

Governing by Proxy: How Regulatory Intermediaries Shape EU AI Act Compliance in Practice

Simon Jarvers, Orestis Papakyriakopoulos

Professorship for Societal Computing, Technical University of Munich, Germany
simon.jarvers@tum.de, orestis.p@tum.de

Abstract

The EU AI Act assigns compliance obligations to defined actor roles along the AI supply chain, but how these obligations are enacted in practice remains understudied. This paper investigates the ecosystem of regulatory intermediaries (consultants, auditors, ethics advisors, certification bodies) through which an AI system provider translates regulatory requirements into organisational practice. Drawing on participant observation and document analysis from an embedded case study at an AI startup, the analysis applies the Regulator-Intermediary-Target (RIT) model to map sixteen intermediaries and trace the compliance pressures that drive pre-enforcement intermediary engagement. Compliance artefacts across risk management, impact assessment, and technical documentation show a consistent pattern: intermediary interpretive choices construct what AI Act compliance means for the regulated organisation. Using Edelman’s legal endogeneity framework, which theorises how organisations’ private compliance constructions can feed back into the meaning of law itself, we discuss how legal requirements are managerialised into procedural deliverables and how compliance forms diffuse without evidence of governance efficacy. The paper contributes on two levels: it extends the RIT model to organisational-level AI governance, demonstrating that compliance intermediaries perform constitutive rather than facilitative functions, and it provides an early empirical signal for legal endogeneity in AI Act compliance. We encourage more guidance from the EU and argue that the compliance ecosystem deserves the same analytical and regulatory attention currently directed at the AI supply chain.

1 Introduction

The EU AI Act (Regulation 2024/1689) mandates compliance obligations from actors occupying defined roles along the AI value chain as Figure 1 illustrates: General Purpose AI (GPAI) model providers must supply documentation and transparency summaries (Art. 53), AI system providers must implement risk management, technical documentation, conformity assessment, and post-market monitoring (Art. 9–17), and deployers must ensure appropriate use and human oversight (Art. 26). What this vertical, role-to-actor mapping leaves unaddressed is *how* compliance is enacted in practice.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: The EU AI Act assigns compliance obligations to defined actor roles along the AI supply chain, but is limited on *how* these obligations are enacted in practice.

In practice, compliance does not stay within the supply chain. For example, AI system providers who face obligations they cannot fulfil alone turn to a lateral ecosystem of compliance consultants, lawyers, ethics advisors, and auditors. These actors are not part of the AI supply chain in the traditional sense: they do not contribute to the AI system’s production. Yet they are integral to how the AI system provider handles its regulatory obligations. A compliance consultant who translates Article 9’s risk management requirements into a standardised process does not only transmit the regulation, but determines what risk categories are enumerable, what assessment methods count as adequate, and what evidence will satisfy an auditor. The intermediary constructs what compliance *is* for the regulated organisation. Whether that construction aligns operational compliance with the Act’s objectives of protecting safety, health, and fundamental rights, or merely satisfies the auditor, is a governance question and not only a descriptive one. Despite this constructive role, existing AI supply chain scholarship centres on production and deployment actors (Cobbe, Veale, and Singh 2023; Hopkins et al. 2025b), and the emerging literature on regulatory intermediaries in AI governance has focused on mandated public-sector actors at the macro level of regulatory transposition (Tervo, Väyrynen, and Iivari 2024, 2026); neither examines the private-market compliance ecosystem through which regulated organisations enact their obligations in practice.

This paper provides an embedded empirical study of AI Act compliance within a multi-actor AI supply chain, documenting both the vertical pressures that trigger compliance action and the horizontal ecosystem through which obligations are enacted. Drawing on participant observation and document analysis from the first author's dual role as researcher and AI Governance Officer at an AI startup, the study asks: **How does an AI system provider's engagement with Regulatory Intermediaries shape the practice and meaning of EU AI Act compliance?**

The embedded case study addresses calls for research from within AI supply chains (Cobbe, Veale, and Singh 2023), documenting compliance dynamics that external researchers cannot observe. On this empirical basis, the paper contributes at two levels. (1) Analytically, it extends the application of the RIT model to AI governance from macro-level role taxonomies (Tervo, Väyrynen, and Iivari 2024) to organisational-level practice, demonstrating how compliance intermediaries perform constitutive functions. (2) Conceptually, it provides an early empirical signal for legal endogeneity dynamics in AI Act compliance, documenting how intermediaries construct compliance meaning under conditions of regulatory ambiguity, procedural regulation with difficult outcome measurement (Matus and Veale 2022), and nascent enforcement.

2 Background

The analysis draws on three bodies of work deployed at different levels. The AI supply chain literature provides the empirical context. The Regulator-Intermediary-Target (RIT) model provides the analytical vocabulary applied in the findings. Edelman's legal endogeneity framework provides the interpretive lens applied in the discussion.

2.1 AI Supply Chains and the Compliance Challenge

Most AI systems are products of multi-actor supply chains, rarely single-organisation outputs. Cobbe, Veale, and Singh (2023) establish this as a governance problem by identifying four structural features of algorithmic supply chains: (1) production and deployment split across organisations, (2) configurations that shift as actors and data flows change, (3) integration across horizontal and vertical market structures, and (4) consolidation around systemically important providers. One central concept, the *accountability horizon*, describes the limit of what any single actor can foresee about how its component will be used downstream. For AI system providers integrating upstream foundation models, this horizon is immediate: they build on components whose training data, architectural decisions, and failure modes they cannot fully inspect (Bommasani et al. 2021). Hopkins et al. (2025a,b) extend this structural analysis with a stakeholder taxonomy and redress typology, demonstrating that market concentration constrains what governance mechanisms downstream actors can feasibly pursue. While Hopkins et al. (2025b) include "Intermediaries & AI Services" as a stakeholder category, this denotes technical production actors—data labelling services, fine-tuning specialists,

MLOps providers—not the regulatory compliance intermediaries analysed in this paper. In one of the few empirical studies of AI supply chains, Widder and Nafus (2023) show that software modularity produces "dislocated accountability," where developers acknowledge potential harms but locate responsibility elsewhere in the chain.

