

State AI Governance Under Constitutional Constraint

An Architecture for Deployer-Centered, Process-Based AI Accountability
Designed to Survive Dormant Commerce Clause and Preemption Scrutiny

Zola Valashiya*

Independent Researcher

LL.M. (Sustainable International Development), University of Washington School of Law

LL.M. (Innovation & Technology Law), Seattle University School of Law

AIGP · CIPT · ISO/IEC 42001 Provisional Implementer

zola.zola.v@gmail.com

Abstract

State AI regulation must govern decisions with local consequences while respecting interstate commerce and federal law. States should begin with the AI-mediated decision: who makes it, which state authority reaches that actor, what evidence accountability requires, and how an affected person can obtain review. Constitutional constraints can guide the choice of duties without dictating a single lawful arrangement.

The argument makes three contributions. It connects Dormant Commerce Clause and preemption analysis to the design of particular deployer duties, comparing their actual burdens and federal overlaps. It develops regulatory resilience as a portfolio of consumer-protection, civil-rights, licensing and procurement channels, while testing whether each can support the same duty and remedy. It then specifies statutory duties, evidence and enforcement consequences informed by the NIST AI Risk Management Framework and ISO/IEC 42001.

Existing state laws already supply important components. The added step is a worked account of how those components fit together, where they conflict with federal requirements, and what remains necessary for implementation. Hospital, hiring and insurance applications test the design against different legal settings. The proposed design makes information access, contestable decisions and enforcement capacity central to accountability. Its durability depends on the duties enacted and the institutions able to enforce them; neither process compliance nor a familiar state authority guarantees protection or relief.

* Zola Valashiya, LL.M. (Sustainable International Development), University of Washington School of Law; LL.M. (Innovation & Technology Law), Seattle University School of Law; AIGP, CIPT, PMI-CPMAI, ISO/IEC 42001 Provisional Implementer, Certified AI Ethicist (CAIE); Oxford AI Ethics, Regulation and Compliance Programme, Saïd Business School, University of Oxford. Artificial intelligence tools were used to support the research for this work: locating and retrieving primary sources, organising and verifying the evidence base, and assisting with formatting and developing the author's original draft. The ideas, thesis, analytical framework, arguments, insights and conclusions are the original work of the author.

Executive Summary

A state choosing to regulate AI should begin with the decision that affects a person. An employer ranking applicants, a hospital using a recommendation in treatment, and an insurer assessing an application occupy different legal settings. Their use of similar technology does not make their duties interchangeable. I propose attaching accountability to the organization using AI in a regulated decision, then specifying what it must assess, what evidence it must retain, who may inspect that evidence, and how an affected person can obtain review and relief.

The first contribution is a method for turning constitutional constraints into drafting choices. Dormant Commerce Clause analysis requires attention to the conduct regulated, its territorial connection, discrimination and burdens on interstate commerce. Preemption requires comparison between a particular state duty and the applicable federal law. The comparison must identify the actor, conduct, sector, territorial connection and proposed remedy, holding the relevant conditions constant and making the tested difference explicit. Locating responsibility at deployment can make that inquiry more concrete; it does not settle the result. A documentation requirement, an audit mandate and a restriction on use can impose different burdens even within the same statute. Duties imposed by several states also require attention to their cumulative operation.

The second contribution is regulatory resilience through a portfolio of consumer protection, civil rights, sectoral licensing and public procurement. These channels reach different actors and authorize different forms of intervention. Their value lies in preserving legally available ways to address harmful conduct when another route narrows. That requires examining each channel's authority, elements, evidence and remedies. Procurement conditions must remain connected to what the state buys and must be tested against federal law. Existing police powers likewise do not ensure that a particular duty survives preemption. Coalition cost identifies a possible political consequence of displacing several institutional arrangements together; it supplies no constitutional barrier to congressional action or demonstrated prediction about legislative difficulty.

The third contribution is a worked statutory translation. The statutory arrangement gives six process obligations a legal setting: documentation, risk assessment, auditability, monitoring, incident response and recourse. The NIST AI Risk Management Framework and ISO/IEC 42001 inform organizational governance and planning. Legislation must select the covered actors, triggers, evidence requirements, reviewing institutions and consequences. The Appendix works through procurement controls, private-deployer assessments, audit trails and dispute pathways. It distinguishes a vendor's submission from the state's review of it, an independent audit from a senior official's attestation, and an administrative complaint from a private action. Each mechanism needs a defined task and a route from evidence to an authorized response.

Existing state practice supplies important components of this design. The New York and California comparisons therefore follow particular duties, institutions and dates. Both states have enacted duties for frontier developers alongside measures directed at particular uses. The comparison distinguishes transparency and incident duties from privacy rights, employment rules and public-sector requirements, and separates enactment from commencement. The historical inventory supports a bounded description of recorded instruments; separately dated comparisons account for selected later developments. Neither supplies an exhaustive current-law census or a causal test of legislative influence.

