



SWISSI INSTITUTE FOR AI

Tiered compliant AI system for regulated financial institutions

Multi agentic execution capable framework with built in DLT audit trails for financial operations in DACH

Walter Kurz¹, Reinhard Magg¹

¹Swissi Institute for AI, office@swissi-ai.institute

August 2025

Abstract

We present a compliance-first architecture for AI in regulated finance that treats regulation as an orientation layer rather than a deterministic ruleset. A matrix of regulatory intent and exposure provides a compact classification handle, which a governed policy compiler then maps into concrete prohibitions, obligations and runtime budgets. Prohibitions constrain feasibility and block externalisation, while obligations extend tasks with artefacts that must meet explicit admissibility criteria. Committee activation remains policy-driven and proportionate, preserving efficiency while ensuring supervisory oversight. Evidence, decisions and reason codes are bound to a permissioned DAG with deterministic timestamping, enabling replay, provenance checks and clear attribution of failure. Clause-level legal indexing with effective dates and capability-based agent routing ensure portability across DACH and the wider EU. The result is assurance by construction: compliance is embedded in execution and verifiable by auditors without sacrificing proportionality or transparency.

Keywords: DACH finance; regulated financial institutions; multi agent expert system; policy compiled orchestration; objective under constraints; permissioned DLT; DAG timestamping; audit trails; EU AI Act; MiFID II; DORA; GDPR; human oversight; execution gating; ESG budgets; verification and assurance

1 Introduction

The financial industry combines data intensity, complex decision processes and strict regulatory oversight, making it a natural environment for the application of artificial intelligence. Yet the adoption of AI in this domain is constrained by the absence of architectures that meet supervisory expectations for accountability, resilience and auditability. Existing implementations often prioritise performance or innovation at the expense of compliance, leaving institutions without frameworks that can be both operationally effective and regulatorily sound. This study addresses that gap by proposing a reference architecture for compliant AI in finance, designed to integrate supervisory requirements into system design from the outset. The analysis focuses on financial institutions in the DACH region. Germany, Austria and Switzerland combine deep capital markets, dense supervisory practice and high cross-border integration [1, 2, 3]. Institutions operate under comparable prudential expectations and reporting cultures: Germany and Austria within the EU prudential framework, Switzerland

through equivalence-based alignment in selected domains. This convergence provides a coherent testing ground, supporting methodological consistency while enabling meaningful comparison across jurisdictions [4, 5, 6, 7, 3]. For the purposes of this study, financial institutions are defined as supervised entities that accept deposits or other repayable funds, extend credit or investment services, insure or reinsure risk, manage collective assets, operate payment or settlement systems, or run market infrastructure. [4, 8, 9, 10, 11, 12, 13, 14, 15]. Within the prudential perimeter, the taxonomy distinguishes credit institutions and banks under CRR/CRD, investment firms under MiFID II, insurance and reinsurance under Solvency II, asset managers under UCITS and AIFMD, and market infrastructures such as trading venues, CCPs and CSDs under MiFIR, EMIR and CSDR. Payment and e-money institutions fall under PSD and EMD [4, 5, 8, 15, 9, 13, 14, 12, 16, 10, 11]. A functional taxonomy complements this perimeter by grouping front-office decisioning and advisory, treasury and risk, payments and settlement, compliance and reporting, and back-office operations. Outsourcing and critical ICT providers enter scope where the ICT risk framework applies [17]. European financial services are governed by an extensive regulatory corpus. The EU *acquis* sets binding requirements on prudential soundness, market conduct and investor protection through CRR/CRD, MiFID II and MiFIR, Solvency II, EMIR and CSDR, PSD and EMD, AML directives, MiCA, DORA and GDPR [4, 5, 8, 15, 9, 13, 14, 12, 16, 18, 19, 20, 17, 21]. These instruments specify licensing, governance, outsourcing, logging, testing, resilience and data protection duties, and enable passporting across the single market [5, 22, 23, 17, 21, 24, 25]. Supervised firms must evidence compliance through internal control systems, audit trails and supervisory reporting [9, 22, 23, 6]. The DACH region presents a coherent yet diverse regulatory landscape. Germany applies EU law through national instruments such as the *Kreditwesengesetz* (KWG) and *Wertpapierhandelsgesetz* (WpHG) under BaFin supervision [26, 27, 28]. Austria implements EU law through the *Bankwesengesetz* (BWG) and *Wertpapiersaufsichtsgesetz* 2018 (WAG 2018) under FMA supervision [29, 30, 31]. Switzerland enforces FinSA and FinIA with FINMA, framed by EU equivalence decisions and bilateral arrangements [32, 33, 34, 7, 3]. The region’s institutional scale and diversity are material: universal banks, savings and cooperative networks, global insurers, specialised asset managers and market utilities operate under comparable supervisory expectations, while preserving sufficient heterogeneity to test generality [1, 35]. Burden asymmetries persist: in Austria, smaller firms bear proportionally higher fixed compliance costs, with studies documenting that regulatory obligations weigh more heavily on them than on larger institutions [36, 37, 38, 39].

2 Related Work

Prior literature on regulated finance in Europe has largely examined prudential soundness, market conduct, and investor protection under the comprehensive EU *acquis*, with attention to national transpositions in Germany and Austria and equivalence-based coordination with Switzerland. Governance frameworks for financial technology have increasingly incorporated operational risk and systemic resilience, yet compliance in practice is often confined to procedural assurances rather than embedded as an architectural invariant. Research on multi-agent systems in finance has advanced considerably, particularly in algorithmic trading, portfolio optimisation, and credit risk modelling. Examples include hierarchical deep reinforcement learning for portfolio optimisation [40], multi-agent dynamic portfolio learning systems [41], and surveys of reinforcement learning applications in finance [42, 43]. These approaches typically optimise task-level performance under static policy assumptions and rarely integrate verifiable execution evidence as part of the architecture itself. Parallel work on enterprise logging and model risk management underscores the need for traceability, oversight, and effective challenge. Surveys of audit log security highlight difficulties in ensuring integrity and immutability [44], while research on BPM systems shows how timestamped, tamper-evident logs can support transparent audit regimes [45]. Yet cryptographic timestamping is seldom treated as a first-class architectural principle, and its integration with multi-agent orchestration remains limited. Distributed ledger research in financial services has concentrated on post-trade settlement, asset tokenisation, and shared record-keeping. Permissioned blockchain ecosystems can improve transparency and enable authorised audit access [46], and recent studies link blockchain with AI

to strengthen audit quality [47]. Still, audit-grade logging via directed acyclic graph (DAG) architectures for AI coordination remains under-explored. Empirical analyses of regulatory burden further demonstrate the case for automation and proportionality. Studies of the Austrian perimeter show that compliance overhead scales disproportionately for SMEs, motivating system designs that embed automated evidence generation and allow for proportional controls [39]. Within ICT risk regulation, the Digital Operational Resilience Act (DORA) stresses robust risk management, third-party oversight, resilience testing, and logging requirements [48, 49, 50]. These reinforce the need for architectures that decouple verifiable orchestration from model training, enforce strict role segregation, and support dual-control governance. Taken together, these strands highlight a persistent gap: no existing enterprise architecture simultaneously treats compliance as both a compile-time and run-time constraint, integrates multi-agent coordination under gated execution, and generates supervision-ready evidence through deterministic timestamping and inclusion proofs. This gap is particularly salient for tiered designs that must serve both institutional and consumer interfaces while remaining aligned with DACH-specific regulatory norms.

