DPAM
INDOSUEZ

October 2026

The AI disclosure gap: what investors still cannot see

An evidence-based assessment of what public markets reveal, aznd fail to reveal, about AI governance

Prepared for investors and stewardship teams
Remy Takang Arrey, CAIO
For DPAM

Disclaimer

Degroof Petercam Asset Management SA/NV (DPAM) I rue Guimard 18, 1040 Brussels, Belgium I RPM/RPR Brussels I TVA BE 0886 223 276 I

For professional investors only.

Investing incurs risks. Past performance does not predict future returns.

All rights remain with DPAM, who is the author of the present document. Unauthorized storage, use or distribution is prohibited. Although this document and its content were prepared with due care and are based on sources and/or third party data providers which DPAM deems reliable, they are provided without any warranty of any kind and without guarantee of correctness, completeness, reliability, timeliness, availability, merchantability, or fitness for a particular purpose. All opinions and estimates are a reflection of the situation at issuance and may change without notice. Changed market circumstance may invalidate statements in this document.

The provided information herein must be considered as having a general nature and does not, under any circumstances, intend to be tailored to your personal situation. Its content does not represent investment advice, nor does it constitute an offer, solicitation, recommendation or invitation to buy, sell, subscribe to or execute any other transaction with financial instruments. This document is not aimed to investors from a jurisdiction where such an offer, solicitation, recommendation or invitation would be illegal. Neither does this document constitute independent or objective investment research or financial analysis or other form of general recommendation on transaction in financial instruments as referred to under Article 2, 2°, 5 of the law of 25 October 2016 relating to the access to the provision of investment services and the status and supervision of portfolio management companies and investment advisors.

Contents

Investor Q&A 5

I. Executive Summary 8

II. Introduction: Why This Report Exists 11
2.1 A governance gap that is now well documented 12
2.2 Why companies and the widerinvestment community cannot treat this as background noise 12
2.3 What this report is not 13

III. The Market Problem: Disclosure Is Structurally Unbalanced 14
3.1 Where disclosure is strongest 15
3.2 Where disclosure is weakest 16
3.3 What this reveals 16

IV. Research Objective and Key Questions 17

V. Methodology and Company Universe 19
5.1 Scope 20
5.2 Assessment basis 20
5.3 Scoring method 20
5.4 Structure assessed 21
5.5 Use of AI in the assessment process 21
5.6 A note on scope and limits 21

VI. The Framework: What It Adds to the Market 22
6.1 Why existing disclosure approaches are insufficient 23
6.2 Why the framework is useful for investors 23
6.3 Why is the framework useful for companies 23

VII. Findings: What the Market Currently Discloses 24
7.1 What companies disclose relatively well 26
7.2 What companies disclose weakly 27
7.3 What companies disclose unevenly 28
7.4 What this reveals about the market 29

VIII. The Disclosure Map 30

IX. Why the Weak-Disclosure Areas Matter for a Sustainable Investor's Fiduciary Duty 33
9.1 Bias-testing outcomes 34
9.2 Remedy and contestability 34
9.3 Audit and assurance 34
9.4 Data governance, lawful basis, and provenance 34
9.5 Impact, social value, and inclusion 34

X. Implications for Investors, and Stewardship Teams 35
10.1 Implications for investors 36
10.2 Implications for stewardship teams 36

XI. Toward a Minimum Investor-Useful Standard for AI Governance Disclosure 37

XII. Conclusion 39

XIII. End Note 41

XIV. Appendix A: The Scoring Model 43

XV. Appendix B: Case Study Snapshots 46

XVI. Appendix C: Company Assessment Summar 52

XVII. Appendix D: Indicator Definitions 55
Section 1: AI Tech Stack 56
Section 2: High Risk AI Application Mapping 56
Section 3: Responsible AI Principles 57



Investor Q&A

Question	Short answer
What is this report about?	It assesses what public markets reveal, and fail to reveal, about corporate AI governance. The core finding is that disclosure is stronger on AI capabilities and intent than on AI accountability in practice.
Why does this matter?	AI governance is now a material investor issue. Litigation, regulatory exposure, and operational failures are already visible, but disclosure remains partial and hard to verify.
What does the report cover?	The report assesses 26 public companies across four AI roles: developers, deployers, platform firms, and infrastructure
How was it assessed?	Companies were scored on a 0-5 evidence scale across three sections: AI Tech Stack, High-Risk AI Application Mapping, and Responsible AI Principles.
How much evidence was reviewed?	The framework scored 793 indicators using three rules: evidence over assertion, outcome over process, and absence is not a pass.
What is the main conclusion?	The market currently discloses more about whether AI is being built at scale than whether it is being well governed.
What are the headline results?	Of 793 indicators, 650 were Red (82.0%), 142 Amber (17.9%), and 1 Green (0.1%). Every company scored Red overall.

Question	Short answer
Where is disclosure strongest?	In familiar reporting areas such as energy use, water, renewable power, data-centre efficiency and board oversight.
Where is disclosure weakest?	In areas that test accountability: audit trails, AI incident reporting, explainability, remedy, contestability and bias-testing outcomes.
Why are these weak areas important?	They are the areas most linked to litigation, compliance, reputational damage, operational risk and long-term stewardship concerns.
Does the report say companies ignore AI governance?	No. It says companies discuss AI governance extensively but provide too little decision-useful evidence to show whether controls work in practice.
What should investors do differently?	The report recommends four shifts: treat policy as a minimum standard, ask for evidence, use a standardised evidence bar and track a minimum disclosure baseline.
What would better disclosure include?	The report proposes a 10-point minimum baseline, including board accountability, risk frameworks, bias-testing results, override metrics, incident reporting, remedy, data provenance, third-party controls, human rights impact assessment, and use-case impact disclosure.

I. Executive Summary

AI governance has moved from a reputational talking point to a material investment issue. Litigation, regulatory exposure, and operational failures linked to AI are already visible in the market. At the same time, investors are being asked to assess these risks using disclosure that remains partial, selective, and difficult to verify. This report puts a number on that gap.



US courts have already ordered nine-figure damages in AI-related cases, while AI discrimination settlements have ranged from single digits to hundreds of millions of dollars . The EU AI Act adds fines of up to EUR 35 million, or 7% of global revenue, for the most serious breaches. Elsewhere, a patchwork of national approaches mean multinational companies face compounding compliance costs rather than a single clear standard to meet.

This report applies a structured, evidence-based AI ESG framework to 26 public companies spanning AI developers, deployers, platform firms, and infrastructure providers, selected for sector and geographic diversity. The assessment draws on public information only. It does not infer compliance from adjacent disclosure, and absence of evidence is scored as absence, not a pass.

Each company was scored on a 0 to 5 evidence scale across the following three part framework:

1. AI Tech Stack
2. High-Risk AI Application Mapping
3. Responsible AI Principles, spanning governance, rights, data, fairness, and impact

Across the full universe, 793 individual indicators were scored. Each indicator was tested against two rules: **evidence over assertion** and **outcome over process**. A policy statement alone cannot score above 1.

Disclosure is strongest where companies already report against established environmental, infrastructure, and governance conventions. These include energy consumption, water use, renewable power sourcing, data-centre efficiency, and broad governance architecture. In the sample assessed here, the AI Tech Stack pillar recorded the highest mean evidence score of any framework component at 1.15 out of 5, and the single strongest individual indicator was AI Strategy & Board Oversight at 1.88.

But naming a board touchpoint is not the same as demonstrating a functioning control. Only 2 of the 26 companies, or 8% of the sample, could evidence board oversight, a documented risk-management framework, and independent audit activity together at even a moderate level.

The weakest disclosure sits precisely where accountability would be tested. Of the 793 indicators scored, 82.0% were rated Red, and only one indicator across the entire sample scored Green. The weakest individual dimensions were Algorithmic Accountability & Audit Trails at 0.38, AI Incident Reporting & Response at 0.54, and Explainability & Interpretability at 0.62. Not one company disclosed AI incident reporting at a moderate level or better.

This is not a gap in AI discussion. Every company in the sample discusses AI extensively. It is a gap in decision-useful evidence. Meta illustrates the pattern directly: it presents a comparatively well-developed Responsible AI communications programme, yet a Senate inquiry and internal document leak in August 2025 revealed that AI Studio chatbots had been permitted to enter into romantic conversations with users identified as minors, a failure that public disclosure alone did not make clear.

Current disclosure practices therefore make it difficult to identify AI-related risks properly for investment purposes. This report supports four immediate shifts for investors and stewardship teams:

- treat policy language as a minimum standard, not a signal;
- ask for evidence, not principles;
- use a standardised evidence bar across the portfolio; and
- track the ten-point disclosure baseline set out in Section 11.

The central conclusion is straightforward:

The market currently discloses more about whether AI is being built at scale than whether it is being governed well.



II. Introduction: why this report exists

Artificial intelligence is a genuine source of economic opportunity. Faster innovation, better productivity, and more efficient use of capital are plausible and evidenced outcomes of AI adoption. Markets have already priced in a considerable share of that promise. What has not kept pace is the governance infrastructure meant to ensure those gains are captured responsibly, without creating downstream liability that erodes the value being built.

This report exists because the gap between what companies say about AI governance and what they can demonstrate is now wide enough to matter for stewardship and company risk profiles.

2.1 A governance gap that is now well documented

The distance between public commitment and demonstrated practice is not a matter of interpretation. It has been measured. Research cited in the report shows that companies are adopting AI faster than their governance frameworks are adapting. Nearly 90% of companies studied had not publicly committed to a named AI governance framework, only 13% disclosed a policy ensuring human oversight of AI systems, and just 2.3% maintained a dedicated complaints mechanism for AI-related harm.

The World Benchmarking Alliance reached a similar conclusion. Of 200 major technology companies assessed, 38% had published AI principles, but only 24 explained their internal AI governance mechanisms in operational detail, and none demonstrated evidence of a comprehensive human rights impact assessment across the AI systems they developed, procured, or deployed.

The problem is compounded by the pace and opacity of AI itself. The 2026 International AI Safety report cites evidence that pre-deployment testing does not reliably predict real-world AI risk, that internal model behaviour remains poorly understood even by the companies that build these systems, and that commercial incentives favour keeping material information proprietary.

Governance is therefore not merely lagging deployment on a fixed target. The target itself keeps moving.

2.2 Why companies and the wider investment community cannot treat this as background noise

AI governance is increasingly treated by institutional investors as a material, investible issue rather than an adjacent ethical topic. It cuts across human rights, climate and nature, governance quality, and operational resilience.

The regulatory backdrop reinforces this. While the current US administration may appear more cautious on AI regulation, other jurisdictions are moving in the opposite direction. Several AI firms have themselves called for clearer boundaries and more explicit guardrails. At the same time, the EU remains highly influential in shaping global standards. As was the case with sustainability regulation, principles established in Europe often diffuse internationally through law, market practice, customer requirements, and investor expectations.

Stewardship activity is also growing. Investors report that engagement and risk assessment are now their most common responses to AI exposure, ahead of formal investment policy . Major asset managers have begun targeted stewardship programmes aimed at closing disclosure gaps ranging from data-centre sustainability to model-level accountability. But these initiatives share a common constraint: investors are asking better questions than the market's current disclosure can answer.