New York City's hiring-audit regime already requires independence and specifies methodology; its implementation record raises questions about enforcement and the use of audit information. That record supports examining additional oversight, without establishing that the proposed additions would improve outcomes. California's public procurement policies and model contract

terms provide another existing sequence of assessment and review. The terms governing an actual purchase depend on their incorporation and the statement of work. The proposed statutory arrangement adds specified responsibilities and evidence channels alongside these practices.

Information access determines whether the duties can be performed. A deployer may depend on a supplier for material evidence about a system. A regulator or auditor needs to distinguish evidence that was obtained from information that remains unavailable, identify the lawful grounds for withholding, and assess what the missing information prevents. A record of requests can expose that limitation; it cannot establish adequate assurance. The design consequently treats access, confidentiality, substitute evidence and evidentiary insufficiency as part of the accountability task. Organizational certification or a completed assessment cannot establish the lawfulness of every consequential decision.

Recourse connects that evidence to the person affected. An explanation must be usable in a challenge, a reviewer must have authority to act, and any available relief must fit the claim. Administrative enforcement, contract remedies and private litigation each have their own predicates and limits. Hospital, hiring and insurance applications test these connections in different legal settings. Federal healthcare requirements, employment protections and insurance supervision alter the duty, the evidence needed and the institution able to respond. A common set of process obligations must be adapted to those differences. The analysis concentrates on structural federalism; a particular enactment also requires its own assessment of speech, due process, equal protection and other applicable constraints.

Enactment requires decisions about capacity as well as legal form. The worked text identifies functions for a state AI oversight office, sectoral regulators, procurement officials and courts, while leaving state-specific funding, appropriations and institutional integration to completion. The annual assessment does not fund every function in the proposal. Smaller organizations, limited public resources, supplier control of information and barriers to individual relief constrain what the arrangement can deliver. States should select duties they can administer, provide the evidence and authority needed to enforce them, and evaluate whether those duties improve protection. The proposal makes those choices explicit. Its durability depends on the law enacted and the institutions able to carry it out.

Reader's Guide

Part I: Introduction begins with the decision that affects a person: who uses AI, what evidence that organization controls, and how the result can be questioned. Recent federal policy and selected state measures establish the setting for the paper's three contributions: constitutional design, an institutional portfolio and worked statutory duties.

Part II: Constitutional Analysis develops the method for comparing state duties. The analysis of *National Pork Producers*, *Pike* and the preemption cases tests discrimination, territorial effects, interstate burdens and federal conflicts. Each comparison specifies the actor, conduct, federal overlap and remedy for the proposed duty.

Part III: The Deployer-Centered Model places responsibility on the organization using AI to make a consequential decision. The historical legal inventory and New York/California comparison distinguish duties, dates and enforcement arrangements. Seven high-impact domains then ground six process obligations: documentation, risk assessment, auditability, monitoring, incident response and recourse.

Part IV: State Authorities tests consumer protection, civil rights, sectoral licensing and procurement as an institutional portfolio. Statutes and cases identify which actors each channel reaches, what must be proved and which remedies it supplies. It asks whether another channel can perform the legal work lost when one route narrows.

Part V: Governance Primitives connects the obligations through procurement controls, impact assessments, audit trails and dispute pathways. California procurement policies and model terms, alongside the NEPA, SOX and FCRA analogies, help specify the tasks of vendors, deployers, reviewers and enforcement bodies. Hospital, hiring and insurance applications test those connections; the final section addresses staffing, funding, supplier evidence and the burdens of administration.

Part VI: Conclusion draws the analysis into a sequence for enactment: identify the decision and legal authority, test the duty, secure usable evidence, and provide institutions and relief. It explains how resources and remaining constitutional questions affect the choices a state must make.

Appendix A: Worked Statutory Text sets out provisions for procurement, impact assessment, audit trails and recourse, followed by general provisions. The clauses and commentaries show how coverage, triggers, records, review and consequences fit together. The Preliminary Note identifies the state-specific institutional, remedial and financial choices needed before enactment.

