mäntymäki (2027), whose conceptual synthesis of 27 academic studies and 27 grey-literature documents asked what governance patterns characterize interorganizational AI governance. Their inductive coding produced 26 first-order concepts, five second-order themes, and two aggregate dimensions, which this report adopts as the two patterns of the ecosystem model. They are summarized in Table 3 and described below.

Table 3. Two patterns of interorganizational AI governance and their components (after Khan & Mäntymäki, 2027).

Pattern	Component	What it consists of
Polycentric governance	External scrutiny	Independent oversight of AI systems across the lifecycle: audits, threat assessment during development, incident investigation after deployment, corrective recommendations, and external assessment of the organization's own evaluation methods
	Collaborative stakeholder engagement	Deliberate involvement of affected groups, academics, domain experts, AI experts, and insurers in benchmarking, risk assessment, feedback collection, and continuous monitoring

Pattern	Component	What it consists of
Value chain governance	Upstream governance	Vendor due diligence and audits, procurement and selection criteria, tailoring procured systems to governance needs, continuous oversight of vendors, limited access as a mitigation for component suppliers
	Downstream governance	A reciprocal relation with customers: the organization prepares customers for AI governance; customers conduct their own risk assessment and feed oversight inputs back
	Risk alleviation	Insuring AI systems as part of the focal organization's responsibility, so that financial risk is reduced for the organization and its customers, while the insurer's assessment and premium act as an incentive to reduce risk

3.3.1 Polycentric governance

Polycentric governance is a concept from the study of the commons (Ostrom, 2010) and denotes a system in which multiple independent actors hold decision-making authority and jointly govern a shared problem through interaction. Applied to AI, it means that governance authority does not sit only with the deploying organization: it is distributed among external actors who scrutinize the organization and collaborate with it. Khan and Mäntymäki (2027) identify two components.

External scrutiny is the involvement of external stakeholders in overseeing AI systems during development and deployment, complementing internal oversight. Its defining feature is independence: multidisciplinary, independent bodies investigate threats during development and incidents after deployment, and can assess not only the AI system but the organization's own evaluation method for it. In banking and insurance this component is unusually well developed, because prudential supervision, conduct supervision, and external audit already exist and are being extended to AI. Section 3.4 names the bodies.

Collaborative stakeholder engagement is the deliberate involvement of diverse external actors (affected groups, academics, domain and AI experts) in governing AI systems. It differs from external scrutiny in who initiates it: scrutiny is initiated by external bodies; engagement is arranged by the organization. Neither is one-sided, however: Khan and Mäntymäki stress that external scrutiny "is conducted in collaboration with the focal organization, not in isolation", and engagement is what gives that collaboration its form. It runs across the lifecycle: consulting experts when benchmarking and assessing risk, collecting feedback on capability and safety, and continuing joint monitoring after deployment. Khan and Mäntymäki's coding also places insurers here: insurers that evaluate the risks posed by an organization's AI systems contribute external assessment, and the insurance offered on that basis gives the organization an incentive to reduce risk.

3.3.2 Value chain governance

Value chain governance embeds governance in the relationships along the chain through which AI systems are produced, deployed, and used, and the focal organization is itself a link in that chain rather than only its object. Khan and Mäntymäki (2027) identify three components, and it is useful to notice that they differ in direction.

In Khan and Mäntymäki's account, *upstream governance* is one-directional. The focal organization oversees its vendors through due diligence, audits, procurement criteria, contractual control, and continuous oversight; the vendor is the object of governance. For a bank or insurer, this is where the sector's existing third-party risk regime meets AI: the same

processes used for outsourcing and ICT third-party risk are extended to cloud providers, foundation-model providers, and application vendors. The evidence in Chapter 4 suggests that, for generative AI in particular, the direction of this relation is less clear-cut than the term implies (Section 4.3).

Downstream governance is reciprocal. The organization prepares its customers for AI governance by disclosing, explaining, and setting expectations, and customers in turn conduct their own assessments and feed oversight inputs back. Khan and Mäntymäki make the important observation that customers hold a *hybrid role*: they transact with the organization and oversee its AI at the same time. In Khan and Mäntymäki's sources the customer is often an organization able to run its own risk assessment; with retail customers the return flow takes more limited forms (feedback, complaints, demands for explanation and human recourse), and Chapter 4 shows what this looks like in retail banking.

Risk alleviation is the third component, and for this project it is the most distinctive. Khan and Mäntymäki treat insuring AI systems as part of the focal organization's responsibility: the insurer, as a value chain stakeholder, reduces the financial risk, so that AI-related risks are alleviated for the organization and its customers, and the premium "serves as a governance mechanism in the value chain that incentivizes the focal organization to reduce AI risks". Insurance here is a governance mechanism, not merely a sector. For a bank, insuring AI risk is a control decision. For an insurer, underwriting AI risk is a product decision, and one that, if done well, makes the insurer a governance actor toward the firms it insures. Khan and Mäntymäki observe that insurers, like customers, play a hybrid role: they assess the organization's AI risks and absorb them partially when necessary. Whether AI risk can actually be transferred or distributed to insurers, and how that would affect governance inside the insured organization, Khan and Mäntymäki leave as a question for future research. This report therefore uses "risk transfer" only for the strategic choice built on this component (Chapter 5, area 5), which is where the report takes the question up.

3.3.3 Why both patterns are needed

The two patterns describe different things. Polycentric governance describes the distribution of *authority*: who, outside the firm, is entitled to look and to judge. Value chain governance describes the embedding of *obligation*: what the firm owes to and is owed by those it buys from, sells to, and insures with. Khan and Mäntymäki (2027) conclude that interorganizational AI governance "is characterized by the synergy of two governance patterns: polycentric and value chain governance — distributing oversight authority across independent external actors and embedding governance obligations in the bi-directional relationship". A bank that attends only to its supervisors has half a picture; so has a bank that attends only to its vendors and customers. The ecosystem model places the two together at the environmental layer for this reason, but draws them differently: polycentric governance acts on the firm from outside, value chain governance runs through it.

3.4 The actor map for banking and insurance

The layers and patterns are general. What makes the model a model of *banking and insurance* is the placement of the sector's actors on it. Table 4 gives the actor map; Appendix B expands each entry into a glossary definition. The map is organized by model component so that the derivation is visible, and it lists the actors relevant to a Finnish bank or insurer operating under EU law. In Finland, the Act on the Supervision of Certain Artificial Intelligence Systems (1377/2025), in force since 1 January 2026, designates FIN-FSA as the market surveillance authority for the high-risk AI systems most relevant to the sector,

including life and health insurance risk assessment and pricing and creditworthiness assessment by supervised entities. It designates the Finnish Transport and Communications Agency (Traficom) as the national single point of contact.

Table 4. Actor map for the responsible AI ecosystem in Finnish and European banking and insurance.

Model component	Actors	Role in governing the focal bank or insurer
External scrutiny	Financial Supervisory Authority (FIN-FSA / Finanssivalvonta); European Central Bank and Single Supervisory Mechanism (for significant banks); European Banking Authority (EBA); European Insurance and Occupational Pensions Authority (EIOPA); national AI Act market surveillance authorities; notified bodies and conformity assessment bodies; external and statutory auditors; audit and certification bodies under ISO/IEC 42006	Supervise, examine, certify, investigate incidents; assess the firm's own evaluation methods; impose sanctions
Collaborative stakeholder engagement	ISO/IEC JTC 1/SC 42 and CEN-CENELEC JTC 21 (standard-setters); industry associations (Finance Finland / Finanssiala; European federations such as the European Banking Federation and Insurance Europe); academic and research institutions, including the FLAIG consortium; civil society and consumer organizations; affected individuals and groups; insurers, as assessors of AI risk	Develop voluntary norms and standards; contribute expertise to benchmarking and risk assessment; give feedback; represent affected interests
Upstream governance	Cloud and infrastructure providers; foundation-model and generative AI providers; core banking, policy administration, and claims system vendors; specialized AI and analytics vendors; consultancies and system integrators; data providers	Supply models, components, data, and services that the firm integrates; carry governance obligations transmitted through procurement and contract
Downstream governance	Retail and corporate customers; policyholders; end users of AI-enabled services; consumer advocacy bodies	Receive and react to AI-enabled services; form expectations; give feedback, complain, demand explanation and human recourse
Risk alleviation	Insurers and reinsurers underwriting operational, cyber, and AI-related risk; for an insurer, its own underwriting function acting toward other firms	Assess and price AI risk; absorb part of the risk, alleviating it for the firm and its customers; discipline the insured through pricing and terms
Cross-layer (translation)	Internal AI governance function or committee; compliance, legal, and risk functions; business units owning AI use cases; developers and data scientists; internal audit	Convert external pressure into practice; act as the organization's stable spokesperson toward the ecosystem

Three observations about the map are worth recording, because they shape the strategy framework.

First, the external scrutiny row is dense. Banking and insurance already possess a mature supervisory and audit apparatus, and the AI Act adds market surveillance and conformity assessment to it. The FLAIG insurance interviews (Section 4.1) confirm that regulation occupies “an especially central position in the insurance sector”, more so than general

ecosystem literature would suggest (Heikkilä, 2026). Supervision of the AI Act itself is divided in Finland among several market surveillance authorities according to the purpose of each AI system, so a bank or insurer may face more than one AI supervisor in addition to its prudential and conduct supervisors. The sector's problem is therefore rarely too little scrutiny; it is the coordination of many scrutinizers, and the question of how much scrutiny a firm should invite beyond what is required.

Second, the upstream row reveals a dependency that the sector has only partly absorbed. Generative AI capabilities in Finnish insurers are substantially bought in rather than built; Heikkilä (2026) cites sector figures indicating that a large share of insurers procure generative AI services externally. The safeguards around those capabilities are partly supplied by model and platform providers. In her words, "responsibility is partially outsourced and co-produced across organizational boundaries" (Heikkilä, 2026, p. 50). The pattern is sector-wide: the supervisor's thematic review found commercial cloud to be the most common deployment environment for AI in Finnish financial firms and listed third-party management and inadequate agreements among the risks they identified (Finanssivalvonta, 2025). Upstream governance is where that co-production is, or is not, made accountable.

Third, the risk alleviation row places the same organization on both sides of the relationship. A non-life insurer is simultaneously a deployer of AI that may seek to transfer its own AI risk and an underwriter that may price other firms' AI risk. No other component of the model has this property, and it is the reason Chapter 5 gives risk transfer and insurability particular attention.

3.5 Institutional translation

The final component of the model is the relation between layers. In Figure 1 the arrow from polycentric governance down to the organization is labeled *institutional translation*, although pressures that arrive through the value chain are translated too. The term is chosen deliberately to reject the view, common in policy documents, that governance requirements are *implemented* in organizations the way a specification is implemented in code.

