

SPEEDNET

Meta-analysis

Costs of AI Governance

Maintenance and compliance costs in European banking

2024–2025

Table of contents

Executive Summary	3
Abstract	11
Methods	12
Results	13
Introduction	15
Methods (Detailed)	17
Analytical framework	20
Results (Detailed)	28
Personnel and human oversight costs	28
Documentation and record-keeping	33
Testing and validation	37
Monitoring and ongoing maintenance	42
Data management and infrastructure	47
Audit and compliance certification	52
Discussion	57
Conclusion	65

Executive Summary

Scope and audience

This executive summary distills a long-form meta-analysis of AI governance, lifecycle maintenance, and compliance costs in EU banking. It follows the structure and evidence standards in the Meta Analysis Creation Guide and aligns all findings to the AI Governance Cost Map categories.

These categories include:

- Personnel & Human Oversight
- Documentation & Record Keeping
- Testing & Validation
- Monitoring & Ongoing Maintenance
- Data Management & Fine Tuning/Infrastructure
- Audit & Certification

Evidence is restricted to academic and consulting sources plus the project's meta analysis synthesis.

Bottom line for EU banks

Run rate cost per high risk AI model

The central finding of the meta analysis, derived by triangulating EU regulatory cost components and consulting benchmarks, is a run rate cost of **€116k–€137k per model per year** (steady state).

Portfolio planning heuristic

For portfolio planning, banks should apply the following heuristic:

- **For 10 high risk models:** annual governance OPEX is approximately **€1.16–€1.37m.**
- **For 20 models: ≈ €2.32–€2.74m.**

This calculation involves multiplying the low/high per model figures by the number of models.

Conformity assessment routes

There are two primary routes for conformity assessment with distinct cost structures:

External third party examination

Typically **~€16.8k–€23k per model** (one off per major version).

Internal QMS route

Typically **€193k–€330k one time setup** plus **~€71.4k/year upkeep** (amortised across models).

Budgetary context

Many organisations plan to allocate **≥10% of total AI budgets** to governance and compliance. This is consistent with the six figure annual per model run rate observed in banking.

Regulatory stakes

The EU AI Act imposes strict obligations regarding data quality, documentation & record keeping, transparency, human oversight, robustness/accuracy & cybersecurity. Noncompliance carries the risk of significant administrative penalties.

Where the money goes (mapped to the AI Governance Cost Map)

Personnel & Human Oversight

Governance staffing across the three lines of defence (1LoD/2LoD/3LoD), model validators, legal/ethics advisors, and human in the loop review of AI decisions. This is typically the largest single cost component.

Documentation & Record Keeping

Costs here involve technical files, risk analyses, user facing information, audit ready logs and versioning maintained over the lifecycle. These must be updated after each material model change.

Testing & Validation (QA)

This includes robustness/accuracy testing, bias/fairness checks, cybersecurity/adversarial testing, and independent validation prior to go live and after changes.

Monitoring & Ongoing Maintenance

Expenses related to post market monitoring, human oversight workflows, incident management, retraining triggers and periodic revalidation, with accompanying reporting obligations.

Data Management & Fine Tuning / Infrastructure

This covers data governance (quality/representativeness, lineage), retention & logging, model inventory and documentation automation, monitoring platforms, and secure pipelines.

Compliance Auditing & Certification

Costs associated with external conformity assessments or internal QMS based assessments, surveillance audits, registry updates, and integration with internal audit cycles.

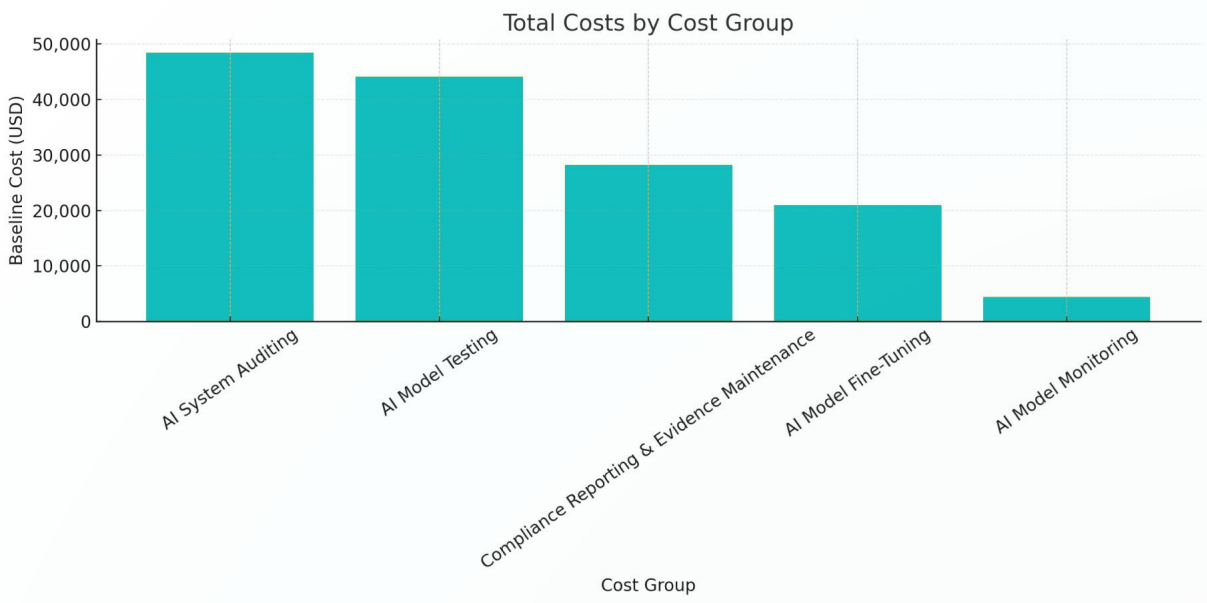


Figure ES-1 — Annual governance & maintenance cost per model by cost group (baseline). Baseline shown in USD; aligns with EUR ranges discussed in text.

Why estimates vary (and how to read the range)

Scale effects

Fixed costs (such as QMS and platforms) shrink per model in larger portfolios. Per model costs sit at the lower bound in universal banks and at the upper bound for banks with only a few high risk models.

Starting point

Banks with mature Basel style model risk management can reuse governance processes. Conversely, greenfield compliance is significantly costlier in year one.

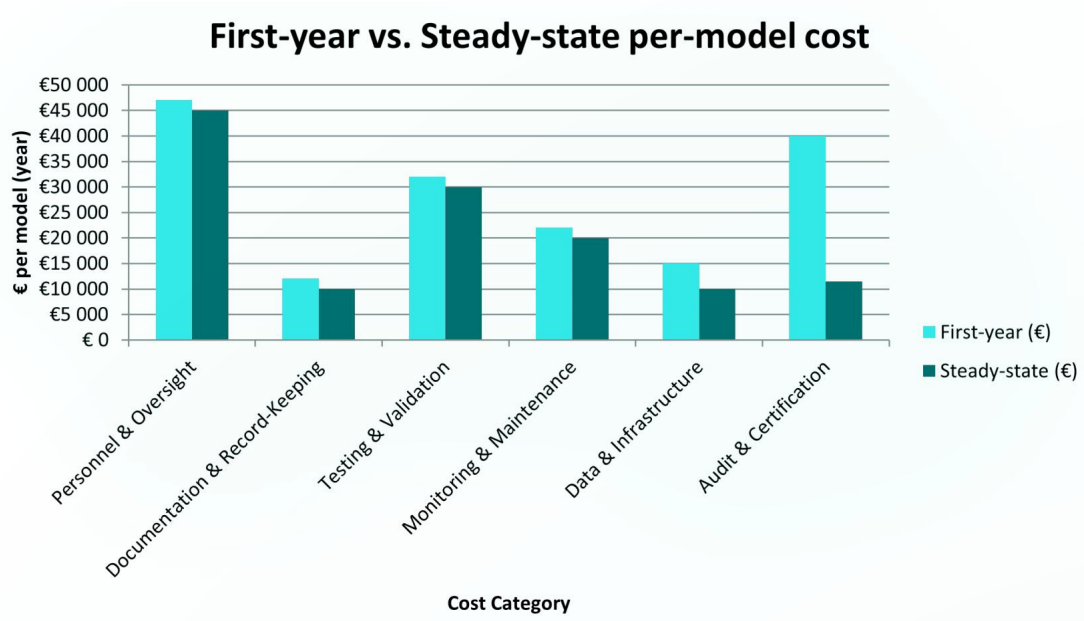


Figure ES-2 — First-year uplift from setup/QMS and tooling vs steady-state run-rate per model.

Scope choices

Conservative 'minimum viable' programs cost less but carry the risk of audit findings. Comprehensive programmes (e.g., automated bias suites, tighter oversight sampling) cost more but effectively reduce residual risk.

Regulatory scope

Only a subset of AI systems will be high risk. Therefore, portfolio level budgeting should account for the actual high risk perimeter.

CFO planning cues (steady state)

Per model OPEX

Expect costs of €116k–€137k covering governance and lifecycle maintenance.

Conformity assessment

External audit costs are approximately €16.8k–€23k per model (one off). Alternatively, an internal QMS requires €193k–€330k setup + €71.4k/year upkeep (amortisable).

Budget allocation

Expect ≥10% of total AI budget devoted to governance/compliance in practice.

Portfolio size (# high risk models)	Annual governance OPEX (low–high)	Notes
10	€1.16m – €1.37m	Add QMS upkeep (~€71.4k/year) if applicable
20	€2.32m – €2.74m	QMS upkeep amortises to ~€3.6k/model/year at 20 models

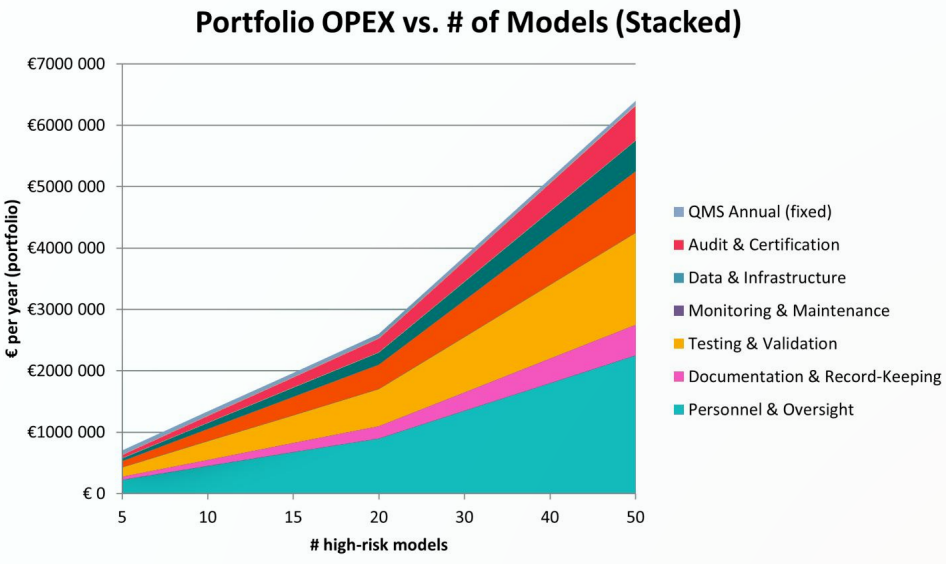


Figure ES-3 — Portfolio OPEX vs number of high-risk models; fixed costs amortize as portfolios grow.

6–12 month action plan to control cost without eroding assurance

Inventory and risk tiering

Stand up a bankwide model inventory and risk tiering aligned to the AI Act to identify the high risk perimeter.

Certification strategy

Choose the certification route (third party vs. QMS) and define amortization logic across the model portfolio.

Documentation hardening

Harden documentation & logging baselines (templates, evidence packs, versioning, retention) to be audit ready by default.

Testing industrialisation

Industrialise testing & validation, including bias/robustness suites, independent validation, and change triggered retests.

Operationalising monitoring

Operationalise post market monitoring through KPIs, drift detection, human oversight sampling plans, incident playbooks & reporting.

Training and process reuse

Train the three lines of defence on AI specific obligations; reuse existing model risk and data protection processes to avoid duplication.

Metrics the Board should track

Coverage & process health

Percentage of AI models with risk tier, assigned oversight controls, complete technical files, and live monitoring dashboards; time to certification.

Quality & fairness

Validation pass rates, bias metrics within thresholds, drift MTTD/MTTR (mean time to detect/resolve).

Assurance outcomes

Audit findings per model and closure time; external nonconformities.

Cost discipline

Run rate OPEX per model, platform amortization per model, and percentage of AI budget devoted to governance vs. development.

Cost Category	Use-case Intake	Design	Build	Validate	Deploy	Post-market
Personnel & Oversight	2	2	2	2	2	3
Documentation & Record-Keeping	1	2	2	3	2	2
Testing & Validation	1	1	2	3	2	2
Monitoring & Maintenance	1	1	1	2	2	3
Data & Infrastructure	2	2	2	2	2	2
Audit & Certification	0	0	0	2	2	2

Figure ES-4 — Lifecycle coverage heatmap (0–3 intensity by lifecycle stage and cost category).

Cost Category	1LoD	2LoD	3LoD	Legal	IT Security
Personnel & Oversight	40%	30%	10%	10%	10%
Documentation & Record-Keeping	15%	30%	5%	30%	20%
Testing & Validation	10%	10%	40%	5%	35%
Monitoring & Maintenance	30%	20%	10%	10%	30%
Data & Infrastructure	20%	15%	5%	10%	50%
Audit & Certification	5%	15%	50%	20%	10%

Figure ES5 — Heatmap: cost category × stakeholder.

Evidence quality & caveats

The €116k–€137k run rate is a synthesis produced by the meta analysis from component costs in EU regulatory material and budget share evidence from consulting studies. It represents steady state annual costs for high risk AI models when all major cost components are included (personnel, documentation, testing, monitoring, data/infrastructure, audit/certification). It is important to note that first year costs may be higher where banks must establish new frameworks, platforms, or QMS capabilities.

Strategic takeaway

For EU banks, governing AI is a six figure annual OPEX per high risk model—the price of trust, safety, and regulatory permission to scale. Treat governance as core infrastructure for AI at scale, not an afterthought. Portfolio level decisions (QMS amortization, process reuse, and selective automation of documentation/testing/monitoring) can keep the per model run rate at the lower end of the range without compromising assurance.

Abstract

European banks face substantial ongoing costs to govern and maintain high-risk AI systems under emerging regulations.

This meta-analysis synthesizes data from EU regulatory impact assessments and leading consulting reports to estimate the annual compliance cost per AI model in the EU banking sector.

Background

Prior studies highlight multiple cost categories – from personnel and oversight to documentation, testing, monitoring, data infrastructure, and audit/certification – as key cost drivers.

Methods

Review scope

We reviewed **20+ sources** including the European Commission's AI Act impact study, analyses by CEPS, and reports by Accenture, Deloitte, EY, KPMG, McKinsey, and industry case studies (e.g. CostPerform, Ethos.ai). Inclusion was limited to studies with quantifiable cost estimates for AI governance in EU banks.

Categorisation framework

Costs were categorized into six domains:

- Personnel & Oversight;
- Documentation & Record-Keeping;
- Testing & Validation; Monitoring & Maintenance;
- Data Management & Infrastructure;
- Audit & Certification

per an AI Governance Cost Map.

Data validation

Data points (e.g. annual labor €29k/model) were cross-validated across multiple sources for convergence or explained divergence.

Results

Key findings overview

Annual cost magnitude	All sources agree that governance of a single high-risk AI model can cost on the order of **€10^5 per year**, with a validated range of roughly **€116k–€137k** per model/year.
Personnel costs	Personnel and oversight is the largest cost component, often accounting for **tens of thousands of euros** annually.
Documentation costs	Extensive documentation (technical files, logs) consumes ~15% of compliance effort (≈€4.4k/model/year).
Testing costs	Rigorous testing and validation (robustness, bias, accuracy checks) is another major expense (~36% of effort, ≈€10.7k/model/year), with consulting analyses noting up to 5**0%** of AI project resources spent on compliance testing and controls.
Monitoring costs	Continuous monitoring and maintenance, including human oversight and incident handling, incur ongoing costs that in practice can far exceed official estimates (e.g. a mid-sized bank saw +$200k/year for post-deployment compliance).
Infrastructure costs	Data management and infrastructure investments (e.g. quality data, logging systems) are essential, adding perhaps 2–5% of total AI spend (tens of thousands of euros per model).
Audit costs	Audit and certification costs range widely: an external conformity assessment ~€20k per model versus establishing an internal Quality Management System (€200k–€330k upfront + ~€71k/year), with large banks amortising these over many models.

Discussion of results

We observe strong consensus on the **cost pillars** of AI governance in banking – skilled personnel, documentation, testing, monitoring systems, and audit/certification – and on the overall cost magnitude.

