Who writes the rules for AI

Following the public record from consultation to final text

Research by The Actual Factual | Brief 004 | Edition 1 | 14 November 2026

Purpose and scope

A documented U.S. case study of the 2021-2023 NIST AI Risk Management Framework process, with a guide to interpreting lobbying records. Reviewed 14 November 2026. This is a historical participation audit, not a comprehensive account of present AI law or every contributor.

Executive summary

NIST published requests for input, drafts, stakeholder submissions and a final voluntary framework. These records make it possible to examine participation and changes in wording. They do not, by themselves, identify which participant caused a particular decision. [1] [2] [6]

A worked comparison finds that Microsoft asked for a consolidated treatment of trustworthiness characteristics; a later draft and the final document organize the characteristics together. This establishes a proposal and a subsequent textual feature, not exclusive authorship or proven influence. [3] [4] [6]

Public comments and lobbying disclosures answer different questions. A comment records a submitted position. Federal lobbying reports identify specified activities, institutions contacted and financial estimates; they are not transcripts of negotiations or records of why an agency accepted a recommendation. [7] [8]

Identify the instrument before asking who shaped it

This investigation follows one voluntary technical framework. NIST's July 2021 notice explicitly described the AI RMF as voluntary and sought information to guide its development. That is the instrument examined here; calling every AI governance document a law obscures who can change it and how it operates. [1]

Our reporting approach is to establish the issuing institution, document version, decision date, participation route and claimed effect before investigating influence. A school board adopting a vendor checklist, a legislature passing a statute and an agency publishing guidance require different records. For this case, the relevant comparison is between public submissions and successive NIST texts. We make no finding about a particular procurement or enforcement decision.

A worked comparison with a corporate submission

On April 29, 2022, Microsoft submitted comments asking NIST to combine separate classes of risks and characteristics into one list. It argued that separate categories could confuse developers and deployers about their responsibilities and recommended explaining their roles across the AI lifecycle. These are Microsoft's expressed arguments, not this publication's endorsement. [3]

The August 18 second draft presents trustworthiness characteristics together in section 4. The January 2023 final framework also presents a combined set in section 3, describing the characteristics as attributes of

sociotechnical systems. The observed structure is consistent with the consolidation Microsoft requested. The documents do not establish that Microsoft originated the idea or uniquely caused its inclusion. [4] [6]

Our analysis: a responsible influence claim needs an additional evidentiary link. That could be an agency explanation assigning a change to a submission, a documented drafting exchange, or minutes recording the decision. Even then, agreement with an external recommendation would not by itself establish improper conduct. This comparison is a method for locating the next question, not a finding of capture.

Other voices are part of the record

The ACLU's September 29, 2022 submission sought stronger roles for affected communities, consideration of non-AI alternatives and attention to institutional incentives. This demonstrates a materially different set of recommendations in the same process. It does not establish whether the final text adopted each proposal. [5]

NIST's second-draft comment directory includes companies, civil-society organizations, researchers and individuals. Among its entries are Microsoft, Google, the ACLU, the Electronic Frontier Foundation and individual contributors. A list of submitters establishes visible participation; it is not a demographic census or a measurement of whose views received the most weight. [7]

Our analysis: counting submissions is useful only after defining the unit. One organization may submit through a trade association and separately in its own name. Several people may sign one submission. A large document is not necessarily more influential than a short one. Any future participation count should publish its classification rules, duplicate-handling method and denominator.

What lobbying disclosures can and cannot establish

The Senate-hosted statutory text lists specific issues and, where practicable, bill numbers and executive branch actions among the required report fields. Our reading approach is to use those fields to distinguish broad themes from named proposals rather than classify every technology reference as AI lobbying. [8]

The Senate-hosted statutory text identifies report fields covering issues, congressional chambers or federal agencies contacted, lobbyists and financial estimates. The joint disclosure guidance explains income for outside firms and expenses for organizations lobbying on their own behalf. These are different reporting measures. [8] [9]

Our analysis: do not label a registrant's entire quarterly total as spending on one AI provision merely because that provision appears among several issues. Do not add a client's reported expenses and its retained firm's income without examining overlap. A report can be a lead for examining participation, but it does not provide a meeting transcript, the position taken in every conversation or an agency's causal reasoning.