2.3 What this report is not

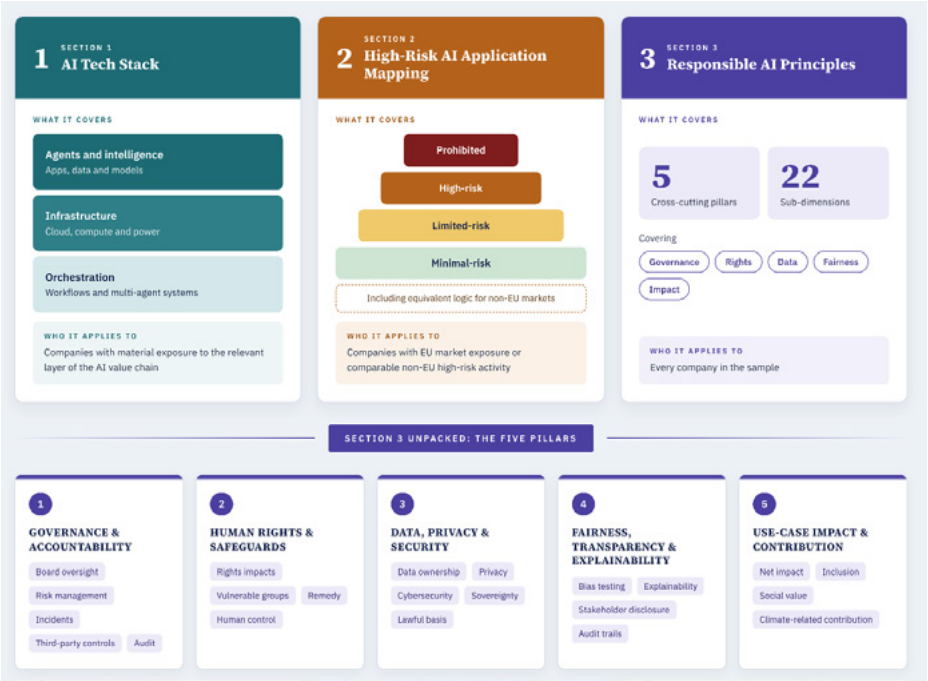
This report does not argue that companies are failing to engage with AI governance or that the topic is being ignored. The evidence gathered across 26 large-cap companies and 793 scored indicators shows the opposite. AI governance is discussed extensively in sustainability reports, responsible AI statements, and investor communications, though this pattern should be interpreted with care given the disclosure bias associated with large-cap companies.

The argument is narrower and more consequential. Companies disclose policies and principles considerably more readily than they disclose evidence that those policies produce outcomes. That evidence would include bias-testing results, override rates, audit findings, or cases where a governance process changed a business decision.

This is familiar from other emerging ESG issues. Climate reporting evolved from greenhouse gas inventories to targets, then to science-based targets, and increasingly to transition plans demonstrating how those targets will be delivered in practice. AI disclosure appears to be at an earlier stage of that maturity curve, with the added complication that the underlying technology is advancing very quickly.

That is not a gap in AI discussion. It is a gap in decision-useful evidence, and it is the gap this report sets out to map, explain, and help close.

Figure 1. Three Sections and Five Responsible AI Pillars



Source: DPAM, reproduced from the report.

III. The market problem: disclosure is structurally unbalanced

Before turning to the company-level findings, it is useful to set out the pattern already visible in the wider public record. Across benchmark studies, sell-side research, and sector reporting reviewed for this report, the same conclusion appears: AI disclosure is structurally uneven.



3.1 Where disclosure is strongest

Disclosure is most developed where AI overlaps with conventions companies already use for corporate reporting. These include:

- energy consumption;
- water use;
- renewable power sourcing;
- data-centre efficiency;
- physical infrastructure scale; and
- broad governance architecture

The reason is straightforward. These areas are comparatively easier to quantify and govern. Companies often already have baseline data, operational KPIs, established methodologies, and reporting routines in place. More novel AI-specific issues, such as model-level societal impact, downstream misuse, or emergent ethical risk, are less standardised and harder to measure, so disclosure is correspondingly less mature.

A survey found that 91% of energy companies are already exploring or actively engaged in data centers and hyperscaler projects, 89% expect AI infrastructure growth to have at least a moderate impact on electricity demand over the next five to ten years, and 78% expect capital expenditure to rise over the next three years. This is a set of metrics companies already know how to count and disclose. AI's footprint within them is captured almost as a by-product of routine reporting.

High-level Responsible AI statements and safety-principle language follow a similar logic. A values statement requires no new measurement infrastructure to produce.

3.2 Where disclosure is weakest

The pattern reverses once the question moves from scale of development, deployment and use to accountability.

The indicators with the weakest scores were the ones that would allow an outsider to verify a governance claim rather than simply read it. These include:

- bias-testing outcomes broken down by protected group;
- AI-specific audit scope and findings;
- remedy and contestability metrics;
- human override and reversal rates;
- dataset lineage and ownership;
- lawful basis and consent architecture behind AI processing;
- AI-specific incident reporting; and
- use-case-level social impact.

The World Benchmarking Alliance found that no company in its 200-company review demonstrated evidence of a comprehensive human rights impact assessment across the AI systems it developed, procured, or deployed, and only 24 explained their AI governance mechanisms in operational detail. Evidence cited in the report points to the same pattern: only 13% of companies disclosed a policy ensuring human oversight of AI systems, and just 2.3% maintained a dedicated AI complaints mechanism.

3.3 What this reveals

Taken together, the picture is not one of AI governance being ignored. Rather, it is being discussed in ways that are comparatively cheap to disclose and difficult for an outside party to verify, while the evidence that would allow an external party to test whether a control works remains largely absent. A more recent study notes that 'AI governance effectiveness is, therefore, not limited to whether companies disclose responsible AI principles; it also depends on whether those principles are embedded in oversight structures, risk management processes, product governance, infrastructure planning, and stakeholder impact assessments .'

The line worth keeping in mind while considering the rest of this report:

The market currently discloses more about whether AI is being built at scale than whether it is being governed well.

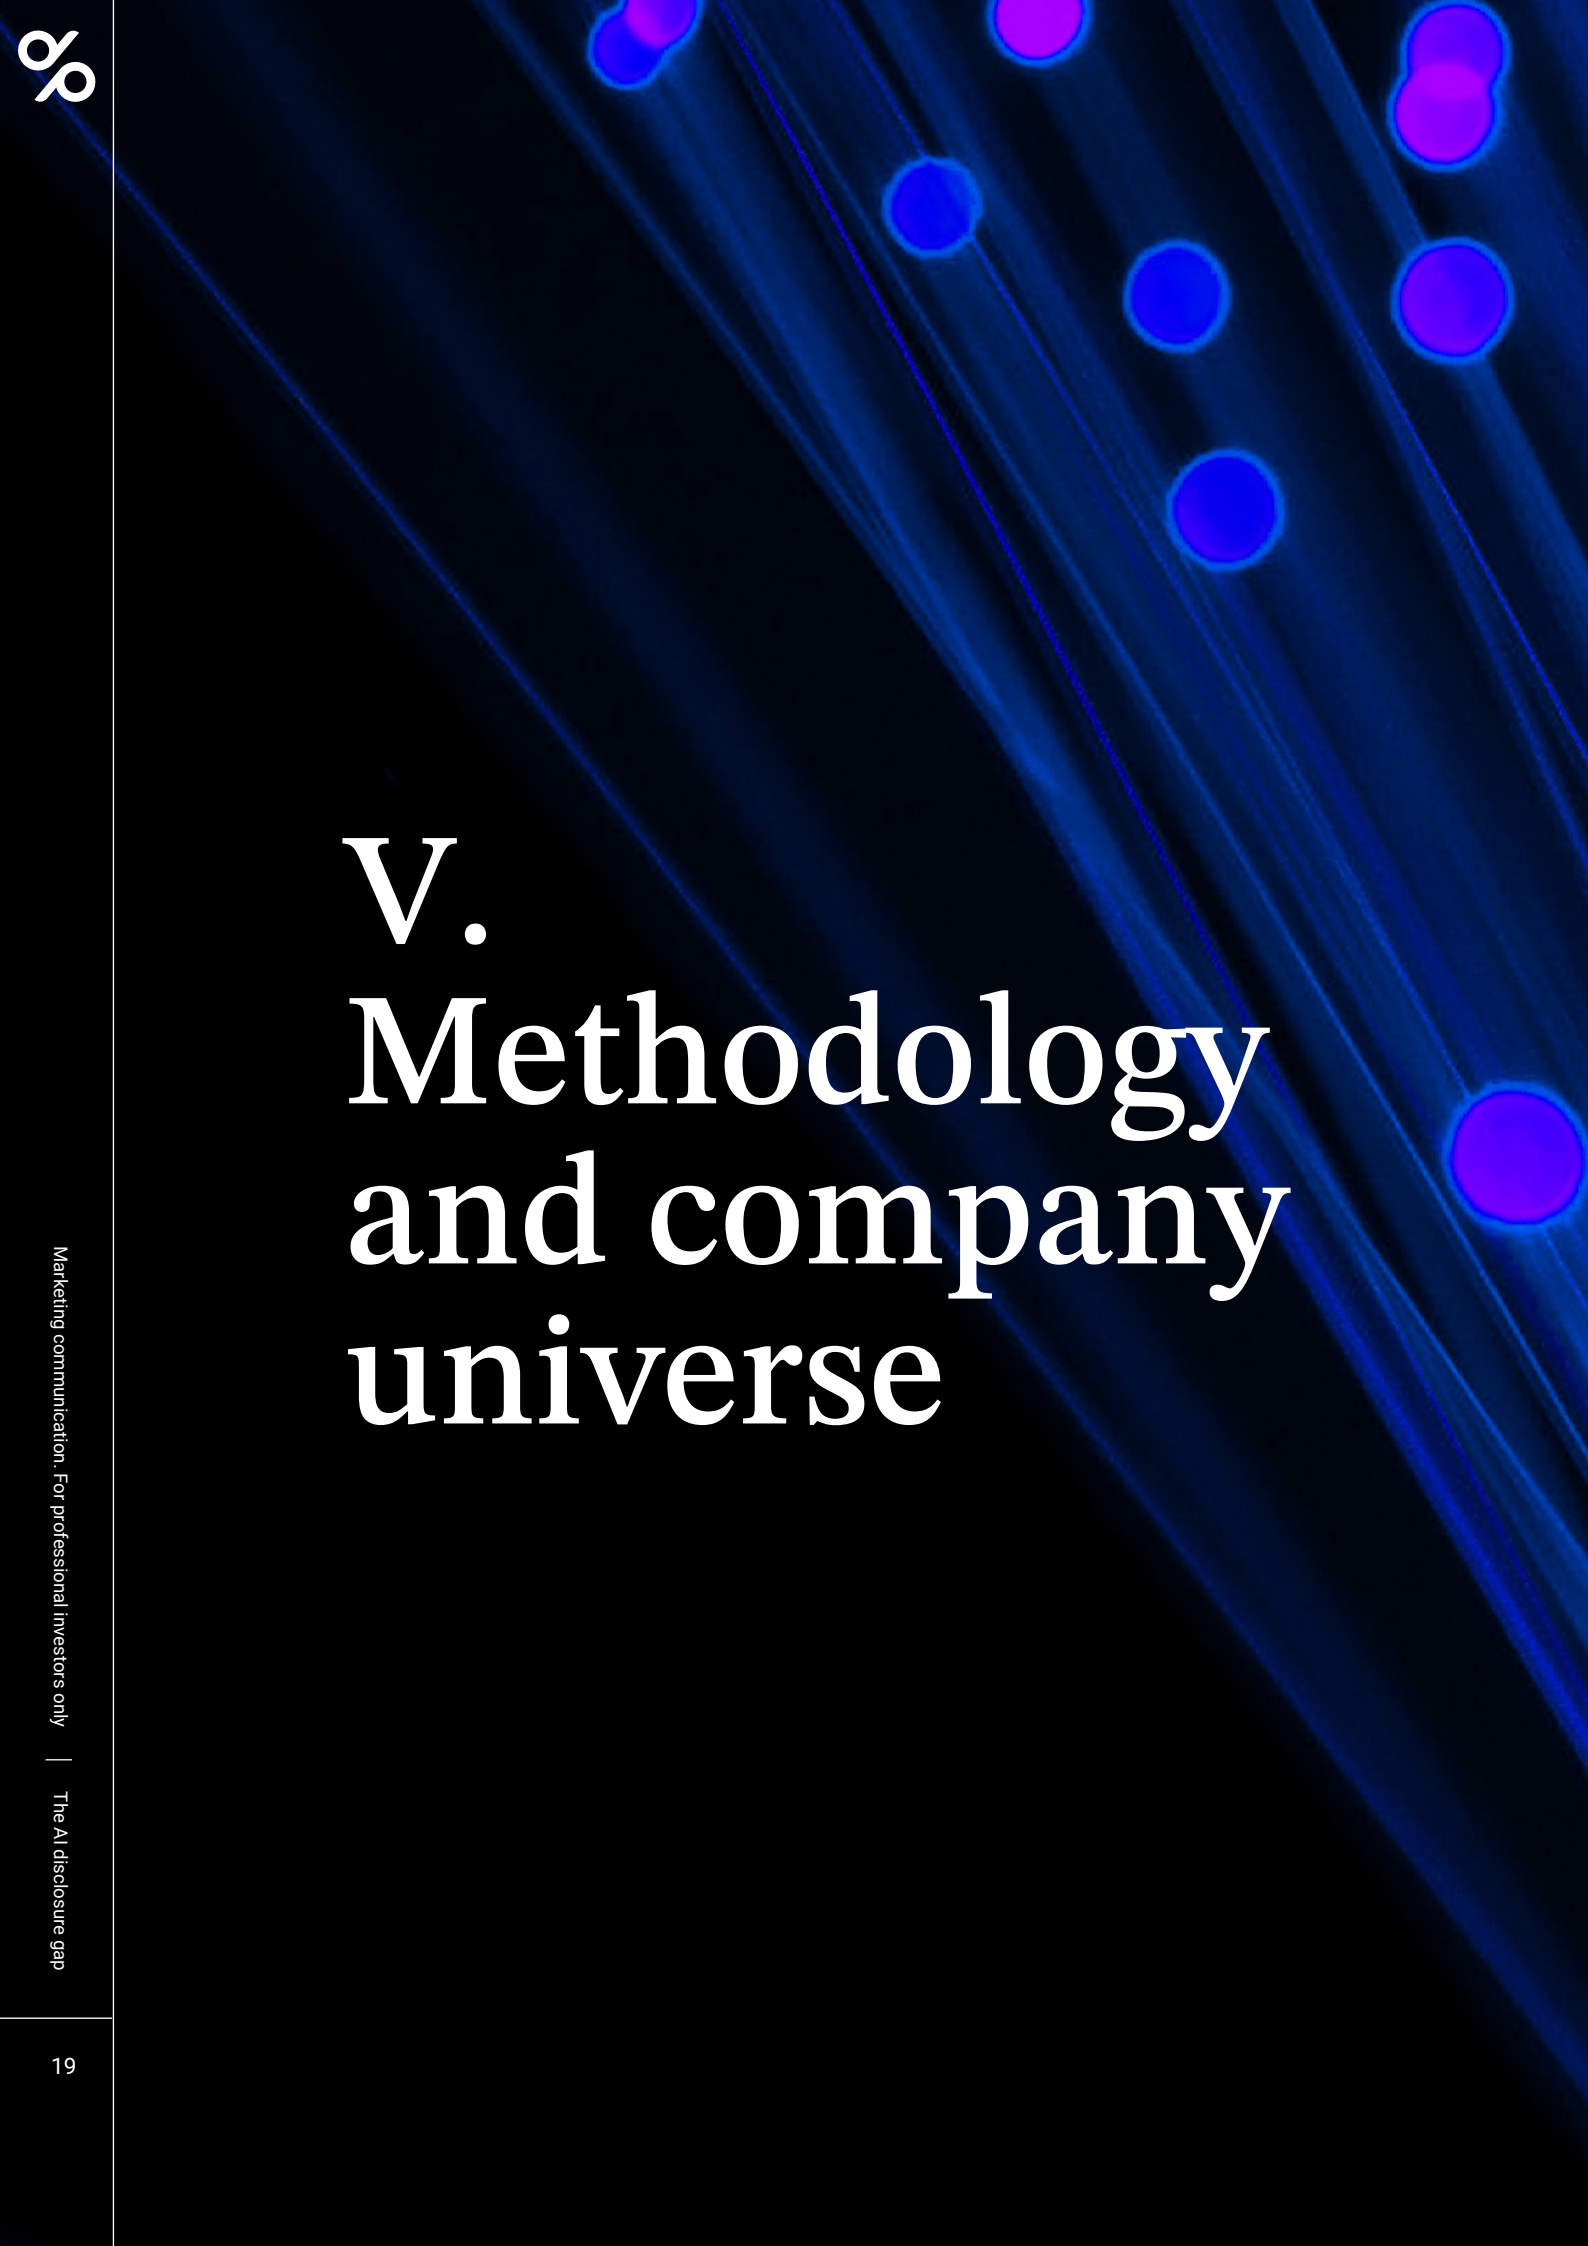
The report's own 26-company assessment tests this pattern directly and confirms it. Every company assessed scored Red overall, with mean scores ranging narrowly between 0.40 and 1.50 out of 5.