Contents

Executive Summary	2
Reader’s Guide	4
Part I: Introduction	7
A. State AI Accountability Must Operate Through Changing Federal Policy	7
B. The Decision Environment Makes the Duty and Its Limits Visible	9
C. Constitutional Constraints Inform the Choice of Duty	9
D. Regulatory Resilience Requires a Usable Portfolio	10
E. Frameworks Need Legal Specification and Evidence	11
F. Accountability Requires Power to Question the Decision	12
Part II: The Constitutional Architecture for State AI Regulation	12
A. The Dormant Commerce Clause as Design Constraint	12
B. Federal Preemption: Compare the Duty with the Federal Rule	26
C. State Police Powers and the Regulatory Heartland	36
Part III: The Deployer-Centered Model: Accountability for AI-Mediated Decisions	39
A. Accountability Begins with the Decision and Its Author	39
B. Existing Laws Supply Components, Not a Single Architectural Verdict	43
C. High-Impact Domains Define Different Accountability Tasks	58
D. Process-Based Obligations as the Regulatory Mode	70
Part IV: Matching State Authorities to Duties, Institutions and Remedies	80
A. Consumer Protection: The Representation and the Available Claim	81
B. Civil Rights: Substantive Protection and the Evidence to Enforce It	83
C. Sectoral Licensing: Authority over the Institution and Its Use	89
D. Public Procurement: Duties Connected to What the State Buys	95
Part V: Governance Primitives: A Modular Architecture for State AI Regulation	98
A. Procurement-First Controls as the Architectural Foundation	98
B. Deployer-Centered Impact Assessments as the NEPA Analogue Adapted to AI	105
C. Audit Trail Mandates as the Backward-Looking Accountability Spine	107
D. Dispute and Remedy Pathways as the Affected Person’s Recourse Architecture	110
E. Three Worked Sector Applications	111
F. Enactment and Enforcement Capacity	121

Part VI: Conclusion: State Governance as Necessary, and the Path Forward	123
A. States Should Build Duties They Can Administer	123
B. Accountability Depends on Access to Evidence and Power to Act	123
C. Enactment Requires Specific Choices About Authority, Evidence and Relief	124
Appendix A: Worked Statutory Text for the Four Governance Primitives	125
Preliminary Note	125
§ 1. Procurement-First Controls (Anchoring Statute)	126
§ 2. Deployer-Centered Impact Assessment	130
§ 3. Audit Trail Mandate	134
§ 4. Dispute and Remedy Architecture	138
§ 5. General Provisions	142

Part I: Introduction

A. State AI Accountability Must Operate Through Changing Federal Policy

An employer can buy a screening system from outside the state and use it to reject a local applicant. A hospital can rely on a vendor's score when allocating treatment. The person affected encounters the decision through the employer or hospital, yet the evidence needed to question it may sit with a supplier. States should make that decision environment the starting point for accountability: identify the responsible organization, specify the information it must retain or obtain, and provide a usable route to contest an adverse result.

Federal policy supplies an unstable setting for that task. Executive Order 14179, issued in January 2025, directed an AI Action Plan and review of actions taken under the earlier, already revoked Executive Order 14110.² Executive Order 14365 later directed a Justice Department litigation task force, an evaluation of state laws and preparation of a legislative recommendation for federal preemption.³ The March 2026 White House framework recommends preemption while reserving specified categories, including generally applicable police-power laws and state procurement and use. These are executive policy recommendations. Their exceptions neither bind Congress nor establish a constitutional boundary.⁴

The distinction matters when policy becomes litigation or agency action. The FTC's July 2026 proposed accuracy statement advanced a preemption position for public comment; a proposal does not itself settle the effect of the FTC Act on a particular state duty.⁵ The April 2026 stipulated enforcement restraint in *X.AI LLC v. Weiser* was procedural relief, not a merits ruling determining which state duties are permissible.⁶ The FTC's December 2025 decision to set aside the *Rytr* order provides another concrete example of a changed administrative position.⁷ Each record supports its particular history. Federal statutes and regulations governing em-

² Exec. Order No. 14179, *Removing Barriers to American Leadership in Artificial Intelligence*, §§ 4–5, 90 Fed. Reg. 8,741 (Jan. 31, 2025) (signed Jan. 23, 2025). Section 5 describes Executive Order 14110 as already revoked and directs review of actions taken under it; section 4 directs development of the Action Plan.

³ Exec. Order No. 14365, *Ensuring a National Policy Framework for Artificial Intelligence*, §§ 3–4, 8, 90 Fed. Reg. 58,499 (Dec. 16, 2025). Section 8 concerns preparation of a legislative recommendation, including its specified exceptions.

⁴ White House, *A National Policy Framework for Artificial Intelligence: Legislative Recommendations* 3–4 (Mar. 2026); compare Exec. Order No. 14365 § 8(b). The proposed exceptions and sectoral allocation are recommendations to Congress.

⁵ Policy Statement Concerning the Suppression of Accuracy in Artificial Intelligence Systems, 91 Fed. Reg. 41,638, 41,639 & n.12 (July 7, 2026) (proposed statement); Fed. Trade Comm'n, *FTC Seeks Public Comment on Policy Statement Addressing AI Accuracy* (July 1, 2026). The dated proposal advances the Commission's preemption position and invites comment.

