

The Classification Study:

How large organisations classify AI across their public policy documents

Richard Foster-Fletcher

MKAI

Correspondence: research@mkai.org

MKAI Working Paper, June 2026

Foster-Fletcher, R. (2026). How Large Organisations Classify AI in Public Governance Documents. MKAI Working Paper, June 2026.

Abstract

This working paper examines where fifteen large publicly listed organisations place artificial intelligence within their public governance architecture. Classification here means placement: which part of the governance structure governs AI, what other domains AI is grouped alongside, and which mechanisms are applied to it. One placement governs software, systems, data, security, acceptable use, product risk, and deployment; the other governs inputs to organisational judgement such as legal advice, external advisers, expert research, conflicts, privilege, and attestation. The study asks whether any public document in the sample governs AI as an input to organisational judgement, in the way professional advice is governed. In the located public record reviewed for this study, no document does so. AI appears primarily within technology, security, product, risk, conduct, and compliance governance, and the mechanisms built for professional sources are absent. Oracle's Artificial Intelligence Terms and Salesforce's AI Acceptable Use Policy name the advice boundary explicitly, which shows that some drafters recognise the analogy without establishing it across the sample.

1. Introduction and Research Question

Large organisations use several established policy mechanisms to manage different kinds of input. Software and technology are usually governed through acceptable use rules, information security policies, access controls, model risk frameworks, data protection requirements, product governance, and audit logs. Professional input is governed through a different set of mechanisms, including legal engagement terms, privilege rules, conflict disclosures, third-party engagement policies, expert reports, board processes, insider-information controls, and attestations.

Classification in this study means placement within a governance architecture. It asks which part of an organisation's governance structure governs AI, what other domains AI is grouped alongside, and which mechanisms are applied to it. This is a different sense from risk classification under the EU AI Act, from technical taxonomies of model type, and from assessments of governance maturity. The concern here is location. When an organisation documents AI governance for the public, where in the governance structure does AI appear, and what does that placement say about the kind of input AI is taken to be.

This study uses that separation as an analytical frame. It does not assume that every organisation operates two cleanly separated systems, or that the public record captures the full internal position. The narrower question is documentary. When large organisations describe AI governance publicly, where does AI appear in their policy materials, and what kind of input is it treated as providing?

The study is concerned with public classification rather than technical capability. A model may generate analysis, recommendations, or drafting support, while the policy document governing it still treats the system as software, a product feature, a data-processing activity, or a security risk. The hosting arrangement is separate from the research question. This inquiry asks whether AI output is publicly governed as an informational input to organisational judgement, and whether any located document applies the mechanisms normally used for professional advice or external expert input.

The study works with two families of governance, separated by the kind of input each manages and by the mechanisms each applies. Systems governance manages software, data, security, and deployment. Its neighbouring domains are acceptable use, information security, model risk, product governance, and privacy, and its mechanisms include access control, usage logging, content restriction, testing, monitoring, and approval. Professional-source governance manages inputs whose authority comes from expertise, accreditation, fiduciary duty, or professional standing, such as legal advice, external counsel, audit opinions, expert reports, and investment research. Its neighbouring domains are engagement and conflicts, disclosure, privilege, and board process, and its mechanisms include engagement terms, conflict declarations, privilege handling, attestation, and professional-judgement requirements. The two families are identified in this study by where a document places AI, what AI is grouped alongside, and which of these mechanisms apply. The meaning of professional-source governance is fixed here before the evidence is read, so that the study measures AI's placement against a standard set in advance rather than one drawn from the findings.

The research question has two parts. Where is AI placed in the public governance architecture of large organisations, and does the located public record show AI being governed as a professional input into organisational judgement?

2. Methodology

The study used a document-first coding approach. Publicly available documents were located through official corporate websites, investor relations pages, regulatory filings repositories, and SEC EDGAR. Each document was read on its own terms before cross-organisational comparison began.

2.1 Sample

This is a purposive sample of fifteen large publicly listed organisations, selected because they publish substantial public documentation relating to AI governance, responsible AI, technology governance, risk management, or corporate conduct. Selection favoured organisations whose public record is detailed enough to code at document level, since the research question concerns what organisations make visible, and an organisation that publishes little cannot show where it places AI. The findings describe this set of large, well-documented, listed firms. They do not generalise to organisations that publish little, to private firms, or to public-sector bodies.

The organisations are JPMorgan Chase, Goldman Sachs, Morgan Stanley, Bank of America, HSBC, Accenture, Microsoft, Salesforce, Adobe, Oracle, Alphabet, Meta, Walmart, Unilever, and Disney.