The evidence for this comes from the actor-network study of Nordic organizations (Birkstedt et al., 2027), which analyzed 27 expert interviews from 23 organizations in two rounds: Finnish frontrunner organizations in 2020, before binding regulation, and Nordic retail banks and insurers in 2022–2023, as they prepared for the EU AI Act. Using Callon's four moments of translation (Callon, 1984), the study identified *organizational practices* as the focal actor that mediates between external demands and internal operations, and "a fit-for-purpose approach to AI governance" as the obligatory passage point through which all actors had to pass. Three of its findings become part of the ecosystem model.

The first is that regulation enters as a non-negotiable force but its operational translation is heterogeneous. One banking informant put it plainly: "Regulatory or administrative pressure blocks quite clearly any black box AI and drives us more towards rule-based solutions." The regulation is the same for all firms; what firms do with it differs, and the difference is settled at the level of individual AI systems and use cases rather than by organization-wide rule-making.

The second is that stable translation requires a stable spokesperson. Organizations that had established "a designated AI governance function or a multi-disciplinary AI governance committee with explicit decision rights" were able to convert external expectations into

durable arrangements; those without such institutional anchoring “experienced repeated re-negotiation of the same governance issues”. The model therefore places the AI governance function at the organizational layer as the point at which the ecosystem meets the firm.

The third is that AI systems themselves shape their own governance. The opacity of certain banking models forced organizations to develop rule-based fallbacks and to redefine which decisions could be delegated at all. This is the bottom-up arrow in Figure 1: properties of the system layer constrain what the organizational layer can promise to the environmental layer.

Read together with the institutional logics of Section 3.2.2, translation explains why the same ecosystem pressure produces uneven organizational responses. A requirement arriving from a supervisor (regulatory oversight) is received by managers (managerial rationalism), data scientists (IT professionalism) and, sometimes, an ethics function (principlism). Whether the result is synthesis, loose coupling, compromise, or polarization depends on the organization’s capacity to negotiate between these logics. Our layered institutional perspective (Minkkinen et al., 2026) calls this “the mediating role of hybrid professional identities”; Birkstedt et al. call the same capacity “translation capabilities” and “productive friction”. In the strategy framework it becomes a decision area in its own right (Chapter 5, area 6). It is also the natural point of contact with WP1: the governance framework developed there is the organizational-layer instrument through which translation is organized.

3.6 Summary of the model

The responsible AI ecosystem model states that a bank’s or insurer’s responsible AI position is a product of three things: where it sits in a structure of external scrutiny, collaborative engagement, upstream, downstream, and risk-alleviation relations; which institutional logics the actors around it and inside it enact; and how well it translates the resulting pressures into practice. The model is deliberately parsimonious. It does not enumerate every organization in the sector, and it does not prescribe governance practices, which is WP1’s task. What it does is make the structure visible, so that the choices a firm has within it can be stated. Chapter 4 checks the model against project evidence; Chapter 5 turns it into those choices.

Two qualifications apply. First, the project evidence refines the model in several respects, which Section 4.3 sets out. Second, the structure the model describes is moving: generative AI is reshaping the upstream component, and the AI Act’s phased application, itself revised in 2026, is reshaping external scrutiny. The model therefore describes the ecosystem at one point in time, and this is why the strategy framework adds to the six decision areas derived from the model’s components a seventh, on anticipating how the structure will change (Chapter 5, area 7).

4 Evidence from the Finnish and Nordic sector

This chapter presents the empirical material produced in the FLAIG project that bears on the ecosystem model. It is organized by what the evidence shows about the model rather than by which study produced it. Two studies are drawn on: an interview study of AI governance in the Finnish insurance sector (Heikkilä, 2026) and a vignette-based interview study of retail banking customers’ expectations of AI (Huopalainen, 2026). The Nordic actor-network study (Birkstedt et al., 2027), which includes interviews with eleven retail banks and insurers, has already been used to build the translation component of the model (Section 3.5) and is not repeated here. Appendix C gives each study’s method, data, and publication status in one

place. The two theses are small qualitative studies; they illustrate and refine the model rather than test it.

4.1 Distributed responsibility and ecosystem functions: the insurance interviews

Heikkilä (2026) asked what functions contribute to the development of responsible AI ecosystems in the insurance sector, and how responsibility is distributed among actors. The study interviewed six professionals in Finnish insurance (an AI solutions architect, a legal counsel, an AI governance specialist, a development manager, a data analytics manager, and a data scientist) between December 2025 and February 2026, and analyzed the transcripts inductively. Its theoretical frame combined responsible AI principles with an ecosystem perspective, using the concept of *functions* to bridge the two: functions are the activities through which principles are operationalized, actors perform functions, and collaboration connects functions across organizational boundaries.

Three findings bear directly on the model. The first concerns the distribution of responsibility. The thesis concludes that “responsibility is therefore not concentrated in a single actor but is distributed across the ecosystem according to actors’ roles, capabilities, and influence over AI development and use” (Heikkilä, 2026, p. 3), and that “responsible AI does not reside in any single tool, policy, or actor, but emerges from the alignment of these elements across the ecosystem” (p. 51). Inside the firm, responsibility “follows the use of AI rather than being centralized in a single function” (p. 60): business units own the AI used in their own area, while developers, compliance and legal, IT, risk management, and management each perform distinct functions. This is the organizational layer of the model seen from inside, and it matches the translation finding of Section 3.5: the informants described cross-functional review forums bringing together contract law, data protection, business law, compliance, and risk management, which is the “stable spokesperson” of Birkstedt et al. (2027) under another name.

The second concerns the ecosystem functions themselves and their provenance. The functions identified (governance models and principles, risk management and compliance, training and AI literacy, human oversight, organizational structures, technical safeguards, documentation and validation) mostly “build on existing governance mechanisms rather than being entirely new” (p. 51). Insurance already has mature risk, compliance, and documentation routines, and AI governance is being grafted onto them. One informant’s formulation captures the tone: “all AI must have its risks assessed and then mitigated if any are identified, and this cannot be bypassed under any circumstances” (I6, p. 50). Validation and human oversight were described as non-negotiable: outputs “must be validated and cannot be blindly trusted”, and legal responsibility “cannot be transferred to any information system” (I1, p. 48). For the model, this means that the external scrutiny and translation components are, in insurance, built on firm foundations, and that the sector’s distinctive challenge is extension rather than construction.

The third concerns the external actors. Regulation and legislation (the EU AI Act, DORA, and FIN-FSA were named) were the most frequently mentioned external actor, and one informant called regulation “most likely the most important external actor”. The informants’ views of regulation were split in an instructive way. On one side, regulation was said to give courage: “Without this kind of EU regulation, no one really dares to do anything here, because it is known that regulation is coming anyway” (I1, p. 55). On the other, the AI Act was described in terms of “misguided interpretations and unnecessary bureaucracy” (I1,

p. 56). DORA was singled out for constraining which partners an insurer may cooperate with, which is upstream governance in the model's terms. Technology and model providers were the second most-mentioned external actor, and were credited with doing "quite a lot of that responsibility work themselves". This is the co-production noted in Section 3.4. Customers and the general public appeared mainly as sources of reputational pressure; one informant observed that transparency about AI use "sometimes causes negative reactions" (I6, p. 57).

Two absences in the insurance evidence are as informative as the findings. Standards did not appear in the interviews at all, although the thesis's literature review lists standard-setters among ecosystem actors. And insurers' role as *underwriters* of other firms' AI risk, the risk alleviation component, was not raised by any informant; insurers spoke of risk only in terms of their own risk management. Both absences mark the distance between the ecosystem model, which is derived from international literature, and current Finnish practice, and both are taken up as positioning questions in Chapter 5.

4.2 Customers as downstream governance actors: the retail banking interviews

Huopalainen (2026) asked what expectations retail banking customers have concerning the use of AI in banking services. The study interviewed ten Finnish retail banking customers in April 2026 using seven vignettes (fictional scenarios ranging from a loan-advice chatbot and a budgeting tool to fully automated credit underwriting, alternative-data credit pricing, a fraud-detection block abroad, and an AI agent managing finances autonomously) and analyzed the transcripts thematically, using the EU High-Level Expert Group's requirements for trustworthy AI as a sensitizing lens. The author is careful to state what the study is not: it does not study "stakeholder pressure mechanisms", and its contribution is "empirical groundwork on what customers themselves expect from AI" (p. 9). Within that limit, it is the project's most direct evidence on the downstream governance component of the model.

The central analytical result is a distinction between two kinds of expectation. *Competence-based expectations* concern whether AI can do the job: its capability and reliability, and its usefulness. *Ethical and relational expectations* concern how the customer is treated: transparency and explanation, human contact and control, privacy and data use, fairness, accountability, and the conditions under which trust is built. A third group, *stakes and context*, does not consist of expectations as such but of conditions that determine how strongly the other two are activated: the higher the financial stakes and the more autonomous the AI, the more the ethical and relational expectations come to the fore.

This distinction does real work for the model. Khan and Mäntymäki (2027) describe downstream governance as reciprocal: the organization prepares customers, and customers feed back. The retail banking evidence shows what customers would feed back, and it is not primarily a judgment of the AI's competence. On competence, customers were skeptical but pragmatic; several thought an AI could be more consistent and less emotional than a human in a credit decision, and reliability was framed as a question of maturity ("banks should wait ... proven in some less critical domain", P7, p. 42). On ethical and relational matters, however, the expectations were specific and firm. Customers wanted to be told that and where AI is used: "the key is that it just should be openly communicated to the customer that in what situations, and how the AI is being used" (P2, p. 45). They wanted explanations after negative outcomes that include instructions on how to proceed. They wanted a human to turn to; the most frequent code in the study was the need for human intervention or service, and one participant asked for a customer "veto-right to demand that a human bank service person

goes through the decision” in automated underwriting (P6, p. 46). They wanted to know where the bank’s data came from, and some wanted the bank’s data governance process to be transparent, “maybe even list all the third-party service providers” (P7, p. 48). And they did not accept legal disclaimers as a legitimate way of transferring liability.

Two features of these findings connect to other components of the model. The demand for transparency about third-party providers links downstream governance to upstream governance: customers are, in effect, asking the bank to account for its vendors. And the demand for a human route of recourse links downstream governance to external scrutiny: the customer who complains to a human is invoking a form of oversight, and the study notes that the right to complain was tied by participants to fairness. The hybrid role that Khan and Mäntymäki attribute to customers (transacting and overseeing at once) is visible here, though in a weaker form than the literature suggests. Customers in the study did not describe themselves as monitors of the bank’s AI; they described the conditions under which they would continue to trust it. For a bank, the practical implication is that downstream governance is largely a matter of *preparing* customers, and that consent and the continued availability of human service were the conditions most often attached to acceptance.