⁶ *X.AI LLC v. Weiser*, No. 1:26-cv-01515, ECF No. 24, at 1–2 (D. Colo. Apr. 27, 2026). The stipulated restraint covers specified violations of SB 24-205 and replacing or amending legislation, through fourteen days after a ruling on the contemplated preliminary-injunction motion. See also United States Complaint in Intervention, ECF No. 17 (Apr. 24, 2026) (pleading federal constitutional objections). The order supplies no merits holding on those objections.

⁷ *In re Rytr LLC*, FTC File No. 232-3052, Order Reopening and Setting Aside Decision and Order (Dec. 22, 2025); Fed. Trade Comm'n, *FTC Reopens and Sets Aside Rytr Final Order in Response to the Trump Administration's AI Action Plan* (Dec. 22, 2025). The Commission identified both inadequacy of the pleaded FTC Act theory and burdens on innovation as grounds for its action.

ployment, medical devices, insurance and other activities remain part of the inquiry when AI is used.

A state response has costs of its own. Different documentation rules, technical requirements and remedial procedures can burden organizations operating across state lines. A national rule may resolve conflicts or allocate responsibility more effectively. The argument here takes a narrower premise: where a state has decided to regulate, it should design duties that address the relevant harm, respect federal requirements and can be administered. That premise leaves the broader case for state regulation open to challenge.⁸

Existing state instruments supply components and competing choices. Colorado’s 2026 legislation changes the earlier allocation of developer and deployer responsibilities; its general operation begins in 2027.⁹ Connecticut’s public-agency framework, Illinois’s employment provisions and New York City’s audit-and-notice law govern different actors through different institutions.¹⁰ California regulates specified privacy processing and employment uses, while New Jersey applies its disparate-impact rules through its own coverage and burdens.¹¹ New York’s RAISE legislation imposes frontier-developer duties under a distinct statutory framework.¹² Healthcare, companion-chatbot and pricing measures add further examples whose commands differ even where their concerns overlap.¹³

These instruments identify legal choices that can be compared duty by duty. The historical inventory and selected current comparisons in Part III distinguish counted enactments from other instruments, dates and institutional developments. The NAIC’s insurance materials, for example, identify model expectations and jurisdictional activity without establishing a uniform operative regime.¹⁴ The constitutional inquiry must follow the actual duty. This Article concentrates on the Dormant Commerce Clause and federal preemption; a defensible position under those doctrines does not dispose of First Amendment, due-process or equal-protection

⁸ The premise is illustrated by the different legislative choices discussed in this Section. The Article addresses the design of state duties where a state chooses to regulate; it does not resolve the broader allocation debate.

⁹ Colorado SB 24-205 (2024 Reg. Sess.); Colorado SB 26-189 (2026 Reg. Sess.), §§ 1–3, 5. The successor generally takes effect January 1, 2027, with specified earlier provisions; *see infra* note 88.

¹⁰ 2023 Conn. Pub. Acts No. 23-16, §§ 1–2 (public-agency inventory and policies); 775 Ill. Comp. Stat. 5/2-102(L) (employment discrimination and notice); N.Y.C. Admin. Code §§ 20-870 to 20-874 (covered employment decision tools). *See infra* notes 196–197, 214.

¹¹ Cal. Code Regs. tit. 11, §§ 7001, 7200–7222; *id.* tit. 2, §§ 11008–11009, 11013(c), 11017; N.J. Admin. Code tit. 13, ch. 16. California’s title 11 privacy regulations took effect January 1, 2026; its title 2 employment regulations took effect October 1, 2025. The ADMT article’s compliance date is January 1, 2027; specified assessment transition rules differ. *See infra* notes 198, 218–219.

¹² 2025 N.Y. Laws ch. 699 (S. 6953-B), as amended by 2026 N.Y. Laws ch. 96 (S. 8828), including amended N.Y. Gen. Bus. Law §§ 1420–1422. Section 1421(1)(a) addresses a frontier developer’s published framework and incorporation of national and international standards and industry consensus practices. *See infra* note 206.

¹³ For the selected healthcare measures, *see infra* notes 297–300; for companion-chatbot measures and their differing coverage, *see infra* note 225; for competition-law pricing controls, *see infra* note 280. Maryland’s enacted HB 895, 2026 Md. Laws ch. 154, § 13-321(B), supplies a separate specified-food-retailer pricing example; its struck-out general merchant disclosure provision was not enacted. These examples do not share a single regulated actor or statutory duty.