2.2 Documents Searched

The inquiry searched for public documents across seven document types. The study does not claim that every publicly available AI-related document for each organisation was located. It reports the evidence found in the public materials listed in the references and appendix.

AI-specific documents included AI usage policies, AI ethics frameworks, AI principles documents, responsible AI materials, responsible AI management systems, AI transparency reports, and AI governance pages.

Technology and security governance documents included acceptable use policies, IT security policies, information security materials, data protection policies, product governance materials, and data classification materials where located.

Professional conduct documents included codes of conduct, codes of ethics, employee standards documents, and public integrity materials where located.

External engagement documents included supplier codes, vendor terms, third-party engagement terms, external counsel materials where available, and procurement-related documents.

Conflict and disclosure documents included conflict of interest policies, insider trading policies, disclosure frameworks, and attestation-related materials where available.

Board governance documents included proxy statements, board governance materials, audit committee materials, and board oversight disclosures where they referenced AI.

Public disclosures included 10-K filings or equivalent annual reports, proxy statements, sustainability reports, ESG reports, topic briefs, and public corporate responsibility materials where AI appeared.

2.3 Coding Framework

For each document located, the coding asked five questions. First, where is AI placed in the policy stack: technology and security, conduct, third-party engagement, AI-specific governance, product governance, disclosure, or another area. Second, what other forms of input, system, risk, or activity is AI grouped alongside. Third, what governance mechanisms are applied to AI, including access control, usage logging, content restrictions, data classification, approval requirements, testing, review, monitoring, contractual limits, attestations, conflict declarations, privilege handling, and professional judgement requirements. Fourth, what vocabulary is used to describe AI. Fifth, what the document says AI is for.

Where internal acceptable-use rules, employee handbooks, executive protocols, board templates, or legal-department guidance were not available publicly, the study treated the absence as missing evidence. It did not treat missing internal materials as evidence that no internal governance exists.

Document coding was conducted using a disciplined, systematic documentary approach. Each public document was evaluated directly against its explicit wording, structural placement, and contextual position within the organisation's overall policy architecture.

3. Findings

3.1 Dominant Public Placement

Across the located public record, AI appears mainly inside technology, security, product, risk, conduct, and compliance governance. The recurring vocabulary is practical and managerial: use, deploy, review, test, monitor, filter, secure, evaluate, approve, comply, and mitigate. The recurring neighbouring concepts include software, cloud services, data handling, cybersecurity, model risk, product risk, content controls, privacy, approved systems, and responsible use.

This supports the central finding. In this sample, public AI governance is most visible through systems, controls, deployment, risk management, product assurance, and compliance. The public documents do not show AI being governed as a professional source of judgement in the way legal advice, external advisers, expert research, or regulated professional opinions are usually governed.

The aggregate findings confirm this system-centric orientation across all fifteen organisations. Ten firms publish a standalone AI-specific public document, seven incorporate AI-specific provisions into their codes of conduct or ethics, and five include explicit AI clauses in supplier terms. Visible board or audit committee oversight appears in two cases. Crucially, across the entire sample, no public document requires formal attestation, applies legal privilege, or mandates engagement-letter controls when AI informs organisational decision-making (see Appendix A for full cross-sample counts).

3.2 Approved-Technology Placement

The clearest support for the technology placement comes from documents that place AI directly inside approved-systems language. Morgan Stanley's 2025 Code of Conduct requires firm information to remain inside "Firm-approved applications and Firm-approved Artificial Intelligence (AI) applications," placing AI within approved-systems control. Goldman Sachs states that AI can only be used in a manner approved by the firm through firm systems. Unilever places AI under Protecting Technology and Information and tells employees to use only approved services, including software as a service and artificial intelligence. Bank of America's code update places AI alongside tools and technology, company-approved applications, AI-enabled search engines, written communications, and insider-trading restrictions.

These examples do not show that AI is treated only as ordinary software. They show a narrower point. In these public documents, AI is governed through the language of approved applications, authorised systems, information protection, employee conduct, and compliance boundaries.

3.3 Professional-Source Governance Is Not Visible in the Located Public Record

The study searched for visible mechanisms that would indicate professional-source treatment. These included conflict declarations triggered by AI-informed judgement, engagement-letter style controls for AI outputs, public disclosure duties where AI materially informed a strategic

recommendation, privilege-style handling of AI-generated analysis, and attestations that identify AI as a source in decision formation.

No located public document in the sample applies those mechanisms to AI. This is an absence in the public record, rather than proof about internal practice. It does, however, show that these organisations have not made professional-source treatment of AI visible in the materials examined.

The same organisations publish extensive governance around legal advice, insider information, supplier terms, disclosure controls, professional conduct, and board oversight. The contrast does not prove institutional motive. It shows that public AI governance and public professional-source governance are documented through different mechanisms in the materials reviewed.