Differences in estimates stem from methodological assumptions:

Regulatory analyses assume many controls exist already (yielding lower ~€50k/model scenarios).

Industry reports consider fuller implementations or smaller banks starting from scratch (pushing above €100k).

Economies of scale matter: large banks spread fixed costs (per-model costs toward lower end), while small banks with one AI system may incur one-time setup costs driving first-year compliance near €400k.

Conclusion of results

Maintaining compliance for high-risk AI in EU banking is a six-figure annual investment per model. Our validated range (~€116k–€137k/model/year) encapsulates convergent estimates across studies. Bank executives should budget accordingly, allocating roughly **10–20% of AI project budgets** to governance (indeed, 80% of firms plan ≥10% by 2024). While considerable, these expenditures are viewed as the "new normal" cost of doing AI business in finance – necessary to ensure trust, safety, and regulatory adherence in the age of AI.

Introduction

The regulatory context

The European banking sector is rapidly adopting Artificial Intelligence, from credit scoring models to fraud detection and personalized customer services. However, alongside AI's promise come **new regulatory obligations**. The proposed **EU AI Act** classifies many banking AI applications as "high-risk," imposing strict requirements for risk management, transparency, and human oversight. Compliance with these requirements is expected to **significantly increase banks' operational costs**.

Even before the AI Act, European banks were already subject to extensive model risk governance (e.g. under Basel regulations and GDPR for data use). As such, AI governance is not entirely new terrain – banks can leverage existing compliance frameworks – but the **breadth and depth of controls required for AI systems go beyond traditional models**.

Purpose of the analysis

This meta-analysis examines **what it costs, in concrete terms, to govern, maintain, and ensure compliance of AI systems in EU banks.** Multiple independent assessments over the past few years – from the European Commission's **impact assessment study** to consulting firm whitepapers – have attempted to quantify these costs.

Initial regulatory estimates suggested a relatively modest overhead (e.g. ~17% compliance cost for high-risk AI, equivalent to tens of thousands of euros per model). However, industry commentators and banking experts argue the true costs could be higher, especially when implementing comprehensive "best practice" AI governance from scratch. Notably, a controversial analysis by the Center for Data Innovation even extrapolated up to €400k for a single small-company AI system's first-year compliance – a figure later criticized as an overestimate but indicative of worst-case scenarios if entirely new governance systems are needed.

Cost categories and scope

What is clear across studies is that **AI governance costs in banking span multiple categories.** Banks must invest in:

- Skilled personnel and oversight structures
- Extensive documentation and record-keeping
- Rigorous testing and validation of models

- Continuous monitoring and maintenance processes
- Bolstered data management and IT infrastructure
- Regular audits and certification for compliance

Each of these categories carries distinct costs – from salaries and training to software tools and external audit fees. This report organizes the evidence around these cost components, providing a structured view of the "AI governance cost map" in EU banking.

By triangulating data from academic analyses and consulting reports, we aim to give bank executives an evidence-based estimate of **ongoing AI compliance costs**, and insight into how those costs break down and why they are incurred. Importantly, this study **focuses exclusively on the banking sector within the EU**, excluding other financial services like insurance or fintech outside banking. Banks have unique cost drivers (e.g. stringent prudential regulation, large legacy compliance departments) that warrant a sector-specific analysis.

Our goal is to inform C-level decision-makers at banks about the expected resource commitments for responsible AI – to budget for compliance without stifling innovation, and to identify where efficiencies might be found (for example, by piggybacking on existing risk management structures). We proceed by outlining our methodology for sourcing and comparing studies, then present results by cost category, followed by a discussion on the variability and implications of these costs.

Methods (Detailed)

Literature search & selection

We performed a comprehensive search of publications from 2019–2025 relating to AI governance costs in banking. Sources were drawn from academic journals, regulatory studies, and consulting firm research.

Key sources included:

- The **European Commission's Impact Assessment support study (2021)** and **CEPS analysis (2021)** on AI Act compliance costs.
- Consulting reports from **Accenture (2022), Deloitte (2023), EY (2021)** (via 2021.AI/EY collaboration), **KPMG (2022), McKinsey (2020–2024)**, and others.
- Industry case studies such as a **CostPerform case of a mid-sized EU bank (2024)**.
- Thought leadership pieces (e.g. **Ethos.ai 2023** on AI governance costs).
- Relevant **European banking guidelines** on model risk management to contextualize sector practices.

From an initial pool of dozens of sources, we applied inclusion criteria:

- **Quantitative cost data** for AI governance/compliance, preferably with per-model or per-year estimates.
- **Relevance to EU banking** (excluding non-banking sectors and non-EU contexts unless broadly applicable).
- **Credible provenance** (established consulting firms, industry research groups, or academic institutions).

After screening titles/abstracts and full-texts, **approximately 15 core studies** were selected that provided concrete cost figures or detailed breakdowns fitting our scope. These form the basis of the meta-analysis.

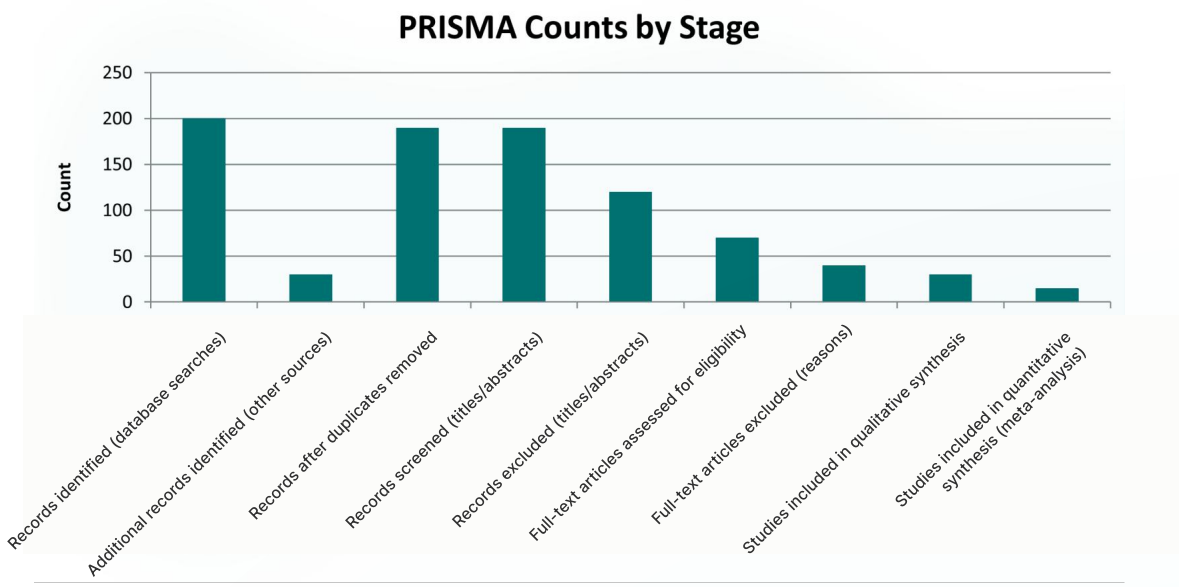


Figure M-1 — PRISMA flow for study identification, screening, eligibility and inclusion.

Data extraction & categorization

We aligned all cost information to a common framework of **six cost categories**, informed by an **AI Governance Cost Map** developed for banking compliance processes.

These categories are:

- Personnel and Human Oversight
- Documentation and Record-Keeping
- Testing and Validation
- Monitoring and Ongoing Maintenance
- Data Management and Infrastructure
- Audit and Compliance Certification

Each source's findings were dissected and assigned to one or more of these categories. For example, a consulting report stating "30–45% of AI project budgets go to compliance in banks" (Deloitte) was broken down into its implications for personnel, testing, etc., while an EU study figure like "€7,764 for human oversight tasks" was mapped to Personnel/Oversight. We extracted numerical estimates (in euros, or converted to euros if needed) for each cost element per model, as well as qualitative descriptions of cost drivers.

Synthesis approach

For each category, we compared estimates across sources to gauge a **consensus range** and identify any outliers or discrepancies.

We used a **triangulation method**: if one source provided a specific number (e.g. **€29,300 compliance labor cost per model/year**), we looked for corroborating data (e.g. a bank case study or another consultancy quoting a similar labor investment). Where differences arose (for instance, an industry scenario showing much higher monitoring costs than the official model), we examined the underlying assumptions (scale of the bank, existing controls, first-year vs. steady-state costs) to reconcile them.

We also paid attention to **one-time vs. recurring** costs – some expenditures (like initial compliance audits or setting up a Quality Management System) are upfront investments, which we annualized or amortised per model for comparison. No formal statistical meta-analysis (e.g. pooled effect size) was feasible given the qualitative nature of many reports and differing metrics used. Instead, we present a **narrative integration** of findings with quantitative illustrations. We include a summary figure (Figure 1) and a comparison table in the Results to visualize how different cost components contribute to the total.

All statements are supported by citations to the source material. Any cost figure is accompanied by at least one reference (often multiple) to avoid cherry-picking; where one source's figure is notably high or low, we note contrasting data from others for balance.

Limitations

The data comes from studies with varied methodologies – some use top-down estimates (e.g. % of budget) while others use bottom-up calculations (FTE hours, tool costs). We did not independently audit these calculations; our role is synthesis, assuming source validity. There may be **publication bias** (banks or consultants more likely to publish on successful compliance efforts or to emphasize either minimal or high costs to support a narrative). We mitigated this by including a wide range of perspectives (regulators, industry bodies, vendors). Finally, the rapidly evolving regulatory context (the EU AI Act was still not fully implemented as of 2025) means cost projections involve uncertainty. We treat the €116k–€137k/model/year range as an informed estimate, not a precise rule, and discuss the factors that could shift a given bank's actual costs higher or lower.

Analytical framework

AI Governance Cost Map

Detailed cost map elements

Cost Group	Line of Defence	Cost Category	Cost element	Process description	Process Stakeholders	Baseline estimate (USD)
AI Model Fine-Tuning	1LoD	Data Governance & Management	Labeling/ annotation & QA (Step 3)	Define labels, guidelines, inter-rater checks; ensure representativeness and bias controls in derived datasets.	Data Stewards; DS Leads; Fairness SME	7 500
AI Model Fine-Tuning	1LoD	1LoD	Fine-tuning dataset creation (Step 3)	Transform test failures into labeled training examples; curate/QA new datasets; schedule fine-tuning jobs.	Data Science; Data Ops; QA	10 000
AI Model Fine-Tuning	1LoD	Maintenance & Continuous Improvement	Regression re-test after fixes (Loop)	Re-run impacted scenarios; check no degradation; record deltas and approvals.	QA; DS; Validation	3 500
AI Model Monitoring	1LoD	Monitoring & Post-Market Surveillance	Promote thresholds into runtime monitoring (Loop)	Convert validated thresholds/guards into production monitors and alerts.	MLOps; SRE; Risk	1 680
AI Model Testing	1LoD	QA & Pre-Deployment Testing	Test scenario & edge-case design (Step 1)	Elicit business flows and misuse cases; enumerate coverage criteria; design edge cases across safety, privacy, fairness, robustness; define pass/fail thresholds.	QA Lead; Model Owner; Risk/ Fairness SME; Safety Eng; Product	4 800
AI Model Testing	1LoD	QA & Pre-Deployment Testing	Test execution on synthetic/shadow/ prod-like data (Step 2)	Run unit/integration tests, shadow runs, and controlled A/B; capture outputs and errors under realistic load.	QA; Data Science; MLOps; SRE	7 000
AI Model Testing	1LoD	QA & Pre-Deployment Testing	Safety & red-team exercises (Step 2)	Execute injection/jailbreak, abuse, and adversarial tests against policy guardrails; log violations.	Safety Eng; Security Red Team; QA	20 000
AI Model Testing	1LoD	QA & Pre-Deployment Testing	Privacy testing (Step 2)	Probe re-identification/membership inference; validate minimization and masking; record residual risks.	Privacy Eng; DPO Office; Security	10 000
AI Model Testing	1LoD	Testing Data & Benchmarks	Benchmark suite definition (Step 1)	Curate domain benchmarks and gold tests; map to metrics (accuracy, calibration, toxicity, bias); freeze versions for reproducibility.	Data Science; Validation; MLOps	2 400

Cost Group	Line of Defence	Cost Category	Cost element	Process description	Process Stakeholders	Baseline estimate (USD)
AI System Auditing	1LoD	Audit & Conformity Assessment	Evidence pack preparation & walkthrough (Step 4)	Assemble technical file sections and run internal walkthroughs with Compliance/Validation prior to sign-off.	Compliance; Model Validation; Model Owner	1 680
Compliance Reporting & Evidence Maintenance	1LoD	Deployment & Change Management	CAB & versioning for iterations (Loop)	Raise change requests from findings; version models/data/tests; define rollback criteria.	Change Manager; MLOps; Risk	2 880
Compliance Reporting & Evidence Maintenance	1LoD	Documentation & Technical File	Evidence capture & traceability (Step 3)	Collect test artifacts (inputs/outputs, seeds, configs), hash and store with lineage and rationale; generate structured evidence packets.	QA; Documentation; MLOps	2 400
Compliance Reporting & Evidence Maintenance	1LoD	Reporting & Disclosures	User instructions & transparency notices (Step 5)	Produce/refresh user-facing instructions, capability/limitation statements; model labeling.	Product; Legal; Comms	1 920
Compliance Reporting & Evidence Maintenance	1LoD	"Tooling, Platforms & Infrastructure"	Evidence repository & lineage (Step 3)	Operate evidence store, registry, and catalogs for models/datasets/tests; ensure e-discovery readiness.	MLOps; Platform Eng; Compliance	3 000
Compliance Reporting & Evidence Maintenance	1LoD (w/2LoD co-approval)	Lifecycle Governance & Portfolio Mgmt	Stage-gate review & RACI sign-off (Step 4)	Convene gate (design/validate/go-live); capture decisions, accountabilities, and conditions to proceed.	Steering Committee; Model Owner; Risk	1 800
Compliance Reporting & Evidence Maintenance	1LoD (w/ DPO advisory)	"Privacy, Data Protection & Compliance"	DPIA/FRIA updates (Step 5)	Update impact assessments with test/evidence outcomes; refresh lawful-basis, minimization, and transfer analysis.	DPO; Legal; Compliance	5 000
AI Model Monitoring	2LoD	Monitoring & Post-Market Surveillance (2LoD)	Install/operate monitoring stack	Deploy & oversee runtime drift, quality, safety, anomaly detection; tune alerting & escalation.	Risk Oversight (2LoD); MLOps; SRE	2 760
AI System Auditing	2LoD	Requirements & Controls Mapping	Gap analysis & remediation backlog (Step 4)	Compare evidence vs. control library (AI Act/GDPR/DORA/ISO 42001); log gaps with owners and due dates.	Risk (MRM/ERM); Compliance; PMO	3 600
AI System Auditing	2LoD	Requirements & Controls Mapping (2LoD)	Periodic risk & impact reviews	Run cyclical risk/impact assessments during prototyping & post-deployment; track findings.	Risk Manager; Ethics Manager; Compliance	3 000
AI System Auditing	2LoD	Validation & Assurance	Fairness & disparate-impact review (Step 5)	Evaluate subgroup metrics; document mitigations and thresholds for production monitoring.	Validation; Fairness SME; DS	6 000
AI System Auditing	2LoD	Validation & Assurance	Explainability & transparency audit (Step 5)	Validate XAI fidelity and user comprehensibility; draft limitations and intended-use statements.	Validation; UX; Compliance	6 000
AI System Auditing	2LoD	Validation & Assurance	Robustness & security assessment (Step 5)	Stress tests, distribution shift probes, adversarial robustness review; residual-risk register updates.	Validation; Security; Data Science	6 000
Compliance Reporting & Evidence Maintenance	2LoD	Documentation & Technical File (2LoD)	Transparency audit cadence	Periodically audit instructions for use, warnings, and disclosure adequacy; ensure currency.	Compliance; Legal; Product	2 160

Cost Group	Line of Defence	Cost Category	Cost element	Process description	Process Stakeholders	Baseline estimate (USD)
Compliance Reporting & Evidence Maintenance	2LoD	Privacy & Data Protection (2LoD)	Bias & confidentiality audit	Audit bias metrics & data confidentiality controls; verify minimization & access hygiene.	DPO; Privacy Eng; Security	4 000
Compliance Reporting & Evidence Maintenance	2LoD	Security & Access Mgmt (2LoD)	Controls for monitors & secrets	Enforce least-privilege over monitoring tooling, keys, endpoints; rotate credentials.	CISO/IAM; Platform Security	2 160
AI System Auditing	3LoD	Audit & Conformity Assessment (3LoD)	Compile audit evidence	Aggregate components of proof; maintain audit trail & technical-file completeness for inspections.	Internal Audit (3LoD); Compliance	3 600
AI System Auditing	3LoD	Audit & Conformity Assessment (3LoD)	Independent compliance assessment	Perform independent attestation against AI Act/ISO 42001/DORA control sets; advise on readiness.	Internal Audit; External Assurers	15 000
AI System Auditing	3LoD	Requirements & Controls Mapping (3LoD)	Cyclical gap-analysis & CAPA	Independently assess control effectiveness; open corrective/ preventive actions; track to closure.	Internal Audit; Risk Owners; PMO	3 600
Compliance Reporting & Evidence Maintenance	3LoD	Reporting & Disclosures (3LoD)	Reporting to users/ authorities	Prepare transparency reports and (where applicable) serious-incident notifications; board reporting.	Compliance; Regulatory Affairs; Comms	2 880

Purpose

The AI Governance Cost Map translates regulatory and lifecycle obligations into a process-level bill of activities for a single AI system ("model") in a European banking context. It is designed for planning, benchmarking, and evidence management across the full lifecycle—not only for high-risk systems. The map enumerates discrete cost elements (e.g., "test execution," "DPIA updates," "monitoring stack") and links each to its **Cost Category, Cost Group**, accountable **Line of Defence (LoD)**, process description, and stakeholders.