The AI Act intervenes in this supply chain through management-based regulation (Coglianese and Lazer 2003): it mandates processes (risk management, documentation, conformity assessment) rather than prescribing technical outcomes. Provisions such as "appropriate" risk management (Art. 9), "relevant" data quality (Art. 10), and "proportionate" human oversight measures (Art. 14) are deliberately open-textured. Interpreting them requires specialised expertise in regulatory analysis, risk methodology, and conformity assessment. AI governance outcomes are, moreover, *credence goods*: their quality cannot be assessed through use or simple inspection (Matus and Veale 2022). These two properties, interpretive ambiguity and outcome opacity, generate structural demand for external expertise that the vertical supply chain alone cannot satisfy.

2.2 Regulatory Intermediaries: The RIT Model

The Regulator-Intermediary-Target (RIT) model reconceptualises regulation as a three-party relationship rather than a bilateral regulator-regulatee dynamic (Abbott, Levi-Faur, and Snidal 2017b). As Figure 2 illustrates, intermediaries (I) stand between regulators (R) and targets of regulations (T), performing functions that include: interpretation (translating ambiguous provisions into actionable terms), information provision (conveying regulatory requirements), facilitation (scaffolding implementation processes), and monitoring (verifying whether targets comply), and standard-setting (establishing what compliance looks like beyond the legal text). The companion enrichment paper (Abbott, Levi-Faur, and Snidal 2017a) extends the model to address intermediary chains and concatenations, where regulation passes through multiple intermediary layers before reaching the target.

Intermediaries possess their own capabilities, interests, and capture risks. Three dynamics matter: *Orientation* distinguishes target-oriented intermediaries (selected and paid by the regulated firm) from regulator-oriented ones (appointed by regulatory bodies). *Selection* determines capture risk: when targets choose their own intermediaries, conditions favour accommodating interpretations. *Functional scope* matters because intermediaries often bundle functions, creating conflicts when the same actor both designs and verifies a compliance process.

Recent work has begun applying the RIT framework to AI governance. Tervo, Väyrynen, and Iivari (2024) identify 15 intermediary roles across AI Act regulatory phases, and their follow-up study (Tervo, Väyrynen, and Iivari 2026) examines how a Finnish ministerial working group translates EU regulation into national implementation. Both contributions study mandated, public-sector intermediaries at the macro level of regulatory transposition. Matus and Veale (2022) apply the RIT model to machine learning certification, drawing structural parallels with sustainability governance. The present paper complements this work by examin-




Figure 2: The Regulator–Intermediary–Target (RIT) model (Abbott, Levi-Faur, and Snidal 2017b). Solid arrows indicate the primary regulatory chain. Dashed arrows indicate return flows.

ing private-market intermediaries at the organisational level: consultants, ethics advisors, and auditors whose role is not mandated by the Act but necessitated by its interpretive demands.

2.3 Legal Endogeneity

Edelman’s legal endogeneity framework (Edelman 1992, 2016), developed through extensive empirical study of US employment law, describes a sequential cycle through which regulated organisations and their professional intermediaries come to define what law means in practice:

(1) Ambiguous law is passed, creating genuine interpretive uncertainty. (2) Organisations face uncertainty about what compliance requires. (3) Intermediaries construct compliance forms: standardised structures that function as evidence of compliance. (4) Organisations adopt these forms as they diffuse through professional networks. (5) Institutions defer: courts or regulators accept the presence of these structures as evidence of compliance. (6) The constructed form becomes the *de facto* legal standard.

Three enabling conditions explain why the cycle initiates (Edelman, Uggen, and Erlanger 1999): (1) ambiguous law that gives organisations interpretive latitude, (2) regulation that constrains procedures more than outcomes, making substantive compliance difficult to verify, and (3) weak or nascent enforcement that creates space for intermediary-constructed forms to fill the interpretive vacuum.

Two concepts from this framework are central to the present analysis. *Managerialisation* (Edelman 2016) describes the recasting of legal requirements as management objectives, a transformation that satisfies formal requirements while potentially altering the law’s substantive intent. The concept of *rational myths* (Edelman, Uggen, and Erlanger 1999) explains why compliance forms diffuse: organisations adopt them not because their efficacy is demonstrated but because they provide legitimacy signals, professions promote them, and external actors accept their appearance as evidence of substance.

3 Case and Method

3.1 Research Design

The study employs an embedded single case study design (Yin 2018). The first author serves as AI Governance Officer at the case organisation while conducting research on the compliance dynamics they participate in, providing direct access to compliance processes, intermediary interactions, and governance decisions in real time. This enables

observation of *enacted* compliance rather than retrospective accounts of it (Jerolmack and Khan 2014).

The case is analytically strategic for two reasons. First, an AI system provider with dependencies in both directions along the supply chain constitutes a productive observation point for intermediary dynamics: the organisation’s compliance capacity is substantially constituted through external relationships, making these dynamics directly observable. Second, the EU AI Act’s substantive obligations for high-risk AI systems are only now coming into force, with provider obligations applicable from December 2027.¹ Empirical research on how organisations navigate these obligations within multi-actor supply chains is effectively absent, so the case captures compliance as it unfolds, favouring open-ended inquiry over hypothesis testing.

3.2 Case Description

The case organisation is a startup that develops AI-powered solutions for training and task execution in industrial settings. It employs approximately fifteen people, is certified under ISO 27001 and ISO 42001, holds an EIC Accelerator grant (a grant under the Horizon Europe programme of the European Union), and is preparing proactively for EU AI Act compliance.

Under the AI Act’s role-based framework the organisation occupies the AI system provider position, integrating an upstream GPAI model (Art. 53) and serving downstream deployers (Art. 26). Its current products are not high-risk, but leadership pursues early compliance certification as a competitive and liability measure, anticipating potential reclassification. Within the RIT framework the organisation occupies the Target position whose compliance practice is shaped by the intermediary relationships this study documents.

3.3 Data Sources

Data collection spans nine months and draws on three sources serving distinct epistemic functions.

Document analysis traces how regulatory requirements are translated across organisational boundaries. The formal corpus comprises internal compliance documents (28 AI-related risks with mitigation strategies, 65 AI-related controls, and 31 AI Management System documents covering policies, procedures, registers, etc.) and the full ISO 42001 audit cycle (initial certification and re-certification after one year, including internal audit, management review, and external audit reports). Other intermediary interactions fall outside this formal corpus and are reflected on through the research journal described below.