4.3 What the evidence adds to the model

The evidence supports the model’s central claim that responsibility for AI is distributed across an ecosystem rather than located in a single function. Both governance patterns are recognizable in Finnish practice: scrutiny by supervisors and the law is the dominant external relation, and upstream and downstream relations are where responsibility is co-produced with providers and negotiated with customers. Small qualitative studies cannot confirm a model, but they can show where it needs refining, and four refinements follow.

Upstream relations involve dependence as well as oversight. Khan and Mäntymäki (2027) describe upstream governance as one-directional oversight of vendors by the focal organization. For generative AI, the insurance interviews describe something closer to the reverse: providers supply part of the safeguards and do “quite a lot of that responsibility work themselves”, and the insurer depends on their practice. The actor-network study makes the underlying point explicit: large technology providers have far greater capacity to shape the network than smaller actors (Birkstedt et al., 2027). The model keeps the upstream component, but its direction depends on bargaining power, which is why collective action is the active shaping option in area 3.

Customers’ overseeing role is latent. The reciprocal relation that Khan and Mäntymäki describe presupposes customers able to run their own assessments. Retail customers instead state conditions for continued trust (disclosure, explanation, human recourse, consent) and would exercise oversight mainly through complaint and exit. In retail banking, downstream governance is therefore mostly a matter of the firm preparing its customers and creating channels through which their expectations can reach it.

AI governance is being grafted onto existing mechanisms. The functions the insurance informants described mostly extend existing risk, compliance, and documentation routines rather than create new ones. In this sector, the model’s components are less often new structures to be built than existing ones to be extended to AI. The strategic question is therefore less whether to take part in the ecosystem than how far to go beyond what existing routines already deliver.

Regulation enables as well as constrains. One insurance informant described EU regulation both as the condition that gives firms the courage to act at all and, in the case of the AI Act, as a source of misguided interpretations and unnecessary bureaucracy. Regulatory oversight thus works not only as a constraint arriving from the polycentric pattern but also as a source of certainty that makes other governance choices possible.

The evidence also shows where Finnish practice currently stops short of the model. Collaborative engagement with standard-setters and civil society was invisible in the insurance interviews; risk alleviation through insurance did not come up in either study; and customers' overseeing role is latent rather than exercised. These are not defects in the model. They are the gaps between the ecosystem as the literature describes it and the ecosystem as the sector currently inhabits it, and they are exactly what a strategy framework should help a firm decide about. Together with the refinements above, they also explain why, as Chapter 5 argues, the live choice for most Finnish banks and insurers in most areas is currently between minimal compliance and managed participation.

5 AI ecosystem strategy framework for banking and insurance

5.1 Purpose and derivation

The ecosystem model describes a structure. The strategy framework turns that structure into choices. Its purpose is to help a bank or insurer decide how it wants to be placed within the responsible AI ecosystem, and to make those decisions explicit enough that they can be discussed, compared across firms, and revisited.

The framework is derived from the model rather than from partner input, and the derivation is visible: each of its seven decision areas corresponds to a component of the model. Five areas map to the five components of the two governance patterns; one maps to the translation relation between layers; one maps to the environmental layer as it changes over time. Table 5 gives the correspondence.

Table 5. The seven decision areas and their derivation from the ecosystem model.

#	Decision area	Model component	Main evidence base
1	External scrutiny posture: how much independent oversight to invite, and from whom	Polycentric: external scrutiny	Khan & Mäntymäki (2027); Heikkilä (2026) on validation and human oversight; standards inventory (ISO/IEC 42001, 42006)
2	Stakeholder engagement design: who is engaged, at which lifecycle stage	Polycentric: collaborative engagement	Khan & Mäntymäki (2027); Heikkilä (2026) on internal and external actors
3	Vendor and upstream governance: due diligence, procurement criteria, contractual control, continuous oversight	Value chain: upstream	Khan & Mäntymäki (2027); standards inventory; the sector's third-party risk regime
4	Customer-facing and downstream governance: disclosure, explanation, recourse, expectation management	Value chain: downstream	Huopalainen (2026); Khan & Mäntymäki (2027); Consumer Credit Directive (2023)
5	Risk transfer and insurability: whether AI risk is insured; for insurers, whether it is underwritten	Value chain: risk alleviation	Khan & Mäntymäki (2027); Heikkilä (2026) on risk assessment and compliance; Eling et al. (2022); insurance market developments

#	Decision area	Model component	Main evidence base
6	Translation capacity: the internal capability to convert external pressure into practice	Cross-layer: institutional translation	Birkstedt et al. (2027); Section 3.2.2 on competing logics
7	Anticipatory capability: a standing routine for scanning what is coming	Environmental layer over time	Jylhä et al. (2026); the scanning pilot in Chapter 6

Each decision area is written the same way: the scope of the decision, its sector context, its research basis, and three positioning options. The options are described next, because they are common to all seven areas.

5.2 Three positioning options

For each decision area the framework offers three positions rather than a maturity score. A maturity scale implies that more is always better; the framework does not assume this. A firm may reasonably choose a minimal position in one area and an active one in another, and the right combination depends on its size, its risk appetite, its dependence on external providers, and its ambitions in the sector. The three positions, which recur under each of the seven areas in Section 5.3, are:

- **Minimal compliance.** The firm does what regulation, supervisory expectation, and contract require, and no more. It treats the ecosystem as a set of fixed rules to be met.
- **Managed participation.** The firm takes part in the ecosystem deliberately: it selects which external relations to invest in, manages them actively, and uses them to reduce its own risk and improve its own governance, without seeking to change the ecosystem's structure.
- **Active shaping.** The firm, alone or with peers, acts on the structure of the ecosystem itself: it helps set a standard, convenes a sector practice, defines what a market for AI assurance or AI insurance looks like, or leads a coalition that others follow.

The first two positions are familiar from compliance and risk management practice. The third is not, and it is the one on which the business ecosystem literature introduced in Section 2.3 has most to say. It is nonetheless consistent with how AI governance is defined more widely: for Wirtz et al. (2022), governance sets rules that promote the use of AI as well as restrict it, and so is not exhausted by the mitigation of risk. The framework offers four tests for choosing a position in each area. Three come from the business ecosystem literature and bear mainly on whether active shaping is realistic: whether shaping can be earned, whether it is worth its cost, and whether the moment is right. The fourth test, drawn from the institutional logics of Chapter 3, bears on all three positions. Where shaping is not realistic, the choice between the other two positions turns mainly on the second and fourth tests. Read against the evidence of Chapter 4, the tests suggest that for most Finnish banks and insurers in most areas the live choice is currently between minimal compliance and managed participation, and that active shaping is realistic in a few areas and mostly collectively, together with peers.

Shaping has to be earned. Active shaping is a bid for ecosystem leadership on a particular issue, and leadership in an ecosystem is conferred rather than declared. Adner (2017) puts it directly: where traditional strategy searches for competitive advantage, ecosystem strategy is “the search for alignment” (p. 49), and the ecosystem leader is simply the firm to whose vision of structure and roles the others defer, so that “successful leadership is thus contingent on willing followership” (p. 48). In the framework's terms, minimal compliance and managed participation are both follower positions, differing in how actively the follower

manages its alignment with others; active shaping is the attempt to become the party others align to. This is a useful corrective. A bank cannot announce itself the shaper of AI assurance practice in Finland; it becomes one only if supervisors, peers, and vendors come to defer to its vision. Active shaping is therefore a position that is earned through the other two, and, as Section 5.4 develops, a move made from managed participation rather than from minimal compliance. Few Finnish banks or insurers can expect that deference on their own, so the realistic form of shaping is often collective. The question is then less whether others will follow the firm's lead than whether they will join it, through a sector association, a joint position toward large providers, or a shared practice convened with peers or research partners. Leadership rests with a coalition rather than a single firm, but the coalition, too, depends on others following it.

Shaping is worth its cost only for some capabilities. Granstrand and Holgersson (2020) describe ecosystem performance as a balance between "growing the pie" across complementary actors and "slicing the pie" among them (p. 9), and their definition of an ecosystem deliberately includes substitute as well as complementary relations. Applied to responsible AI governance, this raises a question firms rarely ask explicitly: is a given governance capability (an audit method, a model-risk practice, a documentation standard) a complement that the sector gains from pooling, or a substitute on which firms compete? Where it is a complement, active shaping grows the pie for everyone, the shaper included, and the cost of building a shared practice is recovered in a better-governed sector. Where it is a substitute, managed participation (building one's own) is the rational choice. The distinction gives the framework a criterion for which areas merit shaping. Most of the active shaping options in Section 5.3 describe a shared practice, standard, or market (sector audit criteria, standardization work, collective procurement positions, joint foresight), where shaping grows the pie. Two are less clear-cut. In customer-facing governance (area 4) and on the insurer's side of risk transfer (area 5), the shaping move is partly a differentiator: a visible practice or a priced product that competes as well as setting a norm. There the firm has to separate the two elements itself: the sector norm or the functioning market is a complement worth building with others, while the firm's own offer within it is a substitute on which it competes.

The moment for shaping is now, but not indefinitely. Tsujimoto et al. (2018) distinguish ecosystems that are designed, where an actor with a strategic intention selects the other participants, from ecosystems that form autonomously (p. 55). FLAIG's earlier finding that no single actor orchestrates the responsible AI ecosystem in Europe (Minkinen et al., 2023) places banking and insurance, for now, in the second category, which is precisely the condition under which a well-placed firm can still shape parts of the structure, before standards, assurance markets, and supervisory expectations settle into a form that later entrants must simply accept. Minkinen et al. (2023) draw the practical conclusion: organizations should "consider this integral role of the emerging RAI ecosystem in their own RAI strategy" (p. 117), deciding which node positions and relationships they seek to obtain, and early adopters "can become RAI champions", gaining through reputation, ecosystem partnerships, and organizational learning. The framework operationalizes that advice by making the choice of position explicit, area by area. Where standards and markets have already settled, the choice falls back to minimal compliance or managed participation.

The fourth test cuts across all three positions and comes not from the ecosystem literature but from the institutional logics of Section 3.2.2: each position is easier to hold from some logics than from others. Minimal compliance is native to regulatory oversight and, in its cost-minimizing form, to managerial rationalism. Managed participation is native to IT

professionalism, which values being in control of one's own systems and relationships. Active shaping draws on AI ethics principlism and on the reputational strand of managerial rationalism. A firm choosing a position should ask which logics dominate its own decision-making, because a position that no internal logic supports is likely to be loosely coupled to practice in the way Berente and Yoo (2012) describe: announced, but not enacted. Area 6 (translation capacity) treats this fit between position and logic as a decision in its own right.