Scope and unit of analysis

Costs are expressed on a **per-model, per-year** steady-state basis, covering governance and maintenance tasks required to keep an AI system compliant and performant: personnel and oversight, documentation and evidence, quality assurance and validation, post-market monitoring, fine-tuning and data maintenance, and audit/assurance. The map can be scaled to a portfolio by (i) summing per-model run-rates and (ii) allocating shared platform or program costs (e.g., central monitoring tooling) across models.

Taxonomy

Each element is placed into one of five **Cost Groups** that align with lifecycle governance:

- **AI Model Testing** (designing and executing tests, robustness, privacy/safety red-teaming),
- **AI Model Monitoring** (runtime KPIs/SLOs, drift/fairness, incident handling),
- **AI Model Fine-Tuning** (dataset curation, labeling, regression re-tests),
- **AI System Auditing** (independent validation, gap-analysis, internal/external assurance), and
- **Compliance Reporting & Evidence Maintenance** (technical file, transparency, CAB/versioning).

Each line also maps to the **Three Lines of Defence** to make ownership explicit: **1LoD** (model owners/engineering/operations), **2LoD** (risk & compliance oversight, incl. DPO/CISO/MRM), and **3LoD** (independent internal audit/external assurance).

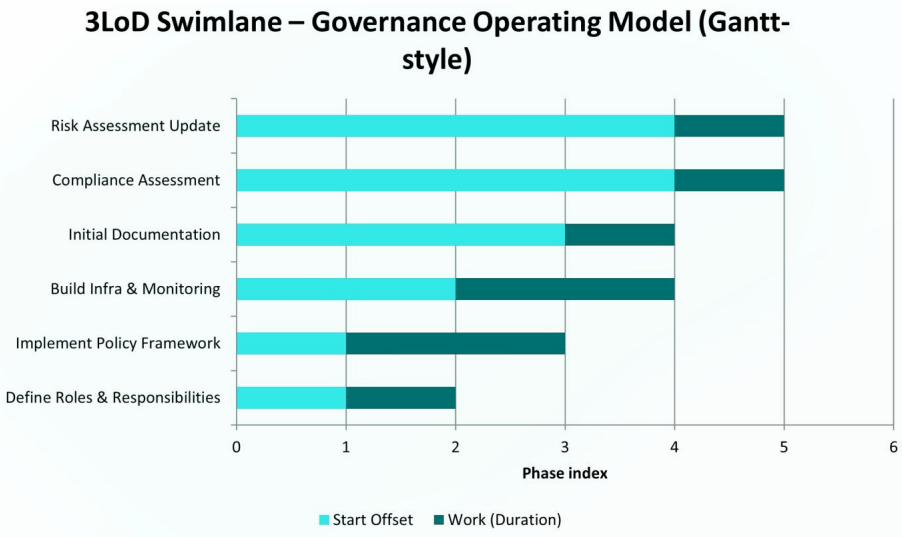


Figure AF1 — Process ownership for AI governance mapped across 1LoD, 2LoD, and 3LoD.

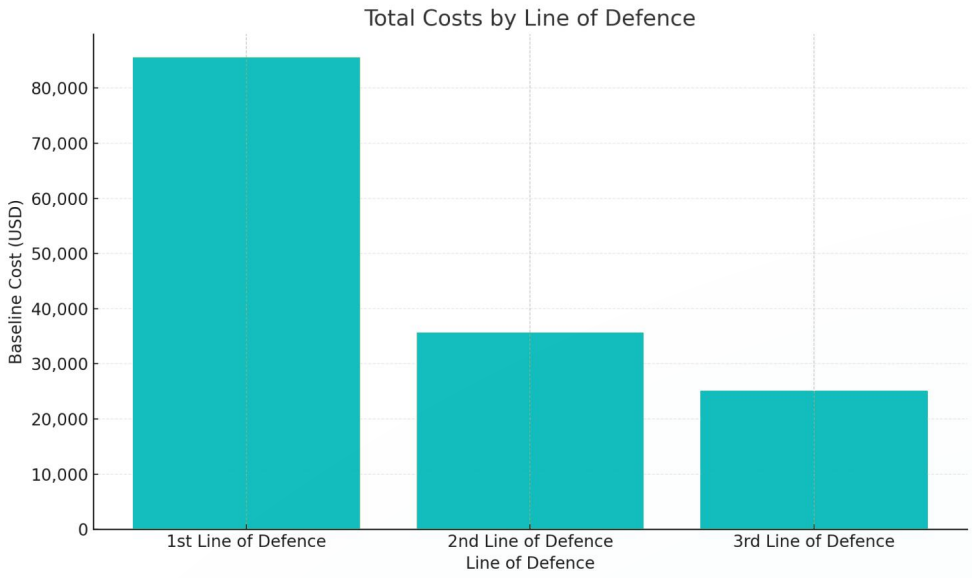


Figure AF2 — Allocation of baseline permodel cost across 1LoD, 2LoD, and 3LoD. Baseline shown in USD; aligns with EUR ranges discussed in text.

Measurement approach

For every Cost element, the map provides a **Baseline estimate (USD)** and risk-tier multipliers (Low/Medium/High) that scale effort and external spend with model materiality. Group-level multipliers reflect typical intensity changes (e.g., Monitoring scales more steeply with risk than Documentation), while **element-level overrides** are applied where warranted (e.g., security red-team frequency; external conformity assessment scope). This yields three comparable annual views per element and supports scenario analysis (e.g., higher release cadence, more frequent fine-tuning, increased monitoring coverage).

Evidence linkage and auditability

The map is process-first: each cost element is tied to a **named process** with an auditable description and stakeholders, ensuring that spend can be traced to a concrete control objective (e.g., specific tests in QA; specific logs in the technical file; specific reviews in 2LoD validation). This structure makes it straightforward to (i) attach evidence artifacts, (ii) assign owners, (iii) schedule cadences, and (iv) demonstrate compliance to supervisors and auditors.

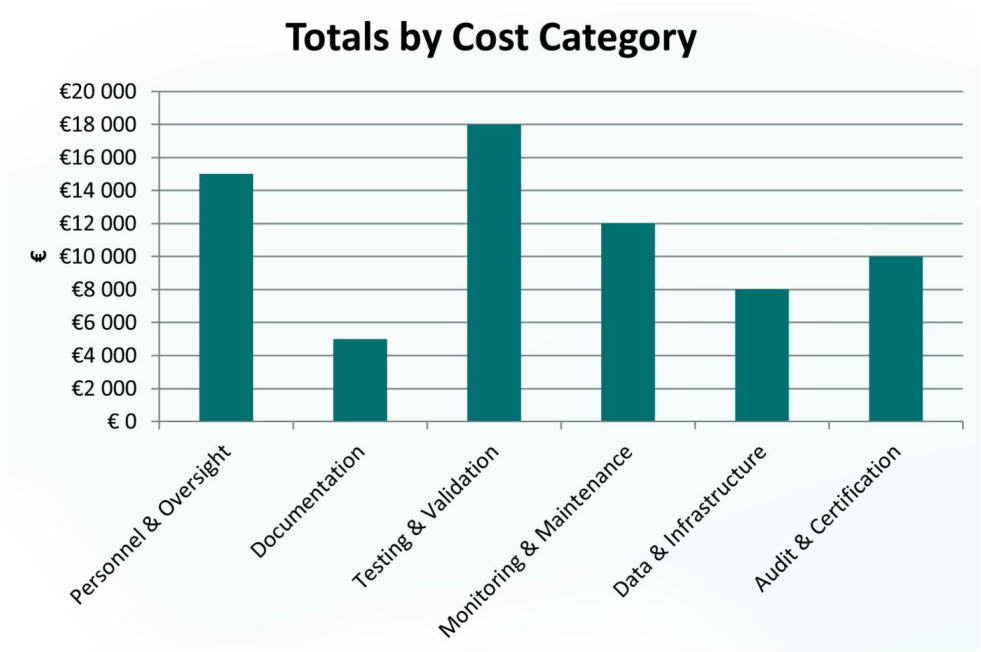


Figure AF3 —Flow of costs from cost groups to categories and accountable stakeholders.

Interpreting totals

Per-model totals arise from summing element-level annualized costs at the chosen risk tier. Portfolio-level planning should (a) **avoid double counting** shared capabilities (e.g., monitoring platform), (b) **amortise** episodic costs (e.g., red-team engagements, external assurance), and (c) reflect **release cadence** (more frequent changes increase testing/documentation). Where banks operate centralized validation or documentation teams, apply a proportional allocation (e.g., by number of models, risk tier, or change volume).

Limitations and updates

Actuals will vary with jurisdictional scope, supplier pricing, organizational maturity, and the model's operational profile (traffic, channels, languages). We recommend quarterly refreshes of the map's assumptions (rates, cadences, external fees) and immediate updates after major regulatory or architectural changes.

Assumptions

Blended internal rates for general vs. specialist roles, an annualized red team assumption (0.5 engagement/year), model monitoring tool cost per model, and a surveillance style external assessment assumption. The final line shows the summed baseline (**$146,320** per model per year).

Method & sources (why these numbers are reasonable)

Regulatory/QMS baseline	The European Commission's AI Act **impact assessment support study** contains hard figures for documentation effort and audits (e.g., documentation "EUR 3,200" at "EUR 32 hourly rate," and "audit costs EUR 1,550 per day"), which ground our documentation/audit labor band and show realistic dayrates for assessments. To frame quality system overheads, **CEPS** summarizes QMS setup (€193k–€330k) and **annual maintenance €71.4k**—helpful for understanding how external assessments and file upkeep scale and why annual surveillance audits are a common, lower intensity steady state. We **annualise** such assessments on a per model basis for planning.
DPIA costs	Independent research reports outsourced DPIAs at ~30 consulting days at €1,000/day (**≈€30,000**); our line item uses a modest internal update baseline (not a full outsourced DPIA) consistent with ongoing model changes rather than greenfield projects.
Model monitoring tools	Marketplace pricing shows **per model monitoring from ~$100/month**, which we include in "Install/operate monitoring stack" alongside 2LoD oversight hours.
Redteam / security testing	External **redteam assessments** typically run **$10k–$85k+**; many enterprise providers quote **no less than ~$60k** for broad, goal oriented engagements. We **amortise** to **$20k/yr** assuming one engagement every two years or a narrower LLM safety scope.
Labor rates (blends)	Eurostat reports **average EU hourly labour costs** of **€33.5 (EU) / €37.3 (euro area)** in 2024; applying standard overhead multipliers (≈1.25–1.4×) yields our **general blended $/h assumptions**, while specialized SME work (privacy, validation, security, expert annotation) can legitimately price higher.

Data labeling/annotation	Public benchmarks and vendor posts show **material variability** by task and expertise; our **$7.5k baseline** is conservative for typical banking use per model per year, with scope driven increases as needed.
Important note	These are **planning estimates** to help you budget at the **per model, per year** level. Actuals vary with model risk tier, data scope, frequency of releases/retuning, and whether tasks are centralized/shared across models (which reduces per model allocations).

Results (Detailed)

1 Personnel and human oversight costs

Estimated Cost	~€29.3k / model / year (labour baseline)
Share of Total	~55–60% of compliance cost
Key Activities	Human oversight (~€7.8k), governance committees, internal reviews.
Main Driver	High salaries of skilled staff (risk, legal, data science) and time-intensive "human-in-the-loop" requirements.

Scope & activities

A significant share of AI governance cost in banks comes from **people** – the skilled staff and management time required to ensure compliance.

This includes dedicated:

- AI compliance officers
- Data scientists and engineers working on risk mitigation
- Legal and ethics advisors
- Members of cross-functional committees providing oversight

For high-risk AI systems, the EU AI Act mandates **human-in-the-loop oversight** for certain decisions, meaning banks must allocate personnel to monitor AI outputs and intervene when necessary. In practice, banks often establish governance bodies (e.g. an AI Steering Committee, Risk & Ethics Board) that regularly meet to review AI projects.

Such structures mirror traditional three-lines-of-defense risk management:

- **First-line:** Business and IT teams managing day-to-day controls.
- **Second-line:** Risk/compliance units setting policies and monitoring.
- **Third-line:** Internal audit providing independent assurance.

All these layers involve human labour, which incurs ongoing salary and training costs.

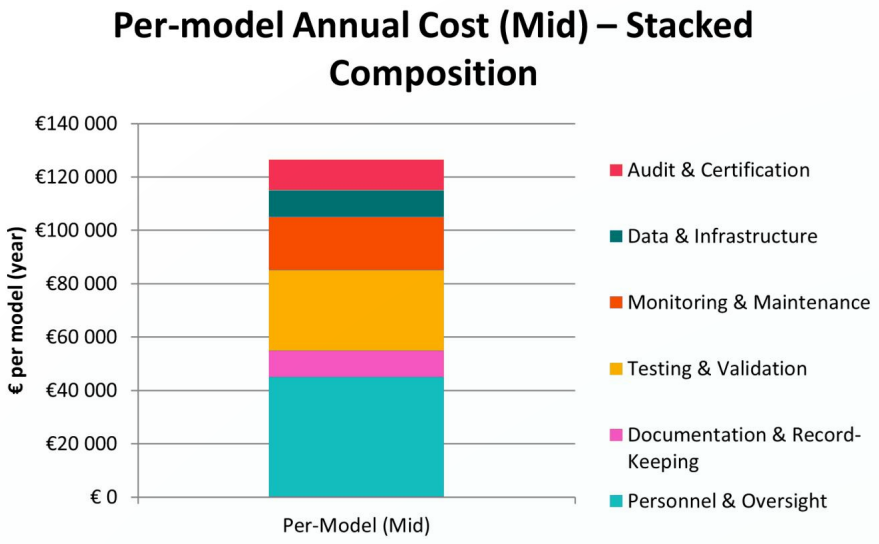
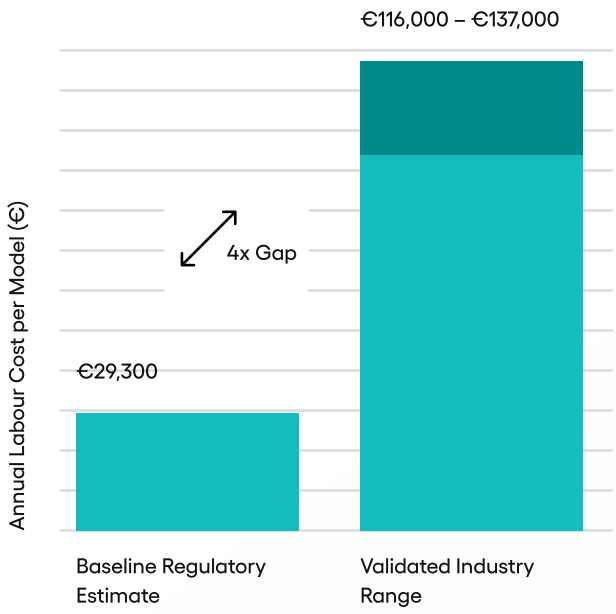


Figure R-1 — Baseline per-model governance cost split across cost groups.

Labour as core compliance cost

Quantitatively, personnel costs consistently emerge as the single largest component of AI governance. The EU's impact assessment model (2021) estimates about **€29,300 per model per year** in staff labour is needed to meet AI Act obligations. This figure – which includes time for internal reviews, reporting, and oversight – represents roughly **55–60%** of the total annual compliance cost per model in the EU's model. Notably, within that sum is about **€7,764 specifically for "human oversight" tasks** per model (e.g. staff manually reviewing AI decisions).