3.4 Oracle and Salesforce as Boundary Cases

Oracle and Salesforce provide the strongest evidence that some policy drafters recognise the professional-advice analogy. Oracle's Artificial Intelligence Terms require users to review and verify AI output, and prohibit its use as a substitute for "any professional (such as healthcare, law, or finance) judgment or advice." Salesforce's AI Acceptable Use Policy prohibits the use of its AI to generate "individualized medical, legal, or financial advice."

These provisions explicitly delineate the advice boundary within public documentation. They demonstrate that AI governance extends beyond technical security and monitoring to address the risk that model output might be received as professional judgement in regulated domains.

Rather than proving a universal rejection of professional-source governance across all corporate functions, these clauses establish clear contractual disclaimers and acceptable-use boundaries. They reflect an explicit policy boundary in the public record regarding professional liability.

3.5 Within-Firm Contrast: Unilever

Unilever illustrates a clear structural contrast by placing different informational inputs into distinct sections of its governance architecture. Under 'Protecting Technology and Information', Unilever instructs employees to use only approved services, explicitly grouping artificial intelligence alongside software-as-a-service. Conversely, professional advice is governed separately under dedicated legal consultation protocols that mandate when senior staff must involve Legal Business Partners on contracts, litigation, and regulatory communications. This within-firm separation demonstrates how public policy frameworks maintain a clear documentary distinction between technology tools and professional advisers.

3.6 Hybrid Cases

Several organisations have public AI materials that go beyond narrow IT-security language. These cases qualify the central finding rather than weakening it.

Disney governs AI through public responsibility materials and an AI topic brief that emphasises human-centred use, creativity, intellectual property, audiences, and information integrity. The mechanism is risk identification and mitigation for proposed AI uses, with executive and board

reporting. This is a wider frame than a simple acceptable-use rule, while still treating AI as a system whose risks must be governed.

Accenture places responsible AI inside ethics, governance, compliance, and audit oversight materials. The mechanisms include risk screening, assessments, standards, controls, mitigation, training, awareness, and monitoring. This resembles an enterprise compliance programme rather than a narrow IT rulebook. The public materials still govern AI as something used across systems and service delivery rather than as an independently disclosed professional source.

HSBC publishes principles for ethical data and AI inside a conduct-oriented architecture. The framing connects data and AI with conduct, risk, uniformity, predictability, and responsible decisions. This is richer than a pure technology control frame, although the public documents still present AI through deployment and governance rather than professional-source treatment.

Meta routes some AI governance through transparency, civil-rights, and human-rights materials, especially around AI-generated content labelling and responsible AI collaboration. This intersects AI with rights, content governance, and public accountability. It remains different from governing AI as an external source of professional judgement.

These hybrid examples reflect the multi-dimensional nature of public AI documentation. Enterprise AI governance frequently extends beyond narrow IT rules to encompass corporate ethics, human rights, product quality, and audit oversight. Nevertheless, even these broader frameworks continue to govern AI as an enterprise capability or system subject to risk mitigation, rather than as an independently disclosed source of professional judgement.

4. What the Study Does Not Show

The study does not show why the organisations have placed AI where they have placed it. The public documents support claims about visible placement, vocabulary, neighbouring policies, and governance mechanisms. They do not support a causal claim that procurement routes produced the classification, or that all organisations made a deliberate strategic decision to exclude professional-source treatment.

The study does not show whether AI is being used in live executive judgement, board paper preparation, strategy work, legal analysis, investment research, or operational decision-making inside these organisations. Public documents are weak evidence for live internal use unless the documents themselves describe that use.

The study does not show that professional-source governance would be better. It only shows that such treatment is not visible in the located public record. Any normative claim about what organisations should do would require a separate argument, especially because AI systems vary widely in use, context, risk, and institutional dependence.

The study does not show that AI is merely software. It shows that public governance materials usually place AI among the mechanisms used for systems, products, risk, security, approved use, responsible deployment, and compliance. That placement may be adequate for some uses and incomplete for others. The present study does not resolve that evaluative question.

5. Methodological Notes and Limitations

The first limitation concerns access. Internal acceptable-use policies, employee handbooks, executive protocols, matter-management systems, legal-privilege procedures, board templates, and AI decision logs are rarely public. Where a document type was not publicly available for an organisation, the study treated this as missing evidence. The findings describe the public record examined, rather than the full internal governance picture.

The second limitation is conceptual. Public documents show how organisations present AI governance to outside readers. They are less able to show how AI is treated in live decision processes. The most probative missing records would include internal generative AI rules, legal-department guidance on prompt confidentiality, board paper templates, executive committee protocols, procurement approval records, model-use registers, and matter-level records for AI use in regulated or strategic decisions.