Participant observation captured through reflexive journaling constitutes the primary record of intermediary engagement. Following established protocols for insider research (Coghlan 2019) and Gibbs’s (1988) reflective cycle, weekly journal entries were maintained over eight months across the study period, beginning in the second month of data collection. Entries explicitly mark whether an observation originates from the practitioner role or the researcher

¹For more information about the Digital Omnibus see https://ec.europa.eu/commission/presscorner/detail/en/ip_26.1024.

role, documenting moments where these perspectives diverge.

Verification workshop with the Chief Technology Officer (CTO). A dedicated session was recorded and transcribed, combining a stakeholder mapping exercise with open-ended questions about compliance pressures, intermediary relationships including financial data, and selection rationales. The workshop functions as a verification instrument rather than a primary data source, triangulating observational and documentary evidence against the CTO's independent perspective.

3.4 Analysis

Analysis proceeds in two stages aligned with the paper's theoretical architecture. Observational and documentary data are coded deductively against RIT categories (Abbott, Levi-Faur, and Snidal 2017b), covering intermediary functions, orientation dynamics, and capture risks, with inductive refinement where empirical patterns did not fit the deductive scheme. The resulting patterns are then interpreted through Edelman's legal endogeneity framework (Edelman 1992, 2016), separating grounded observations of managerialisation and rational myth formation (Edelman, Uggen, and Erlanger 1999) from projective claims about the cycle's later stages.

3.5 Positionality and Limitations

Embedded research generates characteristic tensions between practitioner and researcher roles. Coghlan (2019) specifies three: preunderstanding (organisational familiarity producing blind spots), role duality (serving two functions simultaneously), and organisational politics (confidentiality, career considerations, access negotiation). Independent academic supervision and systematic reflexive journaling provide external review and track role tensions as they occur. Research and organisational objectives are aligned: findings that reveal tensions carry the same analytical weight as those that do not.

The intermediary mapping reflects one organisation's compliance ecosystem at a specific point in time. This means two scope conditions bound the empirical claims. First, the mapping illustrates one ecosystem configuration; it is not a representative sample. The dynamics documented here should be recognisable in comparable contexts where AI system providers engage compliance intermediaries, but whether they manifest in the same form is a question for future comparative work. Second, the study captures compliance before the AI Act's enforcement infrastructure is operational. National competent authorities and market surveillance authorities are being designated but have not yet engaged with AI system providers at the organisational level. The findings and discussion should be read in this pre-enforcement context, which shapes both the composition of the intermediary ecosystem and the governance dynamics it produces. We provide detailed description of context, methods, and findings to enable readers to assess transferability to other organisational settings and regulatory environments.

4 Findings

The findings follow a three-part arc: the compliance pressures that drive the focal organisation toward external engagement (Section 4.1), the intermediary ecosystem that emerges in response (Section 4.2), and how intermediaries construct compliance meaning in practice (Section 4.3). As disclosed in Section 3.5, the mapping reflects a single organisation's ecosystem, and the dynamics unfold in a pre-enforcement regulatory environment where compliance is pursued without active regulatory oversight. Both conditions shape the pressures, composition, and meaning-construction processes documented below.

4.1 Why Compliance Intermediaries Emerge

What drives an organisation to pursue AI Act compliance before enforcement begins is difficult to disentangle. The contributing pressures are interconnected, often implicit, and rarely formally documented, particularly in small and medium-sized enterprises (SMEs) where compliance strategies emerge incrementally. Four contributing factors were observed in the focal organisation (Figure 3).

Customer procurement requirements. Enterprise customers cascade compliance requirements upstream, including a recent request for AI Act compliant instructions for use. A comparable pattern is established in information security: ISO 27001 has shifted from competitive differentiator to market access prerequisite, creating a self-sustaining cascade independent of the underlying regulation's enforcement timeline. Formal certification also functions as a liability shield: a documented compliance record evidences good-faith regulatory engagement, the absence of which would be substantially more damaging in legal proceedings than any shortcoming in its content.

Investor and insurance due diligence. Investor pressure operates through two reinforcing mechanisms. Standard investment terms include representations and warranties with personal founder guarantees of compliance with all applicable laws, concretising AI Act compliance as a personal financial risk to the founders. Investors additionally require Directors and Officers (D&O) insurance, whose providers condition coverage on demonstrated compliance, reinforcing the same exposure through a second instrument. Enterprise customers extend the dynamic by requiring evidence of insurance alongside ISO certification, linking the customer pressure above to the same liability architecture.

EIC Accelerator grant conditions. The grant requires compliance with applicable laws, including the EU AI Act. Non-compliance constitutes a breach of grant conditions, making this the most direct compliance requirement the organisation faces, transmitted through a public funding mechanism. The grant also imposes a formal ethics oversight structure comprising three independent actors operating as a concatenated intermediary chain. The grant body functions as a compliance intermediary in the RIT sense, translating the AI Act's obligations into concrete contractual conditions.

Competitive positioning. Early adoption of AI governance frameworks provides a reputational signal and competitive advantage, particularly in application areas where




Figure 3: Compliance pressures driving AI Act compliance and the horizontal intermediary ecosystem the AI system provider engages in response.

customers are sensitive to regulatory exposure. This advantage may diminish as certification diffuses more broadly, but in the current period it remains a meaningful first-mover differentiator.

Absent pressures. Two pressure sources present in the AI Act’s regulatory design were not observed as active drivers. Upstream GPAI providers impose governance conditions through content filters and acceptable use policies that the organisation does not experience as AI Act compliance pressure, and no distinct internal compliance pressure was observed.

The structural necessity of intermediary engagement. These pressures converge on a compliance challenge the organisation cannot resolve internally. As established in Section 2.1, the AI Act’s open-textured provisions generate structural demand for external expertise, and AI governance outcomes are credence goods whose quality cannot be assessed through inspection (Matus and Veale 2022). This demand exists regardless of organisation size (Costanza-Chock, Raji, and Buolamwini 2022); resource constraints at startup scale intensify the dependence without creating it. Formal engagement also provides accountability signals, certification records that customers, investors, and insurers rely on, that internal processes cannot replicate.