3 Contribution

Current compliance mechanisms in financial institutions remain predominantly ex-post: obligations are verified through periodic audits, reconciliations, and manual attestations, while system execution itself proceeds without embedded regulatory constraints. This approach creates structural gaps, since audit trails can be incomplete, supervisory checks are delayed, and interpretations of legal texts remain siloed across departments. Controls are reactive rather than preventative, and compliance risk accumulates in the time between transaction and inspection. Addressing these deficiencies requires a design in which regulatory duties are formalised as machine-checkable constraints and enforced as part of orchestration and execution rather than appended as a reporting layer. This research develops an enterprise reference architecture for a tiered, compliant AI system in finance, aligned with the regulatory taxonomy and supervisory context of the DACH region yet transferable across the EU. The architecture encodes regulatory duties as machine-checkable constraints and binds them to orchestration and execution. A policy store and rule engine compile EU and DACH requirements into checks enforced both before and during execution. Multi-agent coordination is role scoped across institutional and consumer tiers, with explicit permissions, explainability thresholds and segregation of duties. Gated execution ensures that sensitive actions, such as order placement, occur only under dual control, and every step produces audit-ready evidence. Verification is achieved through a permissioned distributed ledger that employs a directed acyclic graph (DAG) for deterministic timestamping and inclusion proofs. The ledger captures routing instructions, budget usage and control attestations without imposing settlement semantics on application workflows. Decisions are formulated as objectives under constraints, with runtime budgets for risk, latency and ESG, enabling transparent trade-offs at schedule time. The architecture remains technology-agnostic at model level and interoperable with existing systems via supervised queues and APIs. In this way, it provides a reference design that meets supervisory expectations in DACH while supporting broader applicability across EU jurisdictions.

4 System Architecture

The proposed system architecture embeds regulatory compliance as a first-class design primitive rather than an external audit layer added post hoc. Orchestration and execution are bounded by machine-enforceable constraints compiled directly from legal provisions, ensuring that system behaviour remains verifiable at run time. The architecture is tiered across institutional and consumer contexts, portable across jurisdictions, and agnostic to specific model technologies. Its emphasis lies on properties that supervisory authorities recognise as audit-ready: determinism of execution, separation of roles, and transparent evidence generation. A distinctive element of the design is the integration of a permissioned distributed ledger that employs a directed acyclic graph (DAG) structure to anchor audit trails. Each agent action, constraint evaluation, and gating decision is recorded as a

tamper-evident event with deterministic timestamping and inclusion proofs. Unlike traditional logging, which remains vulnerable to post-processing or selective disclosure, the DAG ledger produces an immutable sequence of verifiable attestations that can be inspected by internal control functions and supervisory authorities. The ledger is deliberately lightweight: it does not impose settlement semantics on application workflows but serves solely as a compliance substrate, capturing routing instructions, evidence packs, and oversight signatures. This enables continuous assurance without altering financial transaction flows.

4.1 Agnostic Regulatory Intent and Exposure Matrix

European financial regulation is dense and layered across prudential, conduct, market integrity, data protection and ICT risk instruments. Designing an enterprise system by listing acts and articles provides little direct operational guidance. To address this, we introduce a regulation-agnostic taxonomy that functions as a design primitive. Its purpose is to translate legal text into machine-enforceable control functions that govern orchestration and execution while remaining portable across DACH and the wider EU. We adopt a functional reduction of legal provisions to their operational effect on system behaviour: either a rule forbids an outcome or it demands an outcome with evidence. Compliance therefore reduces to abstention from a prohibited act or to the performance of a positive act that leaves a verifiable trace. This yields a law- and regulation-agnostic orientation in which agents need only decide whether to avoid or to act with evidence, while the specific statutory text is compiled at a later stage into guard families and evidence templates, with exposure handled separately for internal versus external outputs. Practically, this agnostic treatment simplifies handling heterogeneous regulations because the agent can synthesise, per task and context, a checklist of prohibitions and obligations that orchestration executes and tracks, with each item either preventing externalisation or requiring a verifiable artefact with provenance, deterministic timestamping and DAG inclusion proof; alternatively a dedicated compliance agent can apply the same checklist as a pre-output gate, sign the evidence pack, and authorise or withhold release.¹ The taxonomy is formalised as a two-dimensional schema, the *Regulatory Intent and Exposure Matrix*. Each legal provision receives two labels *rule type*, capturing intent as either Prohibition or Obligation, and *exposure*, capturing context as Internal or External.

$$v = (t, e), \quad t \in \{\text{Prohibition, Obligation}\}, \quad e \in \{\text{Internal, External}\} \quad (1)$$

The matrix is an orientation device rather than an execution rule. The label $v = (t, e)$ expresses a stance that guides reasoning and proposes candidate guard families and evidence templates; the selection and activation of concrete gates is compiled later by the rule engine from the policy store and context features, and may be revised by human oversight. In short, v narrows the search space; it does not by itself prescribe what the system must do. A compact label provides the canonical representation. This schema enables candidate guard families and evidence templates to be proposed without first naming a statute, ensuring portability across instruments and jurisdictions. Prohibitions mark contexts where the system should consider blocking externalisation; obligations mark contexts where the system should consider requiring a positive act and admissible evidence. Whether these candidates become active gates depends on compiled policy, context, and approvals, not on the matrix alone.

¹A prohibition can be illustrated with a consumer-tier agent that is not permitted to instruct a client to purchase a particular security or to transmit an order on their behalf without their permission. In this case the guard blocks both recommendation and order externalisation, recording at most a hold or release attestation on the DAG ledger. An obligation is exemplified by onboarding procedures, which require identification and know-your-customer documentation. Here the workflow executes the prescribed checks and generates a signed artefact with provenance metadata, deterministic timestamping, and a DAG inclusion proof before any advisory or execution services can be activated.

This binary mapping is used for orientation in reasoning and as a compact handle for later compilation, as shown in Figure 1.




Figure 1: Taxonomy of regulatory intent. Each legal provision (regulation, law, directive, guidance) maps to Prohibition or Obligation, indicating the primary reasoning objective.

Figure 1 frames two reasoning stances rather than execution rules. A Prohibition prompts scrutiny of potentially disallowed behaviour and a hold on externalisation until resolved. An Obligation prompts identification of duties and anticipation of evidence likely required if the task proceeds. Clauses are stored as atomic statements with mapping and citation to ensure fast reading and consistent interpretation. Concrete gates and logging are derived later once reasoning has reached sufficient confidence.

The second axis classifies exposure. Internal exposure refers to outputs confined to the institution, while external exposure refers to outputs delivered to clients, counterparties, authorities or markets. The distinction is contextual and combines with intent in Figure 2.




Figure 2: Taxonomy of exposure. Internal indicates outputs confined to the institution; external indicates outputs leaving the institutional boundary. Mixed workflows default to external.

This classification orients reasoning by establishing the destination of outputs. Once combined with intent, it guides which questions and candidate actions are relevant. Figure 3 presents the joint view.




Figure 3: Regulatory intent and exposure matrix. Quadrants orient reasoning and suggest candidate guard families and evidence expectations. Mixed workflows default to external.