This brief does not quantify AI lobbying or accuse a named participant of undisclosed activity. A defensible spending analysis would require a defined filing period, identified registrations and amendments, a reproducible search method, and explicit treatment of mixed-issue filings. Current filing obligations should be checked against the applicable version of the law and guidance; the records here are used to explain evidence limits.

Evidence record

The table identifies the record, the finding it supports, and the limit of that evidence. Numbered citations open the original source.

Date or record	Evidence and limits
29 July 2021	**Request for information.** NIST invited input on a voluntary AI risk framework. [1] **Limit:** An invitation does not establish who had the capacity to participate.
17 March 2022	**Initial draft released.** NIST announced the draft, a comment deadline and a March workshop. [2] **Limit:** The announcement describes a process, not its representativeness.
29 April 2022	**Microsoft comment.** The submission requested consolidation of trustworthiness categories. [3] **Limit:** A recommendation is not evidence it was accepted because of that author.
18 August 2022	**Second draft.** Section 4 groups the trustworthiness characteristics together. [4] **Limit:** A textual comparison does not assign authorship.
29 September 2022	**ACLU comment.** The submission argued for a stronger role for impacted communities and consideration of non-AI alternatives. [5] **Limit:** A published position does not prove its weight in the final decision.
26 January 2023	**Framework released.** NIST's development record dates the release of version 1.0. [10] **Limit:** Release does not establish implementation or effectiveness.

Claims findings and unknowns

Claim 1

Claim under examination: A company submitted recommendations, so it wrote the rules.

Finding: The selected comment and later texts establish participation and a comparable structural change. [3] [4] [6]

What remains unknown: The reviewed record does not assign that change exclusively to the company. A decision trail is needed.

Claim 2

Claim under examination: The process was open, so all affected people were represented equally.

Finding: NIST published opportunities for input and a directory of received comments. [2] [7]

What remains unknown: Those records alone do not measure barriers, missing voices, resources or the weight given to each submission.

Claim 3

Claim under examination: A lobbying total measures how much was spent securing one AI rule.

Finding: The reporting structure includes specific issues and overall financial estimates. [8] [9]

What remains unknown: The reviewed guidance does not supply an allocation of each registrant's total to individual provisions or demonstrate that spending secured an outcome.

Questions for decision makers

These are recommended oversight questions. They identify records that could support an answer and do not presume a particular conclusion.

Question 1

Which submissions informed each material revision?

Why this matters: A change log would connect participation with decisions more directly than similarity alone. [2] [3] [4] [6]

Who can answer: The issuing agency and the team responsible for the framework.

Records to request: Versioned drafts, redlines, comment dispositions and explanatory memoranda.

Question 2

Who participated through channels beyond written comments?

Why this matters: A comment directory should not be treated as the complete universe of engagement. [10]

Who can answer: The agency's consultation coordinators.

Records to request: Workshop agendas, attendance records, meeting calendars and publicly releasable meeting summaries, with privacy protections.

Question 3

How were affected communities able to participate and challenge assumptions?

Why this matters: The ACLU's submission makes participation a concrete issue to examine; a directory alone does not resolve it. [5] [7]

Who can answer: Consultation organizers and participating community organizations.

Records to request: Outreach plans, accessible participation formats, support offered, feedback received and explanations of how it was considered.

Question 4

What does each reported lobbying amount actually cover?

Why this matters: Outside-firm income and an organization's expenses are distinct measures, and a filing can cover multiple matters. [8] [9]

Who can answer: The registrant and the organization whose activity is reported.

Records to request: Original and amended reports, reporting-method descriptions and any available issue-level breakdown.

Question 5

What evidence would justify a claim of influence or improper preference?

Why this matters: A public-interest investigation should define the test before selecting evidence that fits a suspicion. [3] [4] [6]

Who can answer: The agency decision-maker and the publication making the claim.