4.2 Intermediary Identification and RIT Function Mapping

Table 1 maps sixteen intermediaries identified in the focal organisation’s compliance ecosystem, coded by RIT function, orientation, payment structure, and engagement type. The intermediaries cluster into five empirically observed groups:

1. **Private compliance service providers:** a compliance-as-a-service provider, a specialist law firm, and a penetration testing partner. All are privately selected and paid by the focal organisation. They perform the core interpretive work of translating ambiguous AI Act provisions into organisational deliverables: templates, assessments, test scopes, and specialised compliance software.

2. **Technical implementation partners:** an infrastructure consulting partner and an external AI development partner. Both contribute to building and configuring technical components of the AI system. Their compliance contribution is indirect: they support technical soundness as a precondition for compliance without explicitly interpreting regulatory requirements.
3. **Open-access knowledge intermediaries:** an AI focused NGO, a national ICT industry association, and academia.² These actors provide information and interpretation through non-market channels: training events, published guidance, and research outputs.
4. **Verification and oversight actors:** an internal and external auditor, an ethics advisor, a university ethics committee, and an EU-level ethics oversight body. This group contains two structurally distinct sub-types: the ISO certification infrastructure (internal and external audit for ISO 42001) and the imposed ethics oversight chain mandated by grant conditions. In an organisation of similar size without the EIC Accelerator grant the second group would most likely not exist.
5. **Standard-setting and regulatory infrastructure actors:** a national information security agency, European and international standardisation bodies (CEN/CEN-ELEC, ISO), and the EU AI Office. None maintain direct relationships with the focal organisation. Their influence operates through circulating artefacts: risk catalogues, emerging harmonised standards, and codes of practice.

Estimated Cost of AI Governance AI governance at the focal organisation costs approximately EUR 25,000–35,000 per year in external relationships: EUR 8,000–12,000 for compliance intermediaries (Group 1), EUR 6,000–10,000 for audit (Group 4), and EUR 10,000–14,000 for the AI-governance-attributable share of technical partners (Group 2). The audit figure covers ISO 42001 certification, which can be viewed as a simpler and probably less expensive proxy in the absence of an AI Act conformity assessment. Internal costs of comparable magnitude cover the AI Governance Officer role, management attention, and compliance-driven development time resulting in a total estimate of EUR 50,000–70,000 per year. Two restrictions apply: the figures count only work that would not be undertaken absent the obligation, and limit scope to AI-specific compliance, excluding ISMS, GDPR, and general legal work. Validated with the CTO and reported as ranges, the figures offer an order-of-magnitude signal only. The EIC ethics oversight is treated separately and amounts to approximately EUR 20,000 over the grant period.

Ecosystem-Level Patterns Three structural observations emerge from the mapping when read through the RIT model.

First, **the ecosystem is overwhelmingly target-oriented.** The majority of actively engaged intermediaries are selected

²The first author’s dual role means the academic institution is also directly connected to the focal organisation; this positionality is addressed in Section 3.5. Here, the entry reflects general research outputs and open-access information provision.

	Intermediary	Description	Interpretation	Info. Prov.	Facilitation	Monitoring	Std.-Setting	Orientation	Payment	Active
Compliance Intermediaries	Compliance as a service	Supplies AI Act governance templates and compliance software whose outputs serve as audit-facing evidence.	✓	✓	✓	×	×	T	✓	✓
	Legal counsel / Law firm	Consulted on request for specific legal questions; not engaged for systematic AI Act compliance guidance.	✓	✓	~	×	×	T	✓	✓
	Penetration testing	Conducts technical vulnerability assessments whose outputs contribute to security hardening and compliance evidence.	✓	✓	~	×	×	T	✓	✓
Technical Partners	Infrastructure consulting	Supports cloud configuration and technical architecture decisions; no regulatory interpretation or verification function.	~	~	✓	×	×	T	✓	✓
	External AI developer	Develops the core AI component; performs AI Act obligations such as documenting, testing, and monitoring.	~	×	✓	×	×	T	✓	✓
Knowledge Intermediaries	Think tank / NGO	Provides training events and publishes AI Act compliance white papers and implementation guidance.	✓	✓	×	×	~	T	×	Ev.
	Industry association	Publishes AI Act position papers, implementation guidance, and standardisation recommendations.	✓	✓	×	×	~	T	×	Ev.
	Academia	Research collaborations and open-access information; paid component reflects the dual practitioner–researcher role.	✓	✓	×	×	~	N	×	✓
Verification & Oversight	Internal auditor	Not an employee of the focal organisation. Conducts the ISO-mandated internal audit preceding external certification.	✓	~	×	✓	×	N	✓	✓
	External auditor	Conducts ISO 42001 audits; closest available proxy for AI Act conformity assessment until harmonised standards are released.	✓	~	×	✓	×	R	✓	✓
	Ethics advisor	Provides independent ethical assessment of the AI governance process; mandated by EIC grant requirements.	✓	~	×	✓	×	N	✓	✓
	Academic ethics committee	Provides independent ethical approval for research purposes; mandated by EIC grant requirements.	✓	×	×	✓	×	N	✓	✓
	EU-level ethics oversight body	The EIC Accelerator grant imposes mandatory ethics review on grant recipients.	✓	×	×	✓	×	R	×	✓
Standards & Regulation	National security agency	Provides the risk catalogue that constitutes the deductive framework for the organisation’s risk management scope.	~	✓	×	×	✓	R	×	×
	Standardisation organisations	ISO 42001 serves as current compliance placeholder; CEN/CENELEC is developing AI Act harmonised standards.	✓	~	×	×	✓	R	×	×
	EU AI Office	Issues interpretive guidelines, codes of practice, and enforcement standards with exclusive EU-level competence.	✓	✓	×	×	✓	R	×	×

Table 1: Compliance intermediaries identified in the case study, organised by group, with their RIT functions in relation to the AI system provider (Target). RIT function cells are shaded: green (✓) = primary function observed; yellow (∼) = limited or indirect; grey (×) = not observed. *Orientation*: T = Target-oriented, R = Regulator-oriented, N = Neutral. *Payment*: Target pays for the relationship. *Active*: ongoing relationship (✓) or event-based only (Ev.). Evidential confidence varies across cells; directly observed relationships carry stronger epistemic weight than inferred ones.

and paid by the focal organisation. Regulator-oriented actors are either imposed through grant conditions or absent because the AI Act’s enforcement architecture is not yet operational at the organisational level. This composition is consistent with the RIT model’s prediction that regulatory targets assemble intermediary relationships that serve their own compliance needs (Abbott, Levi-Faur, and Snidal 2017b).