The third limitation is sampling. The sample comprises fifteen large publicly listed organisations with substantive public documentation. The findings should be read as evidence about this public-document sample. Smaller firms, private firms, public-sector bodies, and organisations in other jurisdictions may publish different governance materials or treat AI differently.

Fourth, document coding was conducted against an explicit, standardized five-question framework. Consistency was maintained by evaluating every document directly on its own verbatim terms and structural placement prior to cross-sample aggregation.

6. Implications for Further Inquiry

The central implication is that public AI governance is currently easier to observe as control governance than as judgement governance. The public documents make visible approved systems, acceptable use, responsible deployment, security, product risk, ethics, conduct, privacy, and board oversight. They rarely make visible the question of whether AI materially informed a professional or strategic judgement.

A stronger claim about organisational intent would require evidence from internal policies, interviews, board-paper templates, procurement approvals, legal guidance, and records of AI use in decisions. The present study can identify an absence in public materials. It cannot explain the internal reasons for that absence.

The Oracle and Salesforce clauses provide a useful starting point for follow-up work. They show that the advice-and-judgement boundary is already visible in some public terms and acceptable-use policies. Future research could test whether similar boundary language appears in internal guidance, and whether it changes how organisations record, disclose, or challenge AI-informed work.

7. Closing Observation

The public record examined here shows AI governed mainly as something to approve, secure, test, monitor, limit, and deploy. Across the fifteen organisations, the governance language for AI is the language of systems and controls, and the mechanisms used for professional sources

remain absent from the located documents. For an outside reader, AI is made legible as infrastructure to be managed. What that record leaves unresolved is what goes unrecorded when organisational judgement is shaped by a source that public governance still describes chiefly as a tool.

References

Corporate documents

JPMorgan Chase. Trustworthy AI Center of Excellence.

<https://www.jpmorganchase.com/about/technology/research/ai/trustai> (accessed 3 June 2026).

JPMorgan Chase. 2025 Proxy Statement (Form DEF 14A). Filed 7 April 2025.

<https://www.sec.gov/Archives/edgar/data/19617/0000019617-25-000321-index.html> (accessed 3 June 2026).

JPMorgan Chase. Letters to Shareholders from Line of Business CEOs (2024 Annual Report).

<https://www.jpmorganchase.com/ir/annual-report> (accessed 3 June 2026). The Consumer and Community Banking letter at <https://www.jpmorganchase.com/ir/annual-report/2024/ar-ceo-letter-marianne-lake> is the primary operating-context AI discussion.

Goldman Sachs. Code of Business Conduct and Ethics. <https://www.goldmansachs.com/our-firm/purpose-and-values/code-of-business-conduct-and-ethics> (accessed 3 June 2026).

Goldman Sachs. Client Security Statement. <https://www.goldmansachs.com/disclosures/client-security-statement.pdf> (accessed 3 June 2026).

Goldman Sachs. Vendor Code of Conduct. <https://www.goldmansachs.com/our-firm/our-vendor-program> (accessed 3 June 2026).

Goldman Sachs. 2025 Annual Report (Form 10-K). Year ended 31 December 2025.

<https://www.goldmansachs.com/investor-relations/financials/current/annual-reports/2025-annual-report> (accessed 3 June 2026). 10-K on SEC EDGAR at <https://www.sec.gov/Archives/edgar/data/0000886982/000088698226000091/gs-20251231.htm>.

Morgan Stanley. Code of Conduct 2025.

https://www.morganstanley.com/content/dam/msdotcom/en/assets/pdfs/Code_of_Conduct_Morgan_Stanley_2025.pdf (accessed 3 June 2026).

Morgan Stanley. General Terms and Conditions (Suppliers). Updated 18 March 2026.

<https://www.morganstanley.com/content/dam/msdotcom/en/disclosures/morgan-stanley-general-terms-conditions.pdf> (accessed 3 June 2026).

Morgan Stanley. International (ex China) Privacy Policy (formerly EMEA Privacy Notice).

<https://www.morganstanley.com/disclaimers/emea-privacy-policy> (accessed 3 June 2026).

Bank of America. Code of Conduct. Updated 12 January 2026.

<https://investor.bankofamerica.com/corporate-governance/governance-library/code-of-conduct> (accessed 3 June 2026).

Bank of America. Annual Report on Form 10-K (fiscal year ended 31 December 2025), Exhibit 19 (Insider Trading Policy). Filed February 2026.