At run time an agent performs a lightweight context check against the proposed matrix. If no mapped provision appears to apply, the agent continues under general IT controls and ordinary policy. If the check suggests applicability, the agent classifies the context using the matrix labels and treats that label as an orientation handle for reasoning rather than as an enforcement command, subject to revision as scope and evidence evolve. The context check can be formalised as a binary predicate that determines whether a task falls under a mapped legal provision:

$$\chi : \mathcal{T} \rightarrow \{0, 1\}, \quad \chi(\tau) = \begin{cases} 1 & \text{if the mapped provision applies to } \tau \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

If $\chi(\tau) = 0$, the task proceeds under baseline IT controls and operational policy. If $\chi(\tau) = 1$, the task is assigned a regulatory label $v = (t, e)$ as defined above. This label serves as an orientation handle that anchors subsequent reasoning.

Once classified, the agent frames the task as a constrained optimisation problem of the form

$$\begin{aligned} \max_{\pi \in \Pi} \quad & U(\pi, \tau) \\ \text{s.t.} \quad & C(\pi, \tau, v) = \text{true}, \end{aligned} \quad (3)$$

where π is a candidate policy, $U(\pi, \tau)$ denotes the utility of executing π on task τ , and $C(\cdot)$ encodes the selected guard families and evidence checks compiled from v and context; the matrix label v orients this compilation but does not, by itself, determine execution. Prohibitions are represented as hard constraints considered for activation to prevent disallowed outcomes, while obligations are represented as candidate duties that, if selected by policy, require admissible evidence. If no provision applies ($\chi(\tau) = 0$), the constraint set C collapses to baseline IT controls and ordinary operational policy. This formulation connects the taxonomy to decision-making without collapsing orientation into determination.

Minting the classification, checklist and gating decisions to the DAG enables ex post assurance and explicit failure analysis: an auditor can reconstruct the task state $\phi(\tau)$ (see (5)), recompute $\chi(\tau)$ (see (2)) and the label $v = (t, e)$ (see (1)) to test whether the agent misunderstood the context, verify that each selected guard $C(\pi, \tau, v)$ in the optimisation constraint (see (3)) was satisfied at the moment of externalisation, and match every obligation item to a concrete artefact with provenance, deterministic timestamping and inclusion proof;

this makes distinct failure modes observable, including (i) misclassification of intent or exposure, (ii) correct classification and checklist generation but incomplete execution, (iii) execution with insufficient evidence quality or missing signatures, and (iv) sequencing or gating errors that allowed externalisation without preconditions, as well as budget breaches for risk, latency or ESG. The proposal turns compliance from after-the-fact attestation into assurance-by-construction, where reasoning, decisions and evidence are bound cryptographically so reviewers can check not only outcomes but also whether the agent’s understanding and follow-through were correct.

We propose a short, curated aid for cases where the context check indicates applicability. A quick reference vector is a concise entry prepared by a qualified oversight committee; it orients reasoning without prescribing execution. Each entry names the action archetype, states the typical matrix label, notes salient triggers and the initial evidence posture, and gives an escalation contact. Entries are versioned and signed so that provenance remains clear; they can be revised as scope or information changes. When a task matches such an entry, the agent adopts the entry as orientation and continues analysing the task toward the decision objective; if no entry fits, the agent proceeds under baseline controls and may flag the gap for later inclusion. The intent is speed and consistency in judgement, not hard rules.

$$r = (a, v, \kappa, \epsilon, \nu, \sigma) \quad (4)$$

where a is the action archetype, $v = (t, e)$ is the matrix label from Equation (1), κ captures salient triggers and context features, ϵ records the initial evidence posture, ν denotes version and effective interval, and σ is the oversight signature.

$$R(\tau) = \{r \in Q \mid \text{sim}(\kappa_r, \phi(\tau)) \geq \theta\} \quad (5)$$

where Q is the approved set, $\phi(\tau)$ is a compact feature map of the task, and θ is a reviewable threshold. If $R(\tau) = \emptyset$, the task proceeds under baseline controls and the context check is recorded for later improvement. If $R(\tau) \neq \emptyset$, the top-ranked entry $\hat{r}$ is adopted as orientation and the analysis continues toward the decision objective in Equation (3).

4.2 Prohibitions as feasibility and obligations as task extension

The matrix provides orientation only; concrete enforcement is selected later by policy compilation and may be revised by human oversight. Given the label $v = (t, e)$ from (1) and the context determined by the predicate $\chi(\tau)$ in (2), the rule engine surfaces candidate prohibitions P and obligations O , then selects active subsets $P^* \subseteq P$ and $O^* \subseteq O$ together with a strict partial order $\prec$ over $P^* \cup O^*$. The order $\prec$ permits short-circuit evaluation where appropriate and defines a narrow top tier $P^\perp \subseteq P^*$ of dominant disqualifiers. Obligations map to mandatory sub-actions and a verifiable evidence artefact; let $E(\cdot)$ denote the admissibility predicate capturing provenance, signatures, deterministic timestamping and DAG inclusion proof. None of these selections is implied by the matrix itself; the matrix narrows the search space, the policy store determines activation.

Prohibitions are encoded as feasibility constraints that prevent disallowed outcomes, while obligations extend the task with preconditions that must hold before optimisation under (3) is meaningful. The feasible policy set for task τ is

$$\Pi_{\text{feas}}(\tau) = \left\{ \pi \in \Pi \mid \bigwedge_{p \in P^*} \neg p(\pi, \tau) \wedge \bigwedge_{o \in O^*} E(\text{artefact}_o(\tau)) \right\}. \quad (6)$$

Optimisation proceeds lexicographically to respect feasibility first and utility second,

$$\pi^* \in \arg \max_{\pi \in \Pi_{\text{feas}}(\tau)} U(\pi, \tau) \quad (7)$$

with optional secondary terms in U for obligation quality such as latency, cost or completeness. This construction preserves the law- and regulation-agnostic stance: prohibitions shape the admissible set, obligations extend the

task graph and demand admissible artefacts, and any burden from satisfying obligations may then be traded within the objective once feasibility is secured.

Externalisation refers to any output that leaves the institutional boundary, including recommendations, orders, client-facing summaries, supervisory reports and API calls, whereas internal artefacts, logs and operator messages remain non-external unless later reused. The gate outcome formalises block versus release for a proposed output $y(\pi, \tau)$,

$$\text{Externalise}(\pi, \tau) = \begin{cases} \text{DENY}(c) & \text{if } \exists p \in P^\star \text{ with } p(\pi, \tau) = 1 \text{ or } \exists o \in O^\star \text{ with } E(\text{artefact}_o(\tau)) = \text{false}, \\ \text{ALLOW}(y) & \text{otherwise,} \end{cases} \quad (8)$$

where c is a machine- and human-readable reason code minted to the DAG with a hold or release attestation. Blocking does not suppress internal reasoning; the system returns an explicit explanation to operators and records the decision and grounds for audit.

Short-circuit evaluation reduces latency and user friction without sacrificing assurance. If there exists a dominant disqualifier $p \in P^\perp$ with $p(\pi, \tau) = 1$, then $\text{Externalise}(\pi, \tau) = \text{DENY}(c)$ by (8) and the remaining lower-tier checks in $P^\star \cup O^\star$ are marked *skipped by policy*. The audit pack minted to the DAG includes the orientation set, the selected $P^\star, O^\star$, the order $\prec$, the firing guard in $P^\perp$, the skipped set with its policy basis, all artefact hashes and admissibility outcomes $E(\cdot)$, and the reason code c . This enables an auditor to reconstruct the state $\phi(\tau)$ used in the quick-reference matching (5), to verify that the selected guards $C(\pi, \tau, v)$ in the optimisation constraint (3) were satisfied at the moment of externalisation, and to distinguish failure modes such as misclassification, incomplete execution of obligations, insufficient evidence quality, or sequencing errors that allowed externalisation without preconditions.