¹⁴ NAIC, *Model Bulletin: Use of Artificial Intelligence Systems by Insurers* (Dec. 4, 2023); NAIC, *Implementation of NAIC Model Bulletin: Use of Artificial Intelligence Systems by Insurers* (status Oct. 8, 2026). The roster places twenty-seven states, the District of Columbia and Puerto Rico in the model-bulletin category and separately identifies four states with other arrangements. *See infra* notes 289–292. A roster category does not establish uniform legal force or implementation.

objections.¹⁵

B. The Decision Environment Makes the Duty and Its Limits Visible

Colorado illustrates why a changed allocation of duties must be read provision by provision. The original SB 24-205 allocated reasonable-care duties to developers and deployers of high-risk systems. Its commencement was postponed before the 2026 legislation replaced the framework.¹⁶ Virginia’s vetoed HB 2094 provides a separate record of objections to a proposed allocation; the veto does not supply an operative statutory regime or establish the constitutional defect of every comparable duty.¹⁷

The Colorado successor changes coverage and obligations while retaining developer-facing provisions alongside deployer duties. Disclosure, explanations, correction and review require separate attention to triggers, exceptions and available relief.¹⁸ Reenactment does not establish why legislators chose each change. Nor does it establish that the amended duties have survived a merits challenge. Legislative learning, litigation pressure and industry opposition are possible explanations to be tested against their own evidence. The sequence does not isolate their causal effects.

The decision-centered method remains useful without that causal claim. Hiring, tenant screening and insurance underwriting each identify an actor making a consequential choice within an existing legal setting.¹⁹ Starting there makes it possible to ask what the actor controls, what it must obtain from a supplier, and what follows when the evidence is inadequate. It also exposes the boundary of the proposal: a duty nominally attached to a deployer can still impose burdens on interstate activity or regulate a federally governed product in operational effect. The organization’s label cannot resolve those questions.

C. Constitutional Constraints Inform the Choice of Duty

The Dormant Commerce Clause directs attention to discrimination, territorial connections, interstate burdens and the benefits claimed for a particular requirement. *National Pork Producers Council v. Ross* rejected the proposed broad rule against laws with out-of-state effects, while its separate opinions require proposition-specific treatment of *Pike* balancing.²⁰ Neither

¹⁵ See *supra* note 6; *infra* Part VI.C. This Article’s structural-federalism focus leaves substantive constitutional questions for particular duties, actors and asserted rights.

¹⁶ Colorado SB 24-205 (2024 Reg. Sess.), former Colo. Rev. Stat. §§ 6-1-1702 to -1703 (developer and deployer duties); Colorado SB 25B-004 (2025 Extraordinary Sess.); Colorado SB 26-189 (2026 Reg. Sess.). See *infra* notes 87–88.

¹⁷ Virginia HB 2094 (2025 Reg. Sess.), enrolled text; Governor of Virginia, Veto Message for HB 2094 (Mar. 24, 2025). The proposal and the objections expressed in the veto record are distinct from an enacted duty or a judicial ruling.

¹⁸ Colorado SB 26-189 (2026 Reg. Sess.), reenacted Colo. Rev. Stat. §§ 6-1-1702 to -1704. The developer-documentation, deployer-disclosure and consumer-procedure provisions have their own triggers, exceptions and timing. See *infra* note 88.

¹⁹ See *infra* Part III.A (decision, actor and deployer-centered design).

²⁰ U.S. Const. art. I, § 8, cl. 3; *Pike v. Bruce Church, Inc.*, 397 U.S. 137, 142 (1970); *Nat’l Pork Producers Council v. Ross*, 598 U.S. 356, 371–76 (2023) (opinion of the Court rejecting the proposed broad extraterritoriality rule); *id.* at 380–83 (opinion of Gorsuch, J., joined by Thomas and Barrett, JJ.); *id.* at 391–93 (Sotomayor, J., joined by Kagan, J., concurring in part). The opinions differ on balancing; the three-Justice discussion of incommensurability is not a holding of the Court. See *infra* Part II.A.

the decision's result nor a general preference for local accountability establishes the validity of every deployer duty. A requirement to retain an existing decision record and a requirement to redesign a nationally supplied system need comparison on their actual facts, and cannot be replaced by assignment to ranked categories.

Preemption asks a different set of questions: what federal law governs, which requirement or remedy it supplies, and how the state obligation interacts with it. Traditional state concern informs the analysis in appropriate settings; it does not override express language or resolve an actual conflict.²¹ State police powers provide an affirmative basis for legislation, while anti-commandeering protects a different constitutional interest. A state may remain free to choose its institutions even when federal law displaces a duty those institutions would otherwise enforce.²²

The first contribution is this connection between doctrine and design. Compare the same proposed duty, actors and facts across available legal arrangements. Specify its territorial basis, the evidence of its benefits and burdens, the federal overlap and the remedy sought. Constitutional constraints can then inform drafting choices without being made to select a single design in advance.

D. Regulatory Resilience Requires a Usable Portfolio

Regulatory resilience describes a strategy of maintaining accountability through several state authority channels, each subject to the federal law that governs its use.²³ Consumer protection, civil rights, licensing and procurement reach different conduct and provide different institutions and remedies. Their diversity creates options, but an option matters only if it can perform the legal work that is needed.