Second, **the organisation’s experience of regulation is multilayered**. Standard-setting actors occupy regulator-adjacent positions but none interact with the focal organisation directly. Their outputs circulate as artefacts that private intermediaries draw upon and repackage. The compliance-as-a-service provider bases its templates on published stan-

dards, auditor feedback, and open-access knowledge from industry associations and NGOs; the organisation receives the integrated product with very little visibility into these upstream inputs. Abbott, Levi-Faur, and Snidal (2017a) describe this concatenated structure as characteristic of complex regulatory environments, with interpretation accumulating at each step. Section 4.3 traces one such chain in detail.

Third, **intermediaries shape each other**. The compliance-as-a-service provider recommended the external certification body that audits the compliance templates the provider itself constructed: the entity that shaped what compliance looks like also influenced who verifies it.

The RIT literature identifies this as a structural capture risk in target-selected intermediary ecosystems (Abbott et al. 2021). When the target selects both the interpretive and the monitoring intermediary, and when one recommends the other, the independence that monitoring requires may become structurally compromised.

4.3 How Intermediaries Shape AI Governance in Practice

This section examines one intermediary relationship in depth, tracing how the compliance-as-a-service provider's interpretation of Article 9's risk management requirement constructed the organisation's risk management system. Two further intermediary dynamics, an AI Impact Assessment template and a technical documentation template, were also observed and analysed, but are not reported in this paper due to brevity and readability; they exhibit a parallel dynamic supporting a pattern claim rather than a single-instance observation. This example follows a three-part structure: empirical description, RIT-informed analysis, and the meaning-construction moment.

The AI Risk Management System Empirical description. The compliance-as-a-service provider (hereafter: compliance partner) is the primary interpretive actor. It translates the risk management requirements of Art. 9 and ISO 42001 into an operational process embedded in its software platform: risk identification, impact and likelihood assessment, treatment selection, control implementation, and residual risk acceptance. The compliance partner's interpretation draws on its engagement with standards bodies, auditors, and client experience, though the organisation has no visibility into these upstream inputs.

The risk identification step requires a methodology for establishing which risks to consider. The organisation independently selected a risk catalogue published by a national information security agency as the deductive framework; the compliance partner did not prescribe it. The selection rationale was auditability-driven: a recognised institutional source provides a defensible position. The catalogue enumerates 78 pre-defined risk categories that the organisation maps onto its AI system before executing the full process within the compliance partner's platform.

The external certification auditor raised one Opportunity for Improvement (OFI): the risk of Cloud Provider Dependence was absent from the register. This risk did not appear in the version of the catalogue used at the time, and the deductive process therefore did not surface it. Following the OFI, the risk was added. Catalogue scope is itself socially constructed and subject to ongoing revision: ISO has, for example, recently mandated climate change considerations across its standards, propagating through audit practice and risk frameworks. What counts as an enumerable risk shifts as standards bodies update their interpretive frames. No further cases of risks identified independently of the catalogue were documented in the observed period.

RIT function analysis. The compliance partner performs three bundled RIT functions. It interprets Art. 9 and ISO 42001, translating open-textured requirements into a

concrete process. It provides information about what a risk management system must contain. It facilitates through the software infrastructure. These functions are tightly coupled: the organisation accesses facilitation and information through the compliance partner's interpretation, with no pathway to the former without accepting the latter.

The risk catalogue performs a standard-setting function without any active relationship to the focal organisation; its influence operates entirely through document circulation. The external auditor's OFI is a monitoring action that produced a constructive corrective effect, surfacing a risk the primary process had not. This corrective function is bounded: it operates after the compliance process has defined the register.

The meaning-construction moment. Meaning-construction operates at two levels. At the first level, the compliance partner's interpretation defines what "AI risk management" means for this organisation: a specific sequence of steps, conducted within a specific software environment. The organisation had no prior operational definition. The compliance partner filled that vacuum entirely, and the organisation adopted the resulting process without evaluating alternatives.

At the second level, the deductive risk catalogue defines which risks are visible at any given time. Risks outside the catalogue's current scope are structurally absent from the identification step, and the catalogue's scope itself shifts as standards bodies revise their frameworks. Cloud Provider Dependence illustrates both: a materially significant risk was not surfaced until external monitoring intervened, and the conditions under which it would have been enumerable through the deductive process are themselves the product of an evolving institutional consensus. Both levels of meaning-construction were shaped by the logic of auditability: the process was adopted because the compliance partner presented it as the operational standard; the catalogue was selected because it offered a defensible institutional anchor. The governing criterion in both cases was what would satisfy an auditor, not what would produce the most complete assessment of risks material to the specific AI system.

Process-Level Patterns A parallel dynamic was observed in two further compliance artefacts: an AI Impact Assessment template and a technical documentation template. Together with the risk management system, four patterns emerge stated here as empirical observations:

First, **the primary operational function of each compliance artefact is verification-facing.** Risk register, AI Impact Assessment, and technical documentation are each structured for audit legibility. This orientation is a structural feature of the audit-certification architecture that frames compliance as a verification problem.

Second, **the scope of compliance processes is defined by an intermediary-constructed artefact.** The compliance partner's process design determines the steps of risk management. The risk catalogue determines which risks are considered. The assessment template determines which impacts are evaluated. The documentation template determines what is documented. The organisation's agency operates within

these architectures.

Third, **the interpretive choices embedded in artefacts are not visible to the organisation.** The organisation receives finished products whose construction history it does not examine; the compliance partner's information ecosystem (standards bodies, auditor feedback, client experience) shapes these choices opaquely.

Fourth, **external monitoring can compensate for scope limitations produced by earlier interpretive choices,** as the auditor's Cloud Provider Dependence OFI demonstrates. This corrective function is bounded: it surfaces risks the auditor's own framework identifies, operating after the primary compliance process has already defined the register.