5.3 The seven decision areas

5.3.1 Area 1: External scrutiny posture

Scope. How much independent oversight of its AI the firm invites beyond what is required, and from whom. Scrutiny is required by supervisors and, for high-risk systems under the AI Act, through conformity assessment and market surveillance. For the credit and insurance-pricing systems typical of the sector, however, the AI Act provides for conformity assessment by the provider itself under internal control, without a notified body. Beyond that, a firm may seek third-party audit of its AI management system, external review of its own evaluation methods, or independent assessment of specific systems.

Sector context. Banking and insurance are already among the most scrutinized sectors. The actor map (Table 4) lists prudential and conduct supervisors, statutory auditors, AI Act market surveillance, and notified bodies. The question is not whether to be scrutinized but how to coordinate scrutiny, and whether to use additional, voluntary scrutiny as a source of assurance and reputation. The standards inventory gives the instruments: ISO/IEC 42001 defines an AI management system that can be certified, and ISO/IEC 42006 sets requirements for the bodies that audit and certify against it (Appendix A). Together they create the possibility of a certified assurance market that is only now taking shape. Because AI Act conformity for the sector's typical systems rests on self-assessment, voluntary third-party assurance is the main route to independent scrutiny of a firm's AI beyond supervision. The quality of that market is not guaranteed. Research on the emerging AI audit ecosystem finds widely varying practices, few agreed standards for what an audit must cover, and limited disclosure of results, and argues for standards and oversight for auditors themselves (Raji et al., 2022; Costanza-Chock et al., 2022); the AI Act's reliance on internal checks has been criticized on similar grounds (Mökander et al., 2022). Shared sector audit criteria are therefore a complement that firms gain from building together, in the sense of Section 5.2.

Research basis. The insurance interviews show validation and human oversight treated as non-negotiable internally (Heikkilä, 2026), but show no use of external, voluntary scrutiny; regulation was the scrutinizer that mattered. Khan and Mäntymäki (2027) note that "mere internal oversight is insufficient for AI governance" and that external scrutiny "can span the entire lifecycle", including assessment of the organization's own evaluation method.

Positions.

- **Minimal compliance:** meet supervisory and AI Act requirements; no voluntary external assessment.
- **Managed participation:** seek ISO/IEC 42001 certification or equivalent third-party review for selected high-risk systems; use external review of internal evaluation methods; treat audit findings as governance inputs.
- **Active shaping:** participate in defining what AI assurance in the sector should consist of (for example, by working with supervisors and audit bodies on sector-specific audit

criteria, or by publishing assurance results) so that the firm's practice becomes a reference for others.

5.3.2 Area 2: Stakeholder engagement design

Scope. Which external stakeholders the firm deliberately engages in governing its AI, at which stages of the lifecycle, and through which channels. Candidates include standard-setting bodies, industry associations, academic partners, consumer organizations, and affected groups.

Sector context. Engagement is the component through which the firm itself gives external scrutiny a collaborative form (Section 3.3.1). In banking and insurance, industry associations and standard-setters are established channels, but engagement with affected groups and civil society on AI is thin. The AI Act's requirements for high-risk systems, including human oversight, transparency to deployers, and, for deployers of creditworthiness assessment and life and health insurance pricing systems specifically, a fundamental rights impact assessment (Article 27), create occasions for engagement that firms can either treat as paperwork or use.

Research basis. In the insurance interviews, external actors beyond regulators and technology providers were mentioned rarely; academia, civil society, and standards bodies did not appear at all (Heikkilä, 2026). Internally, however, engagement was rich: cross-functional forums were the main organizational function identified. Khan and Mäntymäki (2027) list the lifecycle points at which engagement pays: consulting experts when benchmarking and assessing risk, collecting feedback on capability and safety, and joint monitoring after deployment.

Positions.

- **Minimal compliance:** engagement limited to what regulation requires (e.g., consultation obligations in impact assessment).
- **Managed participation:** a defined set of external partners (an academic partner for evaluation, an industry association for shared interpretation of regulation, structured customer or consumer-organization feedback on selected systems) engaged at defined lifecycle stages.
- **Active shaping:** participation in standardization (ISO/IEC JTC 1/SC 42, CEN-CENELEC JTC 21) or in convening a sector forum on responsible AI, so that the firm contributes to the norms it will later be held to.

5.3.3 Area 3: Vendor and upstream governance

Scope. How the firm governs the providers on which its AI depends: cloud and infrastructure providers, foundation-model providers, application vendors, consultancies, and data providers. The instruments are due diligence and procurement criteria, contractual control, tailoring of procured systems to the firm's governance needs, continuous oversight, and limited access as a mitigation for component suppliers.

Sector context. The sector has a mature third-party risk regime, reinforced by DORA for ICT providers, and one informant in the insurance interviews described DORA as constraining which partners an insurer may cooperate with (Heikkilä, 2026). AI adds two difficulties. First, the most consequential upstream actors, foundation-model providers, are few, large, and largely outside the firm's negotiating power; the actor-network study observed that "regulators, large technology providers, and well-resourced organizations have substantially greater capacity to shape the network than smaller actors" (Birkstedt et al., 2027). Second,

generative AI capabilities in the sector are substantially bought in, so that safeguards are partly supplied by providers and “responsibility is partially outsourced and co-produced across organizational boundaries” (Heikkilä, 2026, p. 50).

Research basis. Khan and Mäntymäki (2027) characterize upstream governance as one-directional oversight and stress that pre-procurement due diligence alone is inadequate: continuous oversight after procurement is needed to keep vendor practice aligned with the firm’s standards. The insurance interviews show technology providers credited with doing much of the responsibility work themselves, which is a description of dependence as much as of reassurance. Retail banking customers, for their part, asked banks to be transparent about their third-party providers (Huopalainen, 2026), which connects this area to area 4. The literature on interorganizational governance offers a vocabulary for the instruments involved. It distinguishes contractual mechanisms, meaning enforceable rights and obligations, from relational ones, meaning trust, informal exchange, and shared norms, and finds that the two work as complements rather than substitutes (Poppo & Zenger, 2002). Wirtz et al. (2022) add technological mechanisms for AI: rules built into the systems themselves. The positions below can be read in these terms. Minimal compliance relies on contract. Managed participation adds technological mechanisms, through the tailoring of procured systems. Active shaping is largely relational, because a provider that no single firm can bind by contract has to be moved by other means.

Positions.

- **Minimal compliance:** AI providers handled within the existing outsourcing and ICT third-party risk process; AI-specific criteria limited to what the AI Act requires of deployers.
- **Managed participation:** AI-specific procurement criteria (documentation, testing evidence, data provenance, model update policy), contractual rights to information and audit, tailoring of procured systems to the firm’s governance needs, and continuous oversight of vendor practice.
- **Active shaping:** collective action on upstream terms, through industry associations or joint procurement positions, to obtain from large providers the documentation and assurances that no single firm can obtain alone.

5.3.4 Area 4: Customer-facing and downstream governance

Scope. How the firm prepares its customers for AI-enabled services and what it offers them in return: disclosure of where AI is used, explanation of decisions, routes of recourse, and management of expectations. The reciprocal side of the decision is whether and how the firm collects and uses what customers feed back.

Sector context. These are high-trust services in which customers’ financial lives are at stake, and the AI Act’s high-risk categories cover precisely the customer-facing decisions (creditworthiness, life and health insurance pricing) where expectations are strongest. For consumer credit, the requirement customers voiced most often is also becoming law: under the revised Consumer Credit Directive, applicable from 20 November 2026 and transposed in Finland by amendments to the Consumer Protection Act (31/2026), a consumer whose creditworthiness is assessed through automated processing of personal data has the right to request human intervention, comprising an explanation of the assessment, the opportunity to express their point of view, and a review of the assessment and the credit decision (Directive (EU) 2023/2225, Art. 18(8)). The right is not wholly new: the GDPR already gives a right to human intervention, to express one’s point of view, and to contest a decision based solely on

automated processing (Regulation (EU) 2016/679, Art. 22(3)). What the Directive adds is a credit-specific right that applies whenever the assessment involves automated processing, not only when the decision is fully automated. Reputational risk in this area is immediate: the insurance interviews recorded that customer awareness of AI use is itself a reputational risk and that transparency “sometimes causes negative reactions” (Heikkilä, 2026, p. 57).

Research basis. The retail banking evidence (Section 4.2) is specific. Customers’ ethical and relational expectations were firmer than their competence expectations, and they concentrated on four things: open communication about that and where AI is used; explanations after negative outcomes that include instructions on how to proceed; a human route of recourse, including the option of human service from the outset; and transparency about data use and provenance. Consent and the continued availability of human service were the conditions most often attached to acceptance, and acceptable AI autonomy scaled with stakes. Legal disclaimers were not accepted as a way of transferring liability. Khan and Mäntymäki (2027) frame the relation as reciprocal, with customers as hybrid actors who transact and oversee; the evidence suggests that in Finnish retail banking the overseeing role is latent and is likely to be exercised through complaint and exit unless the firm creates other channels. The demand for human recourse in automated credit decisions broadly anticipates the right that the revised Consumer Credit Directive introduces, so that part of what customers expect is moving from managed participation into minimal compliance, a move the GDPR had already begun for fully automated decisions.

Positions.

- **Minimal compliance:** AI Act transparency obligations and existing consumer-protection duties, including, for consumer credit, the human-intervention right of the revised Consumer Credit Directive; otherwise recourse through the ordinary complaints process.
- **Managed participation:** a customer-facing AI policy that discloses AI use by service, explanation formats designed for negative outcomes, a guaranteed human route with defined response times, stakes-based rules for which decisions are automated, and structured collection of customer feedback on AI-enabled services.
- **Active shaping:** making customer-facing AI practice a visible differentiator by publishing testing results or assurance summaries, offering opt-in human service as a product feature, and working with consumer organizations on sector norms for explanation and recourse.

5.3.5 Area 5: Risk transfer and insurability

Scope. For a bank: whether and how its AI-related risks are insured, and what the insurer’s assessment and pricing tell it about its own governance. For an insurer: whether AI risk is a line of business it will underwrite, on what evidence, and with what effect on the governance of the insured.

Sector context. This is the one area in which banking and insurance meet on the two sides of a single governance mechanism: banks, like other deployers of AI, are potential buyers of AI cover, while non-life insurers are its potential sellers. Khan and Mäntymäki (2027) identify risk alleviation through insurance as a component of value chain governance in its own right, and also place insurers within collaborative stakeholder engagement: insurers “evaluate the potential risks posed by the AI system” and “offer corresponding insurance packages”, and “as a result, organizations are incentivized to reduce risks from their AI development and deployment”. Whether risk can be transferred to insurers, and with what effect on the

insured's internal governance, they leave open for future research; this decision area is where the report takes the question up. The mechanism is familiar from other risk classes: if the insurer prices well, the premium disciplines the insured. What is new is its application to AI, where the underwriter's assessment of a firm's AI governance becomes an external judgment that carries a price.