IV. Research objective and key questions

The objective of this report is to establish and evidence what public markets currently reveal and fail to reveal about corporate AI governance, and to convert the findings into a repeatable tool investors can use for comparison and investor stewardship.

The research is organised around six questions:

1. What does the market currently disclose on AI governance across developers, deployers, platform firms, and infrastructure providers?
2. Which dimensions of disclosure are consistently stronger, and which are consistently weaker, across the assessed universe?
3. Where are disclosure gaps most systematic, meaning present across most companies and sectors rather than specific to one issuer?
4. Which of those gaps are most decision-relevant for investors, in the sense of carrying identifiable financial or impact materiality?
5. What does current disclosure enable, and fail to enable, for investor stewardship?
6. How does a structured, evidence-based framework improve comparability and investor user-friendliness relative to existing high-level AI governance frameworks? Each of these questions is addressed directly in the findings and implications chapters that follow.



V. Methodology and company universe

5.1 Scope

This report assesses 26 companies against a structured AI ESG framework built specifically for investment analysis. The sample spans four functional roles in the AI value chain:

- AI developers, building foundation or fine-tuned models;
- AI deployers, operationalising AI within existing business processes;
- platform firms, distributing AI-enabled products and services at scale; and
- infrastructure providers, supplying cloud, compute and power for AI workloads.

The final universe spans technology and semiconductors, financial services and insurance, pharmaceuticals and healthcare, energy, consumer staples, and diversified holdings across North America and Europe.

The sample was designed to be representative rather than exhaustive. An earlier intention was to create an even split between MSCI World (15 companies) and MSCI Europe (15 companies) but it was adjusted in practice because the final assessment required sufficient public disclosure to score companies meaningfully.

Nine of the 26 companies assessed also appear in the World Bank's list of leading publicly listed AI-related companies: Alphabet, AMD, Apple, ASML, Cisco, Meta, Microsoft, NVIDIA, and Tesla.

5.2 Assessment basis

The assessment relies exclusively on public information, including annual reports, sustainability and responsible AI disclosures, regulatory filings, and other material a company itself has placed in the public domain.

It does not infer compliance or good practice from adjacent disclosure. It does not assume company practice when no evidence is provided. Absence of evidence against an indicator is scored as absence, not as a pass, and is flagged for follow-up rather than left unscored.

5.3 Scoring method

Each indicator is scored on a 0 to 5 evidence scale and converted into a Red, Amber or Green rating at both indicator and company level. The scoring system is governed by three principles:

- evidence over assertion;
- outcome over process; and
- absence is not a pass.

The full scoring methodology, evidence-scale definitions, and Red, Amber, Green (RAG) thresholds are set out in Appendix A.

5.4 Structure assessed

AI Tech Stack

Applied where a company has material exposure to a relevant layer of the AI value chain.

High-Risk AI Application Mapping

Applied where a company has relevant use of AI in high-risk contexts. The framework screens companies against EU AI Act prohibited, high-risk, limited-risk, and minimal-risk categories. Section 2A applies to companies with EU market exposure. Section 2B applies equivalent high-risk logic, benchmarked against comparable global standards, to companies whose AI activity falls outside direct EU market scope.

Responsible AI Principles

Applied to every company in the universe, regardless of sector, geography, or role in the AI ecosystem. This section is organised around five pillars:

- Governance & Accountability
- Human Rights & Fundamental Safeguards
- Data, Privacy & Security
- Fairness, Transparency & Explainability
- Use-Case Impact & Positive Contributions

Every indicator is also classified for both impact materiality and financial materiality, so findings can be read through either lens or both together.

5.5 Use of AI in the assessment process

Claude Sonnet 5.5 was used to support the assessment of the 26 companies using the framework built for this report. All analytical judgments, policy interpretations, and conclusions were developed, reviewed, and verified by the author and the RICC team.

5.6 A note on scope and limits

This report assesses disclosure quality and evidence availability. It does not, and cannot, assess private internal practice that has not been made public. A company that governs AI well but discloses poorly may score similarly to a company that governs AI poorly and discloses nothing. That is not a flaw in the framework. It is the disclosure gap the report is concerned with.

VI. The framework: what it adds to the market

The preceding chapters describe a problem: public disclosure on AI governance is extensive in volume but structurally uneven in substance. It is stronger on scale and stated intent, and weaker on evidence that would allow an outsider to verify a claim.

This chapter explains why the framework applied in this report was built, what it does differently from existing approaches, and why it is useful to both investors and the companies they hold.

6.1 Why existing disclosure approaches are insufficient

Investors consistently report that they need more detailed, company-specific guidance on AI risk management, sector relevance, and comparative analysis than existing high-level frameworks provide. The report cites evidence that investors often read frameworks such as those from the National Institute of Standards and Technology (NIST), the International Organization for Standardization (ISO), the OECD, and UNESCO and then build their own internal versions, which is inefficient and duplicative .

Global AI governance frameworks and voluntary principles are useful for setting direction and establishing a common vocabulary. They were not built to answer the question an investor needs answered: does this specific company have evidence behind its AI governance claims, and how does that evidence compare with its closest peers?

Similarly, sustainability and ESG reporting conventions were developed before AI became a distinct governance category. As a result, AI is often folded into broader technology, data, or innovation disclosure, where AI-specific accountability questions such as bias testing, override rates, and audit scope are diluted or lost altogether.

6.2 Why the framework is useful for investors

For an investor, the framework converts an abstract governance concern into five usable outputs:

1. risk flags identifying where a company's public profile may overstate governance maturity;
2. evidence thresholds distinguishing a policy statement from a verified control;
3. explicit disclosure expectations that can be tracked over time;
4. comparable signals across companies and sectors because every indicator is scored against the same evidence bar; and
5. specific stewardship questions that can be put to management rather than accepting their general reassurance.

The 26-company assessment demonstrates that value directly. Because each company was scored against the same 22 sub-dimensions, the framework makes it possible to say not just that disclosure is weak, but exactly where it is weakest and to what extent.

6.3 Why is the framework useful for companies

For companies, the same framework functions as a specification of what good disclosure and good governance look like in practice rather than in principle. A company that wants to close its own disclosure gap can see in the scoring criteria the kinds of evidence that would move an indicator from Red to Amber or Green: named outcomes, disclosed metrics, independent verification, and traceable governance actions.