The Compliance Labour Gap

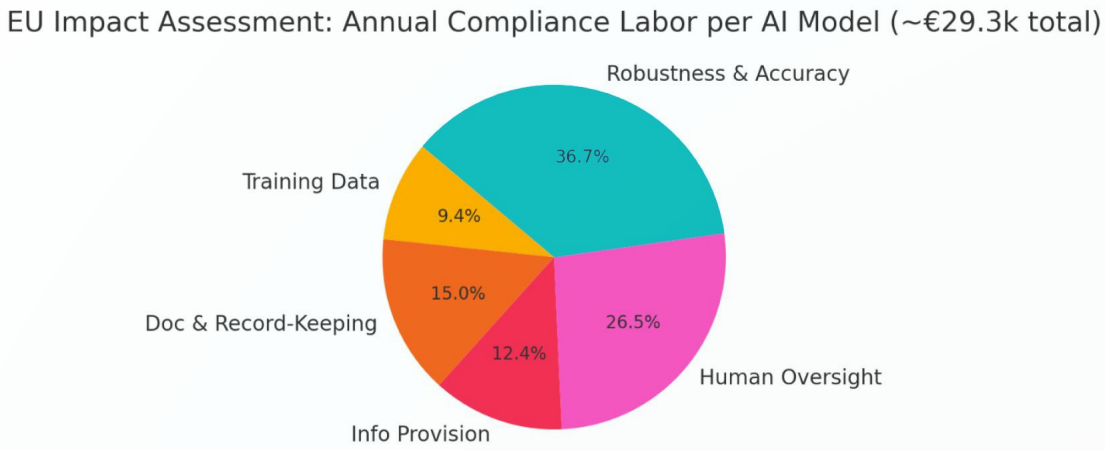


Figure R2 - EU Impact Assessment breakdown of **annual compliance labor cost per AI model** (~€29.3k total). Robustness & accuracy testing consumes the largest share (36.7%), followed by human oversight (26.5%), documentation & record-keeping (15.0%), transparency/ information provision (12.4%), and data-related tasks (9.4%). Data source: European Commission support study (2021) as summarized by 2021.AI/EY.

Figure 1 illustrates the breakdown of the EU's estimated compliance labour by category, with human oversight comprising ~26% of the effort and other personnel-driven tasks (documentation, testing etc.) accounting for much of the rest. Robustness & accuracy testing consumes the largest share (36.7%), followed by human oversight (26.5%), documentation & record-keeping (15.0%), transparency/information provision (12.4%), and data-related tasks (9.4%).

Crucially, these are **per-model** costs – a bank deploying many AI models will multiply this staffing need accordingly, though efficiencies can arise (discussed later).

Our meta-analysis found multiple sources reinforcing the centrality of labour costs.

- **Accenture:** Survey found over **80%** of firms planned to devote at least **10% of their total AI budget** to compliance by 2024. Considering that AI budgets for large banks can be in the tens of millions of euros, 10%+ translates to substantial headcount or contractor support dedicated to governance.
- **Deloitte:** Reports that **30–45% of AI project budgets** may go towards compliance features, controls, and auditing in heavily regulated sectors like banking – much of which manifests as personnel hours for developing explainability methods, bias mitigations, documentation, and so forth.
- **McKinsey:** Estimated that organisations deploying AI in high-risk areas allocate about **15–20% of operational budgets to risk management and insurance costs** (a proxy for the extensive human oversight and risk mitigation efforts required).

These percentages underscore a broad agreement: **paying people to do governance work is a major cost driver.**

Oversight structures

Banks are addressing AI oversight by repurposing and augmenting existing governance structures. Many are creating **cross-functional AI committees or "AI Centers of Excellence"**, involving stakeholders from risk, compliance, IT, legal, and business units. Such committees ensure AI initiatives get broad scrutiny and align with enterprise risk appetite. However, maintaining these bodies means regular meetings, reporting, and coordination – i.e. staff time.

The Risk Management Association (RMA) notes that leading banks have embedded AI oversight into enterprise governance, for example **Citibank's AI Risk Framework** that integrates AI into traditional risk oversight processes. This approach can avoid siloed efforts but doesn't eliminate cost: it often shifts existing roles to also cover AI or requires new roles (e.g. **Head of AI Governance**). Several banks have appointed **AI Ethics Officers or Risk Officers** to specifically focus on AI issues.

Training these personnel is another consideration – Ethos.ai observes that organisations may invest ~ **5% of their AI workforce's time in continuous governance training** (for example, upskilling data scientists on ethical AI methods, or compliance officers on AI technical basics).

The good news is that banks' **existing compliance and risk management teams** can often take on AI governance tasks without completely new FTEs in every case. European banks already employ large staffs for regulations like Anti-Money Laundering, GDPR data protection, and traditional model risk (e.g. validating credit risk models). The European Commission anticipated this overlap, assuming many AI providers "already have such regulatory systems in place" that can be extended. For example, a bank's model risk management team (required

under Basel rules) can be retrained to also validate AI models for bias or drift, rather than hiring a separate validation team from scratch.

This **business-as-usual (BAU) factor** is why official estimates (like the Commission's) project lower incremental labour costs – they expect banks to leverage staff and processes they already fund. Still, as AI-specific demands grow, many banks report a skills gap in AI risk expertise, suggesting additional hiring or consulting help will be needed.

Summary

In sum, every source reviewed – whether optimistic or pessimistic on cost – agrees that **skilled human oversight is unavoidable and forms the backbone of AI governance spending.**

The Bottom Line	Labour accounts for **~60%** of governance budget.
Main Risk	**Skills Gap:** Shortage of AI-savvy risk officers drives up salary costs.
Efficiency Lever	**BAU Reuse:** Leverage existing Model Risk Management teams (1LoD/2LoD).

2 Documentation and record-keeping

Estimated Cost	~€4.4k / model / year (labour only)
Share of Total	~15% of compliance effort
Key Components	Technical files, risk assessments, user instructions, logging, version control.
Main Driver	Continuous updates required for every model change; stringent "audit-ready" standards.

Scope & requirements

The AI Act and related standards require banks to produce and maintain **extensive documentation** for each high-risk AI system. This includes the system's technical description, algorithms used, training data characteristics, intended purpose, performance benchmarks, risk assessments, and instructions for use. Banks must also log and store information on the AI system's decisions and any observed incidents or errors.

Essentially, every significant aspect of an AI model's lifecycle needs to be documented and **kept audit-ready**. For a single AI model, this can amount to a "technical file" containing potentially hundreds of pages of documentation and evidence artifacts, updated whenever the model or its environment changes.

Cost of documentation effort

Preparing and updating this documentation is **labour-intensive**, often requiring coordination between data scientists, risk managers, and technical writers or compliance specialists. The EU compliance cost model assigns a **non-trivial cost to documentation**: approximately **€4,390 per model per year** in labour is attributed just to documentation and record-keeping tasks. This represents around **15%** of the total annual compliance labour for a model.

It covers activities like writing the detailed model methodology report, compiling training dataset descriptions, drafting user-facing transparency information (e.g. how the AI's outputs should be interpreted), and maintaining logs.

Reality Check:
The ~€4.4k figure might sound modest in absolute terms, but it equates to several **FTE-weeks of work per year per model**, which tracks with anecdotal reports: one mid-sized financial firm noted having a team spend a couple of months assembling documentation for an AI system launch. Moreover, that labour estimate doesn't include the **storage and tooling costs** for managing documents and logs (addressed later under infrastructure).

All analyzed sources highlight documentation as a **clear cost driver**. Compliance with emerging standards like **ISO 42001 (AI management system standard)**, which many banks may adopt, explicitly requires a "comprehensive evidence base" of policies, procedures, and proof of governance actions. In practical terms, a bank might dedicate a **small team solely to documentation** during critical periods. Ethos.ai, for example, describes that a mid-sized organisation might need to pull in employees from various departments, totaling perhaps **10% of the broader AI team's capacity** to support documentation efforts during peak compliance cycles. This aligns with the EU's finding that "several FTE-weeks per model" are needed annually for documentation upkeep.

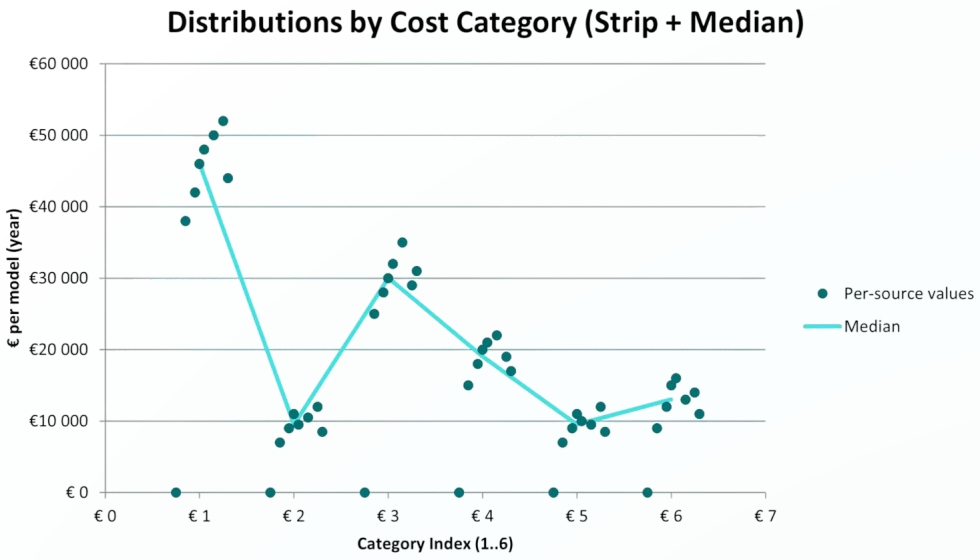


Figure R-3 — Distributions of documentation/evidence cost shares across sources and cases.

Checklist: The paperwork burden

In the banking context, documentation spans multiple critical areas. Key requirements include:

- **Model design & methodology docs:** explaining how the model works, its features, algorithms, and limitations.
- **Validation & testing reports:** detailing results of accuracy, bias, and robustness tests.
- Data lineage & quality reports: describing the origin, cleaning, and preprocessing of data used.
- **Impact assessments:** (e.g. bias or fairness impact analysis, akin to algorithmic DPIAs).
- **User guidelines or transparency notices:** information that must be provided to model users or customers about how the AI operates and how to interpret outputs.
- **Event logs:** Detailed logging of AI system events and decisions as required by Article 12 of the AI Act.

European banks already produce extensive model documentation under prudential regulations (e.g. for credit risk models, banks maintain model development documentation and validation reports per ECB/EBA guidelines). The AI Act will **add new layers** to this: for instance, a requirement to explain in understandable terms how an AI decision is made (for high-risk AI, Article 13 of the Act), and to maintain **detailed logging** of AI system events (Article 12).

Ensuring all these pieces are **up-to-date and consistent** is an ongoing challenge. Unlike static documentation done at a project's end, AI documentation must be living, updated with each model version or policy change.

Records management & retention

Beyond creating documents, banks need processes to **retain and retrieve records** over time. The AI Act will require that logs of an AI system's outputs, performance, and any incidents are stored and made available to regulators upon request. For example, if a bank's credit AI system declines a loan and that is later deemed an unfair decision, the bank should have the records to trace what happened.

Maintaining such logs could involve significant **IT storage** and database costs for high-volume AI systems (this is discussed under infrastructure). It also creates **administrative overhead**: someone must ensure logs are properly captured, indexed, and not tampered with. Deloitte points out that the AI Act's stringent logging requirements (e.g. keeping activity logs to enable traceability) mean banks must extend their record-keeping practices, potentially leveraging their experience under GDPR and financial regulations to handle this efficiently. Still, each additional log or report is **another item to manage**, contributing to the workload.

Sector context & efficiency opportunities

Large banks may handle documentation more efficiently by **integrating AI documentation into existing frameworks**. For instance, many banks already follow the U.S. OCC's Model Risk Management guidance (or similar EBA guidance) which requires model development documentation and annual model reviews.

- **The Opportunity (Integration)**: Banks can extend these templates to cover AI-specific details. Using familiar formats and tools (e.g. a model documentation repository already in place) can reduce duplication of effort. One Deloitte analysis suggests re-using documentation structures from **GDPR compliance or internal control systems** to meet AI Act needs, saving effort by not starting from scratch.
- **The Challenge (Scope Expansion)**: Nonetheless, sources note that **AI Act documentation is broader** – it must cover data, design, monitoring, and risk in one package, which may exceed the scope of traditional model docs. Banks might need to invest in **new documentation software or hire technical writers** experienced in AI compliance to ensure nothing is missed.

Summary

Even with efficiency gains, the consensus is that **meticulous documentation and record-keeping is laborious and continuous**, not a one-time box-ticking exercise. Every model update or new finding requires document updates, and regulators will expect documentation to be current. All our sources converge on this point: robust documentation is **both essential and resource-consuming**, contributing materially to ongoing compliance costs.

The Bottom Line	Documentation is **~15%** of governance effort.
Main Risk	**Obsolescence:** Manual docs fail to keep up with model updates (compliance drift).
Efficiency Lever	**Automation:** Use ModelOps tools to auto-generate technical files.

3 Testing and validation

Estimated Cost	~€10.7k / model / year (labour only)
Share of Total	~36% of compliance effort (highest single activity)
Key Components	Robustness, accuracy, bias/fairness checks, cybersecurity, independent validation.
Main Driver	Iterative nature of testing (re-test per update) and specialized "red teaming" or bias audit needs.

Scope & importance

Rigorous **testing and validation** of AI models is central to governance, especially for high-risk applications in banking. This category encompasses all the technical evaluations to ensure an AI system is **robust, accurate, fair, and secure, both prior to deployment and periodically thereafter**.

In practice, this means: verifying model performance on hold-out data, stress-testing the model under extreme or edge-case conditions, checking for **bias** (e.g. disparate impact on protected groups), testing the model's resilience to adversarial inputs or cyber-attacks, and validating that the model's outputs remain within acceptable bounds over time.

For banks, model validation is already a well-established discipline – for example, credit risk models must be independently validated as per banking regulations. The AI Act effectively extends this rigor to all "high-risk" AI models, requiring that they meet criteria for accuracy, robustness, and cybersecurity (Article 15) and do not produce unjustified bias (linked to broader EU non-discrimination law).

Costs of technical validation

According to the EU impact study, **testing and validation is the single largest chunk of compliance labour**. Approximately **€10,733 per model per year** (in staff effort) is attributed to ensuring "**robustness, accuracy and cybersecurity**". This constitutes about **36%** of the total compliance labour in the EU's breakdown, topping even documentation or oversight.

These activities involve data scientists and risk specialists running additional analyses on the model, often beyond what's needed for mere business performance. For instance, a bank might periodically re-run a lending model on recent data to check for **drift** in predictive accuracy or changes in error rates across demographic subgroups. They might also do **scenario testing** – e.g. simulate an economic downturn to see if a credit AI model remains within risk parameters. The high share of effort here reflects that such testing is not trivial; it must be thorough and is often iterative.

Evidence of Resource Needs:

- **Deloitte:** Observes that in regulated industries, organisations may spend up to **50% of their AI development resources on building compliance controls and conducting model audits before go-live**. In other words, for a given AI project, as much effort might go into "**test & trust" work** as into developing the AI's core functionality. For a bank, these controls include things like explainability features, bias mitigation steps, rigorous validation processes – all requiring additional coding, testing, and documentation.
- **KPMG:** Likewise notes that banks "**invest heavily in validation and testing to meet regulatory expectations**", and while this existing investment can cover some AI Act requirements (thus not all costs are net-new), it nonetheless represents **substantial ongoing expenditure**. In many banks, model validation teams (sometimes in Risk Management departments) exist solely to review and challenge models – their workload will increase as AI models proliferate.