A consumer-protection action must satisfy the relevant statute's elements and enforcement rules; generally applicable form alone supplies no preemption exemption.²⁴ Civil-rights routes depend on the protected interest, covered actor, cause of action and relief, with distinct rules for public enforcement and private claims.²⁵ Banking, insurance and healthcare licensing sit within different federal arrangements.²⁶ Procurement supplies bargaining and oversight opportunities when the state buys or uses AI, subject to downstream, federal-conflict and other limits; its executive-policy treatment is separate from its constitutional treatment.²⁷

²¹ *Rice v. Santa Fe Elevator Corp.*, 331 U.S. 218, 230 (1947); *Medtronic, Inc. v. Lohr*, 518 U.S. 470, 485 (1996); *Wyeth v. Levine*, 555 U.S. 555, 565 (2009). The inquiry remains tied to congressional purpose and the particular federal requirements. See *infra* Part II.B.

²² *Hillsborough County v. Automated Med. Labs., Inc.*, 471 U.S. 707, 718 (1985); *Printz v. United States*, 521 U.S. 898, 935 (1997); *Murphy v. NCAA*, 584 U.S. 453, 477–81 (2018). *Murphy* distinguishes valid preemption governing private conduct from commands to state legislatures. See *infra* Part II.C.

²³ See *infra* Part IV. Compare *Morales v. Trans World Airlines, Inc.*, 504 U.S. 374, 384–86 (1992) (rejecting a general-applicability exception to the express preemption provision there). The portfolio recommendation depends on what each channel can lawfully accomplish.

²⁴ See *infra* Part IV.A (statutory elements, enforcement and federal overlap in consumer protection).

²⁵ See *infra* Part IV.B. *Cuomo v. Clearing House Ass'n, L.L.C.*, 557 U.S. 519, 535–36 (2009), distinguishes judicial enforcement of nonpreempted state law from visitorial supervision; 12 U.S.C. § 25b supplies further rules for national-bank preemption and enforcement. Availability of an enforcement forum does not establish that the underlying state command survives.

²⁶ See *infra* Part IV.C (banking, insurance and healthcare licensing under their distinct federal regimes).

²⁷ See *infra* Part IV.D. *Reeves, Inc. v. Stake*, 447 U.S. 429, 436–39 (1980), addresses market participation; *Crosby v. National Foreign Trade Council*, 530 U.S. 363, 373–74 (2000), supplies a conflict-preemption

The portfolio must therefore be tested against a common task. If one channel cannot require a hospital to provide a particular record, can another law reach the same actor and obtain that record for the same person? A surviving licensing authority is not a substitute for an unavailable damages claim merely because both concern healthcare. Authority, duty, enforceability, relief and better outcomes are separate achievements.

Coalition cost names the proposed political mechanism within resilience. Displacing several established arrangements may require decisions that a narrow preemption measure would avoid. That possibility gives a legislature a reason to consider the breadth of its institutional portfolio. It is not an observed measure of legislative difficulty, a prediction about Congress or a requirement that Congress repeal every underlying authority before displacing AI-specific applications. A single sufficiently broad federal provision may reach several channels. The positive recommendation is to build and test available routes. Their parent statutes can remain in force while a particular duty is displaced.

E. Frameworks Need Legal Specification and Evidence

The NIST AI Risk Management Framework organizes risk work through GOVERN, MAP, MEASURE and MANAGE. Its voluntary guidance supplies a useful vocabulary for documentation, assessment, monitoring and responses to risk.²⁸ ISO/IEC 42001 concerns an organizational AI management system, a different object from the legality or accuracy of an individual decision.²⁹ Neither reference relieves a legislature of deciding what a covered organization must do.

Legal translation requires that decision. The law must specify the regulated actor, required conduct, evidence, reviewer and consequence of noncompliance. Federal cybersecurity procurement illustrates how a technical reference can acquire force through contracts and a defined assessment regime. CMMC's particular levels, assessment methods and contractual predicates matter; certification is not a universal condition attached to every defense acquisition or every system.³⁰

Existing state procurement controls also provide antecedents for AI governance. The contribution here is the worked statutory arrangement in Parts V and the Appendix: procurement-first controls, impact assessment, audit trails and recourse, with explicit connections among duty, evidence and enforcement. A broad reference to a framework cannot settle supplier access, trade-secret protection, burdens on smaller organizations or meaningful review. The state must specify those choices. A fixed technical reference and its later revision raise different questions of legislative control and notice.

The recommendation carries practical objections. Assessments can become ceremonial; credentials can concentrate markets; and legal compliance can reward documented procedures without changing decisions. Evidence of organizational systems must therefore be connected to the

limit in state purchasing. Executive Order 14365 § 8(b)(iii) concerns the executive branch's legislative recommendation. *See supra* notes 3–4.

²⁸ NIST, *AI Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, at 1–5, 20–32 (Jan. 2023); 15 U.S.C. § 278h-1(c). GOVERN, MAP, MEASURE and MANAGE organize voluntary risk work; a legal duty requires its own specification.