Records to request: Dated recommendations, decision explanations, relevant correspondence, alternatives considered and the applicable conflict-of-interest rules. Seek responses from the named parties before making adverse findings.

Connections and their limits

These relationships are a reading aid. A connection does not establish causation, strength of influence, or motive. The interactive version is available on the report landing page.

Consultation invitation and Public submissions

Documented relationship: NIST announced a written-comment process; its directory preserves received submissions. [2] [7]

Limit: An opportunity to comment does not demonstrate equal access or equal consideration.

Microsoft recommendation and Later framework structure

Editorial analysis: The requested consolidation can be compared with the subsequent grouped characteristics. [3] [4] [6]

Limit: Consistency between a proposal and a result is not attribution of causation or motive.

Civil society recommendations and Participation questions

Documented relationship: The ACLU asked for a stronger role for people affected by AI. [5]

Limit: This shows the organization's position, not all affected people's views or its eventual influence.

Lobbying disclosures and Named institutions and issues

Documented relationship: The statutory reporting fields connect registrants with issues and institutions contacted. [8]

Limit: They do not reconstruct each conversation or establish why officials acted.

Textual overlap and Decision trail

Open question: The comparison prompts a request for explanations of drafting choices. [3] [4] [6]

Limit: The link represents an unanswered research question, not an allegation.

Method and corrections history

Edition 1 reviewed 14 November 2026. We examined the original participant submissions, NIST draft and final text, agency process records and official lobbying-disclosure materials linked here. The case selection illustrates a reproducible method; it is not a representative sample of all rulemaking. We did not code the entire comment corpus, audit all meetings, calculate lobbying totals or obtain an agency attribution for the selected textual change. Findings concern the cited historical versions. Analysis and recommended questions are labelled; no finding of misconduct, exclusive authorship or motive is made. Corrections should identify the affected claim, replacement evidence and revision date.

14 November 2026 | Edition 1. Initial publication of this research brief. No corrections recorded in this edition. Future substantive updates should identify what changed and the supporting evidence.

Web report and future updates

Source register

Each numbered citation and source title links to the original record. Source dates identify the cited version, not necessarily its current legal or operational status.

[1] Artificial Intelligence Risk Management Framework request for information

NIST in the Federal Register | 29 July 2021

Primary process record; summary, dates and invitation. Defines voluntary scope.

[2] NIST seeks comments on draft AI Risk Management Framework

NIST | 17 March 2022

Primary agency announcement; initial draft, comment deadline and workshop.

[3] Comments of Microsoft Corporation on the initial draft

Microsoft submission hosted by NIST | 29 April 2022

Primary participant statement; opening recommendations on consolidating categories. Not an independent assessment.

[4] AI Risk Management Framework second draft

NIST | 18 August 2022

Historical draft; section 4, printed pages 10-11. Compared with the participant proposal and final text.

[5] ACLU comments on the second draft

American Civil Liberties Union submission hosted by NIST | 29 September 2022

Primary participant statement; executive summary, printed pages 1-2. Identifies recommendations, not agency adoption.

[6] Artificial Intelligence Risk Management Framework version 1.0

NIST | January 2023

Final version examined; section 3, printed page 12. A voluntary framework, not a comprehensive statement of AI law.

[7] Comments on the second draft of the AI Risk Management Framework

NIST | Created 8 November 2022; updated 1 December 2022

Primary directory of published comments. Not a complete inventory of every engagement channel.

[8] Lobbying Disclosure Act section 5

United States Senate | Statutory compilation reviewed 14 November 2026

Published reporting fields and financial estimates; verify governing law separately for compliance work.

[9] Lobbying Disclosure Act guidance

Clerk of the House and Secretary of the Senate | Version retrieved 14 November 2026

Sections on recording income and expenses. Explains distinct measures and potential overlap.

[10] AI RMF development chronology

NIST | Reviewed 14 November 2026

Agency chronology of the RFI, concept paper, draft consultations and January 2023 release.