A composite example illustrates the semantics. Consider a cross-border retail transaction where the amount exceeds the threshold for retail investors and the client is under age. Orientation via v surfaces multiple candidates; policy selects $P^\star$ that includes an under-age prohibition and places it in $P^\perp$ under $\prec$. The under-age guard fires, so $\text{Externalise}(\pi, \tau) = \text{DENY}(c)$ by (8); cross-border and amount checks are skipped by policy with the dominance certificate recorded in the audit pack. If, in a different context, no dominant disqualifier fires, obligations such as identification and know-your-customer documentation become active preconditions, and externalisation is permitted only once admissibility $E(\cdot)$ holds, after which utility is optimised over $\Pi_{\text{feas}}(\tau)$ according to (7). In both cases the matrix remains an orientation device; selection and activation are determined by compiled policy and oversight, and the full reasoning trail is minted to the DAG for ex post assurance.

We assume only mild regularity, namely that dominant disqualifiers in $P^\perp$ remain policy invariant for a given task state, so that $p(\pi, \tau) = p(\pi', \tau)$ holds for all $\pi, \pi' \in \Pi$; in practice these include conditions such as age thresholds, sanctions hits, or bans tied to specific product tiers. The admissibility predicate $E(\cdot)$ is used to capture evidence quality requirements, including provenance, signatures, deterministic timestamping, and DAG inclusion, and is evaluated directly on the artefacts compiled for τ .

Under these assumptions the feasible set is monotone: if the selected sets expand such that $P_1^\star \subseteq P_2^\star$ and $O_1^\star \subseteq O_2^\star$, then it follows that $\Pi_{\text{feas}}^{(2)}(\tau) \subseteq \Pi_{\text{feas}}^{(1)}(\tau)$, where $\Pi_{\text{feas}}^{(k)}$ is defined by (6) under the selection $P_k^\star, O_k^\star$. The result follows immediately from (6), since adding further prohibitions or obligations introduces additional conjuncts that can only reduce the satisfying set.

Short-circuit evaluation is sound whenever dominant checks are policy invariant: if a single $p \in P^\perp$ evaluates to one for some π , then it evaluates to one for all π' , which implies $\Pi_{\text{feas}}(\tau) = \emptyset$ by (6) and forces $\text{Externalise}(\pi, \tau) = \text{DENY}(c)$ by (8). Conversely, if no element of $P^\perp$ fires and all selected obligations satisfy admissibility so that $E(\cdot) = \text{true}$, then $\Pi_{\text{feas}}(\tau)$ remains nonempty and the optimisation problem in (7) is well defined. Evaluation both terminates and remains auditable. Since the sets $P^\star \cup O^\star$ are finite and $\prec$ is a strict partial order, a topological evaluation sequence necessarily exists, with short-circuiting at $P^\perp$ further bounding runtime. The DAG record

contains the orientation set, the selected P^* , O^* , the order $\prec$, any firing element of $P^\perp$, the admissibility outcomes $E(\cdot)$, and the resulting gate decision (8), which together enable replay and independent re-evaluation using (2), (1), (6) and (3).

If obligation quality is intended to influence the objective once feasibility has been established, the framework admits a straightforward refinement in which $U(\pi, \tau) = U_0(\pi, \tau) - \lambda^\top q(\tau)$, with $q(\tau)$ aggregating latency, cost, and completeness across satisfied obligations and $\lambda \geq 0$ denoting policy weights. The lexicographic regime in (7) guarantees that such trade-offs are only evaluated within the feasible set $\Pi_{\text{feas}}(\tau)$.

4.3 Agent activation by intent and exposure

The orientation matrix provides the initial signal, but the selection and activation of executors is determined by compiled policy and remains subject to oversight revision. Institutions may maintain standing rosters of domain-specific agents and persistent committees for recurring functions, while reserving the possibility of ad hoc committees when context requires. In a banking environment, a payments roster typically spans functions such as wiring funds between accounts, executing foreign-exchange conversions, and arranging treasury transfers. For an investment firm under MiFID II, routine activities encompass the reception and transmission of client orders, order execution on behalf of clients, the provision of investment advice with suitability and appropriateness checks, portfolio management, and post-trade reporting. These rosters remain dormant until activated by the orientation label $v = (t, e)$ from (1) in combination with the context check $\chi(\tau)$ defined in (2). Committee formation follows a policy-driven logic. Where $\chi(\tau) = 0$ or the task is assessed as internal and low risk, such as drafting an internal email or scheduling a meeting, a single role-scoped agent executes under baseline IT controls. When $\chi(\tau) = 1$, activation depends on the exposure level e , the number and dominance of active checks, and the availability of runtime budgets. This is formalised as

$$\text{Committee}(\tau) = \begin{cases} S(\tau, v, \text{context}) & \text{if } \chi(\tau) = 1 \wedge C(\tau), \\ \{\text{primary agent}\} & \text{otherwise,} \end{cases} \quad (9)$$

where

$$C(\tau) := (e = \text{External}) \vee (|P^*| + |O^*| \geq k) \vee (P^\perp \neq \emptyset) \vee \text{budget breach (risk, latency, or cost)}. \quad (10)$$

Here $S(\cdot)$ yields the minimal role set required for the task, P^* , O^* denote the selected checks defined in the previous subsection, $P^\perp$ designates the dominant disqualifiers, and k is a policy-defined threshold. The rule ensures proportionality and avoids unnecessary coordination overhead.

Execution proceeds through a structured handshake. The orchestration agent forwards the task τ together with the orientation v to the designated agent set or committee, which undertakes its domain-specific responsibilities and returns artefacts, rationales, and provisional outcomes. A dedicated compliance agent then applies the pre-output gate defined in (8); if admissibility is confirmed for all selected obligations and no prohibition fires, the output is released, whereas any failure results in blocked externalisation accompanied by a reason code. Each stage mints a committee token to the DAG ledger, recording membership, role scopes, selected checks, reason codes, and inclusion proofs, which collectively support later reconstruction and independent re-evaluation under (3).

Examples illustrate proportionality and dominance. A domestic, low-value internal transfer with $\chi(\tau) = 0$ is executed by a single payments agent without committee escalation, reflecting the proportionality condition that internal exposure with few selected checks below the policy threshold k activates only the primary role. A cross-border retail securities order with $\chi(\tau) = 1$ and $e = \text{External}$ activates the standing “securities orders” set comprising client categorisation and suitability, order execution, best-execution monitoring, and transaction reporting; if a dominant disqualifier in $P^\perp$, such as an under-age client, fires, the committee collapses to the

compliance gate and the externalisation returns $\text{DENY}(c)$ under (8), with all skipped checks logged as *skipped by policy*. Formally,

$$P^\perp \neq \emptyset \wedge \exists p \in P^\perp : p(\pi, \tau) = 1 \implies \text{Committee}(\tau) = \{\text{compliance gate}\}, \quad (11)$$

which ensures that functional agents are excluded by policy while a reason-coded denial is minted under (8). In both directions, whether escalation occurs because exposure or policy thresholds demand it, or collapse follows from dominance being triggered, the orientation matrix identifies relevant agents while compiled policy and supervisory oversight determine actual activation.