<https://www.sec.gov/Archives/edgar/data/0000070858/000007085826000157/bac-20251231.htm> (accessed 3 June 2026).

Bank of America. AI Adoption by BofA's Global Workforce Improves Productivity, Client Service (press release). 8 April 2025. <https://newsroom.bankofamerica.com/content/newsroom/press->

releases/2025/04/ai-adoption-by-bofa-s-global-workforce-improves-productivity--cl.html (accessed 3 June 2026).

HSBC. Principles for the Ethical Use of Data and AI. July 2024. <https://www.hsbc.com/-/files/hsbc/our-approach/risk-and-responsibility/pdfs/240715-hsbc-principles-for-the-ethical-use-of-data-and-ai.pdf> (accessed 3 June 2026).

HSBC. Our Conduct. <https://www.hsbc.com/who-we-are/our-strategy-and-values/our-conduct> (accessed 3 June 2026).

HSBC. Statement of Business Principles and Code of Conduct. 14 May 2025 edition. <https://www.hsbc.com/-/files/hsbc/our-approach/risk-and-responsibility/pdfs/250514-statement-of-business-principles-and-code-of-conduct.pdf> (accessed 3 June 2026).

HSBC. Supplier Code of Conduct. <https://www.hsbc.com/-/files/hsbc/our-approach/risk-and-responsibility/pdfs/hsbc-suppliers-code-of-conduct-english.pdf> (accessed 3 June 2026).

HSBC Holdings plc. Annual Report and Accounts 2025. Year ended 31 December 2025. <https://www.hsbc.com/investors/results-and-announcements/annual-report> (accessed 3 June 2026).

Accenture. Business Sustainability Reporting: Ethics and Governance (Responsible AI Compliance Program). <https://www.accenture.com/us-en/about/company/integrated-reporting-sustainability-ethics-governance> (accessed 3 June 2026).

Accenture. Code of Business Ethics. <https://www.accenture.com/us-en/about/company/business-ethics> (accessed 3 June 2026).

Accenture. 360 Value Report 2025. Published 16 December 2025. <https://www.accenture.com/us-en/about/company/integrated-reporting-sustainability-environment> (accessed 3 June 2026).

Accenture. Annual Report on Form 10-K (fiscal year ended 31 August 2025). Filed 10 October 2025. <https://www.sec.gov/Archives/edgar/data/0001467373/000146737325000217/acn-20250831.htm> (accessed 3 June 2026).

Accenture. Supplier Standards of Conduct. July 2025. <https://www.accenture.com/content/dam/accenture/final/a-com-migration/pdf/pdf-58/accenture-supplier-standards-of-conduct-final-en.pdf> (accessed 3 June 2026).

Microsoft. 2025 Responsible AI Transparency Report. Published 20 June 2025. <https://www.microsoft.com/en-us/corporate-responsibility/responsible-ai-transparency-report/> (accessed 3 June 2026).

Microsoft. Responsible AI. <https://www.microsoft.com/en-us/ai/responsible-ai> (accessed 3 June 2026).

Microsoft. Code of Conduct for Microsoft AI Services. Effective 13 February 2025. <https://learn.microsoft.com/en-us/legal/ai-code-of-conduct> (accessed 3 June 2026).

Microsoft. Annual Report on Form 10-K (fiscal year ended 30 June 2025). <https://www.sec.gov/Archives/edgar/data/0000789019/000095017025100235/> (accessed 3 June 2026).

Microsoft. Standards of Business Conduct (the Trust Code). <https://www.microsoft.com/en-us/legal/compliance/sbc> (accessed 3 June 2026).

Salesforce. Trusted AI. <https://www.salesforce.com/artificial-intelligence/trusted-ai/> (accessed 3 June 2026). Trusted AI principles at <https://www.salesforce.com/blog/meet-salesforces-trusted-ai-principles/>.

Salesforce. Einstein Trust Layer Documentation. <https://www.salesforce.com/artificial-intelligence/trusted-ai/> (accessed 3 June 2026). Technical documentation on help.salesforce.com; Trailhead module at <https://trailhead.salesforce.com/content/learn/modules/trusted-agentic-ai>.

Salesforce. AI Acceptable Use Policy. First published August 2023, updated subsequently. <https://www.salesforce.com/en-us/wp-content/uploads/sites/4/documents/legal/Agreements/policies/ai-acceptable-use-policy.pdf> (accessed 3 June 2026).

Salesforce. Privacy Information. <https://www.salesforce.com/privacy/> (accessed 3 June 2026).

Salesforce. Annual Report on Form 10-K (fiscal year ended 31 January 2026). Filed 2 March 2026. <https://www.sec.gov/Archives/edgar/data/0001108524/000110852426000060/crm-20260131.htm> (accessed 3 June 2026).