This is not a narrow compliance checklist. It is a description of decision-useful evidence that also serves as a roadmap for any company serious about demonstrating, rather than merely stating, that its AI governance works.



VII. Findings: what the market currently discloses

This is the centerpiece of the report. It draws on the full 26-company, 793-indicator assessment and organises the results into patterns rather than company-by-company profiles. Illustrative snapshots appear in Appendix B and the full company summary in Appendix C.

A company like AMD illustrates the exception that proves the rule. Its Services & Devices indicator scored 4.0 out of 5, the only Green score in the entire assessment, supported by supplier audits and responsible-minerals programme coverage. At the other end, Tesla and Eli Lilly both scored zero on AI Risk & Management Framework.

Table 1. Pooled RAG distributionacross all 793 scored indicators

Rating	Count	Share of all indicators
Red	650	82.0%
Amber	142	17.9%
Green	1	0.1%

Source: DPAM based on the 26-company assessment

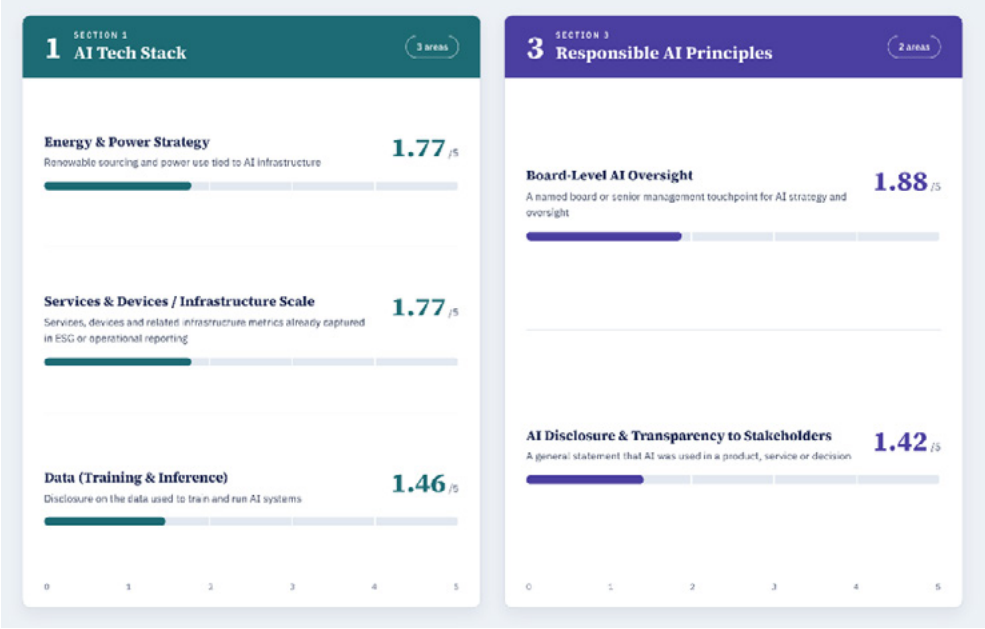
7.1 What companies disclose relatively well

Disclosure is strongest, though still far from mature in absolute terms, in the five areas summarised in Figure 2. All five sit comfortably within reporting conventions companies already use for other purposes, particularly energy, infrastructure, and board governance.

Cisco reports 100% renewable electricity matched to approximately 1.4 million MWh in financial year (FY)25, with Scope 1 and 2 emissions down 90% on FY19. Mastercard and ASML lead on board oversight, each scoring 3.0. Mastercard names a Chief AI and Data Officer, while ASML's Supervisory Board identifies AI strategy as a stated focus. SAP, Novartis, and Microsoft lead on risk frameworks, with SAP classifying every AI use case against seven prohibited red lines.

Several assessed companies, including Apple and Broadcom, also disclose detailed third party-verified environmental metrics, recycled material percentages, emissions reductions, and security-related metrics that fit comfortably into established ESG reporting routines. Across the sample, 21 of the 26 companies were able to name at least some board-level touchpoint for AI oversight, more than for any other indicator in the assessment.

Figure 2. Best-disclosed areas by section



Source: DPAM based on the 26-company assessment

7.2 What companies disclose weakly

This is the strongest and most consequential finding in the report. The indicators with the weakest scores were consistently the ones that would allow an external party to test whether a governance claim is true rather than merely stated.

The weakest dimensions include:

Algorithmic Accountability & Audit Trails

SDG Alignment & Social Value Creation

AI Incident Reporting & Response

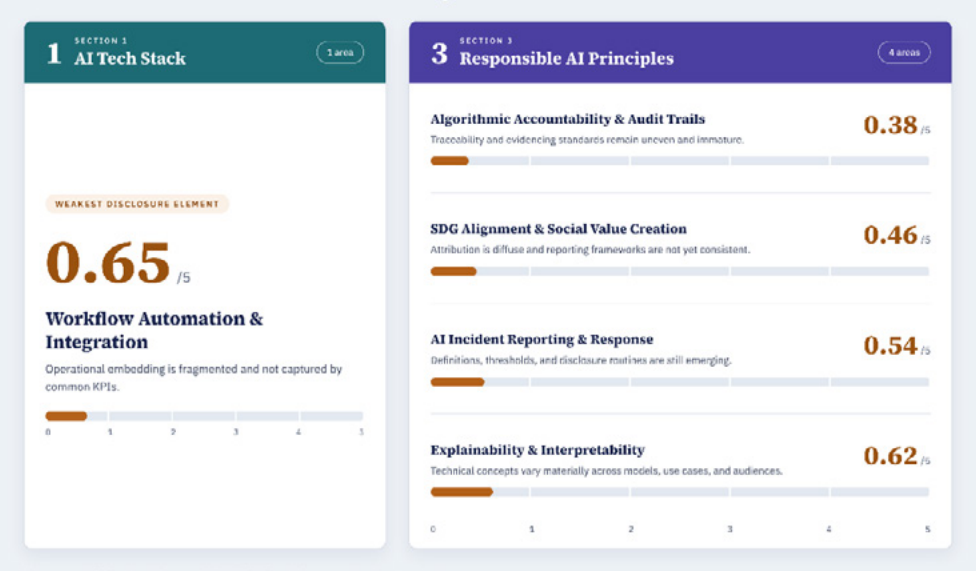
Workflow Automation & Integration

Explainability & Interpretability

Apple is the only company scoring above 2.0 on either audit trails or digital inclusion. Its Private Cloud Compute publishes a tamper-evident transparency log, and its accessibility features also score relatively well. Elsewhere, the picture is thin. Eighteen of the 26 companies scored zero on audit trails, 15 on digital inclusion, and 10 on explainability, including Tesla, UBS, and Shell.

The most striking example is AI incident reporting. Not one of the 26 companies disclosed AI-specific incident reporting at a moderate level or better, and only 3 of the 26 evidenced bias-testing results at that level.

Figure 3. Weakest disclosure areas by section



Source: DPAM based on the 26-company assessment

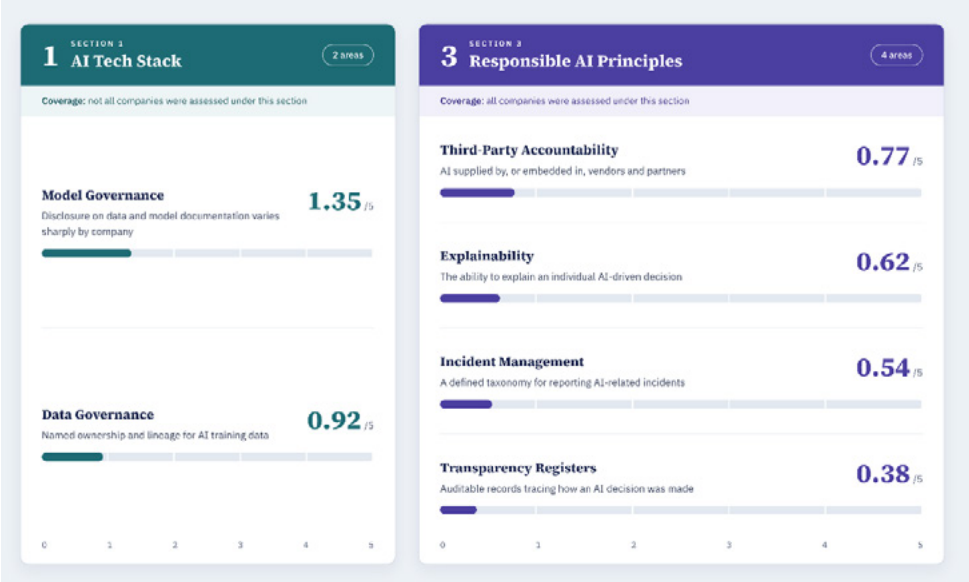
7.3 What companies disclose unevenly

A second tier of disclosure is present but inconsistent. Some companies score reasonably, while others score close to zero on the same indicator. That unevenness is informative, because it shows that detailed, verifiable disclosure is possible when companies choose to provide it.

For example, JPMorgan Chase scored 2.0 on third-party accountability, requiring AI suppliers to provide training data and fairness documentation. By contrast, Tesla, Shell, Berkshire Hathaway, and Allianz scored zero on that same dimension. On governance, Novartis, Eli Lilly, and Cisco reached 2.0, while NVIDIA, Microsoft, Meta, and Broadcom scored zero on some governance-related elements. Tesla's 1.0 on incident management rests on a public federal investigation record rather than its own disclosure, and 12 of the 26 companies scored zero on that dimension.

The constraint is therefore not the technical ability to measure, but rather a reluctance to disclose the specific metrics that would allow independent verification.

Figure 4. Where disclosure is uneven



Source: DPAM based on the 26-company assessment

7.4 What this reveals about the market

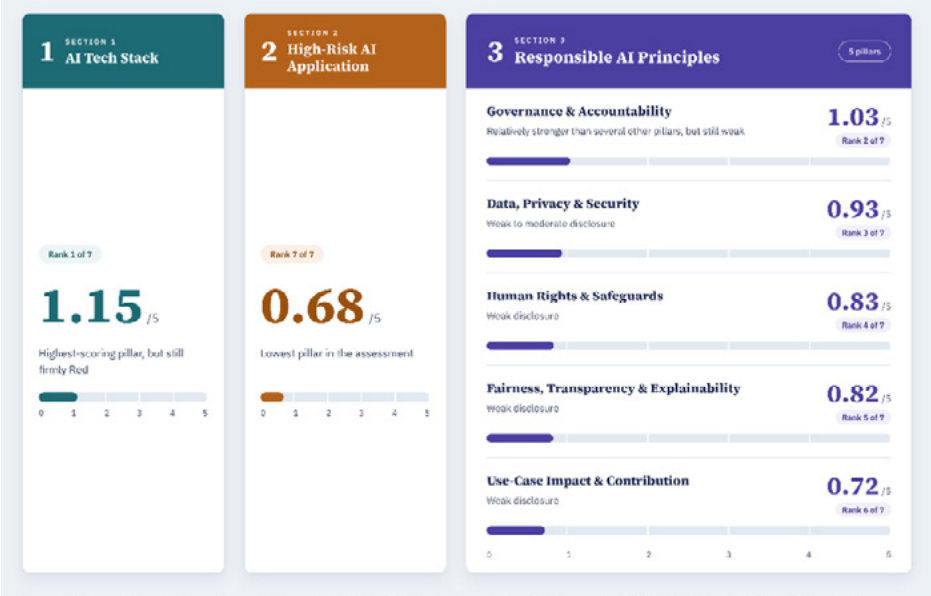
Taken together, the three tiers above point to a clear pattern in how the market discloses AI governance. Figure 5 shows the average score for each framework pillar across the companies assessed where applicable.