4.4 Legal corpus indexing and agent routing

The legal corpus is treated as a structured, versioned asset rather than a flat vector space. Each instrument is normalised, segmented, and cited at clause level, with every chunk l carrying canonical metadata $\text{cite}(l) = \langle \text{jurisdiction, instrument_id, article, paragraph, } \nu \rangle$, where ν encodes version and effective interval, supplemented by tags for instrument class (prudential, conduct, data protection, ICT), intent and exposure labels, and explicit cross-references. Ingestion preserves both the original language and a controlled translation, while recording equivalence mappings between recitals, articles, RTS/ITS, and supervisory Q&A, so that downstream retrieval can resolve different references to a single, time-scoped citation.

Indexing is hybrid in order to preserve meaning as well as citations. Three coordinated views are maintained for the same chunks: a dense semantic index for embeddings $\phi(l)$, a sparse citation and keyword index for exact retrieval, and a citation graph G linking instruments through “refers to,” “implements,” and “interprets” edges. Queries are constructed from task features $\psi(\tau)$, the orientation $v = (t, e)$ from (1), triggers κ from the quick-reference vector (4), and the exposure e . Retrieval proceeds in two stages: a semantic pre-filter selects top- k candidates $\tilde{L}(\tau)$ by similarity $\text{sim}(\psi(\tau), \phi(l))$ under jurisdiction and effective-date filters, and a reranker then fuses sparse scores with graph proximity to produce the final set $L(\tau)$. Every returned item includes $\text{cite}(l)$ together with its text span, ensuring human-readable grounding.

Clustering is dual to reflect legal reading practice. We compute instrument-anchored clusters C_{inst} for each statute or linked set of acts, and aspect clusters C_{asp} derived from the citation graph and tag distributions, covering recurring themes such as suitability, client categorisation, personal recommendations, onboarding, and record-keeping. Chunks may appear in both views. Aspect clusters allow rapid coverage across instruments, while instrument clusters ensure that the exact article and paragraph remain available for citation and audit. The design avoids reliance on a single global cluster.

Routing is performed by capability rather than by statute. Institutions register domain agents with declared capabilities and evidence schemas, covering for example payments, FX, treasury, onboarding, suitability and appropriateness, personal recommendations, order execution, reporting, data protection, and ICT controls. Orientation v , exposure e , and the retrieved set $L(\tau)$ jointly determine which capability bundles are relevant. The router applies the same selection logic as (9) to activate either a single role or a minimal committee. A dedicated citation clerk agent attaches $\text{cite}(\cdot)$ and extracts the governing spans, while a compliance agent executes the pre-output gate (8). This design keeps knowledge current and avoids brittle “one agent per law” silos.

Correlation between chunks and their governing law is guaranteed through metadata rather than embeddings. Each chunk carries its canonical citation, effective dates, and instrument identifiers (e.g. CELEX for EU acts, BaFin circular IDs, FINMA RS numbers). When retrieval proposes candidates, a citation-consistency check eliminates items whose $\text{cite}(\cdot)$ lies outside the jurisdiction or effective interval of the task, and a majority-vote over top- n candidates stabilises the legal basis. The resulting citation pack $\{\text{cite}(l), \text{span}(l)\}_{l \in L(\tau)}$ is signed by the citation clerk and minted to the DAG.

Mapping legal text to matrix intent and exposure constitutes orientation rather than determination. A curator-assisted classifier proposes $m(l) \in \{\text{Prohibition, Obligation}\} \times \{\text{Internal, External}\}$ for each chunk, with all proposals reviewed and versioned. At run time the router applies $m(l)$ together with context to select P^*, O^* and the order $\prec$ (cf. (6), (7)), while concrete guard templates and artefact schemas are compiled from the policy store. The complete retrieval and routing trace $\langle v, e, \psi(\tau), L(\tau), \text{cite}(\cdot) \rangle$ is minted to the DAG alongside the gate result, enabling replay and auditor verification.

This design addresses practical concerns. Legal aspects are not collapsed into a single global cluster but organised into multiple coordinated views. Agents are not tied one-to-one with statutes but are registered by capability and routed according to orientation and retrieved evidence. Correlation between chunks and laws is secured by canonical citations and effective dates rather than embedding heuristics. The outcome is portable across jurisdictions, resilient to textual drift, and ready for supervisory scrutiny.

4.5 Evidence admissibility model

Admissibility must be explicit to ensure that obligation checks in (6)–(7) are verifiable at run time and reproducible on audit. Each obligation $o \in O^*$ is modelled as producing an artefact bundle

$$\text{artefact}_o(\tau) = \langle \text{text}, \text{hash}, \text{sig}, \text{prov}, \text{time}, \text{dag}, \text{ret} \rangle. \quad (12)$$

The admissibility predicate is defined as the conjunction of policy-specified quality dimensions,

$$E(\text{artefact}_o(\tau)) \Leftrightarrow Q_{\text{prov}} \wedge Q_{\text{sig}} \wedge Q_{\text{time}} \wedge Q_{\text{dag}} \wedge Q_{\text{ret}} \wedge Q_{\text{access}}, \quad (13)$$

where Q_{prov} establishes provenance and chain of custody, Q_{sig} validates authorised signatures and key material, Q_{time} confirms trusted time binding, Q_{dag} attests DAG inclusion through Merkle proof and index, Q_{ret} enforces retention horizon and lawful basis, and Q_{access} guarantees role-scoped readability with appropriate redaction of personal data.

Let $\text{Hash}(\cdot)$ denote a collision-resistant digest and $\text{TS}(\cdot)$ a trusted timestamp. A minimal admissible bundle must then satisfy

$$\text{hash} = \text{Hash}(\text{text}) \wedge \text{sig} = \text{Sign}_{k_{\text{oversight}}}(\text{hash}) \wedge \text{time} = \text{TS}(\text{hash}) \wedge \text{dag} = \text{Include}(\text{hash}, \text{time}), \quad (14)$$

with prov linking to registered source identifiers and ret demonstrating that the object is scheduled for retention in accordance with policy and applicable law. Privacy is preserved by anchoring digests and metadata on the DAG while raw artefacts remain in controlled storage, accessible only through capability- and role-scoped URIs. This definition operationalises the obligation term in the feasible set (6) and ensures that evidence produced by obligation checks is admissible by construction and audit-ready for supervisory verification.

4.6 Policy compilation and governance

The orientation matrix provides the initial signal, while policy compilation determines the active checks. Given the label $v = (t, e)$ from (1) and the task context $\chi(\tau)$ from (2), a compiler maps orientation and context into selected guards and their evaluation order, subject to human oversight and revision. The policy store is a versioned, time-scoped rule base with effective intervals, reviewer sign-off, and rollback support.