²⁹ ISO/IEC 42001:2023, *Information technology — Artificial intelligence — Management system* (Dec. 2023); ISO, *ISO/IEC 42001:2023: AI Management Systems*, <https://www.iso.org/standard/42001> (describing the organizational management system and continual-improvement methodology).

³⁰ 32 C.F.R. pt. 170; 48 C.F.R. §§ 204.7504, 252.204-7021. The CMMC level, required assessment and phased contract application govern the actual obligation. *See infra* Part V.A; Appendix commentary following § 1.

disputed use, while affected persons and regulators retain ways to test it. Statutory translation provides legal machinery for that inquiry. It does not establish its cost or effectiveness.³¹

F. Accountability Requires Power to Question the Decision

The throughline is the asymmetry between the organization using AI and the person subject to its decision. An applicant may see only a rejection. A patient may receive a denial whose reasoning and inputs are held elsewhere. Accountability requires more than identifying which organization used the system: it requires evidence that can be obtained, someone able to examine it, and a consequence when a duty is breached. The six process duties seek to make that power usable without assuming that a completed assessment or an explanation supplies relief.³²

The paper moves from constitutional analysis (II), through the decision, existing instruments, domains and obligations (III), to the authority portfolio (IV). Implementation and sector applications follow in V, with worked text in the Appendix and next steps in VI. The recommendation is concrete: attach responsibility to consequential uses of AI, then provide the evidence, institutions and remedies that responsibility needs.

Part II: The Constitutional Architecture for State AI Regulation

A. The Dormant Commerce Clause as Design Constraint

1. Discrimination and *Pike*: The Questions a State Duty Must Answer

An AI governance duty needs a constitutional account of what it requires, whom it reaches, and what burden it places on interstate commerce. *Pike v. Bruce Church, Inc.* states the familiar rule for an evenhanded law serving a legitimate local interest with only incidental effects on interstate commerce: it stands unless the burden is “clearly excessive in relation to the putative local benefits.” The nature of the interest and the possibility of achieving it with less impact on interstate activity matter to that inquiry.³³ An in-state duty bearer and a legitimate state purpose identify parts of the comparison. They do not decide it. Arizona lost in *Pike* even though its order applied to a grower operating in the State and pursued the reputation of Arizona produce.

National Pork Producers Council v. Ross preserves that distinction between a legal question and its answer. The Court treated antidiscrimination as central to the Dormant Commerce

³¹ See *infra* Parts V.B–D and the Appendix explanations (incorporation authority, costs, evidence, institutions and assessment design). NIST’s voluntary mandate, 15 U.S.C. § 278h-1(c), does not itself settle a state’s authority to impose a particular duty. The objections in the text are design questions; the proposed arrangement’s effectiveness requires evidence.

³² See *infra* Part III.D (the six obligations), Part V.D (recourse) and Appendix §§ 2–3 and their explanations. The intended connection among information, review and power to act must be supplied through the operative duties and institutions.

³³ *Pike v. Bruce Church, Inc.*, 397 U.S. 137, 142–46 (1970) (general rule; legitimate local interest; application requiring an unnecessary in-state packing facility), <https://www.govinfo.gov/content/pkg/USREPORTS-397/pdf/USREPORTS-397-137.pdf>.

Clause and left a route for challenges to genuinely nondiscriminatory burdens.³⁴ It rejected the producers' proposed "almost *per se*" rule against laws with the practical effect of controlling commerce outside the State.³⁵ Its common opinion also reaffirmed that a State may not discriminate purposefully against out-of-state economic interests.³⁶ The Justices divided on why the pleaded *Pike* claim failed. Justice Gorsuch's warning against a judicial "roving license" belongs to the three-Justice portions identified in Section II.A.3, which also explains the separate opinions and their different grounds.³⁷ Neither the rejection of an effects-alone rule nor the judgment on that complaint determines the validity of an AI assessment, disclosure, or design requirement.

The older transportation cases make the burden question concrete. In *Southern Pacific Co. v. Arizona*, the Court invalidated a limit on the length of interstate trains.³⁸ A railroad could break and remake trains at the border or conform a longer route to Arizona's limit. The record included the extra trains, equipment, labor, and operating costs that compliance would require.³⁹ The Court examined the nature and extent of the interference with interstate operation against the safety interest, which it found slight and doubtful on that record.⁴⁰ The case gives an AI legislator a reason to examine the actual consequences of interacting state commands. It supplies no measurement of those consequences for AI infrastructure.