Salesforce. Code of Conduct. Effective May 2025. <https://www.salesforce.com/company/legal/compliance/> (accessed 3 June 2026).

Adobe. Code of Business Conduct. April 2025 edition. <https://www.adobe.com/content/dam/cc/en/corporate-responsibility/pdfs/code-of-conduct-ext.pdf> (accessed 3 June 2026). The AI Ethics section is at Section 6.

Adobe. Annual Report on Form 10-K (fiscal year ended 29 November 2024). Filed 13 January 2025. <https://www.sec.gov/Archives/edgar/data/0000796343/000079634325000004/adbe-20241129.htm> (accessed 3 June 2026).

Adobe. Annual Report on Form 10-K (fiscal year ended 28 November 2025). Filed 15 January 2026. <https://www.sec.gov/Archives/edgar/data/0000796343/000079634326000003/adbe-20251128.htm> (accessed 3 June 2026).

Adobe. Ethics and Integrity. <https://www.adobe.com/about-adobe/integrity.html> (accessed 3 June 2026).

Oracle. Oracle Artificial Intelligence Terms. Version effective 26 September 2025. <https://www.oracle.com/contracts/docs/oracle-ai-terms.pdf> (accessed 3 June 2026).

Oracle. ISO/IEC 42001 Certification: Oracle's Responsible AI Management System. <https://www.oracle.com/europe/corporate/responsible-ai-iso-42001/> (accessed 3 June 2026).

Oracle. Code of Ethics and Business Conduct. <https://www.oracle.com/webfolder/assets/ebook/employee-code-of-conduct-and-ethics/index.html> (accessed 3 June 2026).

Oracle. Supplier Information and Physical Security Standards. <https://www.oracle.com/a/ocom/docs/oracle-supplier-contractor-security.pdf> (accessed 3 June 2026).

Alphabet (Google). Responsible AI Progress Report (published February 2025, covering 2024). <https://ai.google/static/documents/ai-responsibility-update-published-february-2025.pdf> (accessed 3 June 2026).

Alphabet (Google). Google AI Principles. Restructured February 2025 around three tenets. <https://ai.google/principles> (accessed 3 June 2026).

Alphabet (Google Cloud). Responsible AI. <https://cloud.google.com/responsible-ai> (accessed 3 June 2026).

Alphabet. Annual Report on Form 10-K (fiscal year ended 31 December 2025). Filed 4 February 2026. <https://www.sec.gov/Archives/edgar/data/0001652044/000165204426000018/goog-20251231.htm> (accessed 3 June 2026).

Meta. Our Approach to Labeling AI-Generated Content and Manipulated Media. Published 5 April 2024, updated 1 July 2024. <https://about.fb.com/news/2024/04/metasp-approach-to-labeling-ai-generated-content-and-manipulated-media/> (accessed 3 June 2026).

Meta. Human Rights at Meta. <https://humanrights.fb.com/> (accessed 3 June 2026). The 2022 Meta Human Rights Report at <https://humanrights.fb.com/wp-content/uploads/2023/09/2022-Meta-Human->

Rights-Report.pdf. An Update on Meta’s Civil Rights Progress (March 2023) at <https://about.fb.com/news/2023/03/an-update-on-metas-civil-rights-progress/>. The AI-fairness companion post at <https://ai.meta.com/blog/responsible-ai-progress-meta-2022/>.

Meta. Annual Report on Form 10-K (fiscal year ended 31 December 2025). Filed 28 January 2026. <https://www.sec.gov/Archives/edgar/data/1326801/000162828026003942/meta-20251231.htm> (accessed 3 June 2026).

Walmart. Our Responsible AI Pledge: Setting the Bar for Ethical AI. 17 October 2023. <https://corporate.walmart.com/news/2023/10/17/our-responsible-ai-pledge-setting-the-bar-for-ethical-ai> (accessed 3 June 2026).

Walmart. Digital Trust Commitments. https://tech.walmart.com/content/walmart-global-tech/en_us/blog/post/responsible-use-of-data-and-technology.html (accessed 3 June 2026). The October 2024 scaling release at <https://corporate.walmart.com/news/2024/10/09/walmart-reveals-plan-for-scaling-artificial-intelligence-generative-ai-augmented-reality-and-immersive-commerce-experiences>.

Walmart. Ethics and Integrity. <https://corporate.walmart.com/purpose/ethics-integrity> (accessed 3 June 2026).

Walmart. FY2025 ESG Report (Delivering Shared Value). Published September 2025. <https://corporate.walmart.com/content/dam/corporate/documents/esgreport/2025/FY2025-Walmart-ESG-Report.pdf> (accessed 3 June 2026).