Three conclusions follow:

- disclosure is strongest where firms can rely on established ESG or infrastructure reporting traditions;
- disclosure is weakest precisely where it would expose operational accountability, such as bias results, audit findings, incident counts, and override rates; and
- the market still reports governance intent considerably better than governance outcomes.

The mean score of every pillar remains below the 2.00 Red/Amber threshold. But AI Tech Stack and Governance & Accountability, both associated more closely with stated intent and inherited reporting structures, score meaningfully higher than Use-Case Impact & Contribution and High-Risk AI Application Mapping, which are more closely associated with demonstrated outcomes .

Figure 5. Average disclosure by framework pillar



Source: DPAM based on the 26-company assessment

VIII. The disclosure map

The findings above can be summarised as a simple disclosure map: which dimensions of AI governance are relatively better disclosed, which are partially visible, and which remain systematically weak, set against how much investors need that evidence to assess risk.

Table 2. Disclosure gap map: typical disclosure strength versus investor usability, by dimension

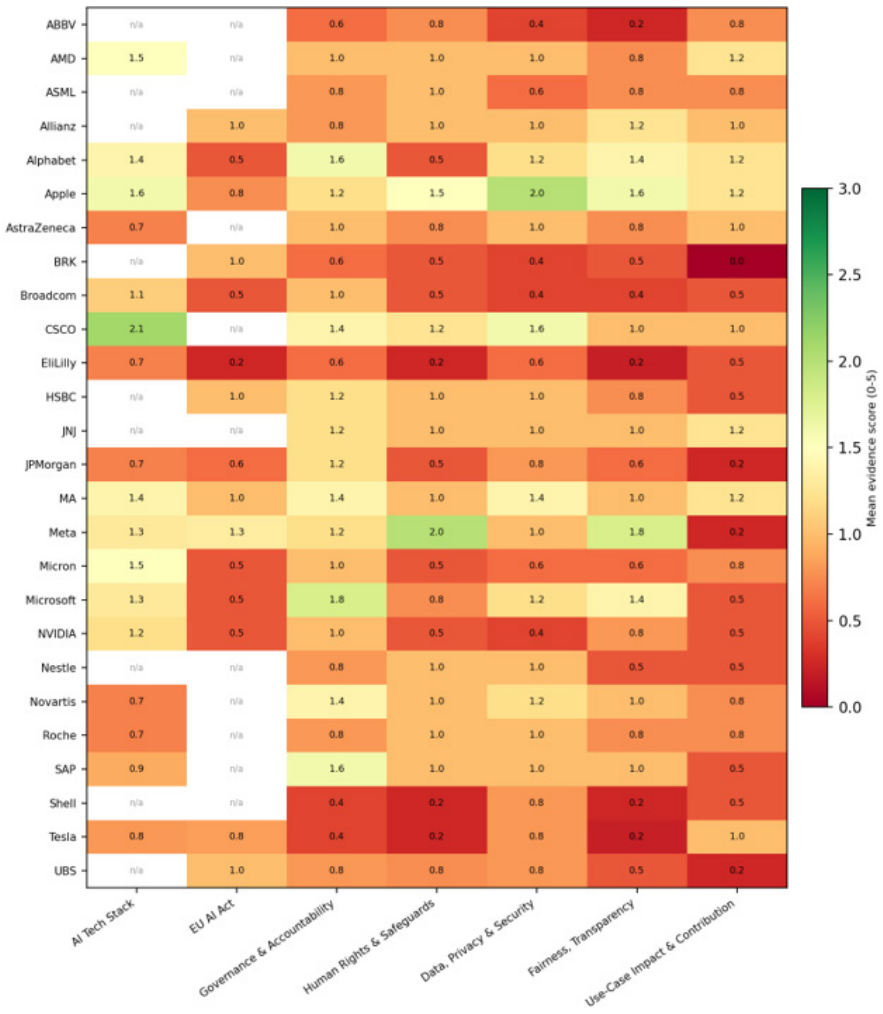
Dimension	Typical disclosure strength	Mean score (/5)	Investor usability
Energy, power & climate footprint	Stronger	1.5–1.9	Moderate to high
Board-level governance structures	Moderate	1.3–1.9	Moderate
Safety principles & architecture language	Moderate	1.0–1.4	Low to moderate
Data governance & cybersecurity	Weak to moderate	0.9–1.3	Moderate to high
Bias testing outcomes	Weak	0.4–0.9	High need, low visibility
Remedy & contestability	Weak	0.8–1.0	Very high need, low visibility
Audit & assurance	Weak	0.4–0.7	Very high need, low visibility
Data lineage & lawful basis	Weak	0.7–0.9	High need, low visibility
Societal impact & vulnerable groups	Weak	0.5–0.8	High need, low visibility
AI incident reporting	Systematically weak	0.54	Very high need, near-zero visibility

Source: DPAM based on the 26-company assessment

This synthesis is designed to travel. An investor reviewing a single company's disclosure against this map can immediately see whether that company's own pattern, is strong on energy and governance architecture, silent on audit and remedy, matches or diverges from the market-wide norm.

At the time of writing, no reactions or comments were sought from the companies assessed in this report.

Figure 6. Disclosure score heatmap by company and framework pillar



Source: DPAM based on the 26-company assessment

IX.
Why the weak-disclosure areas matter for a sustainable investor's fiduciary duty

AI's societal impact is moving to the forefront as concerns grow around safety, accountability, bias, and real-world harm. This is accelerating regulatory attention due to political and public pressure. Yet legal responsibility between AI developers and deployers remains only partially defined, as a result, governance strength and risk-management capability are likely to be important determinants of which companies can monetise AI safely and sustainably over the coming decade.

Tesla provides a practical illustration. The company has faced complaints, regulatory investigations, and legal penalties linked to Autopilot and Full Self-Driving systems, while still providing very limited public evidence on fairness, transparency, or explainability in those use cases.

The weak disclosure areas identified in this report are not equally consequential, but each map onto a distinct risk channel.

9.1 Bias-testing outcomes

Weak bias disclosure creates exposure to litigation, regulatory enforcement, and reputational damage. When testing is not disclosed, investors cannot distinguish between a company that tested for bias and found no material issue and a company that never tested at all. Only 3 of the 26 companies in the sample evidenced bias-testing outcomes at a moderate level or better.

9.2 Remedy and contestability

Weak remedy and contestability disclosure signals a governance weakness. Where no contestation mechanism has been demonstrated, unresolved harm and compliance risk tend to surface later and at greater cost. The mean score for Rights to Remedy & Contestability across the sample was 0.81.The 26-company assessment demonstrates that value directly. Because each company was scored against the same 22 sub-dimensions, the framework makes it possible to say not just that disclosure is weak, but exactly where it is weakest and to what extent.

9.3 Audit and assurance

Weak audit and assurance disclosure removes one of the main levers and stewardship teams must verify whether a stated control works. Only 5 of the 26 companies evidenced any independent AI audit or assurance activity at a moderate level or better, and the mean score for this area was 0.65.

9.4 Data governance, lawful basis, and provenance

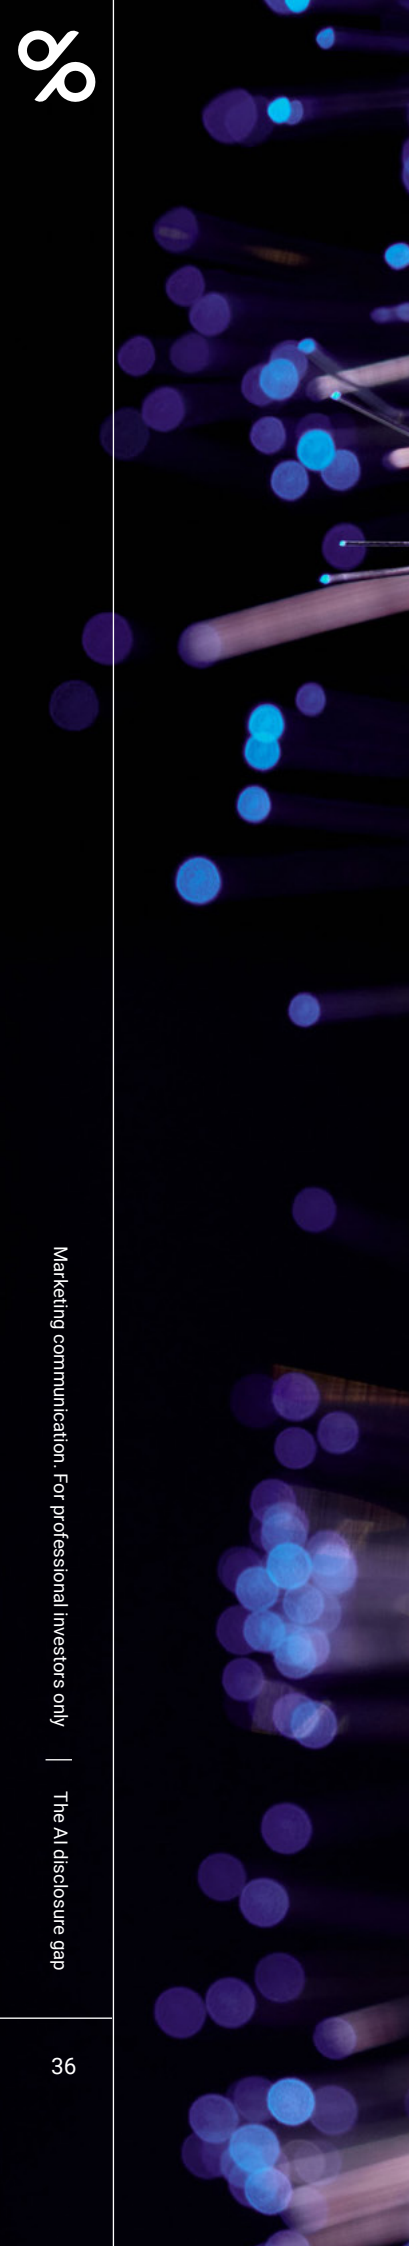
Weak data governance disclosure carries direct General Data Protection Regulation (GDPR), EU AI Act, and cross-border data risk. It is also increasingly relevant to M&A and operational due diligence. Consent Architecture & Lawful Basis for AI Processing scored 0.67 across the sample.

9.5 Impact, social value, and inclusion

Weak impact disclosure makes it difficult to distinguish AI deployment that produces genuine, verifiable social or environmental benefit from AI deployment that simply externalises cost and risk onto users, workers, or society. SDG Alignment & Social Value Creation scored 0.48, while Digital Inclusion & Access to AI scored 0.58.