Compilation is modelled as a deterministic function of the effective policy version ν at decision time t :

$$\text{Compile}_{\nu, t}(v, \text{context}, \tau) \rightarrow (P^*, O^*, \prec, P^\perp, B), \quad (15)$$

where $P^\star \subseteq P$ and $O^\star \subseteq O$ are the selected prohibitions and obligations, $\prec$ is a strict partial order over $P^\star \cup O^\star$, $P^\perp \subseteq P^\star$ are dominant disqualifiers, and $B = \langle B_{\text{risk}}, B_{\text{lat}}, B_{\text{cost}} \rangle$ are the active runtime budgets. Selection is constrained by jurisdiction and effective dates from the legal index, by the exposure label e , and by any institutional policies specific to product or client tier. The matrix narrows the candidate set but does not enforce activation.

Budgets operate both at compile time and at run time. A budget-breach predicate informs committee activation in (9) and the externalisation gate in (8):

$$\text{Breach}(\tau) := (\text{risk}(\tau) > B_{\text{risk}}) \vee (\text{latency}(\tau) > B_{\text{lat}}) \vee (\text{cost}(\tau) > B_{\text{cost}}). \quad (16)$$

Here risk may be a calibrated model score, while latency and cost are measured against defined service tiers. If the predicate fires, the compiler may escalate to a committee under (9) or deny externalisation under (8).

Obligation templates are bound to admissibility during compilation. For each $o \in O^\star$, the compiler selects an evidence schema and admissibility test $E(\cdot)$ from (13), producing a concrete template for $\text{artefact}_o(\tau)$ with required provenance, signatures, trusted time, and DAG inclusion. The feasible set (6) and optimisation (7) then operate over these compiled guards rather than the full candidate space.

Governance ensures reproducibility and institutional control. Each compilation yields a certificate

$$\text{Cert}_{\text{comp}} = \langle v, \text{hash}(\text{context}), P^\star, O^\star, \prec, P^\perp, B, \nu, \sigma_{\text{compiler}}, \sigma_{\text{oversight}} \rangle, \quad (17)$$

which is signed both by the compiler and by an oversight key, and anchored on the DAG with a reference to the policy store version. Replays reconstruct $\text{Cert}_{\text{comp}}$ and verify that identical inputs under the same ν yield identical selections and order.

Overrides are permitted under controlled conditions. An authorised reviewer may issue $\text{Override}(\text{Cert}_{\text{comp}}, \Delta)$ that modifies $(P^\star, O^\star, \prec)$ with reason code and dual control. Safety monotonicity applies to dominant prohibitions: if $p \in P^\perp$ and $p(\pi, \tau) = 1$, the override cannot deactivate p when exposure is external ($e = \text{External}$); any attempted weakening is logged and forces $\text{DENY}(e)$ under (8). Strengthening is always admissible, while removal of non-dominant checks requires reviewer signatures and the issuance of a new certificate version ν' .

Determinism and time anchoring close the loop. The compiler is deterministic given $(v, \text{context}, \nu, t)$, and any change in output is attributable to context updates, policy version changes, or time-scoped rule amendments. Trusted time binds both ν and $\text{Cert}_{\text{comp}}$ to replayable evidence via (14). The result is a governance process that integrates admissibility, feasibility, committee formation, and externalisation, while keeping the orientation matrix as a guidance device rather than a source of hard-coded execution rules.

4.7 Stylised simulation pseudocode

The following pseudocode illustrates a pass through context detection, policy compilation, dominant disqualifier short-circuiting, admissibility checks, and the externalisation gate defined in Equation (8).

Algorithm 1 Execution flow for task τ (Part I)²

```

1: Input: task  $\tau$ , context, candidate policy set  $\Pi$ , policy version  $\nu$ , time  $t$ 
2: if  $\chi(\tau) = 0$  then
3:    $v \leftarrow (\text{None}, \text{Internal})$ 
4:    $\text{MINT\_DAG}(\tau, \text{decision} = \text{"allow\_internal"})$ 
5:   return allow\_internal
6: else
7:    $v \leftarrow \text{ORIENT}(\text{context})$ 
8:    $(P^*, O^*, \prec, P^\perp, B) \leftarrow \text{Compile}_{\nu, t}(v, \text{context}, \tau)$ 
9:   if  $\text{Breach}(\tau, B)$  then
10:     $\text{MINT\_DAG}(\tau, \text{decision} = \text{"deny"}, \text{reason} = \text{"budget breach"})$ 
11:    return deny(reason = "budget breach")
12:   end if
13:   for all  $p \in P^\perp$  do
14:     if  $\exists \pi \in \Pi$  such that  $p(\pi, \tau) = 1$  then
15:        $\text{MINT\_DAG}(\tau, \text{decision} = \text{"deny"}, \text{reason} = \text{"dominant prohibition"})$ 
16:       return deny(reason = "dominant prohibition")
17:     end if
18:   end for
19: end if

```

(continued on next page)

Algorithm 2 Execution flow for task τ (Part II, continued)

```

20: if  $\chi(\tau) \neq 0$  then
21:   for all  $o \in O^*$  in order  $\prec$  do
22:      $a \leftarrow \text{EXECUTE}(o, \tau, \text{context})$ 
23:     if  $E(a) = \text{false}$  then
24:        $\text{MINT\_DAG}(\tau, \text{decision} = \text{"deny"}, \text{reason} = \text{"inadmissible obligation"}, \text{failed} = o)$ 
25:       return  $\text{deny}(\text{reason} = \text{"inadmissible obligation"})$ 
26:     end if
27:   end for
28:    $\Pi_{\text{feas}} \leftarrow \{\pi \in \Pi \mid \forall p \in P^* : p(\pi, \tau) = 0, \forall o \in O^* : E(a) = \text{true}\}$ 
29:   if  $\Pi_{\text{feas}} = \emptyset$  then
30:      $\text{MINT\_DAG}(\tau, \text{decision} = \text{"deny"}, \text{reason} = \text{"no feasible policy"})$ 
31:     return  $\text{deny}(\text{reason} = \text{"no feasible policy"})$ 
32:   end if
33:    $\pi^* \leftarrow \arg \max_{\pi \in \Pi_{\text{feas}}} U(\pi, \tau)$ 
34:    $\text{MINT\_DAG}(\tau, \text{decision} = \text{"allow"}, \text{output} = \pi^*)$ 
35:   return  $\text{allow}(\text{output from } \pi^*)$ 
36: end if

```

This pseudocode shows how the system detects applicability, assigns the orientation label, compiles policy deterministically, short-circuits on dominant disqualifiers, evaluates admissibility for obligation artefacts, optimises over the feasible set defined by (6)–(7), executes the externalisation gate (8), and records every decision path by minting a replayable, time-bound audit pack.

5 Discussion

The architecture treats regulatory constraints as first-class objects while preserving the distinction between orientation and determination. The matrix provides a compact label $v = (t, e)$ for classification and routing, while policy compilation selects and orders prohibitions and obligations subject to possible revision through oversight. Compliance is operationalised through feasibility and admissibility as defined in (6), (7), and (8), with every obligation bound to trusted time and anchored on a permissioned DAG to ensure replayability, provenance validation, and clear attribution of failure. A central benefit arises from the separation of concerns. Orientation remains stable and portable across jurisdictions, whereas activation is institution specific through the policy store and its governance processes. Prohibitions do not trade against utility, and obligations extend the task with artefacts that must satisfy the admissibility predicate $E(\cdot)$ defined in (13). Committee activation under (9) ensures proportionality, while short-circuiting of dominant checks formalised in (11) reduces latency without compromising the completeness of audit trails. There are trade-offs. Misclassification or an excessively broad dominance set may lead to false blocks, while incomplete policy stores can result in false releases. The legal index and the intent–exposure mapping $m(l)$ therefore require continuous curation with versioning and reviewer sign-off. Privacy constraints further limit what can be anchored on the ledger; hashes and metadata are sufficient for attestation, but disciplined key management and controlled storage are required for raw artefacts. Determinism also relies on stable time sources and reproducible compilation given $(v, \text{context}, \nu, t)$. The scope of the design is clearly delimited. It does not replace legal interpretation or product governance but instead provides an executable substrate that is law- and regulation-agnostic at the orientation layer and institution specific at the activation layer. Portability across DACH and the wider EU follows directly from this separation, since only

²Initial pseudocode generated using Claude Code in response to structured system prompts [51].

compilation rules and evidence templates vary with local implementation.