Walmart. FY2025 Annual Report. Released 24 April 2025. https://stock.walmart.com/_assets/_81906fd8150d99dd80e92168ae1aff5d/walmart/db/950/9949/annual_report/Walmart+2025+Annual+Report.pdf (accessed 3 June 2026).

Walmart. 2025 Proxy Statement (Form DEF 14A). Filed 24 April 2025. <https://www.sec.gov/Archives/edgar/data/0000104169/000010416925000055/wmt-20250424.htm> (accessed 3 June 2026).

Unilever. Code of Business Principles and Code Policies. <https://www.unilever.com/files/code-of-business-principles-and-code-policies-english.pdf> (accessed 3 June 2026). The Legal Consultation section is within this combined document.

Unilever. Responsible Partner Policy. <https://www.unilever.com/files/56455dcf-0f8f-41cb-a3a9-2f77bf8a6291/unilever-responsible-partner-policy.pdf> (accessed 3 June 2026).

Unilever. The EU AI Act Has Arrived: How Unilever Is Preparing. November 2024. <https://www.unilever.com/news/news-search/2024/the-eu-ai-act-has-arrived-how-unilever-is-preparing/> (accessed 3 June 2026).

Unilever. Annual Report and Accounts 2025. Year ended 31 December 2025. <https://www.unilever.com/files/unilever-annual-report-and-accounts-2025.pdf> (accessed 3 June 2026).

Disney. Operating Responsibly (Artificial Intelligence section). <https://impact.disney.com/operating-responsibly/> (accessed 3 June 2026).

Disney. Artificial Intelligence Topic Brief. As of 23 May 2025. <https://impact.disney.com/app/uploads/2025/05/Artificial-Intelligence-Topic-Brief.pdf> (accessed 3 June 2026).

Disney. Annual Report on Form 10-K (fiscal year ended 27 September 2025). Filed 13 November 2025. <https://www.sec.gov/Archives/edgar/data/0001744489/000174448925000155/dis-20250927.htm> (accessed 3 June 2026).

Disney. 2026 Proxy Statement (Form DEF 14A). Filed 22 January 2026.
<https://www.sec.gov/Archives/edgar/data/0001744489/000174448926000013/dis-20260122.htm>
 (accessed 3 June 2026).

Disney. Standards of Business Conduct. <https://impact.disney.com/app/uploads/Current/TWDC-Standards-of-Business-Conduct.pdf> (accessed 3 June 2026).

Disney. Journalistic Integrity Topic Brief. As of 23 May 2025.
<https://impact.disney.com/app/uploads/2025/05/Journalistic-Integrity-Topic-Brief.pdf> (accessed 3 June 2026).

Regulatory and standards documents

National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. Published 26 January 2023. <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10> (accessed 3 June 2026). DOI: <https://doi.org/10.6028/NIST.AI.100-1>.

European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, OJ L 2024/1689, 12 July 2024. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng> (accessed 3 June 2026).

International Organization for Standardization. ISO/IEC 42001:2023 Information Technology, Artificial Intelligence, Management System. Published December 2023. <https://www.iso.org/standard/42001> (accessed 3 June 2026).

Appendix A: Coded Inventory by Organisation

The table below sets out the public documents located and coded for each organisation, together with their placement in the public policy stack. Citations and URLs for each document appear in the references section.

Organisation	Core public documents located and coded	Policy-stack location(s)
JPMorgan Chase	Trustworthy AI Center of Excellence; 2025 Proxy Statement; 2024 Annual Report (Line-of-Business letters).	AI-specific research and public governance; board and public filing; operating model disclosure.
Goldman Sachs	Code of Business Conduct and Ethics; Client Security Statement; Vendor Code of Conduct; 2025 Annual Report.	Conduct and ethics; technology and security; supplier governance; annual report.
Morgan Stanley	Code of Conduct 2025; General Terms and Conditions for Suppliers; International (ex China) Privacy Policy.	Conduct and ethics; external engagement and supplier terms; privacy and compliance.
Bank of America	Code of Conduct; 2025 10-K Exhibit 19 (Insider Trading Policy); AI workforce use press release.	Conduct and ethics; conflict and disclosure exhibit; operational disclosure.
HSBC	Principles for the Ethical Use of Data and AI; Our Conduct page; Statement of Business Principles and Code of Conduct; Supplier Code; 2025 Annual Report.	AI-specific; conduct and ethics; supplier governance; annual report.
Accenture	Business Sustainability Reporting (Responsible AI section); Code of Business Ethics; 360 Value Report 2025; Annual Report (FY2025 10-K); Supplier Standards of Conduct.	AI-specific and compliance; conduct and ethics; board and audit oversight; supplier governance; annual report.
Microsoft	Responsible AI Transparency Report 2025; Responsible AI corporate pages; Microsoft AI Services Code of Conduct; FY2025 10-K; Standards of Business Conduct.	AI-specific; technology and service governance; public filing and annual report; general conduct.
Salesforce	Trusted AI page; Einstein Trust Layer documentation; AI Acceptable Use Policy; Privacy Information; FY2026 10-K; Code of Conduct.	AI-specific; product governance; privacy and compliance; annual report and investor disclosure; conduct.
Adobe	Code of Business Conduct (AI Ethics section); FY2024 and FY2025 10-Ks; Ethics and Integrity page.	Conduct and ethics; annual report; integrity and compliance.
Oracle	Oracle AI Terms; ISO/IEC 42001 responsible AI page; Code of Ethics and Business Conduct; Supplier Information and Physical Security Standards.	AI-specific contractual terms; technology and compliance; conduct and ethics; supplier and security.