Research basis. The FLAIG evidence is thin. The insurance interviews show a strong internal risk-assessment culture ("all AI must have its risks assessed and then mitigated"), but no informant discussed underwriting other firms' AI risk (Heikkilä, 2026). Khan and Mäntymäki (2027) themselves list interviews with insurance companies among the primary data most needed, and observe that insurers play a hybrid role, assessing risk and partially absorbing it. The framework therefore presents this area as the one where its positions rest on the thinnest evidence.

Outside the project, the market is moving in two directions at once. Some insurers have begun to introduce exclusions for AI-related losses into existing policies, so that AI risk that used to sit silently within existing lines is being carved out; at the same time, affirmative cover for AI risk, while still sparse, exists (Munich Re, 2026). Munich Re has insured the performance risk of AI models since 2018, and a specialist Lloyd's coverholder now offers standalone AI liability policies with limits of up to USD 25 million (Armilla AI, 2026). Insurance research gives reasons to expect underwriting to remain difficult: AI may shift some risks from high-frequency, low-severity to low-frequency, high-severity events, which requires insurers to rethink traditional cover (Eling et al., 2022). For the framework, this has two consequences. For a bank, the minimal position of leaving AI risk within existing cover is becoming less tenable as exclusions spread. For an insurer, the emerging affirmative market shows that underwriting AI risk on the basis of governance evidence is feasible, but also that it is at an early stage.

Positions.

- **Bank, minimal compliance:** AI risk left within existing operational and cyber cover, without AI-specific assessment, a position that holds only as long as that cover does not exclude AI.
- **Bank, managed participation:** AI-related exposures identified and discussed with insurers; the insurer's assessment used as an external check on governance.
- **Bank, active shaping:** working with insurers to define what evidence of AI governance should support insurability, so that the bank's practice sets the reference.
- **Insurer, minimal compliance:** AI risk excluded or treated silently within existing lines.
- **Insurer, managed participation:** AI risk underwritten selectively, with underwriting criteria drawn from governance evidence such as management-system certification and documentation.
- **Insurer, active shaping:** developing AI risk as a product with explicit governance-based pricing, thereby becoming a governance actor for the sector.

5.3.6 Area 6: Translation capacity

Scope. How the firm organizes its capacity to convert external pressure into internal practice: whether it has a designated AI governance function or committee with decision rights, how it handles conflicts between logics, and whether it manages compliance at the level of individual systems or through organization-wide rules.

Sector context. Every other decision area generates external relations; this one determines whether those relations change anything inside the firm. The actor-network study of Nordic banks and insurers found that organizations with “a designated AI governance function or a multi-disciplinary AI governance committee with explicit decision rights” achieved durable arrangements, while those without “experienced repeated re-negotiation of the same governance issues” (Birkstedt et al., 2027). It also found that compliance is managed more effectively “at the level of individual AI systems or use cases than through organization-wide rule-making”. Our layered institutional perspective adds that the outcome of translation depends on who enacts which logic, how well roles are performed, whether conflicts resolve into polarization, compromise, or synthesis, and whether hybrid professionals are available to broker (Section 3.5).

Research basis. The insurance interviews describe exactly the anchoring structure the model calls for: cross-functional review forums combining legal, compliance, risk, and business expertise, with responsibility following AI use into the business units (Heikkilä, 2026). They also describe the friction: regulation-driven decision-making was perceived as rigid and slow, and one informant saw EU regulation as both enabling and, in the AI Act’s case, bureaucratic. In the model’s terms this is regulatory oversight meeting managerial rationalism, and the forum is where the negotiation happens.

Positions.

- **Minimal compliance:** AI governance assigned to an existing function (compliance, risk, or IT) without dedicated decision rights; conflicts escalated case by case.
- **Managed participation:** a designated AI governance function or committee with explicit decision rights, system-level compliance management, and deliberate staffing with hybrid profiles able to translate between technical, legal, and business logics.
- **Active shaping:** the translation capacity made visible outside the firm, with the governance function acting as the firm’s spokesperson toward supervisors, standard-setters, and peers, and the firm’s translation practice offered as a model through a sector forum or the WP1 governance framework.

5.3.7 Area 7: Anticipatory capability

Scope. Whether the firm maintains a standing routine for scanning developments in AI technology, regulation, and the ecosystem, and how the outputs of that routine enter governance decisions.

Sector context. The ecosystem drawn in Figure 1 is moving. Generative AI has changed the upstream component within three years; AI agents are changing the AI system layer; the AI Act’s phased application changes the scrutiny component each year through 2028, and its timetable has already been revised once; and standards are published in a steady stream (Appendix A lists seven under development). The expert-interview study found that generative AI “has outpaced governance” and that experts expected governance and technology to continue co-evolving (Jylhä et al., 2026). This is the pacing problem familiar from the governance of emerging technologies. The response proposed for it at the policy level is a recursive governance process, restarted whenever monitoring shows a need to reframe (Wirtz et al., 2022); this decision area asks whether the firm has an equivalent of its own. A study of nine Finnish public sector organizations that prepared for the AI Act while it was still being negotiated shows what this looks like from inside an organization: regulation reshaped what the organizations saw as possible long before it applied, and their uncertainty did not disappear once the text was final (Vainionpää et al., 2026). Their responses ranged

from informal watch by an individual, through a designated person or team following the process, to attempts to influence the outcome through ministries and system providers, a gradient close to the three positions below. The authors also conclude that waiting is not neutral: anticipatory capability is cumulative, and wait-and-see approaches can narrow later options by holding back capability development. Banks and insurers are in a comparable period now, with the AI Act in force but its high-risk obligations postponed.

Research basis. Chapter 6 describes the horizon-scanning process piloted in the project and its outputs. The insurance interviews, meanwhile, showed a gap between experimenting with AI and implementing it, and a noted lack of management strategy in at least one organization (Heikkilä, 2026). These are conditions under which scanning tends to happen informally, through vendors and consultants, rather than as a governed routine.

Positions.

- **Minimal compliance:** regulatory watch through compliance and legal functions; technology watch left to vendors.
- **Managed participation:** an annual, lightweight horizon-scanning routine of the kind described in Chapter 6, with a defined owner, a defined set of internal and external informants, and a defined path from scanning output to governance and strategy decisions.
- **Active shaping:** participation in shared sector foresight (joint scanning with peers, associations, or research partners) so that the firm helps to define what the sector treats as emerging and how it prepares.

5.4 Using the framework

The framework is used by working through the seven areas and recording, for each, the firm's current position, its intended position, and the reasons. Appendix D provides this as a one-page worksheet. Three points of guidance apply.

First, positions are independent across areas. It is coherent for a mid-sized insurer to hold a managed position on scrutiny and translation and a minimal position on stakeholder engagement while considering active shaping on risk transfer, because that is where its capabilities and the ecosystem's openness coincide. The framework is designed to make such profiles visible, not to push every area to the right-hand column.

Second, the positions are cumulative in the sense that Adner's leadership argument implies. Active shaping in an area presupposes that the firm already manages its participation there; a firm cannot credibly help define sector audit criteria while its own external scrutiny is minimal. In practice this means that active shaping should be planned as a move from managed participation, not from minimal compliance.

Third, the worksheet's most useful output is often the *reasons* column. When a firm records that its position on upstream governance is minimal because its foundation-model provider will not negotiate, it has identified a structural feature of the ecosystem (a power asymmetry) that individual action will not change and that collective action (area 3, active shaping) might. When it records that its position on risk transfer is minimal because no insurer offers a product, it has identified a gap in the ecosystem that a domestic non-life insurer may be positioned to fill.

6 Piloting a horizon-scanning process

The model in Chapter 3 describes the ecosystem at one point in time, and Section 3.6 noted that its structure is moving. This chapter turns to how a firm can keep track of that movement. It documents the horizon-scanning process piloted in the project to enable continuous foresight of the responsible AI ecosystem. The pilot was conducted as a research study rather than as a corporate routine, and this chapter sets out its steps (how informants were identified and recruited, what was asked, how the material was analyzed, and what came out) as a repeatable routine, so that a bank or insurer can run a lighter version of it annually, at the managed-participation position in decision area 7 (Section 5.3.7). The full findings of the pilot are reported in a paper accepted for presentation at ISPIM Connects Sendai 2026 (Jylhä et al., 2026); here they are summarized only as the pilot's output.

6.1 What was piloted

The pilot took the form of an expert-elicitation study on the future of generative AI and its governance. Horizon scanning, as the term is used in futures research, is the systematic examination of potential developments at the margins of current thinking, with the aim of informing decisions before those developments become pressing. Expert interviews are one established method for it, and they were chosen here because the developments of interest (generative AI, AI agents, and the governance responses to them) were moving too fast in 2025 for published documents alone to keep up with. Vainionpää et al. (2026) identify four sets of practices through which organizations anticipate regulation in the making; horizon scanning formalizes the first of these, forming expectations, which organizations otherwise carry out informally by reading drafts, following news and peers, and comparing new rules with familiar ones.

Informants. Twenty-three professionals in AI development, policy, and management were interviewed between March and October 2025. They were recruited purposively from the researchers' networks and LinkedIn and by snowball sampling; 36 were approached, and recruitment continued until thematic saturation was reached. All had more than six years of professional experience with AI, generative AI, or machine learning and were professionally active in Europe. Roles ranged from data analysts and AI engineers to heads of AI, chief analytics officers, directors of government relations, and chief executives of AI companies and consultancies. Informants were not drawn specifically from banking and insurance: the pilot scanned the AI and governance environment the sector shares with other industries, and Section 6.2 reads its outputs for what they imply for the sector's decision areas.

Protocol. Interviews were semi-structured, lasted about an hour, and were conducted remotely in English or Finnish. The guide covered the shift from narrow AI to generative AI and beyond, AI governance frameworks and mechanisms at industry level, the limitations, concerns, and future potential of generative AI, and societal and ethical consequences.

Analysis. Interviews were recorded, transcribed, anonymized, and analyzed thematically in six phases. Coding combined inductive and deductive elements and produced 139 hierarchical codes, reduced to 17 and then to five main themes with fourteen subthemes. Coding was carried out by one researcher and reviewed with three others, with disagreements resolved in discussion.

6.2 What came out

The pilot's output is organized by the five themes. Three of them describe *limitations* of generative AI as experts currently perceive them: limited validity (applicability and knowledge constraints on organizational use), limited utility (governance, regulatory, and technological constraints), and limited choice (fairness, bias, and geopolitical constraints). The fourth describes *concerns* that these limitations give rise to: for work and the economy, for cognition, and for technological dependence. The fifth describes *potential future benefits* and the directions in which experts expect technology and governance to move.