X. Implications for investors, and stewardship teams





10.1 Implications for investors

Current market disclosure may:

- overstate governance maturity, since board oversight and principle-level language are common while evidenced controls are rare;
- understate operational AI risk, since the indicators most likely to reveal risk are also the least disclosed;
- impair comparability across issuers, because companies disclose selectively and inconsistently;
- make it difficult to price governance quality separately from AI capability or scale; and
- bias analysis toward visible infrastructure metrics over less visible but potentially more consequential accountability failures

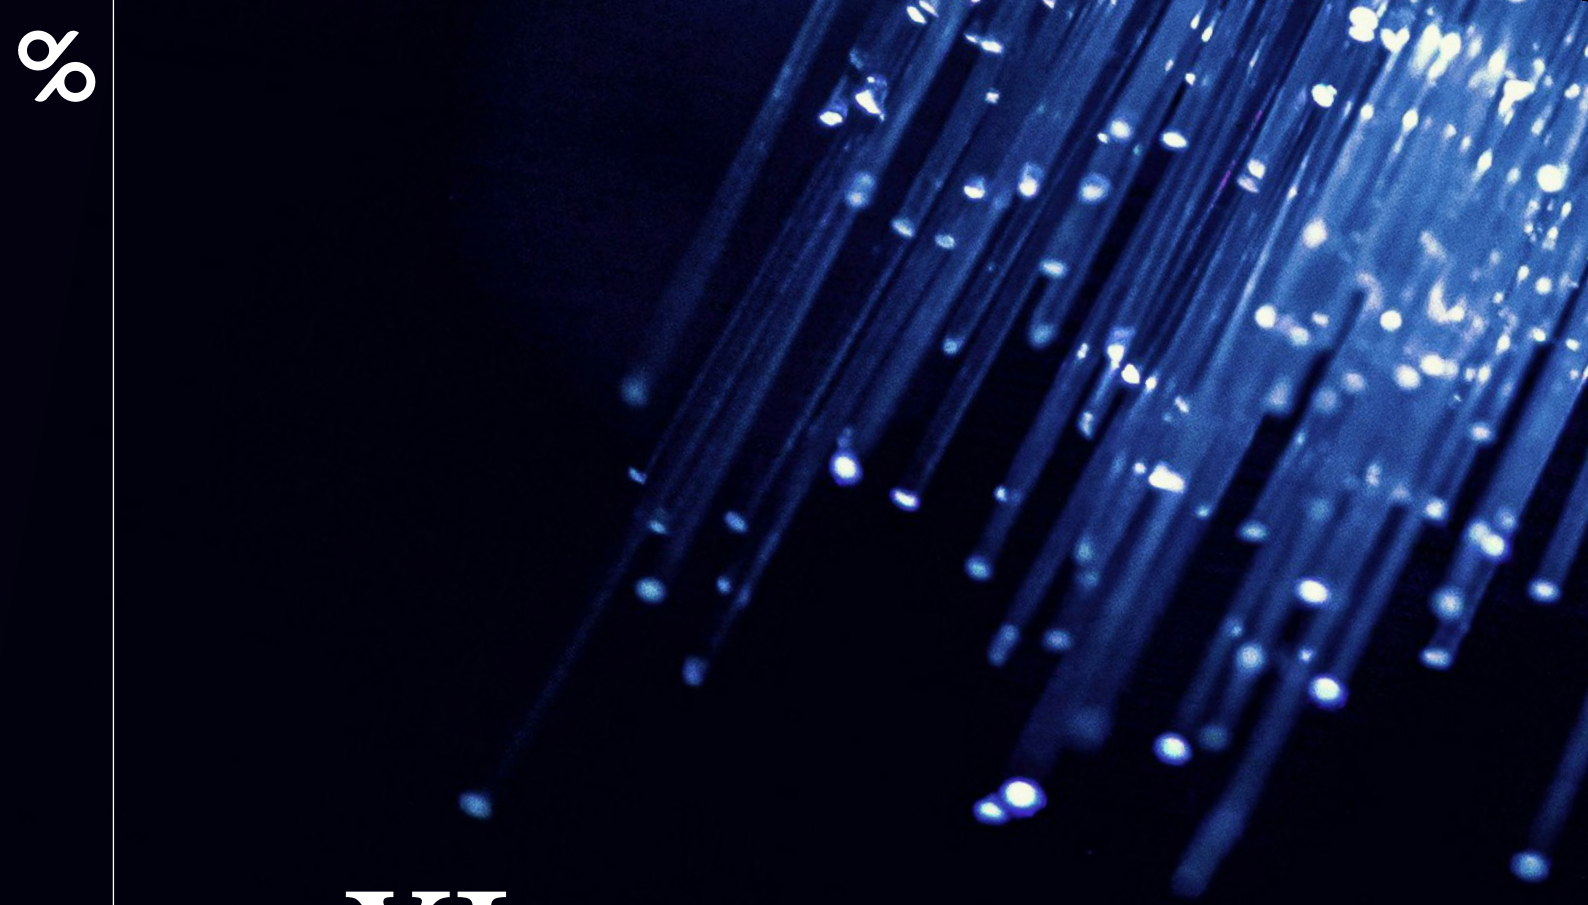
10.2 Implications for stewardship teams

Stewardship teams are already leading engagement on AI governance for many institutional investors, but often lack practical tools calibrated to the evidence a company can reasonably be expected to produce.

The findings support three conclusions:

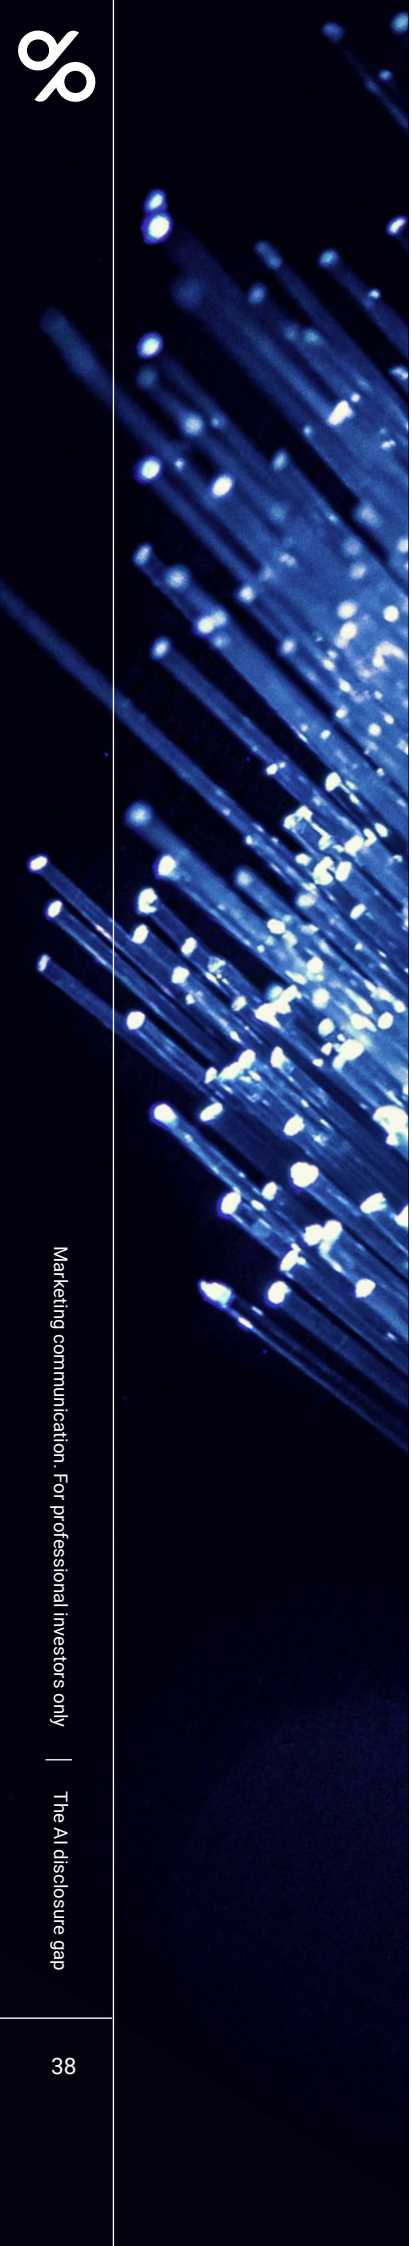
- current disclosure is insufficient to support high-quality engagement on its own;
- engagement needs to move beyond principles to specific, evidenced asks; and
- nvestors collectively need standardised asks across companies and sectors

The framework set out in this report is designed to translate broad concern about AI governance into the specific, comparable stewardship questions this requires.



XI.
Toward a minimum investor-useful standard for AI governance disclosure

The preceding chapters have described what is missing. This chapter is deliberately constructive. Rather than only critiquing current disclosure, it proposes a concise minimum baseline for investor-useful AI disclosure.



Minimum investor-useful AI disclosure should include:

01 **Board oversight and named accountability**
A named individual or committee with publicly defined AI governance responsibility.

02 **An AI risk-management framework and clear risk ownership**
Covering identification, assessment, and mitigation of AI-specific risk.

03 **Bias-testing outcomes, not just fairness principles**
Disclosed results, by protected characteristic where relevant, with a stated remediation threshold and process.

04 **Human oversight and override metrics**
Disclosed rates at which human review overrides or reverses an AI-driven decision.

05 **Incident reporting and remediation**
A defined AI incident taxonomy, disclosed incident counts, and named remediation actions.

06 **Contestation and remedy performance**
A functioning mechanism by which an affected party can contest an AI-driven decision, with disclosed usage and outcome data.