5 Discussion

The findings document what compliance intermediaries do. This section asks what these dynamics mean for AI governance by reading them through Edelman's legal endogeneity framework (Edelman 1992, 2016). The discussion proceeds along an epistemic gradient: Section 5.1 interprets dynamics directly grounded in the findings, Section 5.2 connects observed patterns to mechanisms documented in adjacent literatures, Section 5.3 projects where these dynamics may lead, and Section 5.4 returns to the AI supply chain literature to locate the paper's contribution.

5.1 The Managerialisation of AI Act Obligations

Edelman (2016) uses *managerialisation* to describe how legal requirements are recast as management objectives. Compliance ceases to be a substantive achievement and becomes a procedural activity conducted through standardised deliverables. The findings document this transformation across the three compliance artefacts examined in Section 4.3. Article 9's requirement for "appropriate" risk management arrives at the focal organisation as an operational process embedded in compliance software. The AI Impact Assessment and technical documentation template undergo analogous conversions, reducing open-textured legal provisions to structured forms whose primary audience is the verification infrastructure.

Management-based regulation (Coglianese and Lazer 2003) deliberately mandates processes rather than outcomes, so this conversion is expected by regulatory design. Managerialisation becomes analytically significant when it narrows the space of possible compliance responses without explicit justification for the narrowing. The compliance partner's risk management process is one defensible interpretation of Article 9, yet it forecloses alternatives the organisation never evaluates. The risk catalogue's 78 categories define the universe of risks considered. Both narrowing choices are embedded in the deliverable's design and received as operational standards.

Power's constitutive auditing thesis (Power 1997) identifies the mechanism driving this narrowing. Audit reshapes organisational realities: organisations restructure to become "auditable," and the demands of the audit architecture come to define governance practice. The process-level patterns in Section 4.3 provide direct evidence. Across the three arte-

facts, the governing criterion was what would satisfy an auditor. The Cloud Provider Dependence example illustrates the consequence: a materially significant risk was absent from the register because the deductive methodology did not enumerate it, and was surfaced only when external monitoring intervened. Managerialisation privileges procedural form over substantive inquiry without eliminating governance substance entirely.

Two sources of the documented dynamics need to be kept apart. The case observes an organisation that delegates interpretation to its compliance partner, in a setting where doing so is a defensible allocation of scarce expertise; the delegation is not treated here as a failure of diligence. One source is structural: because governance quality is a credence property (Matus and Veale 2022), whether a given risk methodology produces better outcomes cannot be verified at any organisational scale, so management-based regulation over such unobservable outcomes generates the demand for external interpretation independently of firm size. The other source is scale-dependent. What size changes is interpretive bargaining power: the capacity to commission alternative interpretations or to build the internal expertise needed to evaluate what an intermediary supplies. Talesh's (2015) account of insurers constructing compliance meaning for small businesses without in-house legal capacity describes this same small-firm condition. A single small-firm case cannot establish how the balance shifts across the size distribution. That the structural demand persists while bargaining power grows is a scope condition on the constitutive claim.

5.2 Compliance Forms as Rational Myths

The compliance artefacts documented in the findings are standardised products deployed across the compliance partner's client base, connecting the single-case findings to a diffusion dynamic at ecosystem scale. Edelman, Uggem, and Erlanger (1999) explain such diffusion through the concept of *rational myths*: organisations adopt compliance forms whose governance efficacy has never been demonstrated, yet which provide legitimacy signals that external audiences accept. The clearest instance is the organisation's ISO 42001 certification, adopted as a proxy for the AI Act conformity assessment that does not yet exist: a form acquired for the legitimacy it signals to customers, investors, and insurers before any demonstration of governance efficacy. The compliance partner's templates exhibit the same structure. The organisation adopts them because the partner presents them as the operational standard and because evaluating alternatives requires expertise the organisation does not have. Auditors familiar with the templates then review them, reinforcing their authority. Templates designed for auditor legibility are verified by auditors who accept them on those terms.

Eckstein and Shapira's (2024) analysis of compliance gatekeepers identifies a reinforcing mechanism. In a *two-sided reputation market*, intermediaries and clients jointly benefit from the appearance of rigour, with no structural incentive on either side to expose weaknesses in the compliance artefacts they produce together. The findings corroborate this: the compliance partner recommended the certification body that audits the compliance partner's own tem-

plates, creating a closed loop where the entity that shaped what compliance looks like also influenced who verifies it. Costanza-Chock, Raji, and Buolamwini (2022) document a critical absence in the algorithmic auditing ecosystem that sustains this dynamic: no systematic mechanism exists to assess whether audits achieve their governance objectives.

Bromley and Powell's (2012) distinction between two forms of decoupling clarifies the governance risk. *Policy-practice decoupling* describes policies adopted on paper but never implemented. *Means-ends decoupling* describes a different failure: policies are genuinely implemented yet remain disconnected from their intended outcomes. The process-level patterns suggest the focal organisation is firmly in the second category. It implements compliance processes rigorously, yet the artefacts produced are optimised for verification. AI Act compliance may be closing the gap between policy and practice while opening a new gap between practice and governance objectives.

5.3 Conditions for Legal Endogeneity

Edelman's endogeneity framework (Edelman 1992, 2016) identifies three conditions under which regulated organisations and their intermediaries come to define what law means in practice: ambiguous law, procedural regulation with difficult outcome verification, and nascent enforcement. All three are structurally present in the AI Act context. The Act's open-textured provisions create genuine interpretive latitude. AI governance outcomes are credence goods (Matus and Veale 2022), making procedural compliance the primary observable signal. Enforcement is nascent: market surveillance authorities are being designated, harmonised standards remain under development, and the compliance consulting ecosystem is actively constructing what AI Act compliance looks like before any regulatory authority has assessed the results.

The case study's empirical contribution is locatable within the cycle's early stages. Stage 1, the passage of ambiguous law, is asserted as a structural precondition from the Act's documented regulatory design rather than examined directly here. Stages 2 and 3, organisational uncertainty and intermediary meaning-construction, are directly observed. Stage 4, form adoption and diffusion, is beginning: the compliance partner's templates are deployed across its client base, and the certification body's acceptance reinforces their authority. Stages 5 and 6, institutional deference and full institutionalisation, are projected from Edelman's evidence in other regulatory domains.