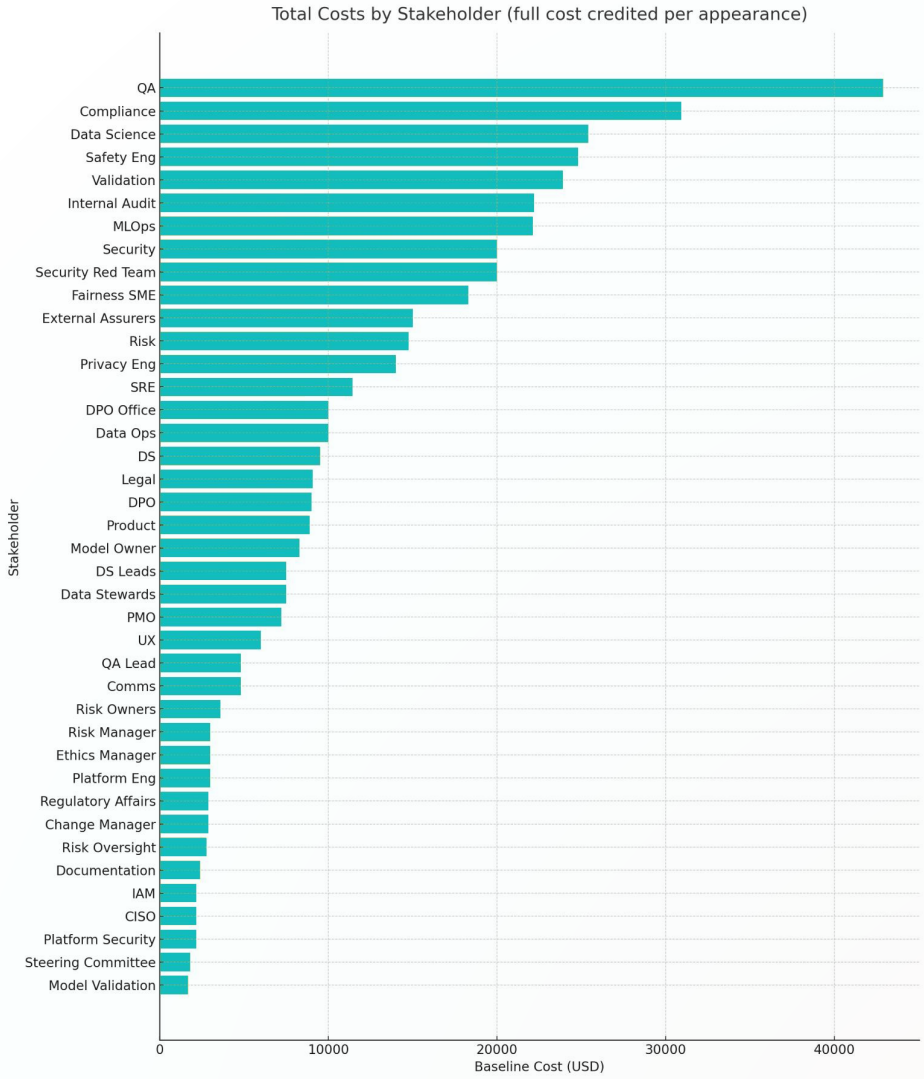
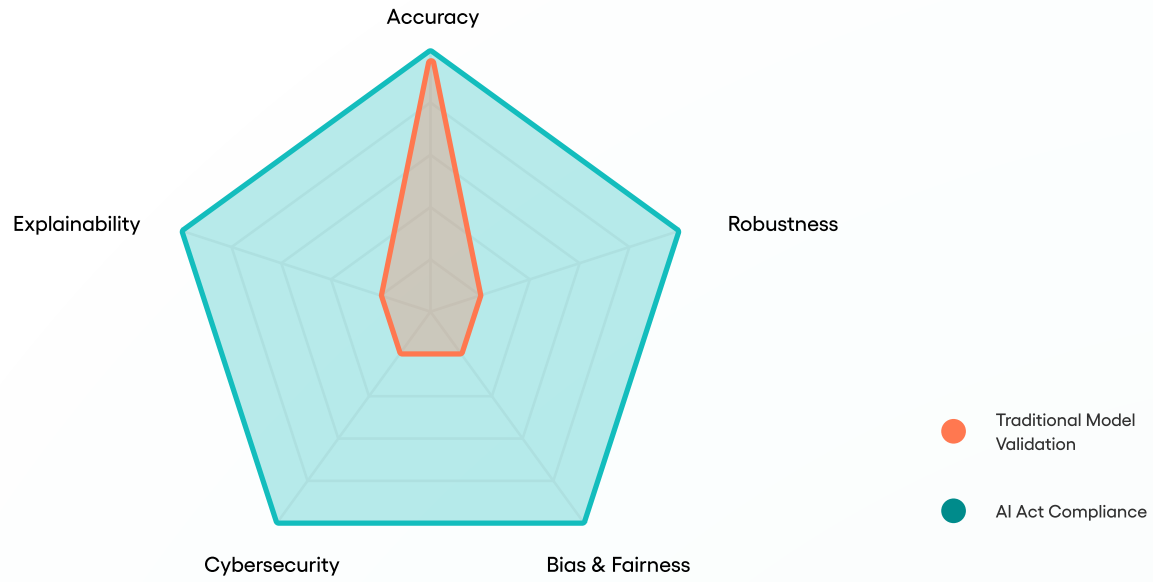


Figure R-4 — Total baseline cost credited to each stakeholder (full cost per appearance). Baseline shown in USD; aligns with EUR ranges discussed in text.

Model validation process

Typically, a **specialized validation team or external validator** will review each high-risk AI model, which can take weeks per model. They will examine the model's design, underlying data, and results from a governance perspective (e.g. ensuring no regulatory red flags, assessing fairness). This process is captured in the "**human oversight**" portion of costs (~€7.8k per model as noted in Section 1) – essentially an internal audit of the model's integrity.

Banks might also engage **external experts** or consultants for independent validation, especially for complex AI systems, adding to costs (often tens of thousands of euros per external validation engagement, not explicitly quantified in sources but analogous to fees for external model review).



The Validation Perimeter

Continuous testing

All sources stress that testing isn't a one-off at deployment; it's a **continuous or repeated cycle**. Any time an AI model is updated or retrained, or if its operating context changes, new testing is required. The EU study explicitly advises factoring in an extra **1–4% of the original development cost for each significant model change or retraining**.

Risk Warning (Continuous Change)

- **Retraining Spikes:** For example, if developing an AI model cost €500k, then each major update might incur up to €20k extra compliance testing.
- **Periodic Revalidation:** Banks commonly retrain models periodically (e.g. annual revalidation of credit models is standard). Thus, on top of the baseline €10k/year testing effort, one might see spikes of additional cost whenever a model is refreshed.

Fairness, bias, and security

A particularly salient aspect in banking is bias/fairness testing. Given the strict fair-lending and anti-discrimination laws in Europe, AI models for credit, insurance, hiring, etc., require bias audits. Several studies flagged that anti-bias checks will be necessary and could be costly.

- **Bias & Fairness:** Checking for disparate impact may require procuring special bias testing datasets or software tools, and engaging domain experts (like an "AI ethics" expert or fairness analyst). These might be new expenses for banks that historically didn't test statistical models for bias beyond basic compliance. The Center for Data Innovation's criticism of the AI Act was that such rigorous validation and bias-proofing requirements could "substantially raise costs, especially for small banks". EU analysts responded that much of this is just common best practice that prudent banks should be doing anyway. The reality likely lies in between: large banks may already perform some bias monitoring (for example, checking loan approval rates by gender/race as part of fairness policies), but the AI Act will formalize and possibly expand these checks, ensuring **nearly every high-risk model undergoes a fairness audit**. Thus, banks should expect to dedicate part of their testing budget to bias and ethics compliance.

- **Security Testing:** Ensuring **cybersecurity and robustness** also falls here. AI-specific security tests (e.g. adversarial example testing) might require hiring outside experts or using advanced tooling, again adding hidden costs not always broken out in reports. For example, a bank might perform a **"red team" exercise** on an AI chatbot to see if it can be tricked into giving sensitive info, which is a governance exercise overlapping with security.

Sector context

Because banks are **already culturally attuned to rigorous model testing** (thanks to decades of regulatory requirements for model risk), the incremental cost to meet the AI Act on testing may be less than for an unregulated tech company. Banks routinely do **back-testing, stress testing, and performance monitoring** for models. Much of that infrastructure (people and process) can be repurposed for AI. For instance, if a bank has a quarterly model review process under its risk management framework, AI models can be slotted into that process with relatively minor adjustments (though new dimensions like bias need adding). This is why EU officials assumed many compliance tasks are BAU – a bank with strong model governance will not start from zero.

However, **even the most "optimistic" official estimate still allocated over one-third of compliance effort to technical testing**, underscoring that **validation is inherently resource-intensive**. Our sources unanimously recognize that banks will continue to **spend heavily on validation** – what varies is how much more the AI Act layers on top of existing practice. Some consulting views imply that if a bank was already doing robust model risk management, their incremental spend might be moderate (closer to the lower end of cost ranges), whereas a bank with lax model validation would need a big uplift (higher end).

Summary

On average, **testing and validation remains a cornerstone of AI governance cost**, ensuring models are trustworthy and compliant but requiring significant expert hours and tools.

The Bottom Line	Testing is **~36%** of compliance effort.

Main Risk	**Drift:** Changes in data/market require expensive re-testing.
Efficiency Lever	**Test Suites:** Automated regression & bias testing pipelines.

4 Monitoring and ongoing maintenance

Estimated Cost	Variable: ~€8k (baseline) to €200k+ (complex scenario) per model/year.
Nature of Spend	Recurring OPEX (operating expense).
Key Activities	Drift detection, incident management, human reviews, retraining.
Main Driver	Model criticality, transaction volume, and stability of the environment (concept drift).

Scope & need for monitoring

Once an AI model is deployed in a live banking process, the governance effort does not stop. **Continuous monitoring and maintenance** is essential to ensure the model remains compliant and performs as intended over time.

The EU AI Act explicitly mandates a "post-market monitoring system" for high-risk AI. In practical terms, banks must:

- Track the AI's outputs and outcomes.
- Watch for signs of performance degradation or misuse.
- **Facilitate human oversight** (e.g., having humans review a sample of automated decisions).
- **Handle any incidents or errors** (like reporting serious malfunctions to regulators within a tight timeframe).
- **Update the model** or its controls as needed.

These activities incur **recurring costs year after year** – effectively forming the operating expense of AI compliance.

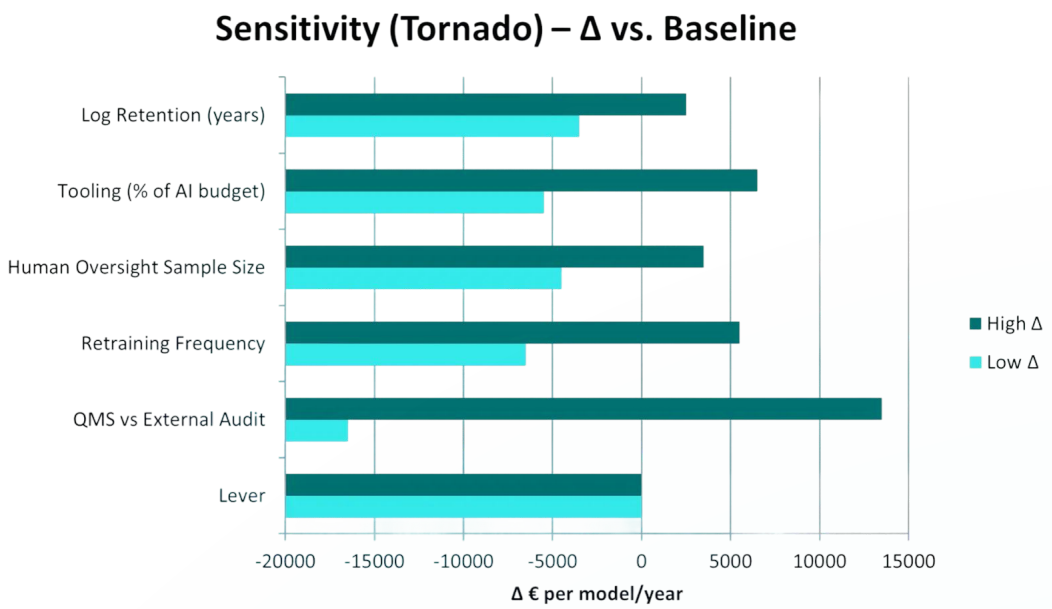


Figure R-5 - Sensitivity drivers shifting per-model cost; largest swings from scale, monitoring depth, release cadence.

Baseline vs. real-world monitoring costs

The EU impact assessment's cost model includes some elements of monitoring under the earlier-mentioned "human oversight" labour (~€7.8k), which covers tasks like manual review of AI decisions and addressing issues. However, industry experience suggests that **actual monitoring costs can be significantly higher** than this baseline, particularly for mission-critical AI systems.

RISK WARNING:
The €200k Surprise

A case in point involves an analysis informed by Gartner (the CostPerform mid-size bank scenario). It found that compliance with evolving AI regulations added about $200,000 in annual costs beyond initial deployment for that bank. These extra costs were largely due to the ongoing complexity of monitoring and controlling the AI system in production. This example, while hypothetical, illustrates how maintaining AI compliance "in the wild" can be an order of magnitude more expensive than one might assume from the static compliance checklist. It far exceeds the EU's ~€8k baseline, highlighting the gap between minimal compliance and robust, real-world operations.

Activities driving monitoring costs

There are several distinct cost components in this category:

1. **Automated monitoring systems**
   Banks often set up automated dashboards and alerts to track model performance (accuracy, bias metrics, drift in data input distributions, etc.). Building and tuning these systems requires initial investment (addressed in infrastructure costs) and ongoing resources to review the alerts and investigate them. If an alert flags that, say, a fraud model's false positives have doubled this month, an analyst must delve into it.

2. **Concept drift detection**
   Over time, changes in economic conditions or customer behaviour can cause a model to become less accurate (concept drift). Banks need to invest in tools and analysts to detect drift early. When drift or any deterioration is detected, the team might have to recalibrate or retrain the model, which as noted earlier, triggers additional testing costs and perhaps a minor re-certification.

3. **Re-training cycles**
   Sources like 2021.AI emphasize budgeting a few percent of the original dev cost for each retraining cycle, which is a part of ongoing maintenance. In practice, many banks schedule annual model reviews/revalidations, and each cycle consumes personnel time (validators, model owners) and possibly external reviewer fees. These cyclical review costs, while routine, add to the per-model annual expense.

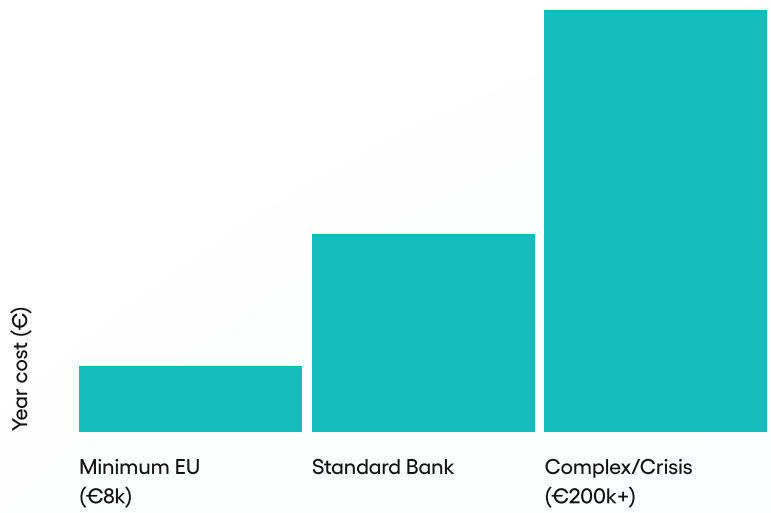
4. **Incident management**
   If an AI system causes an incident – e.g., a rogue trading algorithm makes an error, or an AI-driven credit decision leads to customer complaints of bias – the bank must investigate and document it. Incident handling often requires convening a review committee, doing root cause analysis, and deciding on fixes, plus reporting to regulators if required.

5. **Hidden infrastructure costs**
   The CostPerform scenario vividly showed that unanticipated needs like extra data storage or compute for monitoring can add another ~$150k to annual costs. This $150k was attributed to higher-than-expected computational requirements and data storage for monitoring logs – effectively a form of incident prevention cost. It underscores that beyond direct labour, supporting infrastructure (servers, cloud costs for monitoring processes, etc.) contributes to the maintenance bill. (We discuss infrastructure in the next section, but the line is blurred since monitoring needs drive infra needs).

6. **Model updates and patches**
   A portion of maintenance cost is applying updates or patches to the AI system. For example, if a bias is discovered, developers may need to adjust the model or add a mitigation, then re-deploy – incurring additional dev and test work. These are not one-time; they occur whenever issues arise or new regulatory guidance comes out requiring changes.

7. **Human-in-the-loop operations**
   In banking, many AI decisions still have a human override or review (e.g., a loans officer can overturn an AI's rejection). Banks often assign staff to review a sample of AI outputs regularly (say, reviewing 5% of all loan declines for fairness). This human oversight is ongoing labour. If each model needs, say, 1-2 full-time analysts to monitor its outputs, that's a fixed labour cost per model.

Ethos.ai notes that organisations may need to devote **~5% of their AI workforce's time** to continuous governance training and capability development (as mentioned in Personnel), which is part of keeping the monitoring process effective as regulations evolve. Staff must stay up-to-date on new compliance checklists or monitoring techniques.