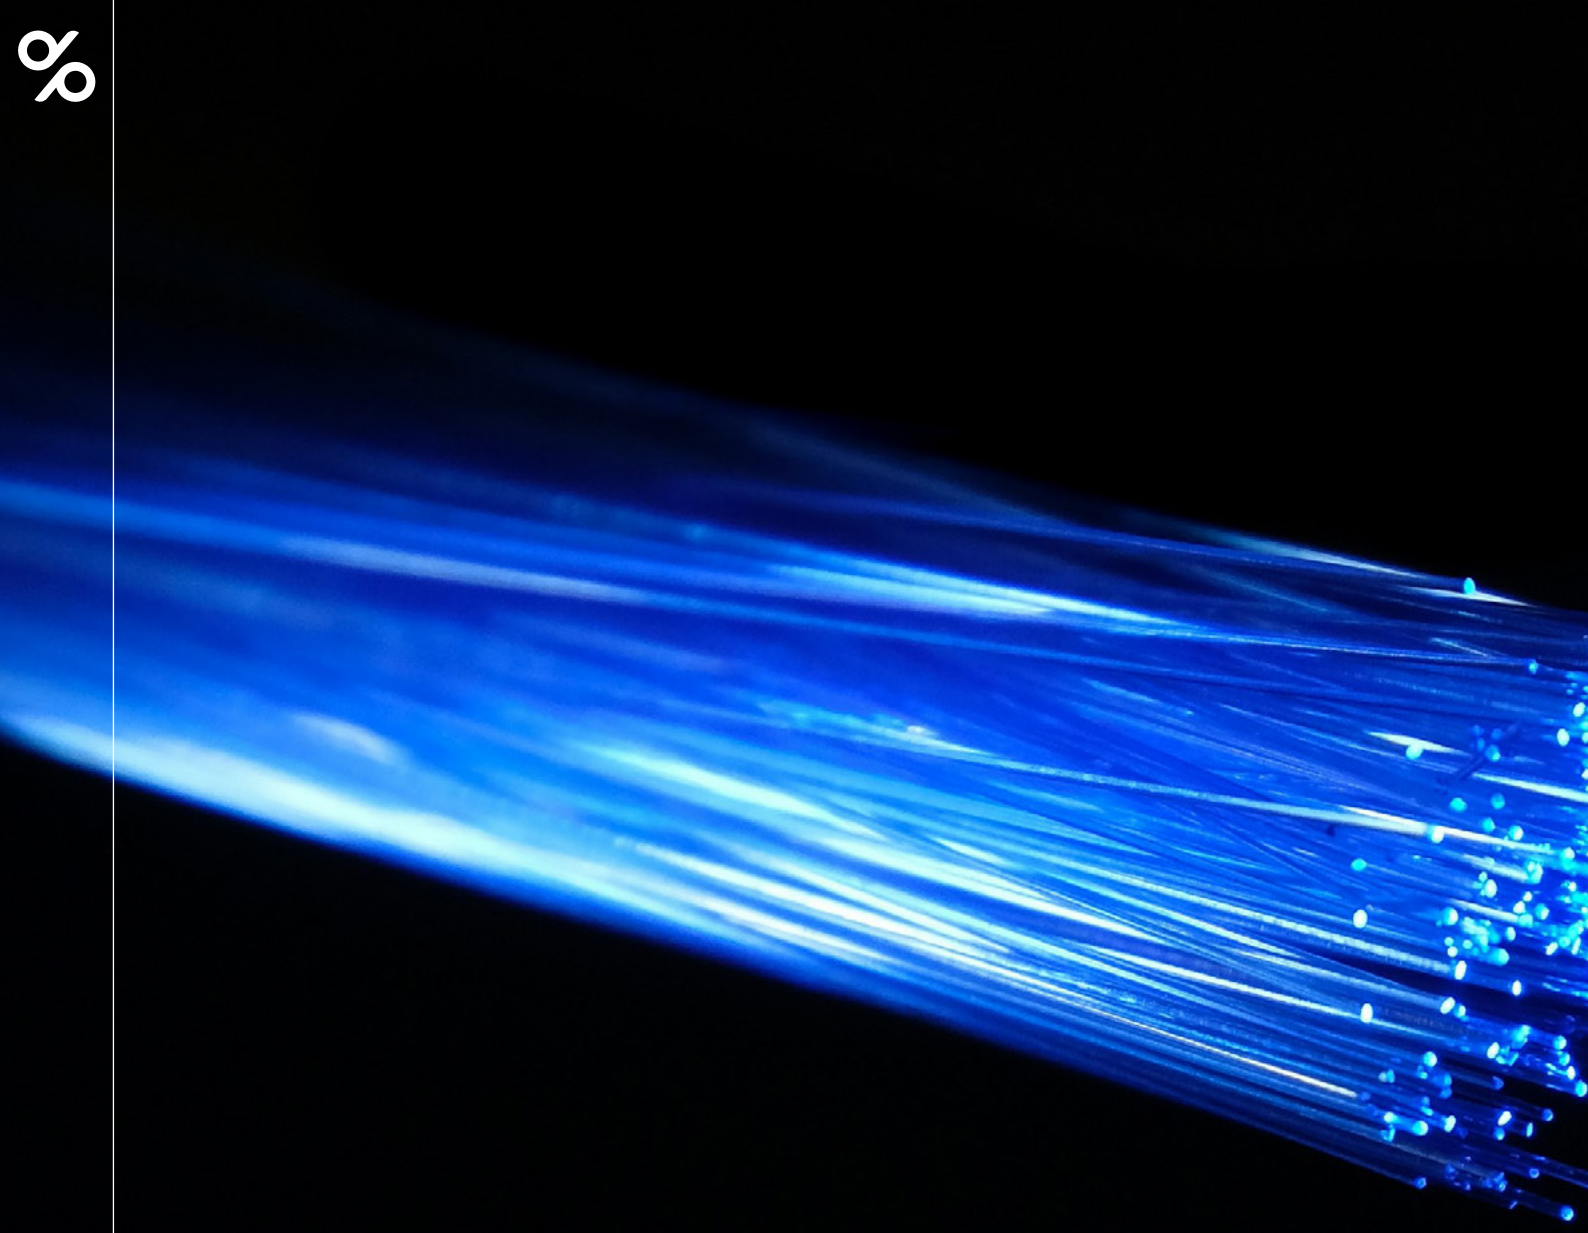
07 **Data provenance, lawful basis, and privacy controls**
Named data lineage, ownership, and the lawful basis relied upon for AI training and inference, with evidence of verification.

08 **Third-party AI accountability**
Disclosure of due diligence and accountability mechanisms covering AI supplied by, or embedded within, third-party products, services, clients, and supply chains

09 **Human Rights Impact Assessment for high-risk AI applications**
Conducted before design and deployment, with clear disclosure of scope, rights considered, affected stakeholders, and how findings are incorporated into governance, testing, and oversight.

10 **Use-case-level societal impact disclosure**
Quantified, verifiable evidence of the real-world effects of specific AI applications, whether positive or negative

None of these ten elements requires new technology to produce. Each requires a disclosure decision a company can already make. That is precisely why the gap documented in this report is a disclosure gap.



XII. Conclusion



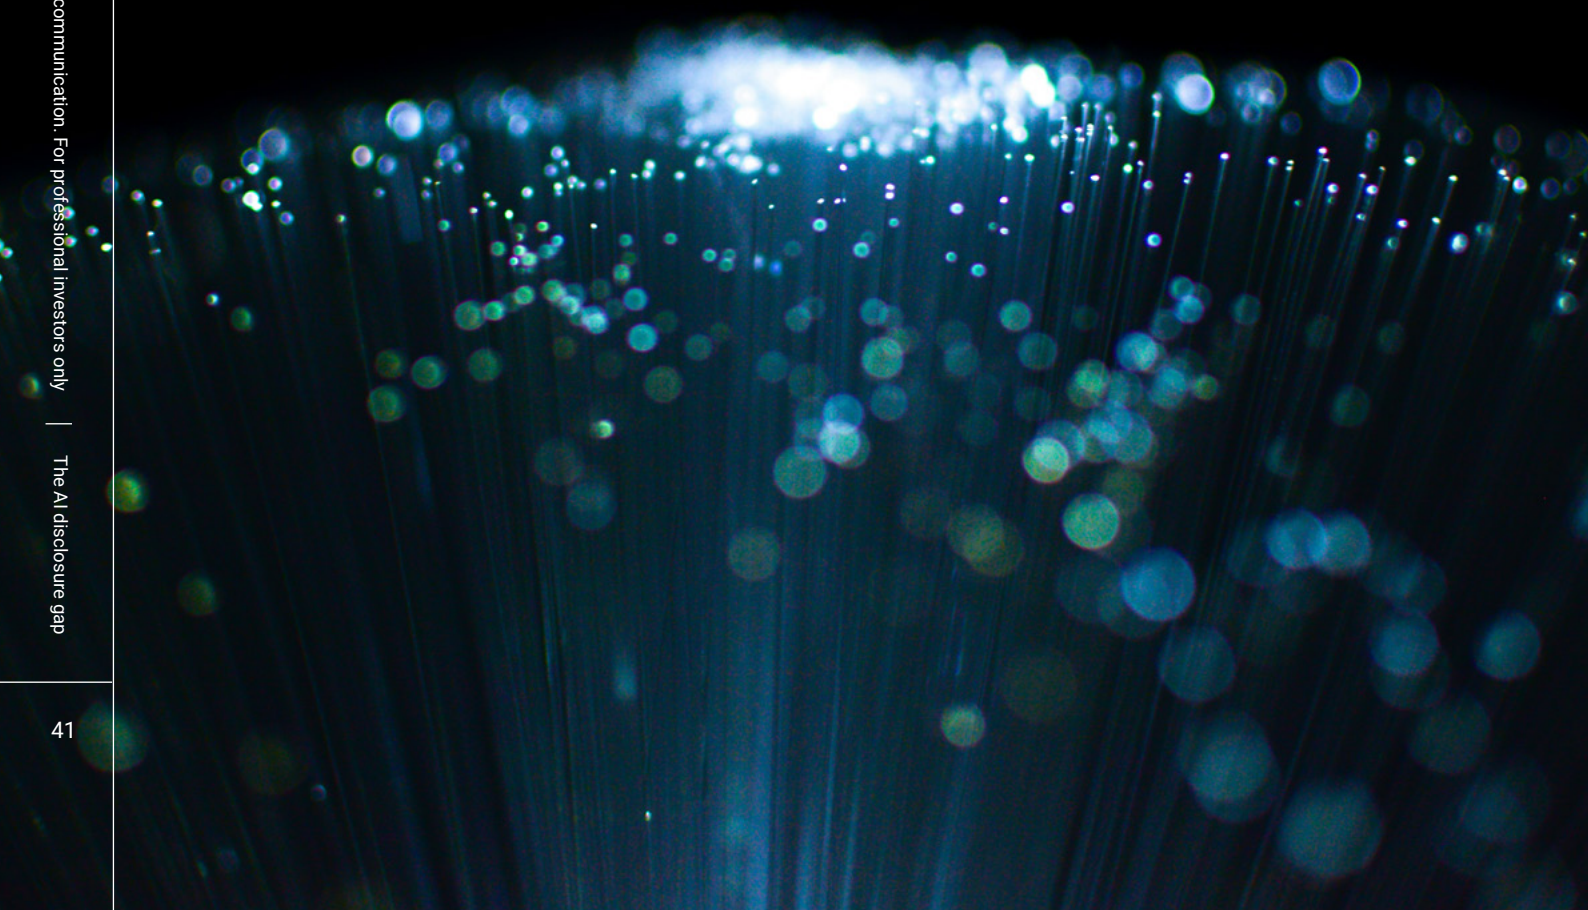
AI governance is now a material issue for investors, not an adjacent ethical issue to be addressed after financial questions are settled. Litigation, regulatory enforcement, and reputational exposure linked to AI are already visible in the market, while the pace of AI deployment continues to outstrip the governance infrastructure meant to contain its risks.

Investors need decision-useful evidence, not policy abundance. This report's assessment of 26 companies and 793 individually scored governance indicators shows that the market currently discloses selectively: relatively strongly on infrastructure, energy, and high-level governance architecture, and weakly—sometimes at zero—on bias testing, audit findings, remedy mechanisms, incident reporting, and use-case-level impact.

Every company assessed scored Red overall, with mean scores clustered narrowly between 0.40 and 1.50 out of 5. That range is too tight, and too uniformly weak, to be explained by company-specific circumstances alone.

The weakest disclosures identified in this report are consistently the ones most relevant to accountability and long-term risk. This creates a practical gap for stewardship and for capital allocation. The framework applied throughout the report is designed to close that gap by converting AI governance into structured, comparable, investment-grade analysis and by giving investors and companies alike a shared evidence-based standard for what good AI governance disclosure should look like.

XIII.
End Notes



Executive Summary

[1]See Anthropic's $1.5 Billion Settlement in July 2026; also, the $250 Million Apple Settlement in May 2026.

Introduction: why this report exists

[1]AI Company Data Initiative (AICDI) / Thomson Reuters Foundation, Corporate AI Governance Report 2025, produced in partnership with UNESCO; see also AICDI, Responsible AI in Practice (2026). https://www.trust.org/resource/ai-company-data-initiative-2025-insights/

[2]World Benchmarking Alliance, "Tech sector progress on AI accountability threatens to stall" (2026), based on an assessment of 200 major technology companies. https://www.worldbenchmarkingalliance.org/tech-sector-progress-ai-accountability-threatens-stall

[3]International AI Safety Report 2026, International AI Safety Report panel/AI Safety Institute. https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026

[4]Fitzjohn-Sykes, L. (2026, April). The role of investors in AI governance [Whitepaper]. Oxford Martin AI Governance Initiative, University of Oxford. See pp. 3, 4, 7 and 12. The shareholder proposal figures are reported on p. 7, citing Mishra (2025), and the coalition figures citing World Benchmarking Alliance (2025). The participants are predominantly stewardship professionals, so the views reported are investors' perceptions rather than measured financial impacts. https://aigi.ox.ac.uk/wp-content/uploads/2026/04/The-Role-of-Investors-in-AI-Governance-Oxford-AIGI-FINAL-130426.pdf

The market problem: disclosure is structurally unbalanced

[1]Reuters Insights, in cooperation with Rystad Energy, 'Powering the AI Era: The global energy industry's capacity, constraints and choices' (2026), survey of 148 senior energy executives across North America, Europe, the Middle East, Asia and Latin America. https://www.rystadenergy.com/insights/reuters-insights-rystad-energy-powering-the-ai-era

[2]World Benchmarking Alliance (2026), op. cit.