Four findings stand out as most relevant to banking and insurance. First, generative AI has outpaced governance: experts described governance mechanisms as fragmented, in some organizations well developed and in others absent, and current regulatory guidance as insufficient for generative AI. Second, generative AI is nonetheless viewed as a net organizational benefit, with automation and job transformation seen as inevitable but not inherently threatening; experts expected human-machine collaboration and the need for humans in the loop to increase, not decrease. Third, adoption remains early, constrained by limited knowledge, governance uncertainty, and unresolved fairness and geopolitical questions. Fourth, experts expected further regulation on transparency, data ownership, and privacy, expected model providers to have to disclose more about how systems are developed and trained, and described a movement “from compliance to a more all-encompassing, mindful GenAI use” in which AI literacy becomes part of governance.

The study also proposes a process model of anticipatory governance in which perceived limitations give rise to concerns, concerns trigger anticipatory sensemaking about future development, and that sensemaking shapes governance that is proactive and adaptive rather than reactive. For the strategy framework, the model supplies the argument for treating anticipatory capability as a decision area in its own right: if proactive governance depends on concerns being turned into anticipatory sensemaking, that step cannot be left to informal watch by compliance functions and vendors, as the minimal-compliance position in Section 5.3.7 does. The findings also inform three other decision areas. Insufficient regulatory guidance bears on external scrutiny posture (area 1); expected disclosure duties for model providers bear on vendor and upstream governance (area 3); and fragmented governance mechanisms, together with the turn to AI literacy, bear on translation capacity (area 6). Mapping scanning output onto decision areas in this way is how it enters the framework, and it is the step that the routine in Section 6.3 makes regular.

Read against the model, the findings also connect the pilot to the sector evidence of Chapter 4. The expected disclosure duties for model providers bear directly on the first refinement in Section 4.3, that upstream relations involve dependence as well as oversight: if such duties materialize, they would give banks and insurers more of the information on which oversight of providers depends. The turn to AI literacy matches the insurance informants' inclusion of training and AI literacy among the ecosystem functions (Section 4.1). And the finding that generative AI has outpaced governance is the pacing problem behind the second qualification in Section 3.6: the ecosystem the model describes is changing faster than the governance arrangements within it.

6.3 Running a lighter version annually

The pilot also showed what the full research design demands. Recruiting 23 informants from 36 approached, interviewing them over eight months, and reducing 139 codes to five themes across four researchers produced rich material, but on a timescale and with resources that few

firms could sustain every year. And because the informants were not drawn from the sector, their outputs had to be read for banking and insurance after the fact. The lighter version below addresses both points.

A bank or insurer does not need 23 interviews and a research team to keep a scanning routine alive. A lighter version that preserves the pilot's structure would consist of five steps:

1. An owner is named: typically the AI governance function, which at the managed-participation position in decision area 6 is the firm's translation point; where area 6 is held at minimal compliance and no such function exists, the routine falls to the existing function to which AI governance is assigned.
2. A panel of ten to fifteen informants is assembled, mixing internal experts (data science, compliance, product, IT) with external ones (a vendor, an academic partner, a supervisor contact, a peer).
3. A short protocol is used, condensed from the pilot's interview guide into four questions: what has changed in AI capability, what has changed in governance and regulation, what limitations and concerns are emerging, and what developments the informant expects within two years.
4. The material is analyzed by a small group using a stable coding frame derived from the pilot's themes, so that years can be compared, plus an open category for material that fits none of them, so that the frame does not filter out the marginal developments that scanning exists to detect.
5. The output is written as a one-page brief that enters the firm's governance and strategy cycle at a defined point, for example the firm's periodic review of the AI risk assessments that one insurance informant described as mandatory for all AI (Section 4.1).

Run once a year and compared with the previous year, this routine discharges decision area 7 at the managed-participation position. Run jointly with peers or a research partner, it moves toward active shaping, subject to the tests set out in Section 5.2. The pilot's own 2025 findings are the natural baseline for the first comparison, read with care: they come from a larger, cross-industry panel and a fuller analysis than the lighter routine will produce.

7 Implications and limitations

7.1 Implications for banks and insurers

The report's central implication for the sector can be stated simply. A bank's or insurer's responsible AI position is not fully under its own control, because it depends on supervisors, vendors, insurers, customers, and standard-setters; but it is not outside its control either, because the firm can choose how it is placed among those actors. The ecosystem model shows the structure; the strategy framework turns it into seven decisions. Three of those decisions deserve emphasis for the Finnish sector in particular. Upstream governance is where dependence on a few large model providers is greatest and where individual firms have least leverage, which makes it the strongest case for collective action through associations and joint positions. Downstream governance is where customer expectations are firmest and most specific, and the evidence suggests that meeting them is largely a matter of disclosure, explanation, recourse, and consent rather than of technical performance; for consumer credit, it is partly a matter of law from November 2026. Risk transfer is where banking and insurance meet: as exclusions for AI spread through existing policies, a Finnish

non-life insurer that underwrites AI risk on the basis of governance evidence would become a governance actor for the sector. In all three areas the structure is still moving, so positions chosen today will need revisiting. The annual scanning routine in Chapter 6 is a lightweight way of knowing when, and it gives the seventh decision area, anticipatory capability, a concrete form.

7.2 Implications for the FLAIG project

For Work Package 1, the model locates the governance framework and auditing guidelines at the organizational layer and identifies translation capacity (decision area 6) as the point where WP1's instruments and WP2's ecosystem meet. The layered institutional perspective paper's observation that ecosystem pressure produces uneven organizational responses is one reason WP1's framework needs to be fit for purpose rather than uniform, and Birkstedt et al.'s finding that compliance is best managed at system level is directly usable in the auditing guidelines. For Work Package 3, the AI system layer is where quantitative measures of trustworthiness enter the ecosystem: a measure of explainability or stability is, in the model's terms, an artifact through which the system layer reports upward to the organization and, through it, to scrutinizers and customers. The two deliverables of this report therefore connect to the other work packages at specific, nameable points rather than in general.

7.3 Implications for research

The report contributes to IS research on AI governance in three ways. It gives the interorganizational patterns identified by Khan and Mäntymäki (2027) an initial sector-level illustration, drawing on primary data of the kind their paper calls for, and uses that evidence to refine them (Section 4.3). It combines those patterns with the layered institutional perspective developed in the project into a single model that has both a vertical and a horizontal axis, which neither source has alone. And it derives a strategy framework from that model using business ecosystem theory, so that ecosystem positioning becomes a set of explicit choices rather than an aspiration. Three research directions follow. The risk alleviation component, and whether it amounts to risk transfer, needs empirical study on the insurer side, where the project's own studies found no evidence. The hybrid, overseeing role of customers needs study at the organizational level, to see whether and how customer expectations become governance pressure, a question the retail banking thesis explicitly leaves open. And the framework itself, including the annual scanning routine in Chapter 6, needs evaluation in use, which falls outside the scope of this project.

7.4 Limitations

The report has five limitations that should be read alongside its claims. First, the evidence is small and qualitative: six insurance professionals, ten retail banking customers, 27 Nordic experts on governance, and 23 European experts on generative AI. It illustrates the model and refines it; it does not test it. Second, the model is a synthesis of existing conceptual work, and it inherits the limitations of its sources: the interorganizational patterns were derived from literature published between 2022 and March 2025, and the layered institutional perspective is conceptual, untested, and not yet peer-reviewed. Third, the strategy framework has not been validated by practitioner use, and the lighter scanning routine in Chapter 6 is derived from a research study whose informants were not drawn from the sector, rather than tested inside a bank or insurer. Fourth, the actor map reflects the regulatory position in autumn 2026, during the AI Act's phased application; supervisory competences and the standards landscape will continue to change, and the map should be read as dated. Fifth, the model and

framework were assembled by a single author from the project's studies, and the business ecosystem concepts used in Chapters 2 and 5 were developed for commercial value propositions and are applied here by extension (Section 2.3); other syntheses of the same material are possible.

References

Act on the Supervision of Certain Artificial Intelligence Systems (1377/2025) [Laki eräiden tekoälyjärjestelmien valvonnasta]. *Suomen säädöskokoelma*, 29 December 2025. <https://www.finlex.fi/fi/lainsaadanto/saadskokoelma/2025/1377>

Adner, R. (2017). Ecosystem as structure: An actionable construct for strategy. *Journal of Management*, 43(1), 39–58. <https://doi.org/10.1177/0149206316678451>

Armilla AI. (2026, January 21). *Armilla AI raises Lloyd's-backed coverage to \$25M as traditional insurers retreat from AI risk* [Press release]. <https://www.armilla.ai/resources/armilla-ai-raises-lloyds-backed-coverage-to-25m-as-traditional-insurers-retreat-from-ai-risk>

Berente, N., & Yoo, Y. (2012). Institutional contradictions and loose coupling: Postimplementation of NASA's enterprise information system. *Information Systems Research*, 23(2), 376–396. <https://doi.org/10.1287/isre.1110.0373>

Bernardi, R., & Exworthy, M. (2020). Clinical managers' identity at the crossroad of multiple institutional logics in IT innovation: The case study of a health care organization in England. *Information Systems Journal*, 30(3), 566–595. <https://doi.org/10.1111/isj.12267>

Birkstedt, T., Minkkinen, M., & Mäntymäki, M. (2027). Governance by negotiation: How organizations translate AI ethics into practice. In *Work, life, and society in a sustainable, intelligent future* (Lecture Notes in Computer Science, Vol. 16912, pp. 301–312). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-38014-2_22

Birkstedt, T., Minkkinen, M., Tandon, A., & Mäntymäki, M. (2023). AI governance: Themes, knowledge gaps and future agendas. *Internet Research*, 33(7), 133–167. <https://doi.org/10.1108/INTR-01-2022-0042>

Boonstra, A., Eseryel, U. Y., & van Offenbeek, M. A. G. (2018). Stakeholders' enactment of competing logics in IT governance: Polarization, compromise or synthesis? *European Journal of Information Systems*, 27(4), 415–433. <https://doi.org/10.1057/s41303-017-0055-0>

Callon, M. (1984). Some elements of a sociology of translation: Domestication of the scallops and the fishermen of St Briec Bay. *The Sociological Review*, 32(1_suppl), 196–233. <https://doi.org/10.1111/j.1467-954X.1984.tb00113.x>

Costanza-Chock, S., Raji, I. D., & Buolamwini, J. (2022). Who audits the auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1571–1583). <https://doi.org/10.1145/3531146.3533213>

Eling, M., Nuessle, D., & Staubli, J. (2022). The impact of artificial intelligence along the insurance value chain and on the insurability of risks. *The Geneva Papers on Risk and Insurance — Issues and Practice*, 47(2), 205–241. <https://doi.org/10.1057/s41288-020-00201-7>

European Commission. (2020). *White paper on artificial intelligence: A European approach to excellence and trust* (COM(2020) 65 final). European Commission.

European Parliament & Council of the European Union. (2016). Regulation (EU) 2016/679 of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union*, L 119, 1–88.