Where Edelman's courts close the loop by deferring to compliance structures as evidence of legal compliance, market surveillance authorities occupy that position in the AI Act context. Assessing compliance at scale with limited resources and without the adjudicative adversarialism that can slow judicial deference, they face structural incentives to accept standardised intermediary-constructed forms as adequate. One structural disanalogy tempers this projection. Edelman's cases offered no formal public channel for constructing compliance standards, and courts extended deference to privately constructed forms because nothing else occupied that space. The AI Act scripts deference toward pub-




Figure 4: Edelman's (1992; 2016) legal endogeneity cycle, adapted to the EU AI Act context.

lic forms instead: harmonised standards carry a presumption of conformity (Section 5.4). This turns the projection into a question of timing, namely whether CEN/CENELEC output arrives before private templates accumulate de facto authority.

This trajectory is contestable. Roberts and Ziosi (2026) document the politics of AI standards-making, showing that standardisation processes are sites of active negotiation. Whether intermediary-constructed compliance forms become institutionalised depends on whether countervailing forces, particularly regulatory capacity-building and empirical research on compliance effectiveness, intervene before the cycle consolidates. The present analysis identifies the structural conditions for endogeneity. Whether the cycle completes is a question for longitudinal research.

5.4 Implications for AI Governance

The sixteen intermediaries documented here populate a horizontal ecosystem whose interpretive work constitutes what AI Act compliance means in practice, a dimension the vertical supply chain literature leaves underexplored (Cobbe, Veale, and Singh 2023; Hopkins et al. 2025a). Cobbe, Veale, and Singh (2023) introduce the concept of the *accountability horizon* along the AI supply chain that upstream AI developers have limited understanding and control how their general purpose systems will be used downstream, and that downstream providers have limited insights into the training data and assumptions that shaped the upstream model. We argue this concept extends horizontally: the focal organisation observes intermediary outputs but has very little insight into the reasoning embedded at each step that produced them.

The endogeneity dynamic Sections 5.1 to 5.3 document is best read as a structural condition rather than a failure internal to the intermediary market. The AI Act's open-textured provisions and the credence-good character of governance outcomes (Matus and Veale 2022) leave substantial interpretive latitude, while regulator-side resources such as harmonised standards and AI Office guidance are maturing slowly. Private compliance intermediaries fill the resulting

vacuum because nothing else does on the same operational timeline. Intermediary engagement is structurally necessary (Section 4.1); what is at issue is the conditions under which it gives rise to rational myth formation rather than substantive compliance.

The intermediaries in Table 1 do not face uniform incentives. Profit-driven compliance services operate within the two-sided reputation market identified above, where client retention and audit success are the dominant performance signals (Eckstein and Shapira 2024). Non-market actors (NGOs, think tanks, academia, standards bodies) face different incentives, with authority resting on perceived independence; emerging open-source compliance solutions may extend this category. In the case ecosystem, profit-driven intermediaries dominate active compliance practice while non-market actors appear mainly through event-based relationships, strengthening the conditions for rational myth formation.

A structural lever exists upstream of the intermediary market: where regulators or regulator-adjacent bodies produce interpretive resources directly, the space within which private intermediaries operate narrows, and the dynamic Edelman's framework describes can be slowed. Harmonised European Standards developed by CEN/CENELEC are the AI Act's primary mechanism for this, and their continued delay extends the period during which intermediary-constructed forms accumulate *de facto* authority. The AI Office's emerging guidance and information platforms signal a similar regulatory intent; whether such initiatives mature into operationally usable resources for SME providers will determine how meaningfully they reduce private interpretive latitude.³ The implication concerns regulator-side action rather than further regulation of intermediaries themselves: clearer guidance reduces legal uncertainty for targets and helps align compliance with the Act's objectives of protecting safety, health, and fundamental rights. Whether the trajectory in Section 5.3 consolidates depends on the pace at which regulator-side resources mature, an empirical question this single, pre-enforcement case cannot resolve.

6 Future Work

The single-case design bounds the empirical claims. Comparative multi-case research across AI system providers differing in size, risk classification, sector, and intermediary composition would test whether the patterns documented here are structural features of the AI Act's compliance architecture or artefacts of one configuration. Cases where organisations build internal compliance capacity rather than outsourcing it would be particularly informative for the constitutive claim.

Longitudinal research is necessary to track whether the endogeneity cycle progresses beyond Stage 4 (form diffusion): whether market surveillance authorities, once operational, accept intermediary-constructed compliance forms as evidence of regulatory conformity or develop independent

assessment capacity. The compliance partner's client base offers a natural site for studying template diffusion across organisational contexts, and the auditor's corrective function documented in Section 4.3 suggests that monitoring intermediaries merit dedicated study.

7 Conclusion

This paper asked how an AI system provider's engagement with regulatory intermediaries shapes the practice and meaning of EU AI Act compliance. The embedded case study documented an ecosystem of sixteen intermediaries whose interpretive work constructs what compliance means for the regulated organisation. Regulatory intermediaries are constitutive of AI governance rather than exogenous to it. A compliance partner who translates Article 9 into a structured risk management process does not relay the regulation; the partner determines which risks are enumerable and what evidence will satisfy an auditor. The risk catalogue, the impact-assessment template, and the technical documentation framework each embed interpretive choices the organisation receives as operational standards without examining the reasoning behind them.

These findings carry implications at two levels. Analytically, the RIT model proves productive at the organisational level, revealing dynamics that macro-level analyses of regulatory transposition cannot capture: bundled intermediary functions, target-oriented selection, and structural capture risks where the regulated entity selects both the interpretive and the monitoring intermediary. Conceptually, Edelman's legal endogeneity framework identifies a governance trajectory worth monitoring. The AI Act's structural conditions, ambiguous provisions, procedural regulation, and nascent enforcement, match those under which compliance forms have historically become *de facto* legal standards in other regulatory domains. Whether this trajectory materialises depends on countervailing forces that remain uncertain: regulatory capacity-building, empirical research on compliance effectiveness, and the contested politics of AI standardisation.