6 Conclusion

The paper advances a compliance-first architecture that positions regulation as an orientation layer rather than a deterministic ruleset. The Regulatory Intent and Exposure Matrix provides the compact handle $v = (t, e)$ for classification and routing, while execution is delegated to a governed policy compiler. Prohibitions constrain feasibility through the admissible set $\Pi_{\text{feas}}(\tau)$ in (6) and block externalisation via (8). Obligations extend tasks with admissible artefacts that must satisfy the predicate $E(\cdot)$ in (13). Optimisation then proceeds over $\Pi_{\text{feas}}(\tau)$ under lexicographic semantics in (7). Proportionality follows from policy-driven committee activation in (9), while dominance rules in (11) enable early termination without compromising audit completeness. Evidence, decisions and reason codes are bound to a permissioned DAG with deterministic timestamping, ensuring replayability, provenance checks and precise failure attribution. Portability across DACH and the wider EU results from the separation between orientation and activation, a clause-level indexed legal corpus, and capability-based agent routing. The architecture transforms compliance from retrospective attestation into assurance by construction, offering a generalisable design pattern that remains law- and regulation-agnostic at the orientation layer while retaining institution-specific control at activation.

7 Limitations and Future Research

Several limitations remain. The design depends on curated mappings from legal text to orientation $m(l) \rightarrow v$ and on the integrity of policy compilation; drift or gaps can cause false blocks or unintended releases. Formal guarantees for compiler determinism and reproducibility under $(v, \text{context}, \nu, t)$ require explicit proof obligations and differential testing. The admissibility predicate $E(\cdot)$ must be calibrated to institutional standards for provenance, signatures, trusted time and retention, and cryptographic choices and key management directly affect audit reliability. Ledger privacy remains constrained by linkage risks across hashes and metadata, and selective disclosure, salting and zero-knowledge proofs warrant further evaluation. Institutional dependencies also arise. Externalisation boundaries can shift when internal artefacts are reused in external contexts, which necessitates systematic exposure re-checks and reuse monitors. Committee thresholds k and runtime budgets B determine latency and cost, so empirical calibration, red-teaming and scenario replay are needed to quantify trade-offs and client impact. Retrieval robustness relies on similarity functions, feature maps and citation consistency, suggesting the need for adversarial tests and benchmarks specific to legal retrieval. Human and governance factors remain material. Curation of the legal index and the intent–exposure mapping requires ongoing reviewer effort with versioning and sign-off. Reviewer workload, override workflows and reason-code clarity all influence governance quality and practical adoption. Future research will extend evaluation on institution-specific scenarios in DACH, develop auditor-facing query suites tied to (2), (1), (6), (13) and (8), and investigate privacy-preserving ledger designs and time-binding mechanisms that improve verifiability without undermining data minimisation.

It should also be noted that empirical evaluation at meaningful scale would require full system deployment across institutional and regulatory contexts, which entails investments beyond the scope of academic research. Smaller prototypes would not capture the systemic properties of budget enforcement, admissibility, or externalisation, and could therefore provide misleading evidence. For this reason, the present contribution is deliberately theoretical and algorithmic: it establishes the formal and operational basis on which future large-scale empirical work can be responsibly pursued.

References

1. European Central Bank. Financial integration and structure in the euro area. Report. ECB, 2024. Available from: <https://www.ecb.europa.eu/pub/pdf/fie/ecb.fie202406.en.pdf>
2. Bank for International Settlements. BIS international banking statistics and global liquidity indicators at end-September 2024. Report. BIS, 2025. Available from: <https://www.bis.org/statistics/rppb2501.pdf>
3. Swiss State Secretariat for International Finance (SIF). European Union: Market access and equivalence. 2025. Available from: <https://www.sif.admin.ch/en/european-union-eu>
4. Regulation (EU) No 575/2013 on prudential requirements for credit institutions and investment firms (CRR). Official Journal of the European Union, 2013. Available from: <https://eur-lex.europa.eu/eli/reg/2013/575/oj>
5. Directive 2013/36/EU on access to the activity of credit institutions and the prudential supervision of credit institutions (CRD). Official Journal of the European Union, 2013. Available from: <https://eur-lex.europa.eu/eli/dir/2013/36/oj>
6. Commission Implementing Regulation (EU) 2021/451 laying down ITS for supervisory reporting under the CRR. Official Journal of the European Union, 2021. Available from: https://eur-lex.europa.eu/eli/reg_impl/2021/451/oj
7. European Commission. Equivalence of non-EU financial frameworks: overview and list of decisions. 2025. Available from: https://finance.ec.europa.eu/eu-and-world/equivalence-non-eu-financial-frameworks_en
8. Directive 2014/65/EU on markets in financial instruments (MiFID II). Official Journal of the European Union, 2014. Available from: <https://eur-lex.europa.eu/eli/dir/2014/65/oj>
9. Directive 2009/138/EC on the taking-up and pursuit of the business of insurance and reinsurance (Solvency II). Official Journal of the European Union, 2009. Available from: <https://eur-lex.europa.eu/eli/dir/2009/138/oj>
10. Directive 2009/65/EC on undertakings for collective investment in transferable securities (UCITS). Official Journal of the European Union, 2009. Available from: <https://eur-lex.europa.eu/eli/dir/2009/65/oj>
11. Directive 2011/61/EU on Alternative Investment Fund Managers (AIFMD). Official Journal of the European Union, 2011. Available from: <https://eur-lex.europa.eu/eli/dir/2011/61/oj>
12. Directive (EU) 2015/2366 on payment services in the internal market (PSD2). Official Journal of the European Union, 2015. Available from: <https://eur-lex.europa.eu/eli/dir/2015/2366/oj>
13. Regulation (EU) No 648/2012 on OTC derivatives, central counterparties and trade repositories (EMIR). Official Journal of the European Union, 2012. Available from: <https://eur-lex.europa.eu/eli/reg/2012/648/oj>
14. Regulation (EU) No 909/2014 on improving securities settlement in the European Union and on central securities depositories (CSDR). Official Journal of the European Union, 2014. Available from: <https://eur-lex.europa.eu/eli/reg/2014/909/oj>
15. Regulation (EU) No 600/2014 on markets in financial instruments (MiFIR). Official Journal of the European Union, 2014. Available from: <https://eur-lex.europa.eu/eli/reg/2014/600/oj>
16. Directive 2009/110/EC on the taking up, pursuit and prudential supervision of the business of electronic money institutions (EMD II). Official Journal of the European Union, 2009. Available from: <https://eur-lex.europa.eu/eli/dir/2009/110/oj>
17. Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA). Official Journal of the European Union, 2022. Available from: <https://eur-lex.europa.eu/eli/reg/2022/2554/oj>
18. Directive (EU) 2015/849 on the prevention of the use of the financial system for the purposes of money laundering or terrorist financing (4AMLD). Official Journal of the European Union, 2015. Available from: <https://eur-lex.europa.eu/eli/dir/2015/849/oj>