Research objective and key questions

[1]Mishra, S. (2026, October). Governing AI: Corporate oversight and shareholder engagement. Harvard Law School Forum on Corporate Governance. https://corpgov.law.harvard.edu/2026/10/06

Methodology and company universe

[1]World Bank. (2026). World development report 2026: The promise of artificial intelligence. World Bank. https://doi.org/10.1596/978-1-4648-2331-2. See Table O.1, p. 9. Firms are defined as leading based on supply chain presence and valuation.

The framework: what it adds to the market

[1]Fitzjohn-Sykes (2026, pp. 12–13).



XIV. Appendix A: The Scoring Model

Table A1. The 0-5 evidence scale applied to every indicator across all three framework sections.

Evidence scale

Score	Label	Criteria
0	No evidence	No verifiable evidence was provided against the indicator. Absence is scored as absence, not as a pass.
1	Minimal	A principle or policy is stated but untested, unexercised or unmeasured. A policy document alone cannot score above 1.
2-3	Moderate (developing)	Practice is developing towards a defined standard; controls exist with some independent verification or partial metric disclosure.
4	Moderate-to-strong	Controls are in place with credible independent verification and a disclosed metric, but the outcome trail is not fully closed.
5	Comprehensive	Independently verified evidence, operational controls, disclosed metrics and a named outcome traceable to a governance action.

Table A2. RAG rating thresholds, applied consistently at both indicator and company overall level.

RAG thresholds

Score	Label
Red	Below 2.00
Amber	2.00 to 3.99
Green	4.00 and above



Three governing principles

- Evidence over assertion: a policy alone cannot score above 1; evidence requires a named artefact, disclosed metric, independent confirmation, or traceable decision.
- Outcome over process: a control that exists on paper but has never been tested, exercised or shown to produce a result does not score above the moderate range.
- Absence is not a pass: a company that cannot show evidence against an indicator is scored 0 and flagged for follow-up, rather than given the benefit of the doubt



XV. Appendix B: Case Study Snapshots

Meta Platforms

Category	Summary
Overall score	1.30 / 5
Overall rating	RED
Risk level tied to Sector 2 on high-risk AI application	HIGH
Main strength	Meta shows stronger governance structures than many peers. It carried out independent human-rights reviews of AI and has the Meta Oversight Board, which can make binding decisions on some content cases.
Main weakness	Meta gives very little hard evidence about the real-world impact of its AI systems on users. Its Use-Case Impact & Positive Contributions score was only 0.25.
Why it matters	Strong policies and transparency structures did not prevent a serious operational failure. In August 2025, a Senate inquiry and internal document leak showed that AI Studio chatbots were allowed to have romantic conversations with users identified as minors
Takeaway	Meta talks more clearly than many peers about governance but gives limited proof that its AI is producing safe and responsible outcomes in practice.

Apple

Category	Summary
Overall score	1.50 / 5
Overall rating	RED
Risk level tied to Sector 2 on high-risk AI application	MEDIUM
Main strength	Apple had one of the highest overall scores in the sample and was especially strong on privacy and security. The report highlights Private Cloud Compute and Apple's security bounty programme
Main weakness	Apple gives less evidence on how its AI products perform in practice, including bias testing, accessibility results, consent review, named data owners and dataset lineage.
Why it matters	Even the strongest company in the sample remained in red territory, showing that strong privacy and infrastructure disclosure do not fully close the accountability gap.
Takeaway	Apple is better than most peers on privacy and security but still does not show enough evidence about how its AI products are tested and governed in practice.

Broadcom

Category	Summary
Overall score	0.70 / 5
Overall rating	RED
Risk level tied to Sector 2 on high-risk AI application	MEDIUM
Main strength	Broadcom gives strong environmental and infrastructure disclosure, including detailed greenhouse-gas reporting. This shows a solid culture of measuring established business metrics.
Main weakness	The same level of detail does not appear in its AI governance disclosure. The report found almost no evidence on AI bias mitigation, explainability or AI data governance.
Why it matters	Broadcom reports well where clear reporting systems already exist but says very little where new AI-specific disclosures would be needed.
Takeaway	Broadcom reports strongly on traditional business and environmental metrics but says very little about AI accountability.

Tesla

Category	Summary
Overall score	0.60 / 5
Overall rating	RED
Risk level tied to Sector 2 on high-risk AI application	HIGH
Main strength	Tesla shows at least some sign of engagement with AI risk categories, including EU AI Act-related disclosure.
Main weakness	Tesla gives almost no meaningful public evidence on fairness, transparency or explainability for Autopilot and Full Self-Driving.
Why it matters	Tesla uses AI in a high-stakes real-world setting, driver assistance and autonomous driving, but gives very little public proof that these systems are being tested and governed properly.
Takeaway	Tesla has one of the highest-risk AI uses in the sample, but one of the weakest levels of public accountability evidence.

Berkshire Hathaway

Category	Summary
Overall score	0.40 / 5
Overall rating	RED
Risk level tied to Sector 2 on high-risk AI application	MEDIUM
Main strength	Some subsidiaries show limited governance activity. For example, Gen Re maintained a Responsible AI Committee chaired by its CTO.
Main weakness	The report found no clear group-wide AI strategy, risk-management framework or incident-reporting mechanism across Berkshire as a whole.
Why it matters	The main issue is not uneven disclosure in one area, but the absence of a visible AI governance system across the group.
Takeaway	Berkshire shows little evidence of group-level AI oversight, which creates a distinct governance risk for investors.



XVI. Appendix C: Company Assessment Summary

Summary overall results for all 26 companies assessed for this report, sorted by overall score. Full underlying assessments, including indicator-level scores, evidence and gap notes, are available on request.

Company	Overall score (/5)	Overall RAG	Inherent business activity risk tied to likely AI high-risk use
Berkshire Hathaway	0.40	RED	MEDIUM
Shell	0.45	RED	MEDIUM
AbbVie	0.50	RED	MEDIUM
Eli Lilly	0.50	RED	MEDIUM
Tesla	0.60	RED	HIGH
UBS Group	0.65	RED	HIGH
Broadcom	0.70	RED	MEDIUM
JPMorgan Chase	0.70	RED	HIGH
ASML Holding	0.77	RED	LOW
Nestlé	0.77	RED	LOW
NVIDIA	0.80	RED	MEDIUM
Roche Holding	0.81	RED	MEDIUM
AstraZeneca	0.84	RED	MEDIUM
Micron Technology	0.90	RED	MEDIUM

Company	Overall score (/5)	Overall RAG	Inherent business activity risk tied to likely AI high-risk use
HSBC Holdings	0.91	RED	HIGH
Novartis	0.97	RED	MEDIUM
Allianz SE	1.00	RED	HIGH
SAP	1.00	RED	MEDIUM
Johnson & Johnson	1.10	RED	MEDIUM
Microsoft	1.10	RED	HIGH
Advanced Micro Devices (AMD)	1.20	RED	MEDIUM
Alphabet (Google)	1.20	RED	HIGH
Mastercard	1.30	RED	HIGH
Meta Platforms	1.30	RED	HIGH
Apple	1.50	RED	MEDIUM
Cisco Systems	1.50	RED	MEDIUM



XVII. Appendix D: Indicator Definitions



Section 1: AI Tech Stack

Assessed for companies with material exposure to the relevant layer of the AI value chain.

- **Agents, Services & Devices:** AI embedded in consumer- or enterprise-facing products and hardware.
- **Agents, Applications:** end-user-facing AI applications and features.
- **Agents, Data Estates:** the datasets an organisation holds and manages that feed AI systems.
- **Intelligence, Data (Training & Inference):** the data used to train and run AI models.
- **Intelligence, Model (Foundations & Fine-tuned):** the AI models themselves, foundation and fine-tuned.
- **Infrastructure, Cloud Infrastructure:** the cloud computing environment supporting AI workloads.
- **Infrastructure, Accelerated Computing (GPUs/TPUs):** specialised compute hardware powering AI training and inference.
- **Infrastructure, Power:** the energy supply and strategy underpinning AI infrastructure.
- **Orchestration, Workflow Automation & Integration:** how AI is integrated into and automates business workflows.
- **Orchestration, Multi-agent & Agentic Systems:** coordinated systems of multiple AI agents acting with a degree of autonomy.

Section 2: High Risk AI Application Mapping

Every company is screened against the EU AI Act's prohibited, high-risk, limited-risk and minimal-risk categories, including an exclusion check against Article 5 prohibited practices applied to all companies regardless of sector. Section 2A applies to companies with EU market exposure; Section 2B applies the equivalent high-risk logic, benchmarked against comparable global standards, to companies whose AI activity falls outside EU market scope.

Section 3: Responsible AI Principles (22 sub-dimensions across 5 pillars)



Pillar 1: Governance & Accountability

- **AI Strategy & Board Oversight:** named board-level or senior-management accountability for AI strategy and risk.
- **AI Risk & Management Framework:** a documented framework for identifying, assessing and mitigating AI-specific risk.
- **AI Incident Reporting & Response:** a defined AI incident taxonomy, with disclosed counts and remediation actions.
- **Third-Party & Supply Chain AI Accountability:** due diligence and accountability mechanisms for AI supplied by, or embedded in, third parties.
- **AI Audit & Assurance:** independent audit or assurance activity covering AI systems, with disclosed scope and findings.



Pillar 2: Human Rights & Fundamental Safeguards

- **Human Rights Impact Assessment (HRIA) for AI:** a structured, independently attributed assessment of AI's human rights impact.
- **Protection of Vulnerable Groups:** specific safeguards for children and other vulnerable users of AI systems.
- **Rights to Remedy & Contestability:** a functioning mechanism for affected parties to contest an AI-driven decision.
- **Human Oversight & Meaningful Control:** disclosed human review and override capability over AI-driven decisions.

Pillar 3:
Data, Privacy & Security

- **Data Governance & Lifecycle Management:** named data ownership, lineage and lifecycle controls for AI training and inference data.
- **Privacy by Design & Data Minimisation:** technical and organisational measures limiting personal data use in AI systems.
- **AI Cybersecurity & Adversarial Robustness:** security controls, including bug-bounty and adversarial-testing programmes, protecting AI systems.
- **Data Sovereignty & Cross-Border Data Flow:** controls governing where AI training and inference data is stored and processed.
- **Consent Architecture & Lawful Basis for AI Processing:** the disclosed lawful basis and consent mechanism underlying AI data processing.

Pillar 4:
Fairness, Transparency & Explainability

- **Algorithmic Bias Detection & Mitigation:** testing for bias by protected characteristics, with a defined remediation threshold.
- **Explainability & Interpretability (XAI):** the ability to explain individual AI-driven decisions to affected parties.
- **AI Disclosure & Transparency to Stakeholders:** disclosure that content or a decision was AI-generated or AI-assisted.
- **Algorithmic Accountability & Audit Trails:** retained, auditable records of AI decision-making for accountability purposes.



Pillar 5:
Use-Case Impact & Positive Contributions

- **Net Positive Impact Assessment:** a structured, quantified, independently corroborated assessment of AI's net real-world impact.
- **Digital Inclusion & Access to AI:** measures ensuring equitable access to the benefits of AI systems.
- **SDG Alignment & Social Value Creation:** quantified mapping of AI use cases to recognised social value or development frameworks.
- **AI for Climate Change & Environmental Solutions:** AI applied to external climate or environmental benefit, distinct from a company's own operational footprint.



Contact details

Responsible Investment Competence Center

ri.competencecenter@degroofpetercam.com

+ 32 2 287 97 01

dpam@degroofpetercam.com

dpaminvestments.com

dpaminvestments.com/blog

/company/dpam