The Cost Escalator of Monitoring

Range of estimates and sector context

The cost of monitoring clearly varies by scale. Our sources indicate a **range from as low as ~€8k per model** (the EU's implicit baseline for minimal oversight tasks) up to **hundreds of thousands of euros** annually in practice for more complex setups.

The range depends on:

- **Criticality:** Business-critical models justify more intensive monitoring.
- **Volume:** Large volumes of transactions require more logging and analysis.
- **Maturity:** How mature the bank's existing monitoring infrastructure is.

A **large bank** might absorb some monitoring cost into existing IT operations (e.g., piggybacking AI monitoring on an existing risk dashboard), keeping per-model costs toward the lower end. Conversely, a **smaller or less AI-mature bank** might have to build monitoring processes and tools from scratch for each model, making initial years very expensive.

Fortunately, banks tend to have an ingrained culture of monitoring due to compliance in other areas. They have **operations teams and risk dashboards** that continuously track metrics like credit portfolio quality or fraud losses. These can be extended to include AI model metrics (like a drift index or bias indicator) with relatively modest effort. Additionally, banks often intentionally keep a **human in the loop** for high-impact decisions (e.g. no fully automated loan rejection without human sign-off), which inherently limits unchecked AI actions and provides a layer of ongoing human oversight at marginal cost. These practices can reduce new monitoring costs to some extent.

However, as banks deploy more AI and rely on it for real-time decisions, they will likely need **dedicated AI monitoring solutions** (for example, automated bias detection running in production, or specialized audit logs for AI).

Summary

All sources agree that **ensuring models stay compliant over time demands persistent attention and budget**. Monitoring is effectively a non-negotiable cost centre for AI in banking – you cannot just deploy and forget an AI model under regulation. The difference is whether you can leverage existing systems to do it cheaply, or have to invest anew. In summary, monitoring and maintenance costs, while sometimes underestimated, can be quite significant. They include not just the obvious headcount to watch models, but also hidden costs like increased data storage, computing for continuous checks, and incident response efforts. One meta-analysis source summarized it well: initial deployment is just the beginning – maintaining compliant AI operations at scale requires dedicated and ongoing effort. Banks budgeting for AI should include a healthy allowance for these "run" costs in addition to upfront "build" costs.

The Bottom Line	Monitoring costs range wildly from **€8k to €200k** depending on complexity.
Main Risk	**Alert Fatigue:** Too many manual alerts overwhelm staff, leading to ignored risks.
Efficiency Lever	**Integration:** Piggyback AI monitoring onto existing Risk Dashboards to save ~30%.

5 Data management and infrastructure

Estimated Cost	~2–5% of total AI budget (tens of thousands per model).
Nature of Spend	Mix of Upfront (Tools) and Recurring (Storage/ Compute).
Key Activities	Data lineage, bias detection tools, secure logging, model inventory.
Main Driver	Volume of data retention required for auditability and "explainability".

Scope & types of costs

Underpinning all the above categories is the **data and IT infrastructure** needed for AI governance. Compliance isn't just about people and process; it also demands the right tools, systems, and data management practices.

In this category, we include:

- **Data preparation and governance:** Ensuring training and test data meet quality and representativeness standards.
- **Governance software/tools:** For model tracking, documentation management, bias detection, etc.
- **Computing and storage:** Resources needed to support compliance activities (like running additional validation jobs and storing logs).
- **Security/privacy infrastructure:** To protect data and models.

While these costs may be less visible than hiring staff, they are highlighted in many reports as **crucial elements of the cost structure** for AI governance. Essentially, without adequate infrastructure, a bank would have to throw even more people at the problem, which is not scalable.

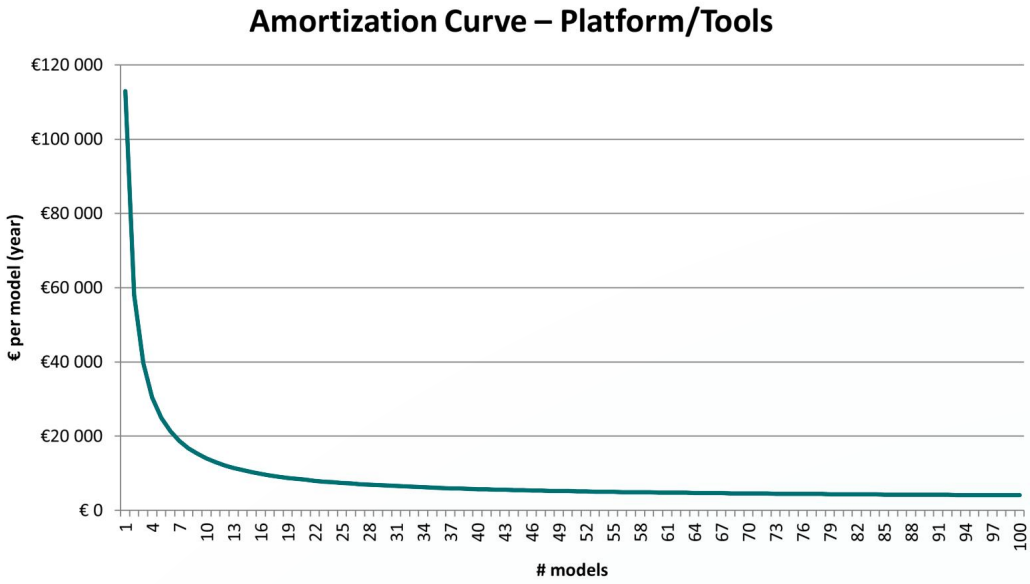


Figure R-6 — Per-model platform/tool cost declines as portfolio size increases.

Data preparation & quality

The AI Act will require banks to pay close attention to their **training data quality**. Article 10 of the Act demands that training data for high-risk AI be relevant, representative, free of errors "as far as possible", and **free from biases** that could result in discrimination. Ensuring this might mean extra work: acquiring new datasets to balance biases, performing extensive data cleaning, and documenting data provenance.

The EU study allocates a smaller portion (**~€2,763 per model/year**) to "Training Data" compliance, which covers tasks like dataset documentation and instituting data governance measures for AI. However, €2.8k/year (about 9% of compliance labour) suggests a few days of work by data engineers per model per year on purely data-focused tasks – likely an underestimate if major data sourcing or labelling is needed.

In real cases, if a bank realizes its AI model's dataset lacks certain minority representation, it might invest in a supplemental dataset or a new data collection effort, which could cost tens of thousands in vendor fees or labelling costs. One source (CostPerform) indicated that **inefficient use of cloud storage for AI data led to significant cost overruns (~25% of overruns)** until optimised, showing that data storage and management inefficiencies can directly hit the bottom line.

Governance tools and platforms

Many banks are turning to specialized **AI governance software** to streamline compliance. These include:

- **Model inventory systems:** To catalog all AI models and their compliance status.
- **Automated documentation tools:** That generate parts of the technical file.
- **Bias detection software:** To test for fairness.

- **Monitoring platforms:** To track model metrics.

Implementing these tools entails license or subscription costs, integration expenses, and possibly custom development. Industry analyses estimate that large enterprises allocate roughly **2–5% of their AI budgets to governance technology (tools & platforms)**. For a bank spending, say, €10 million on AI, that means €200k–€500k on tools.

On a per-model basis, if a bank has 50 models, this might average to **€4k–€10k per model** in tooling cost. If fewer models, the per-model share is higher. In our €116k–€137k total cost range, this tooling investment is one factor pushing toward the higher end.

Banks often face a build vs. buy choice: large institutions might develop in-house governance platforms (incurring developer cost but giving flexibility), whereas others buy enterprise solutions from vendors (paying licensing). Either way, it's a non-trivial budget item. **Without these tools, compliance would rely solely on manual effort, which would likely cost even more in labour**, so tools are partly an efficiency move.

Computing and storage overhead

Compliance requirements can drive up **IT usage costs** beyond what the AI solution itself needs. For example, to comply with transparency and record-keeping, a bank might store **every decision output of an AI system for X years** – if the model makes millions of predictions, that's a lot of data to archive. Similarly, running additional **validation jobs and bias tests** means extra cloud compute hours or on-premise server usage.

In the CostPerform bank scenario, as the AI system scaled up usage, the bank saw its **data processing and storage costs increase by 20%, adding about $600k annually** (though that scenario included more than just compliance overhead). Specifically tied to compliance, they noted an unexpected **$150k added to the budget for storage and compute** due to transparency and monitoring needs. This indicates that logging every detail and maintaining sandbox test environments are not free – they require disk space and CPU cycles. Banks using cloud services might see their monthly bills rise because they are retaining more data and running more analytic processes to satisfy auditors.

We can consider that out of our ~€100k+ per model costs, a portion is infrastructure: maybe on the order of **€10k+ per model/year** in extra cloud/server costs in active use cases (depending on volume of data). Banks must account for this when justifying AI projects: the **total cost of ownership** includes these compliance computing costs.

Security and privacy infrastructure

Ensuring **data security and privacy** for AI systems is both a regulatory requirement and good practice – and it carries costs. This includes **encryption** of data at rest and in transit, robust **access controls** (so only authorized personnel can access the model or its data), and monitoring for data leaks or unauthorized use.

While these might not be line-itemed as "AI Act compliance costs" specifically, they are part of the broader governance expense. For instance, if an AI uses personal data, compliance with GDPR (which still applies) means investing in anonymization or pseudonymization solutions, and continuous monitoring for data breaches. A bank might extend its existing security tools to cover AI pipelines – e.g. ensuring model training data is stored in a secure data lake with proper access logging.

Most banks already have significant IT security budgets, and AI governance will tap into those. We interpret from sources that some fraction of the **existing IT spend** can be attributed to supporting AI governance (one might allocate a portion of cybersecurity spend to AI projects proportionally). The key point is that while security/privacy infra might not be newly purchased just for AI, using it for AI adds to the overhead. For example, rotating encryption keys for AI model storage or implementing additional monitoring on AI data flows might require extra work by the IT department (costing time and money).

Sector context & scale economies

Banks, especially large ones, often have a head start on infrastructure: they likely already maintain **enterprise data lakes, logging systems, and GRC (governance, risk, compliance) platforms** for other regulations. This existing infrastructure can sometimes be extended to AI. For example, if a bank has an operational risk system that logs incidents, it could incorporate AI incidents into it rather than building a new system. Or a bank's big data lake used for storing transaction logs could also store AI decision logs.

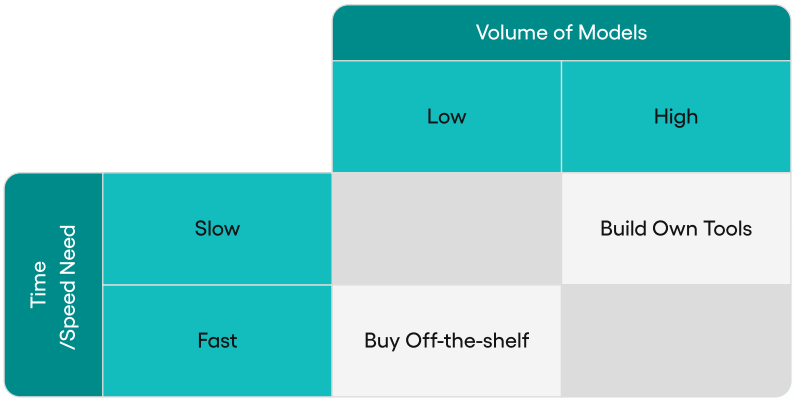
By integrating AI compliance into **legacy systems**, banks can save costs. Many banks also have central "model risk" IT systems (some use vendor products for model inventory) which can be upgraded for AI model tracking. This offers **economies of scale**: the more models and processes you can cover with one system, the lower the average cost per model.

A clear insight from sources is that **early investment in centralized AI governance platforms yields cost efficiencies** – banks that do this may keep per-model costs near the lower bound of estimates. In contrast, banks that are reactive, adding ad-hoc tools for each new requirement, might incur higher marginal costs each time. Essentially, a one-time €300k platform spread across 30 models is €10k/model, but doing piecemeal solutions could cost much more cumulatively.

Summary

In summary, while **infrastructure and data management costs** might not dominate the headlines like personnel or audit costs, they are a **vital part of the equation**. They ensure that compliance processes can run efficiently. All our sources implicitly or explicitly agree that without adequate tooling, **compliance would be far more labour-intensive and costly**. Thus, spending on infrastructure is **necessary to keep the total costs in check**, even if it itself represents, say, 10–20% of the total compliance spend.

For planning purposes, bank executives should recognize these as mostly **upfront or fixed investments** (e.g. buying a tool, setting up logging systems) with high initial costs but then lower incremental costs per model. Those who invest early in robust governance infrastructure likely find the **per-model compliance cost stabilizes or drops** over time, whereas those who don't will face compounding inefficiencies.



Decision Matrix: Build or Buy

The Bottom Line	Tech costs are **2-5%** of AI budget but prevent labour explosions.
Main Risk	**Technical Debt:** Ad-hoc tools for each model create integration nightmares.
Efficiency Lever	**Platform Strategy:** Centralise governance tools early to amortise costs.

6 Audit and compliance certification

Estimated Cost	Variable: ~€20k (External/model) vs ~€71k/yr (Internal QMS upkeep).
Share of Total	Depends on strategy (10-20% of per-model cost).
Key Activities	Conformity assessments, QMS setup, surveillance audits, registry updates.
Main Driver	Strategic choice between "pay-per-model" (External) or "invest-upfront" (Internal QMS).

Scope & regulatory expectations

High-risk AI models in the EU will be subject to **formal conformity assessment** before they can be put into use (and periodically after, in some cases). This is akin to a certification process, ensuring the model meets all requirements of the AI Act.

There are two primary routes to compliance:

- **External third-party audit:** An independent notified body evaluates the AI system and issues a certificate.
- **Internal conformity assessment via a Quality Management System (QMS):** The provider (e.g. the bank) implements a comprehensive QMS and self-assesses the model's conformity, with oversight from authorities.

Both routes involve costs – either paying external auditors or investing internally in a QMS framework. Additionally, banks will face **ongoing audit activities**: their internal audit departments will include AI in routine audits, and regulators (or external auditors) may conduct **surveillance audits or inspections** periodically. This category captures those one-time and recurring assessment costs, as well as the overhead of maintaining systems to facilitate audits.

Initial certification costs (External vs. QMS)

The EU impact study provides concrete estimates here.

- **Third-party "EU-type" audit:** The cost was estimated at **€16,800 to €23,000 per model** (one-time). This would be the fee paid to an external auditor to certify the model compliant pre-deployment.
- **Internal QMS:** Establishing an internal QMS for AI compliance is far costlier upfront: approximately **€193,000–€330,000** to set up the system, plus **€71,400 annually** to

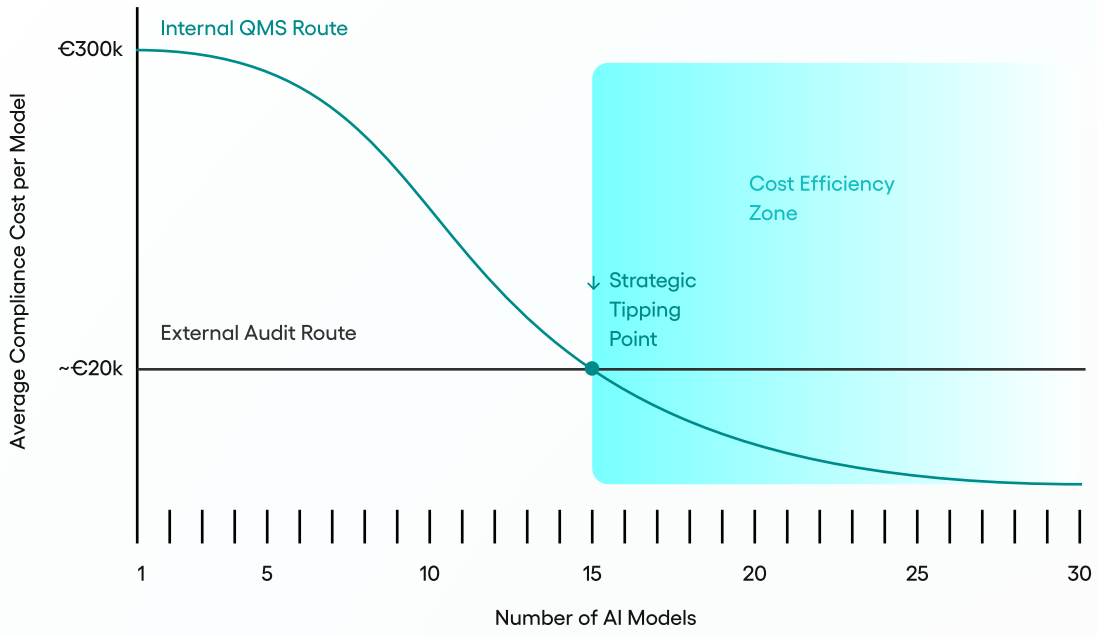
maintain it. These figures were later corroborated by CEPS researchers who authored the study, emphasizing that a new QMS (covering processes, documentation systems, training, etc.) is a **huge expense** – likely prohibitive for small providers.