19. Directive (EU) 2018/843 amending Directive (EU) 2015/849 (5AMLD). Official Journal of the European Union, 2018. Available from: <https://eur-lex.europa.eu/eli/dir/2018/843/oj>
20. Regulation (EU) 2023/1114 on markets in crypto-assets (MiCA). Official Journal of the European Union, 2023. Available from: <https://eur-lex.europa.eu/eli/reg/2023/1114/oj>
21. Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data (GDPR). Official Journal of the European Union, 2016. Available from: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
22. European Securities and Markets Authority. MiFID II, Article 16: Organisational requirements (Interactive Single Rulebook). 2023. Available from: <https://www.esma.europa.eu/publications-and-data/interactive-single-rulebook/mifid-ii/article-16-organisational-requirements>
23. Commission Delegated Regulation (EU) 2017/565 supplementing Directive 2014/65/EU. Official Journal of the European Union, 2017. Available from: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32017R0565>
24. European Securities and Markets Authority. MiFID II, Article 34: Freedom to provide investment services and activities (passporting). 2023. Available from: <https://www.esma.europa.eu/publications-and-data/interactive-single-rulebook/mifid-ii/article-34-freedom-provide-investment>
25. European Banking Authority. PSD2, Article 28: Application to exercise the right of establishment and freedom to provide services. 2024. Available from: <https://www.eba.europa.eu/regulation-and-policy/single-rulebook/interactive-single-rulebook/16236>
26. Gesetz über das Kreditwesen (KWG). Gesetze im Internet (BMJ/Bundesanzeiger Verlag), 2025. Available from: <https://www.gesetze-im-internet.de/kredwg/>
27. Wertpapierhandelsgesetz (WpHG). Gesetze im Internet (BMJ/Bundesanzeiger Verlag), 2025. Available from: <https://www.gesetze-im-internet.de/wphg/>
28. BaFin. Securities Trading Act (WpHG)—overview. 2019. Available from: https://www.bafin.de/SharedDocs/Veroeffentlichungen/EN/Aufsichtsrecht/Gesetz/WpHG_en.html
29. Bankwesengesetz (BWG). RIS Bundeskanzleramt, 2025. Available from: <https://www.ris.bka.gv.at/GeltendeFassung.wxe?Abfrage=Bundesnormen&Gesetzesnummer=10004827>
30. Finanzmarktaufsicht (FMA). Securities Supervision Act 2018 (WAG 2018) [convenience translation]. Check against official RIS version. 2024. Available from: <https://www.fma.gv.at/wp-content/plugins/dw-fma/download.php?d=2085>
31. Finanzmarktaufsicht (FMA). Financial market supervision in Austria (integrated model). 2025. Available from: <https://www.fma.gv.at/en/financial-market-supervision-in-austria/>
32. Federal Act on Financial Services (FinSA). Fedlex, 2024. Available from: <https://www.fedlex.admin.ch/eli/cc/2019/758/en>
33. Federal Act on Financial Institutions (FinIA). Fedlex, 2024. Available from: <https://www.fedlex.admin.ch/eli/cc/2018/801/en>
34. FINMA. FINMA’s legal basis (Financial Market Supervision Act and sectoral laws). 2025. Available from: <https://www.finma.ch/en/documentation/finma-s-legal-basis/>
35. EFAMA. Trends in European investment funds: Fact Book 2025. Report. European Fund and Asset Management Association, 2025. Available from: https://www.efama.org/sites/default/files/fact-book-2025_lowres.pdf
36. European Investment Bank. EIB Investment Survey 2024: Country overview—Austria. Report. EIB, 2025. Available from: https://www.eib.org/files/documents/lucalli/20240238_econ_eibis_2024_austria_en.pdf
37. Financial Stability Board. Evaluation of the effects of financial regulatory reforms on SME financing. Report. FSB, 2019. Available from: <https://www.fsb.org/wp-content/uploads/P291119-1.pdf>

38. Financial Stability Institute. Proportionality in banking regulation: a cross-country comparison. Report. Bank for International Settlements, 2018. Available from: <https://www.bis.org/fsi/publ/insights1.pdf>
39. Kurz W, Magg R, and Stromeyer K. Financial and Operational Impacts of Regulatory Compliance on the Austrian Securities Industry. English. Journal of Next-Generation Research 5.0 2025; 1:137. DOI: [10.70792/jngr5.0.v1i5.137](http://dx.doi.org/10.70792/jngr5.0.v1i5.137). Available from: <http://dx.doi.org/10.70792/jngr5.0.v1i5.137>
40. Sun R. A novel multi-agent hierarchical deep reinforcement learning framework for portfolio optimization. Journal of Finance and Computation 2025
41. Cheng LC. Multi-agent based deep reinforcement learning framework for trading strategy and portfolio management. Applied Artificial Intelligence in Finance 2024
42. Hambly B, Xu R, and Yang H. Recent advances in reinforcement learning in finance. Mathematical Finance 2021; 31:782–813. DOI: [10.1111/mafi.12315](https://doi.org/10.1111/mafi.12315)
43. Malibari N, Katib I, and Mehmood R. Systematic review on reinforcement learning in the field of FinTech. Journal of Cloud Computing 2023; 12:1–34. DOI: [10.1186/s13677-023-00449-1](https://doi.org/10.1186/s13677-023-00449-1)
44. Ali A, Ahmed M, and Khan A. Audit logs management and security: a survey. Security and Privacy 2021; 4:e155. DOI: [10.1002/spy2.155](https://doi.org/10.1002/spy2.155)
45. Singh P and Kumar A. Ensuring data integrity and auditability in BPM systems: A blockchain-based approach. International Journal of Information Technology & Management Information Systems 2025
46. Faccia A. Is permissioned blockchain the key to support the external audit process? International Journal of Disclosure and Governance 2022; 19:354–64. DOI: [10.1057/s41310-022-00153-2](https://doi.org/10.1057/s41310-022-00153-2)
47. Arham MW. Transforming auditing through AI and blockchain. American Journal of Industrial and Business Management 2025; 15:225–41. DOI: [10.4236/ajibm.2025.152011](https://doi.org/10.4236/ajibm.2025.152011)
48. EIOPA. Digital Operational Resilience Act (DORA). https://www.eiopa.europa.eu/digital-operational-resilience-act-dora_en. Overview of DORA measures and scope. 2025
49. Skadden. The EU’s Digital Operational Resilience Act (DORA). <https://www.skadden.com/insights/publications/2024/07/the-eus-digital-operational-resilience-act>. Analysis of ICT risk requirements under DORA. 2024
50. European Banking Authority. Regulatory Technical Standards on ICT services supporting critical or important functions. <https://www.eba.europa.eu/activities/single-rulebook/regulatory-activities/operational-resilience/regulatory-technical-standards-policy-ict-services-supporting-critical-or-important-functions>. DORA RTS on ICT third-party risk governance. 2024
51. Anthropic. Claude Code. <https://www.anthropic.com/index/claude-code>. Accessed 17 August 2025. 2024