European Parliament & Council of the European Union. (2023). Directive (EU) 2023/2225 of 18 October 2023 on credit agreements for consumers and repealing Directive 2008/48/EC. *Official Journal of the European Union*, L 2023/2225.

European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*.

European Parliament & Council of the European Union. (2026). Regulation (EU) 2026/1744 amending Regulation (EU) 2024/1689 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI). *Official Journal of the European Union*, 24 July 2026.

Faik, I., Barrett, M., & Oborn, E. (2020). How information technology matters in societal change: An affordance-based institutional logics perspective. *MIS Quarterly*, 44(3), 1359–1390. <https://doi.org/10.25300/MISQ/2020/14193>

Finanssivalvonta [Financial Supervisory Authority]. (2025). *The use of artificial intelligence by financial sector actors* [Thematic review, 27 June 2025]. https://www.finanssivalvonta.fi/globalassets/en/publications/supervision-releases/2025/the-use-of-artificial-intelligence-by-financial-sector-actors_en_.pdf

Gasser, U., & Almeida, V. A. F. (2017). A layered model for AI governance. *IEEE Internet Computing*, 21(6), 58–62. <https://doi.org/10.1109/MIC.2017.4180835>

Granstrand, O., & Holgersson, M. (2020). Innovation ecosystems: A conceptual review and a new definition. *Technovation*, 90–91, 102098. <https://doi.org/10.1016/j.technovation.2019.102098>

Hansen, S., & Baroody, A. J. (2020). Electronic health records and the logics of care: Complementarity and conflict in the U.S. healthcare system. *Information Systems Research*, 31(1), 57–75. <https://doi.org/10.1287/isre.2019.0875>

Heikkilä, S. (2026). *Responsible AI ecosystems: Insights from the insurance industry* [Master's thesis, Information Systems Science, University of Turku].

Huopalainen, J. (2026). *Customer expectations of AI in retail banking: A vignette-based qualitative study* [Master's thesis, Information Systems Science, University of Turku].

Jylhä, H., Minkinen, M., & Mäntymäki, M. (2026). Rage against or with the machine? From perceived limitations, concerns, and future potential of GenAI to anticipatory governance [Paper accepted for presentation]. ISPIIM Connects Sendai: Innovating for a Resilient Future, Sendai, Japan, 30 November–2 December 2026. To appear in the conference proceedings.

Khan, F., & Mäntymäki, M. (2027). Interorganizational AI governance in academic and grey literature. In *Work, life, and society in a sustainable, intelligent future* (Lecture Notes in

Computer Science, Vol. 16912, pp. 326–336). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-38014-2_24

Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. (2022). Defining organizational AI governance. *AI and Ethics*, 2, 603–609. <https://doi.org/10.1007/s43681-022-00143-x>

Minkkinen, M., Jylhä, H., & Mäntymäki, M. (2026). AI governance: A layered institutional perspective [Unpublished working paper]. FLAIG project, University of Turku.

Minkkinen, M., & Mäntymäki, M. (2023). The institutional logics underpinning organizational AI governance practices. In *Proceedings of the 14th Scandinavian Conference on Information Systems (SCIS 2023)*, paper 12. <https://aisel.aisnet.org/scis2023/12>

Minkkinen, M., Zimmer, M. P., & Mäntymäki, M. (2023). Co-shaping an ecosystem for responsible AI: Five types of expectation work in response to a technological frame. *Information Systems Frontiers*, 25(1), 103–121. <https://doi.org/10.1007/s10796-022-10269-2>

Mökander, J., Axente, M., Casolari, F., & Floridi, L. (2022). Conformity assessments and post-market monitoring: A guide to the role of auditing in the proposed European AI regulation. *Minds and Machines*, 32(2), 241–268. <https://doi.org/10.1007/s11023-021-09577-4>

Munich Re. (2026). *AI risks and Munich Re's AI insurance solutions* [Client seminar programme, 14 October 2026]. <https://www.munichre.com/en/company/seminars/client-programme/2026/casualty-ai-risks-and-munichre-ai-insurance-solutions-14oct-am.html>

Ostrom, E. (2010). Beyond markets and states: Polycentric governance of complex economic systems. *American Economic Review*, 100(3), 641–672. <https://doi.org/10.1257/aer.100.3.641>

Poppo, L., & Zenger, T. (2002). Do formal contracts and relational governance function as substitutes or complements? *Strategic Management Journal*, 23(8), 707–725. <https://doi.org/10.1002/smj.249>

Qiu, Y., Gopal, A., & Hann, I.-H. (2017). Logic pluralism in mobile platform ecosystems: A study of indie app developers on the iOS App Store. *Information Systems Research*, 28(2), 225–249. <https://doi.org/10.1287/isre.2016.0664>

Raji, I. D., Xu, P., Honigsberg, C., & Ho, D. (2022). Outsider oversight: Designing a third party audit ecosystem for AI governance. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 557–571). <https://doi.org/10.1145/3514094.3534181>

Shneiderman, B. (2020). Bridging the gap between ethics and practice: Guidelines for reliable, safe, and trustworthy human-centered AI systems. *ACM Transactions on Interactive Intelligent Systems*, 10(4), 1–31. <https://doi.org/10.1145/3419764>

Stahl, B. C. (2022). Responsible innovation ecosystems: Ethical implications of the application of the ecosystem concept to artificial intelligence. *International Journal of Information Management*, 62, 102441. <https://doi.org/10.1016/j.ijinfomgt.2021.102441>

Stahl, B. C. (2023). Embedding responsibility in intelligent systems: From AI ethics to responsible AI ecosystems. *Scientific Reports*, 13, 7586. <https://doi.org/10.1038/s41598-023-34622-w>

- Thornton, P. H., & Ocasio, W. (1999). Institutional logics and the historical contingency of power in organizations: Executive succession in the higher education publishing industry, 1958–1990. *American Journal of Sociology*, 105(3), 801–843. <https://doi.org/10.1086/210361>
- Tsujimoto, M., Kajikawa, Y., Tomita, J., & Matsumoto, Y. (2018). A review of the ecosystem concept — Towards coherent ecosystem design. *Technological Forecasting and Social Change*, 136, 49–58. <https://doi.org/10.1016/j.techfore.2017.06.032>
- Vainionpää, F., Väyrynen, K., Lanamäki, A., Parmiggiani, E., & Hietala, H. (2026). Anticipation practices of public sector organizations in liminal AI design spaces – the case of the EU AI Act. *European Journal of Information Systems*. Advance online publication. <https://doi.org/10.1080/0960085X.2026.2651227>
- Wirtz, B. W., Langer, P. F., & Weyerer, J. C. (2022). An ecosystem framework of AI governance. In J. B. Bullock et al. (Eds.), *The Oxford handbook of AI governance*. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780197579329.013.24>

Appendix A. Inventory of relevant AI and IT governance standards

This inventory lists ISO/IEC standards relevant to the governance of AI in banking and insurance, as compiled in the FLAIG project (list updated 12 February 2026). It is presented here as an *actor mapping*: ISO/IEC JTC 1/SC 42 and the related committees are participants in the ecosystem (Table 4, collaborative stakeholder engagement), and these documents are the artifacts through which they act. Publication status was re-checked against the ISO catalog in October 2026.

Table A1. Published and active standards.

Standard	Title	Relevance to the ecosystem model
ISO/IEC 38500:2024	Governance of IT for the organization	Organizational layer: the IT governance into which AI governance is embedded
ISO/IEC 42001	Artificial intelligence management system	Organizational layer; certifiable: the object of external scrutiny (area 1)
ISO/IEC 38507	Governance implications of the use of AI by organizations	Organizational layer; board-level governance of AI
ISO/IEC 5338	AI system life cycle processes	AI system layer; lifecycle points for scrutiny and engagement
ISO/IEC 5339	Guidance for AI applications	AI system layer
ISO/IEC 12792	Transparency taxonomy of AI systems	Downstream governance: disclosure and explanation (area 4)
ISO/IEC 22989	AI concepts and terminology	Shared vocabulary across actors
ISO/IEC 23894	Guidance on AI risk management	Organizational layer; risk transfer evidence (area 5)
ISO/IEC 42005	AI system impact assessment	Collaborative engagement: consultation of affected groups (area 2)
ISO/IEC 42006	Requirements for bodies providing audit and certification of AI management systems	External scrutiny: constitutes the assurance market (area 1)
ISO/IEC TR 24027	Bias in AI systems and AI-aided decision making	AI system layer; fairness in credit and pricing
ISO/IEC TR 24028	Overview of trustworthiness in AI	Shared vocabulary
ISO/IEC TR 24368	Overview of ethical and societal concerns	AI ethics principlism logic in standard form
ISO/IEC TS 6254	Objectives and approaches for explainability	Downstream governance; link to WP3
ISO/IEC TS 12791	Treatment of unwanted bias in classification and regression	AI system layer
ISO/IEC 5259-1 to 5259-5; TR 5259-6	Data quality for analytics and machine learning (Parts 1–6)	Upstream governance: data providers (area 3)
ISO/IEC 8183	Data life cycle framework	Upstream governance
ISO/IEC 23053	Framework for AI systems using machine learning	AI system layer

Standard	Title	Relevance to the ecosystem model
ISO/IEC 24668	Process management framework for big data analytics	Organizational layer
ISO/IEC 25059	SQuaRE quality measures for AI	AI system layer
ISO/IEC TR 5469	Functional safety and AI systems	AI system layer
ISO/IEC TR 17903	Overview of machine learning computing devices	AI system layer
ISO/IEC TR 20226	Environmental sustainability aspects of AI systems	Collaborative engagement; societal concerns
ISO/IEC TR 21221	Beneficial AI systems	AI ethics principlism
ISO/IEC TR 24030	AI use cases	Shared vocabulary
ISO/IEC TS 8200	Controllability of automated AI	AI system layer; human oversight
ISO/IEC TS 25058	SQuaRE (AI)	AI system layer
ISO/IEC TS 42119-2	Testing of AI — Overview of testing AI systems	External scrutiny; validation
ISO/IEC 5392	Reference architecture of knowledge engineering	AI system layer
ISO/IEC 20546	Big data overview and vocabulary	Shared vocabulary
ISO/IEC 20547-3; TR 20547-1, -2, -5	Big data reference architecture (Parts 1, 2, 3, 5)	AI system layer
ISO/IEC 24029-2; TR 24029-1	Assessment of the robustness of neural networks (Parts 1–2)	AI system layer
ISO/IEC TR 24372	Overview of computational approaches for AI	Shared vocabulary
ISO/IEC TS 4213	Assessment of machine learning classification performance	AI system layer

Table A2. Standards under development or in draft (status as of October 2026).