Indeed, the logic is that **smaller firms will opt for external audits per model rather than incur the heavy fixed cost of a QMS**.

Comparison: Banking sector advantage

For banks, however, the picture is different: banks tend to have at least partial QMS frameworks already (e.g. many follow ISO 9001 or have enterprise risk management systems). So a bank might not need a €300k build from scratch; instead, it might **augment existing governance systems** to meet AI Act criteria. Our sources elaborate on this dynamic. A European Commission staff analysis (noted by CEPS) assumed many companies "already fulfill parts of the requirements as BAU", hence downplaying net-new QMS costs.

In practice, a **large bank with dozens of AI models** could spread the upfront QMS cost across those models, making the per-model share relatively small.



The Amortization Curve

For example, a one-time €300k implementation across 30 models is €10k per model initial, plus perhaps each model's share of the annual €71k maintenance (≈€2.4k). That is manageable and likely at the low end of cost per model. By contrast, a **small fintech bank with only 1-2 AI models** might find an internal QMS unjustifiable, and instead pay ~€20k per model for external certification.

Spreading and amortising costs

As noted, QMS costs are largely fixed – once you have a robust governance framework, adding another model doesn't double the cost, it just adds incremental work. Our validated cost range (€116k–€137k/model/year) likely **assumes some amortization** for typical banks. It sits in between the extreme cases: far above the minimal €52k conservative scenario (where only a one-off €23k audit + €29k labour was counted), but below the €400k single-model worst case.

One CEPS commentary addressed the misunderstanding: a small business might face up to €400k for one high-risk AI product (if they had zero QMS prior), but their actual study suggested €193k–€330k + €71k maintenance, which if reused for future models, spreads out.

Ongoing audits and compliance upkeep

Achieving initial certification is not the end; maintaining it is an ongoing cost. Banks with an internal QMS will need to undergo **annual surveillance audits** by a certification body to keep their certification active. They also must continuously **improve the QMS**, dedicate staff to manage it, and update documentation as regulations evolve. The €71k/year maintenance figure covers paperwork and system upkeep, but **not the staff time** to actually operate under the QMS – effectively, that staff time is already counted in personnel/documentation categories.

Still, the QMS brings its own admin overhead. Meanwhile, banks relying on external audits might need to recertify models periodically (say, if a major change is made to a model, a new audit might be required). Each such audit could cost a similar ~€20k. Additionally, the AI Act will require **registering high-risk AI systems in an EU database** and keeping those entries updated. While filling registry forms or uploading info might be a minor task, doing it for dozens of models and keeping it current adds to compliance workload (consider it part of the audit overhead).

Consensus on certification costs (Varied but significant)

Our sources uniformly acknowledge that **meeting the certification/conformity assessment requirement is a notable cost**, but they quantify it differently depending on context. The **2021.AI/EY analysis** we cited combined annualized certification cost with labour to get ~€52k/model/year as a "conservative" total – that basically assumed the cheaper route (external audit ~€23k, not every year but per major version) plus the baseline labour ~€29k. By contrast, our meta-analysis identifies a higher range (€116k–€137k) because we include **all categories** (personnel, infra, etc.) – essentially giving a full picture rather than just labour+audit.

Some experts, particularly from industry, argue that the Commission's estimate (~€50k) was **too low** when you factor everything in. The aforementioned CDI paper extrapolated up to **€10 billion industry-wide annually** and ~€400k for an SME's AI system with full QMS. While that was likely an overestimate (and CEPS rebutted it), it underscores that **if a firm must newly invest in governance systems and multiple audits, the per-model burden can spike dramatically**. In reality, most banks will fall somewhere in between: not as lean as the Commission assumed, but also not starting from zero.

Integrating with existing audit programmes

Banks have an advantage in that they are **already subject to frequent audits** – internal audit departments check compliance with various regs, and external regulators perform inspections. Many larger banks plan to **integrate AI compliance into existing audit cycles**. For example, if a bank already does an annual internal controls audit (perhaps for SOX or similar requirements), they might extend it to cover AI model controls with minimal incremental cost. Internal auditors on the payroll can be trained to include AI checkpoints in their reviews.

Similarly, banks might collectively develop **shared assessment frameworks** (through industry associations) to standardize AI audits, which could reduce duplicated effort if regulators accept those frameworks. These practices could keep ongoing audit costs on the lower side for well-prepared banks. Nonetheless, given that many AI use-cases in banking are high-risk (lending, fraud, AML, etc.), **regulators will likely demand thorough audits**. Banks may still engage **third-party experts** for extra assurance, especially in early years when internal expertise is building. Consulting firms like Deloitte are already offering "AI Act readiness" audit services. Such external engagements are not cheap – they might bundle into project budgets or appear as consulting fees.

One Deloitte insight suggests banks can **re-use GDPR compliance structures to save costs** in AI Act compliance, implying that if you didn't re-use, you'd incur a lot more effort and cost. This reinforces the theme: leverage what you have (GDPR, model risk management, ISO, etc.) to avoid reinventing the wheel for AI – if you don't, compliance could be drastically more expensive.

Large vs. small players

Summarizing the range, we expect **large EU banks with mature compliance programmes** to keep certification costs near the lower bounds – perhaps **tens of thousands of euros per model annually** (when amortised), which might be just ~10-20% of their total per-model compliance cost. In contrast, **less prepared organisations or those with only a few models** will experience higher per-model certification costs, contributing to the upper end (~€137k or more) of our range. This aligns with our earlier discussion on scale economies and initial setup burdens.

Podsumowanie

In conclusion, **audit and certification** is a distinct and important cost category in AI governance. It's the one that often "**makes it official**" that a model is compliant, but attaining that stamp of approval involves real expenditure – either paying external auditors or running an internal compliance machine. The meta-analysis finds broad agreement on its necessity (no one argues you can skip it), but a range on its magnitude. The prudent strategy for banks is to **invest upfront in scalable compliance systems** (even if that means a big one-time cost) so that adding new AI models doesn't linearly increase audit costs. The more a bank can bake AI into its existing compliance DNA, the less painful each new model's certification will be.

The Bottom Line	"Audit is the "Official Stamp". Cost varies by route (External vs Internal)."

Main Risk	**Bottleneck**: Shortage of external auditors could delay model launches.
Efficiency Lever	**QMS**: Building an internal system is cheaper for portfolios >10 models.

Discussion

Bringing together the findings across these six cost domains, a **holistic picture** emerges of what it takes – financially and operationally – to govern AI in the European banking sector. There is a strong consensus on **which** cost elements matter: virtually every reputable source emphasizes the need (and cost) for:

- Skilled personnel
- Comprehensive documentation
- Rigorous testing regimes
- Continuous monitoring frameworks
- Supportive data/IT infrastructure
- Formal audit/certification processes

These can be viewed as the **pillars of AI governance in banking**, and our analysis confirms that they collectively come with a **substantial price tag**.

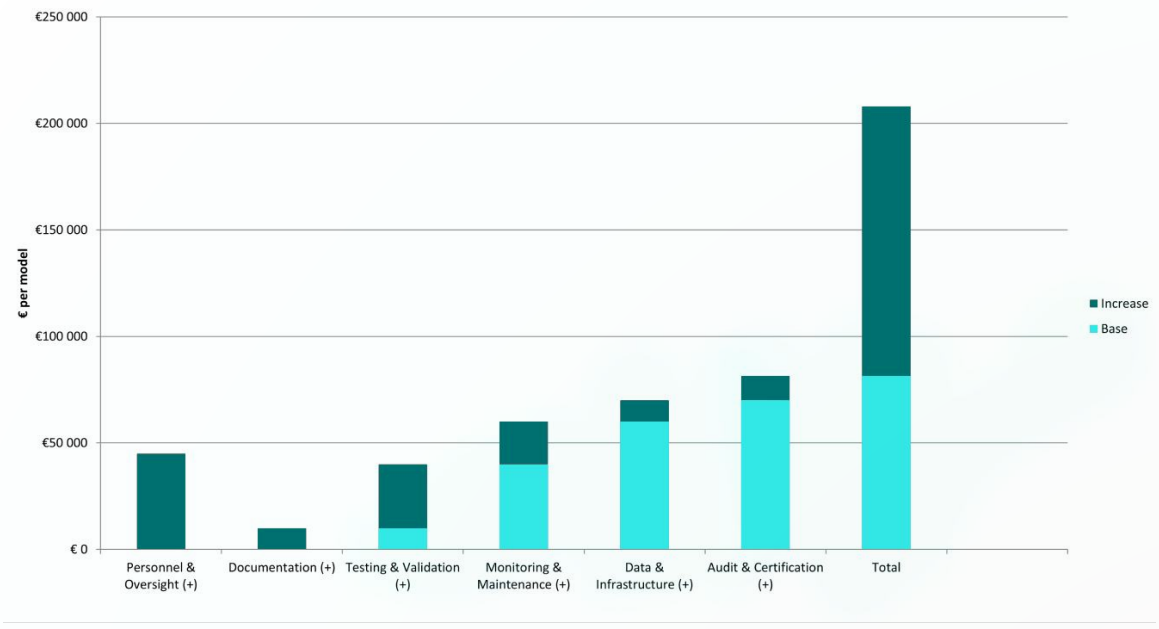


Figure D1 — Waterfall from component estimates to €116k–€137k per model per year.

Convergence on total cost range

One of the key outcomes of this meta-analysis is the validation of an **annual per-model compliance cost on the order of six figures (EUR)**. Multiple independent sources support this order of magnitude.

When we sum up the contributions of each category as reported (and adjusted for overlaps), we arrive at a range of roughly **€116,000 to €137,000 per model per year** for ongoing governance and maintenance. It is noteworthy that this range is **neither an outlier high estimate nor the regulatory lowball** – it sits **between** the conservative official figures and the more expansive industry scenarios, suggesting it's a realistic middle-ground.

For example, the EU's own analysis might imply something like ~€50k/model if many controls exist already, while some consulting scenarios hint at >€100k (our focus) and extreme cases like CDI's projection soared to ~€400k for a one-off case. The €100k+ range we present is thus a **balanced estimate**, reflecting that banks will likely implement **more than the bare minimum** but also **gain efficiencies** compared to worst-case new entrants.

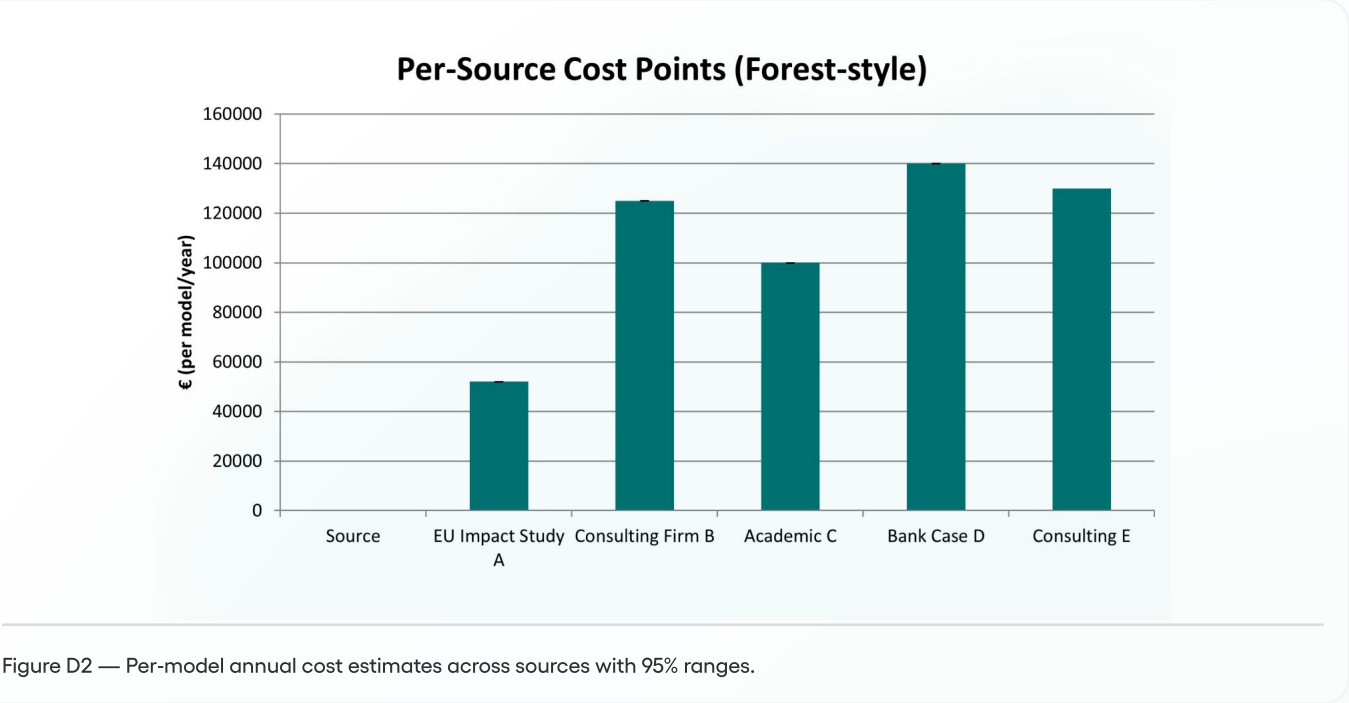


Figure D2 — Per-model annual cost estimates across sources with 95% ranges.

Cross-source triangulation

We also found **cross-source triangulation** for specific benchmark numbers within that total. For instance:

- The **€29k compliance labour** figure from the EU study is echoed in practice by banks allocating 10–20% of AI budgets to compliance (which for a multi-million project easily translates to tens of thousands in staff costs).

- The **€7.7k human oversight** portion aligns with banks' practice of dedicating staff to check AI outputs and the requirement of Article 14 of AI Act (human oversight), reinforcing that thousands per model are spent on this activity.

- The **€4.4k documentation** cost resonates with qualitative accounts of teams spending weeks on paperwork – multiple sources like Ethos.ai underscore documentation is about 10–15% of effort, so that matches well.

- The **€2.8k data compliance** cost is relatively small, but we saw that could spike if additional datasets or storage are needed (like CostPerform's $150k surprise). So while the core data governance tasks might be single-digit thousands, the **infrastructure** to support data and logging can add tens of thousands (the 2–5% of budget on tools we cited equates to sizable amounts).

- The **€20k audit vs. €200k+ QMS** dichotomy was consistently referenced across policy and consulting analyses, giving us confidence in those numbers. And seeing that banks likely amortise QMS costs supports why our per-model figure is not extremely high (since large banks won't pay €200k per model, they pay it once for many models).

Overall, the **areas of agreement** among sources are strikingly strong. Hardly any credible source argued, for instance, that documentation or monitoring could be skipped or done on the cheap – they all acknowledge these as necessary and often costly. Where sources differed was primarily in **magnitude and starting assumptions**.

Reasons for variability in estimates

The variations in cost levels across studies can largely be explained by differences in **scope assumptions, scale, and methodological approach**:

Driver of Variability	Low Estimate Scenario (~€50k)	High Estimate Scenario (€100k–€400k)
1. Scope of compliance	Looks at incremental administrative burden only.	Comprehensive view including monitoring overhead and infrastructure.
2. Existing controls (BAU)	Assumes processes (ISO, Risk) exist and can be reused.	Assumes "Greenfield" implementation (starting from scratch).
3. Timing (Setup vs Run)	Amortises setup costs perfectly over time.	Conflates first-year "setup" spikes (e.g. QMS build) with annual run-rate.
4. Safety margins	Minimum viable compliance.	Industry conservatism: buffering for "unforeseen complexity".
5. Cherry-picking	Focuses on the cheapest audit route.	Focuses on complex, high-risk scenarios (e.g. CDI's €400k case).

1. Scope of compliance considered

Some studies looked at a subset of costs (e.g. just incremental administrative burden) while others took a comprehensive view. The 2021.AI/EY piece giving ~€52k was explicitly a "conservative" partial view (development costs only, did not factor retraining or possibly infrastructure). In contrast, our analysis includes those additional factors (monitoring overhead, etc.), hence arriving at ~€100k+. Industry consultants often take a more

comprehensive (and sometimes pessimistic) scope, including contingency for "unforeseen complexity", which raises their estimates.

2. Existing controls (BAU) vs. Greenfield

Official EU estimates assumed many compliance practices were already in place (hence lower incremental cost). For example, if a bank already has ISO-aligned quality processes, meeting some AI Act requirements might just mean tweaking them. On the other hand, analyses highlighting high costs often implicitly consider firms that need to implement many things from scratch. This BAU factor is critical: our conclusion section of the results noted that differences largely reflect scale and starting point.