The vertical AI supply chain has attracted sustained scholarly attention as an accountability problem; the horizontal compliance ecosystem through which organisations enact their obligations has not. This paper provides initial evidence that its intermediaries exercise constructive authority over what the Act's requirements mean in practice, with direct consequences for whether AI regulation achieves its objective of protecting safety, health, and fundamental rights. The argument is diagnostic, not a case against intermediary involvement: management-based regulation of credence goods generates structural demand for external expertise that resource-constrained organisations cannot meet internally, which makes intermediaries a necessary feature of AI Act implementation at this stage. What remains open is whether regulator-side resources, harmonised standards and operationally usable AI Office guidance, mature quickly enough to narrow the private interpretive latitude these intermediaries currently fill.

³See for example the AI Office's initiative on clear guidelines for AI Act implementation: <https://digital-strategy.ec.europa.eu/en/news/supporting-implementation-ai-act-clear-guidelines>.

Positionality Statement

The first author conducted this research from within the organisation it examines, occupying simultaneously the role of AI Governance Officer responsible for implementing compliance work and the role of researcher studying that same work. This position enabled a depth of access that external researchers rarely obtain, including direct participation in intermediary interactions, internal compliance deliberations, and the production of the artefacts the paper analyses. It also shaped what could be seen. The first author cannot fully separate the interpretation of intermediary dynamics from the investment of having worked alongside many of the intermediaries described, or from the professional stake in governance work succeeding at the organisation. Coghlan (2019) calls this pre-understanding: the accumulated assumptions that organisational membership produces, which can obscure patterns an external researcher might more readily notice.

Having no affiliation with the studied organisation, the second author provided an external academic perspective throughout. Regular supervisory meetings served as a recurring check on insider framings, surfacing questions the first author had not considered and pushing back on claims that risked overstating the constructive role of intermediaries or the consequences of their interpretive choices. This arrangement does not dissolve the tensions inherent to insider research, but it introduces a vantage point outside the organisational culture and commercial context of the research site. Our findings should be read with this positioning in mind.

Generative AI Usage Statement

Generative AI tools were used during the preparation of this manuscript. Specifically, Claude (Anthropic, versions Sonnet 4.6 and Opus 4.6/4.7) assisted with grammar and style editing to improve clarity and fluency, as well as formatting of figures and tables. All AI-generated suggestions were critically evaluated by the authors, and the authors retain full responsibility for the originality, accuracy, and integrity of all content in this manuscript.

References

- Abbott, K. W.; Genschel, P.; Snidal, D.; and Zangl, B. 2021. Beyond opportunism: Intermediary loyalty in regulation and governance. *Regulation & Governance*, 15: S83–S101.
- Abbott, K. W.; Levi-Faur, D.; and Snidal, D. 2017a. Enriching the RIT framework. *The Annals of the American Academy of Political and Social Science*, 670(1): 280–288.
- Abbott, K. W.; Levi-Faur, D.; and Snidal, D. 2017b. Theorizing regulatory intermediaries: The RIT model. *The ANNALS of the American academy of political and social science*, 670(1): 14–35.
- Bommasani, R.; Hudson, D. A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M. S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Bromley, P.; and Powell, W. W. 2012. From smoke and mirrors to walking the talk: Decoupling in the contemporary world. *Academy of Management annals*, 6(1): 483–530.
- Cobbe, J.; Veale, M.; and Singh, J. 2023. Understanding accountability in algorithmic supply chains. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 1186–1197.
- Coghlan, D. 2019. *Doing action research in your own organization*. Sage Publications Ltd.
- Coglianesse, C.; and Lazer, D. 2003. Management-based regulation: Prescribing private management to achieve public goals. *Law & Society Review*, 37(4): 691–730.
- Costanza-Chock, S.; Raji, I. D.; and 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*, 1571–1583.
- Eckstein, A.; and Shapira, R. 2024. Compliance Gatekeepers. *Yale J. on Reg.*, 41: 469.
- Edelman, L. B. 1992. Legal ambiguity and symbolic structures: Organizational mediation of civil rights law. *American journal of Sociology*, 97(6): 1531–1576.
- Edelman, L. B. 2016. *Working Law: Courts, Corporations, and Symbolic Civil Rights*. University of Chicago Press.
- Edelman, L. B.; Uggen, C.; and Erlanger, H. S. 1999. The endogeneity of legal regulation: Grievance procedures as rational myth. *American journal of Sociology*, 105(2): 406–454.
- Gibbs, G. 1988. Learning by doing: A guide to teaching and learning methods. *Further Education Unit*.
- Hopkins, A.; Cen, S. H.; Struckman, I.; Ilyas, A.; Videgaray, L.; and Madry, A. 2025a. AI supply chains: An emerging ecosystem of AI actors, products, and services. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 8, 1266–1277.
- Hopkins, A.; Struckman, I.; Klyman, K.; and Silbey, S. S. 2025b. Recourse, repair, reparation, & prevention: A stakeholder analysis of AI supply chains. In *Proceedings of the 2025 ACM conference on fairness, accountability, and transparency*, 209–227.
- Jerolmack, C.; and Khan, S. 2014. Talk is cheap: Ethnography and the attitudinal fallacy. *Sociological methods & research*, 43(2): 178–209.
- Matus, K. J.; and Veale, M. 2022. Certification systems for machine learning: Lessons from sustainability. *Regulation & Governance*, 16(1): 177–196.
- Power, M. 1997. *The audit society: Rituals of verification*. OUP Oxford.
- Roberts, H.; and Ziosi, M. 2026. The Inter-and Intra-Institutional Politics of AI Standards-Making. *Available at SSRN 6302398*.
- Talesh, S. 2015. Legal intermediaries: How insurance companies construct the meaning of compliance with antidiscrimination laws. *Law & Policy*, 37(3): 209–239.
- Tervo, E.; Väyrynen, K.; and Iivari, N. 2024. Increasing understanding about the role of regulatory intermediaries in regulation—a scoping review and implications for the European AI Act. *MCIS 2024 Proceedings*.

Tervo, E.; Väyrynen, K.; and Iivari, N. 2026. Regulatory Intermediaries and Chains—the Complex Process of Sense-making for AI Act Implementation in Finland. In *Proceedings of the 59th Hawaii International Conference on System Sciences*. University of Hawaii press.

Widder, D. G.; and Nafus, D. 2023. Dislocated accountabilities in the “AI supply chain”: Modularity and developers’ notions of responsibility. *Big Data & Society*, 10(1): 20539517231177620.

Yin, R. K. 2018. *Case study research and applications*, volume 6. Sage Thousand Oaks, CA.