Organisation	Core public documents located and coded	Policy-stack location(s)
Alphabet	2024 Responsible AI Progress Report; Google AI Principles; Google Cloud Responsible AI page; FY2025 10-K.	AI-specific; product development governance; public filing.
Meta	AI-generated content labelling documents; Civil-rights and human-rights materials referencing responsible AI; FY2025 10-K.	AI-specific transparency; human-rights and civil-rights governance; public filing.
Walmart	Responsible AI Pledge; Digital Trust Commitments; Ethics and Integrity page; FY2025 ESG Report; FY2025 Annual Report; 2025 Proxy Statement.	AI-specific pledge; ethics and integrity; annual report; proxy and public disclosure.
Unilever	Code of Business Principles and Code Policies (including Legal Consultation section); Responsible Partner Policy; EU AI Act preparation note; Annual Report 2025.	Technology and information governance; legal-consultation comparator; supplier governance; AI-assurance note; annual report.
Disney	Operating Responsibly (AI section); AI Topic Brief; FY2025 10-K; 2026 Proxy Statement; Standards of Business Conduct; Journalistic Integrity Topic Brief.	AI-specific governance; public filing and proxy; conduct and ethics; journalistic integrity.

Binary Indicator Matrix

The matrix below records five binary indicators for each organisation, coded Y where the indicator is present in the located public documents and N where it is not. Primary placement within technology, security, product, risk, conduct, and compliance governance applies to all fifteen organisations and is recorded in the cross-sample counts. An AI-specific supplier clause means an AI-specific obligation, restriction, standard, or governance requirement in supplier or vendor terms, rather than the presence of a general vendor code alone. Professional-source mechanisms means engagement-letter treatment, conflict disclosure, privilege handling, or public attestation applied to AI because it informed organisational judgement.

Organisation	Standalone AI document	AI in code of conduct	AI-specific supplier clause	Board oversight	Professional-source mechanisms visible
JPMorgan Chase	Y	N	N	N	N
Goldman Sachs	N	Y	N	N	N
Morgan Stanley	N	Y	Y	N	N
Bank of America	N	Y	N	N	N
HSBC	Y	Y	Y	N	N
Accenture	Y	Y	Y	Y	N
Microsoft	Y	N	N	N	N
Salesforce	Y	N	N	N	N

Organisation	Standalone AI document	AI in code of conduct	AI-specific supplier clause	Board oversight	Professional-source mechanisms visible
Adobe	N	Y	N	N	N
Oracle	Y	N	Y	N	N
Alphabet	Y	N	N	N	N
Meta	Y	N	N	N	N
Walmart	Y	N	N	N	N
Unilever	N	Y	Y	N	N
Disney	Y	N	N	Y	N
Total recorded (Y)	10	7	5	2	0

Board oversight is recorded Y only where oversight of AI is attributed to the board or audit committee in the public record. HSBC is recorded N on this indicator because its public materials attribute AI oversight to executive committees, with board-level oversight of conduct and risk operating more generally.

Cross-Sample Counts

Count	Finding in this public-document sample
15 of 15	AI appears primarily within technology, security, product, risk, conduct, and compliance governance.
10 of 15	A standalone AI-specific public document was located.
7 of 15	AI-specific wording appears in a code of conduct or ethics document.
5 of 15	An AI-specific supplier clause was located.
At least 2 of 15	Board or audit committee oversight of AI is visible in the public record.
0 of 15	Professional-source mechanisms are visible for AI in the located documents.
0 of 15	No located public document requires public attestation solely because AI informed a strategic position.
0 of 15	No located public document applies legal-privilege or engagement-letter style treatment to AI outputs as such.