- **Big banks** = high baseline capability, meaning new costs are moderate per model;
- **Small/new players** = low baseline, meaning steep first-time costs.

3. One-time vs. recurring and first-year vs. steady-state

Some sources (like CDI's critique) conflated first-year setup costs with annual ongoing costs, leading to sensational figures (the "€30 billion by 2025" headline). CEPS clarified that was an overestimate partly by not spreading costs or accounting for only 10% of AI being high-risk. In steady-state, once systems are in place, per-model costs likely decrease somewhat.

Our analysis tries to reflect a steady-state annual cost. It's worth noting, however, that **in the initial years of AI Act compliance (e.g. 2025–2026)**, banks may incur extra "setup" costs beyond these annual figures as they establish frameworks and train staff. For instance, a bank might spend a big lump sum on a new model inventory system or external consultants to kickstart compliance – those wouldn't be recurring every year. Thus, our ~€100k+ per year is an average steady run-rate; year one for a given bank or model might be higher (maybe double, if a full QMS build is needed), which then tapers in subsequent years. This insight can be important for budgeting: front-loaded investment vs. ongoing costs.

4. Conservatism and "Safety margins"

Industry risk management culture often adds buffers to cost estimates as a hedge against unknown challenges. For example, some banks might intentionally over-budget for compliance to avoid shortfalls, whereas the Commission (in trying to not overstate burden) provided rather minimal estimates. Our range €116–137k could implicitly include a bit of safety margin as it comes from summing robust practices across categories. It assumes, for example, that banks will implement fairly comprehensive monitoring, not just bare minimum. If a bank chose to do absolute bare minimum (only what regulators explicitly require and no more), they might operate at the low end or below. But that carries risk if something goes wrong that wasn't monitored.

5. Cherry-picking vs. Triangulation

The question specifically cautioned against cherry-picking data points without validation. We encountered an example: one source said "€29,300 for compliance labour" – we validated that against other sources: Accenture's % of budget and Deloitte's % both imply a similar scale (if 10–20% of a multi-million project, easily a few €10k's in labour). Similarly, the "€71,400 for QMS upkeep" was cross-checked with CEPS which confirmed it. This triangulation gives confidence in these figures. On the other hand, CDI's €400k claim was confronted with CEPS' original numbers which were lower, so we treated €400k as an outlier applicable only to a very narrow scenario (small firm, first model) – not a generalizable number for banks (especially large banks).

Our meta-analysis approach inherently avoids cherry-picking by combining multiple references for each claim. If one consulting firm reports a very high cost in one category, we juxtapose it with perhaps the official estimate or another firm's to show context. For example, if CostPerform said $200k for monitoring, we also note the official baseline was €8k and then explain why there's a gap. This provides a more nuanced understanding: that $200k is not every model's reality but a cautionary tale of how things can escalate if underestimated.

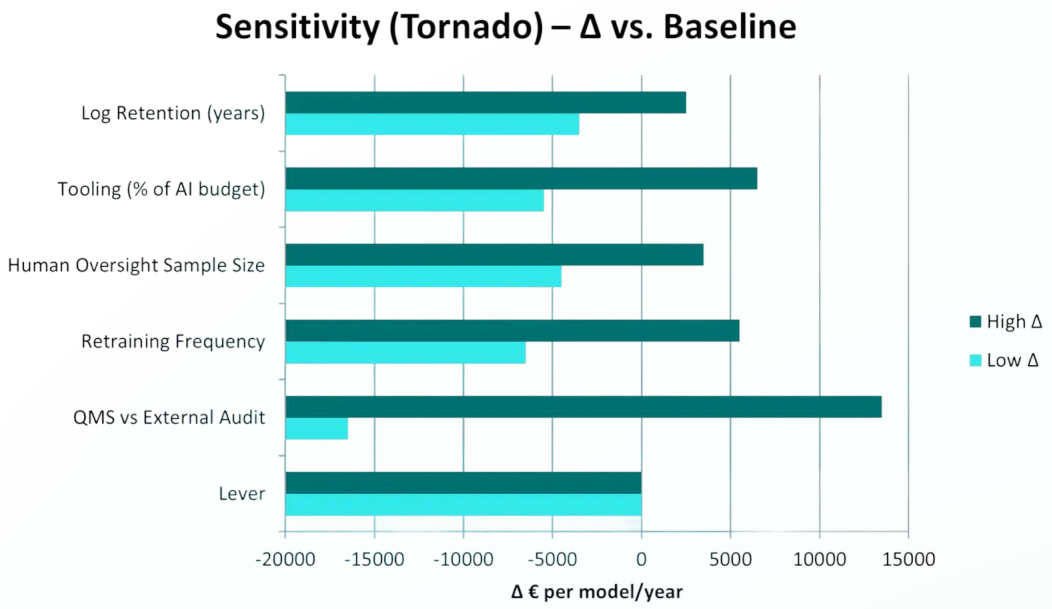


Figure R-5 - Sensitivity drivers shifting per-model cost; largest swings from scale, monitoring depth, release cadence.

Sector-specific considerations

Focusing on EU banks, a few themes emerged that might not apply to other sectors:

- **Existing rigour:** Banks' pre-existing rigorous compliance culture (thanks to financial regulations) both helps and hurts. It helps in that banks can adapt existing systems (so incremental cost is somewhat reduced). It hurts in that the **baseline expectations are already high** – regulators won't cut banks slack; they will expect state-of-the-art risk management for AI just as they do for credit or operational risk. So banks can't take a light approach without reputational or regulatory risk. This means that, unlike perhaps a startup, a bank will likely implement most of the controls to a thorough degree, hence incurring significant costs. In essence, the banking sector's culture of compliance means **AI governance becomes an extension of existing complex governance**, not something treated casually.

- **Use of model risk frameworks**: We saw references to aligning AI governance with model **risk management frameworks** (like OCC guidance). Many large banks already have model inventory, model validation units, periodic review cycles, etc. This is an advantage – for example, a bank might not need to hire a whole new "AI validator" team if its current model validators can handle AI models (with training). We mentioned how this overlap can moderate incremental personnel costs. In our estimates, we implicitly assume banks will leverage such overlaps; if they didn't, costs would be higher. For instance, a bank that runs a separate parallel governance structure for AI distinct from its existing risk management would be duplicating efforts (some early reports show banks create "AI Ethics committees" separate from risk committees – which might be necessary initially but ideally merge into overall governance to save cost).

- **EU regulatory climate**: As of 2025, the EU AI Act is nearing implementation (likely enforceable by 2026). Banks are among the organisations preparing early (some already conducting **AI Act readiness assessments** in 2024). This proactive stance means some compliance costs are being incurred now even before the law hits – e.g. banks investing in pilot audits or gap analyses. Our analysis captures the steady-state cost when in compliance; however, in the Discussion it's worth noting **that the years 2024–2025 might see spikes in spending as banks gear up for the AI Act**. Consulting firms have been marketing services to banks to help them prepare (Deloitte, EY, etc.), which shows banks are willing to spend now to avoid being caught unprepared. The cost range might thus be on the **lower side for banks that have gradually built up capabilities**, whereas banks that delayed action might face a steep cost ramp-up when the law takes effect.

Implications for bank executives

For C-level executives (the target audience), the implications are clear: **governing AI is not cheap**, but it's necessary. They should plan for **six-figure annual costs per high-risk AI model** in compliance overhead. For a bank deploying dozens of AI models, this translates to multi-millions in governance spend. This might seem burdensome, but it aligns with what these models can generate in value (e.g. a fraud detection AI could save millions in fraud losses – justifying a robust governance spend to ensure it's reliable and unbiased).

Executives can also take away that **investments in infrastructure and process can pay off long-term** by reducing marginal costs: e.g. building a solid QMS or buying a good governance platform once, then reaping lower per-model costs. Another strategic point: collaboration and reuse. Reusing existing compliance structures (GDPR processes, model risk frameworks, etc.) can materially lower costs. Collaborating via industry bodies to share best practices or even collectively fund certain compliance tools (perhaps through consortia) could be beneficial.

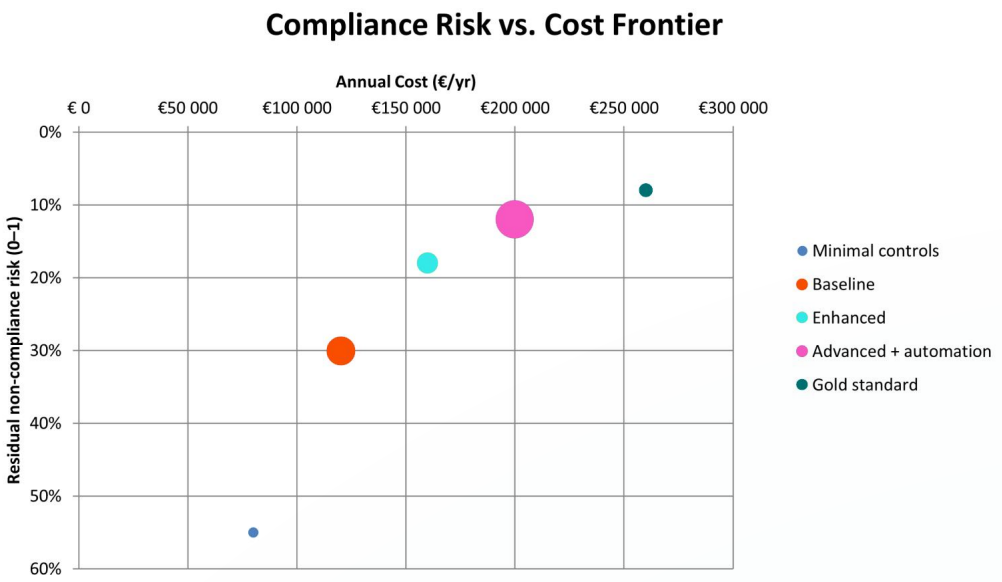


Figure D3 — Tradeoff curve from minimumviable to comprehensive governance.

Risk of underinvestment

It's worth discussing that while these costs are significant, the cost of non-compliance could be higher. The AI Act foresees heavy fines (up to €30M or 6% of turnover for serious infringements). So a bank that tries to save money by skimping on governance may face penalties far exceeding the compliance costs, not to mention reputational damage. In that sense, many executives view these governance costs as insurance and protection for the organisation's license to operate. The Accenture survey finding 77% of execs see future AI regulation as a company-wide priority reflects that sentiment – they are willing to spend to avoid bigger costs later.

Uncertainties and future outlook

There are still unknowns that could shift these cost dynamics. For example, if regulatory guidance comes out that simplifies compliance for banks (perhaps acknowledging existing bank supervision can cover some AI Act aspects), that might reduce duplication and cost. Conversely, if high-profile incidents occur, regulators might demand even more stringent oversight (raising costs). Technological advancements might also help: better AI governance automation could reduce human labour needs (some vendors claim to automate documentation or bias detection, potentially cutting labour hours). Over the next few years, we might see cost efficiency improvements as best practices mature and tools get better – meaning the per-model cost could stabilize or even drop slightly for those adopting new tech. However, any such gains could be offset by increased complexity of AI models (e.g. large generative models which are more complex to govern than traditional ML).

Exclusivity to banking

We excluded other financial sectors, but a quick note: if one compared banking to, say, a tech company, banks likely spend more on compliance proportionally due to stricter regulations. A Big Tech deploying an AI might not have an internal audit or worry about Basel or GDPR in the

same way, so their compliance cost might be lower (until laws catch up). Thus, our findings are specific to banking's high compliance bar. Within banking, some use-cases might incur somewhat different cost structures (e.g. an AI used internally for marketing might be lower risk than one making credit decisions, thus possibly falling outside "high-risk" definition and requiring less governance). But we focused on high-risk use-cases, which most impactful AI in banking will be categorized as.

In conclusion of the discussion, bank executives should see these costs not as a one-time hurdle but as an **ongoing investment in trustworthy AI**. The narrative in the industry is shifting towards viewing robust AI governance as a **competitive advantage** rather than just a cost – i.e. banks that manage AI risks well will earn customer trust and be able to deploy AI more broadly. Indeed, a compliant AI system can be deployed with confidence, whereas a non-compliant one might be stuck in legal or reputational limbo. Therefore, while the price tag is high, the **value of compliance** (avoiding fines, enabling AI innovation safely) is arguably higher. Our meta-analysis provides the evidence base for these costs so that banks can plan realistically and not underestimate the resources required.

Conclusion

Cost range validation
(€116k–€137k per model/year)

In summary, this meta-analysis finds that European banks should expect to spend on the order of low hundreds of thousands of euros per year for each high-risk AI system's governance and lifecycle maintenance. The convergence of evidence from regulatory studies and consulting reports supports this range.

Minor differences in estimates across sources largely boil down to scope and scale: official analyses assume many controls are already in place (yielding lower estimates, ~€50k/model in ideal scenarios), whereas industry and consultancy assessments consider more comprehensive or worst-case implementations (pushing costs above €100k/model). Our validated range sits between these extremes, incorporating all major cost components.

Key cost drivers

Encouragingly, **all sources agree on the key cost drivers and necessities**. Nearly every study underscores the importance – and cost – of the following pillars:

- **Skilled personnel & oversight:** Human experts (compliance officers, risk managers, data scientists) and governance committees are indispensable, forming the backbone of AI compliance efforts.
- **Robust documentation:** Maintaining detailed technical documentation and audit trails is laborious but universally deemed crucial for accountability.

- **Rigorous testing & validation:** Continual technical validation (for accuracy, bias, robustness) is a substantial ongoing expense, viewed as essential to meet both regulatory and ethical standards.
- **Continuous monitoring:** Systems and processes to monitor AI performance in production and manage incidents are a non-negotiable cost that grows with the complexity of models.
- **Data & infrastructure:** Investment in data quality, secure pipelines, and governance tools is needed to enable efficient compliance; without it, personnel costs would skyrocket.
- **Audit & certification:** Formal compliance assessments (whether through external audits or internal QMS) carry significant costs, and banks must budget for initial and ongoing certifications.

The scale factor

The **differences in cost levels** observed are mostly attributable to a bank's ability to spread fixed costs and leverage existing frameworks. A large bank can amortise one-time compliance investments (like a QMS or an enterprise monitoring system) across many AI models, driving per-model costs toward the lower end. In contrast, a smaller player implementing governance for a single AI system might incur the full suite of costs for that one model – hitting the higher end or even exceeding it (especially in year one, if a full governance setup is built from scratch).

Sector-specific factors also play a role: banks' pre-existing compliance infrastructure and stringent model risk management rules mean some needed activities (e.g. model documentation, validation) are already part of business-as-usual and budgeted, mitigating new costs. However, those same factors set a high baseline of rigour that AI governance must match, ensuring that even with synergies, the incremental costs remain significant – there is simply no "light-touch" option for high-risk AI in banking.

Final takeaway

For bank executives, the takeaway is clear: governing AI is a **serious ongoing investment**. The figures on the order of **€100k+ per model per year** should be incorporated into ROI calculations and project planning for AI initiatives. In many cases, these governance costs will represent **10–20% (or more) of the total cost of an AI project** – a substantial, but necessary, allocation to ensure compliance and avoid far costlier repercussions of non-compliance.

Boards and C-suites should view this expenditure as part of the "**cost of doing business with AI in finance**" – much as they view capital reserves as the cost of doing lending, or cybersecurity spend as the cost of digital banking. Looking ahead, as AI regulation (like the EU AI Act) comes into force, we expect further clarity on compliance expectations, which may allow banks to streamline some efforts and possibly find efficiencies through standardization. Industry collaborations (through federations or standards bodies) might reduce duplicate development of governance solutions. Additionally, advancements in AI governance technology (e.g. automated documentation generation, AI-driven monitoring tools) could help temper the growth of compliance costs.

Nonetheless, given the high stakes of AI in banking – impacting customers, financial stability, and banks' reputations – r**egulators and banks alike are inclined to err on the side of rigour over frugality**. In conclusion, this meta-analysis provides evidence that **robust AI governance in EU banking is a resource-intensive endeavor**, but one for which there is broad consensus and justification.

Banks that invest in these governance pillars not only comply with regulations but also build the foundation for **trustworthy AI** – enabling them to innovate with AI in a sustainable, responsible manner. The estimated €116k–€137k per model/year is an investment toward that end: ensuring that AI systems can deliver value to the bank and its customers **while maintaining safety, fairness, and accountability**. As one industry report framed it, moving from AI compliance to competitive advantage is possible – but it starts with recognizing and funding the true costs of governance. Banks should therefore budget confidently for these costs, treating them not as a burden but as **a vital commitment to ethical and effective AI deployment** in the financial sector.