Reference	Title
ISO/IEC DIS 4213	Performance measurement for AI classification, regression, clustering and recommendation tasks
ISO/IEC AWI 22989-2	AI concepts and terminology — Part 2: Healthcare
ISO/IEC CD TR 24030	AI use cases (revision)
ISO/IEC AWI 25058	SQuaRE (revision)
ISO/IEC FDIS 24029-3	Assessment of the robustness of neural networks — Part 3
ISO/IEC TS 42119-3 (under publication)	Testing of AI — Part 3: Verification and validation analysis
ISO/IEC FDIS 25059	SQuaRE quality measures (revision)

Appendix B. Glossary of ecosystem actors

The glossary defines the actor terms used in Figure 1 and Table 4. Actors are grouped by the model component under which they primarily act; several act under more than one.

Supervisors and regulators. Public authorities with binding powers over the focal organization. In Finland, the Financial Supervisory Authority (FIN-FSA, Finanssivalvonta) supervises banks and insurers; significant banks are supervised directly by the European Central Bank within the Single Supervisory Mechanism. The European Banking Authority and the European Insurance and Occupational Pensions Authority issue guidelines and technical standards. Under the AI Act, national market surveillance authorities oversee high-risk systems; in Finland, Act 1377/2025 assigns this role for the sector's high-risk systems to FIN-FSA, with Traficom as the national single point of contact. Characteristically enact regulatory oversight.

Conformity assessment and audit bodies. Organizations that assess whether a system or management system meets a requirement. They include notified bodies under the AI Act, statutory auditors, and bodies accredited to audit and certify AI management systems under ISO/IEC 42006. For the credit and insurance-pricing systems typical of the sector, the AI Act provides for conformity assessment under internal control, so notified bodies play little role there and voluntary certification carries more of the weight. Their independence from the focal organization is what makes their scrutiny *external*.

Standard-setting bodies. Organizations that develop voluntary technical and management standards. ISO/IEC JTC 1/SC 42 is the international committee for AI; CEN-CENELEC JTC 21 develops the European standards that will support AI Act conformity. Their outputs are listed in Appendix A. Characteristically enact IT professionalism.

Industry associations. Bodies representing the sector's firms collectively, such as Finance Finland (Finanssiala ry) nationally and, at European level, the European Banking Federation and Insurance Europe. They are channels for shared interpretation of regulation and for collective positions toward providers and standard-setters.

Academic and research institutions. Universities and research organizations that generate knowledge about AI impacts and governance and can act as evaluation partners. The FLAIG consortium is an instance.

Civil society and consumer organizations. Organizations representing affected communities and the public, including consumer advocacy bodies. Characteristically enact AI ethics principlism.

Affected individuals and groups. People who bear the consequences of AI-enabled decisions (loan applicants, policyholders, claimants), whether or not they are the firm's customers.

Cloud and infrastructure providers. Providers of the computing and storage on which AI systems run. Governed upstream through outsourcing and ICT third-party risk processes, reinforced by DORA.

Foundation-model and generative AI providers. Providers of large pre-trained models that the focal organization licenses and adapts. Under the AI Act they carry obligations as

providers of general-purpose AI models. They are the upstream actors over which a single bank or insurer has least negotiating power.

Application and system vendors. Providers of core banking, policy administration, claims, fraud-detection, and analytics software into which AI is embedded. Governed upstream through procurement and contract.

Consultancies and system integrators. Firms that advise on or build AI systems and governance for the focal organization; frequently the practical channel through which standards and regulation are interpreted.

Data providers. Organizations supplying data used to train or run AI systems, including credit bureaus and external data services. Governed upstream, with data provenance as the central concern.

Customers, policyholders, and end users. Individuals and firms who receive AI-enabled services. In the model they hold a hybrid role, transacting with the organization and, at least latently, overseeing its AI through expectation, complaint, and exit.

Insurers and reinsurers (as risk-alleviation actors). Firms that assess, price, and partly absorb the AI-related risk of a deploying organization, alleviating it for that organization and its customers. Through their risk assessment they also take part in collaborative stakeholder engagement. In this role an insurer is a governance actor for other firms; the same insurer is also a deployer governed by everything else in the model.

Internal governance actors. The AI governance function or committee, compliance, legal, and risk functions, business units owning AI use cases, developers and data scientists, and internal audit. Not ecosystem actors in the strict sense, but the point at which the ecosystem meets the firm; listed because the translation component depends on them.

Appendix C. Sources for the synthesis in the report

Minkkinen, M., Jylhä, H., & Mäntymäki, M. (2026). “AI Governance: A Layered Institutional Perspective.” Conceptual paper using a theory-synthesis approach grounded in systems theory. Develops the three-layer model (societal/ecosystem, organizational, AI system; the first is called the environmental layer in this report) and identifies four institutional logics (AI ethics principlism, managerial rationalism, IT professionalism, regulatory oversight) and four institutional dynamics. No primary data. Unpublished project working paper, not peer-reviewed. Used in Chapter 3 for the layers and logics.

Khan, F., & Mäntymäki, M. (2027). “Interorganizational AI Governance in Academic and Grey Literature.” Conceptual synthesis of 27 academic studies (from the Senior Scholars’ journal set, January 2022–March 2025, with backward and forward search) and 27 grey-literature documents (from 26 organizations identified by expert consultation), coded inductively in the Gioia style into 26 first-order concepts, five second-order themes, and two aggregate dimensions. Presented at I3E 2026 and published in Lecture Notes in Computer Science. Used in Chapter 3 for the two governance patterns and throughout Chapter 5.

Birkstedt, T., Minkkinen, M., & Mäntymäki, M. (2027). “Governance by Negotiation: How Organizations Translate AI Ethics into Practice.” Interpretive study applying actor-network theory to 27 semi-structured interviews from 23 Nordic organizations, in two rounds: 13 interviews with 12 Finnish frontrunner organizations (October–November 2020) and 14 interviews with 11 retail banks and insurers in Finland, Norway, and Sweden (August 2022–December 2023). Triangulated with public AI policy documents and the development of the EU AI Act. Presented at I3E 2026 and published in Lecture Notes in Computer Science. Used in Section 3.5 and decision area 6.

Jylhä, H., Minkkinen, M., & Mäntymäki, M. (2026). “Rage Against or With the Machine? From Perceived Limitations, Concerns, and Future Potential of GenAI to Anticipatory Governance.” Qualitative study of 23 semi-structured interviews (March–October 2025) with AI professionals with more than six years’ experience, active in Europe, analyzed thematically (139 codes → 17 → 5 themes, 14 subthemes). Accepted for presentation at ISPIM Connects Sendai (Sendai, Japan, 30 November–2 December 2026) and for publication in the conference proceedings. Used as the output of the horizon-scanning pilot in Chapter 6.

Heikkilä, S. (2026). “Responsible AI Ecosystems: Insights from the Insurance Industry.” Master’s thesis, Information Systems Science, Turku School of Economics, University of Turku; supervised by Matti Minkkinen; conducted within FLAIG. Five semi-structured interviews with six professionals in Finnish insurance (December 2025–February 2026), coded inductively with NVivo in the Gioia style. Used in Section 4.1 and across Chapter 5.

Huopalainen, J. (2026). “Customer Expectations of AI in Retail Banking: A Vignette-Based Qualitative Study.” Master’s thesis, Information Systems Science, Turku School of Economics, University of Turku; supervised by Matti Minkkinen; conducted within FLAIG. Ten vignette-based semi-structured interviews with Finnish retail banking customers (April 2026), analyzed thematically with the EU HLEG requirements as a sensitizing lens. Used in Section 4.2 and decision area 4.

Minkkinen, M., Zimmer, M. P., & Mäntymäki, M. (2023). “Co-shaping an Ecosystem for Responsible AI: Five Types of Expectation Work in Response to a Technological Frame.” *Information Systems Frontiers*, 25(1), 103–121. Published article from the predecessor AIGA project; five EU documents and 15 expert interviews on the European responsible AI ecosystem. Used in Chapter 2 and Section 5.2.

Minkkinen, M., & Mäntymäki, M. (2023). “The Institutional Logics Underpinning Organizational AI Governance Practices.” SCIS 2023, paper 12. Published conceptual paper identifying the four institutional logics adopted in Minkkinen et al. (2026). Used in Section 3.2.

FLAIG standards list (updated 12 February 2026). Internal inventory of ISO/IEC standards compiled in the project. Reproduced as Appendix A.

Appendix D. Framework self-assessment worksheet

This worksheet reproduces the seven decision areas of the AI ecosystem strategy framework (Chapter 5) as a grid that a bank or insurer can complete. For each area, mark the current position and the intended position (minimal compliance / managed participation / active shaping) and record the main reason. The reasons column is often the most useful output (Section 5.4).

Table D1. AI ecosystem strategy self-assessment.

#	Decision area	Minimal compliance	Managed participation	Active shaping
1	External scrutiny posture	Supervisory and AI Act requirements only	Voluntary certification / third-party review of selected systems; external review of evaluation methods	Co-defining sector assurance criteria; publishing assurance results
2	Stakeholder engagement design	Engagement only where regulation requires	Defined partners engaged at defined lifecycle stages	Participation in standardization; convening a sector forum
3	Vendor and upstream governance	AI providers within existing third-party risk process	AI-specific procurement criteria, contractual audit rights, continuous oversight	Collective action on terms with large providers
4	Customer-facing and downstream governance	AI Act transparency and consumer-protection duties; human-intervention right in consumer credit	AI use disclosed by service; explanation formats; guaranteed human route; stakes-based automation rules; feedback collection	Customer-facing AI practice as visible differentiator; work with consumer bodies on sector norms
5	Risk transfer and insurability	AI risk within existing cover (bank) / excluded or silent (insurer)	Exposures identified and discussed with insurers (bank) / selective underwriting on governance evidence (insurer)	Defining insurability evidence with insurers (bank) / AI risk as a product with governance-based pricing (insurer)
6	Translation capacity	AI governance within an existing function, no dedicated decision rights	Designated function or committee with decision rights; system-level compliance; hybrid staffing	Governance function as spokesperson toward ecosystem; practice offered as a sector model
7	Anticipatory capability	Regulatory watch via compliance; technology watch via vendors	Annual horizon-scanning routine with owner, panel, and decision path (Chapter 6)	Joint sector foresight with peers or research partners

Table D2. Record of positions.

#	Decision area	Current position	Intended position	Main reason
1	External scrutiny posture			
2	Stakeholder engagement design			
3	Vendor and upstream governance			
4	Customer-facing and downstream governance			
5	Risk transfer and insurability			

#	Decision area	Current position	Intended position	Main reason
6	Translation capacity			
7	Anticipatory capability			

Guidance for completion: positions are independent across areas; active shaping presupposes managed participation in the same area; and where the reason for a minimal position is a structural feature of the ecosystem (for example, a provider that will not negotiate, or the absence of an insurance product), note it explicitly, since it identifies where collective action or a partner's product decision would change the position.