A safety label likewise supplies no answer. *Kassel v. Consolidated Freightways Corp.* invalidated Iowa's restriction on sixty-five-foot double-trailer trucks.⁴¹ Justice Powell's plurality weighed the safety justification and the burden, drawing on *Raymond Motor Transportation's* recognition that genuine safety judgments deserve respect.⁴² The plurality found Iowa's asserted safety benefit illusory and described the routing and trailer-separation costs imposed on carriers.⁴³ Consumer protection, antidiscrimination, and safety remain reasons to regulate AI. Their invocation must be connected to the benefit of the actual duty and the burden it creates.

I therefore use the doctrine as a design discipline. A legislator should define the covered activity, identify the local harm, examine discriminatory treatment and territorial reach, and specify

³⁴ *Nat'l Pork Producers Council v. Ross*, 598 U.S. 356, 377–80 & n.2 (2023) (Part IV-A, opinion of the Court; antidiscrimination and nondiscriminatory burdens); see *infra* Section II.A.3.

³⁵ *Pork Producers*, 598 U.S. at 371–77 & n.1 (Part III, opinion of the Court; proposed effects-alone rule and reserved territorial questions).

³⁶ *Id.* at 364, 369–71 (opinion of the Court).

³⁷ *Pork Producers*, 598 U.S. at 361 (opinion-part lineup), 380–83, 387–89 (Gorsuch, J., joined by Thomas and Barrett, JJ.); *United Haulers Ass'n v. Oneida-Herkimer Solid Waste Mgmt. Auth.*, 550 U.S. 330, 343 (2007). Section II.A.3 distinguishes this reasoning from the Court's common portions and the other opinions.

³⁸ *S. Pac. Co. v. Arizona ex rel. Sullivan*, 325 U.S. 761, 763–64, 781–83 (1945), <https://www.govinfo.gov/content/pkg/USREPORTS-325/pdf/USREPORTS-325-761.pdf>.

³⁹ *Id.* at 771–76 (operating costs, border adjustment, route-wide conformity, and interference with interstate operation).

⁴⁰ *Id.* at 770–71, 779, 781–82 (burden inquiry and safety evidence). The application to an AI duty is my inference, not a holding about AI.

⁴¹ *Kassel v. Consol. Freightways Corp.*, 450 U.S. 662, 665–67, 679 (1981) (judgment; Powell, J., plurality opinion).

⁴² *Id.* at 670–71 (plurality opinion); *Raymond Motor Transp., Inc. v. Rice*, 434 U.S. 429, 441–44 (1978) (opinion of the Court), 449–50 (Blackmun, J., concurring). The opinions' deference and burden statements are not a categorical exemption for laws described as safety measures.

⁴³ *Kassel*, 450 U.S. at 671–74 (plurality opinion). The AI comparison preserves the need for evidence of the actual benefit and burden; it does not equate all documentation costs with a restriction on interstate trucks.

the evidence needed to assess the interstate burden. I favor duties directed at the decisions for which a deployer is accountable because that choice identifies an actor, an affected person, and a regulatory task. This preference informs the design developed in Part III. Whether a particular version satisfies the Dormant Commerce Clause remains a duty-specific inquiry, including when the duty is imposed on a developer.

2. Territorial Limits After *Pork Producers*

Healy v. Beer Institute used broad territorial language while invalidating Connecticut's beer price-affirmation law. Its opinion said that a statute directly controlling commerce wholly outside the State exceeds the State's authority, and directed attention to the regulation's practical effect.⁴⁴ It also examined how the law interacted with other States' legitimate rules. Connecticut's linkage of its prices to neighboring prices could prevent competitive pricing elsewhere and produce interlocking restrictions if other States adopted similar laws.⁴⁵ Five Justices joined that reasoning; Justice Scalia joined only Parts I and IV. The historical language cannot be dismissed as a plurality simply because an earlier case had fewer votes.

That earlier line included *Baldwin v. G.A.F. Seelig, Inc.*, which invalidated New York's restriction on sales of milk bought outside the State below New York's minimum producer price.⁴⁶ In *Edgar v. MITE Corp.*, Justice White's Part V-A objected to Illinois's direct regulation of transactions wholly outside the State, including transactions involving persons with no connection to Illinois.⁴⁷ Its analogy to the territorial limits on state judicial power drew four votes.⁴⁸ Part V-B, however, separately applied *Pike* and had five, including Justice Powell. Treating all of Part V as a plurality would obscure a ground on which the Court invalidated the Illinois law.⁴⁹ *CTS Corp. v. Dynamics Corp. of America* later upheld an Indiana voting-rights law for corporations the State had created, distinguishing the Illinois law's reach.⁵⁰

Brown-Forman Distillers Corp. v. New York State Liquor Authority and *Healy* involved a particular mechanism: a State linked the prices charged within its borders to prices charged elsewhere.⁵¹ In *Healy*, that linkage operated alongside other States' pricing rules and constrained transactions beyond Connecticut.⁵² *Pork Producers* reads this lineage in its discrimination context and rejects the proposed general rule against extrastate effects alone. Its common opinion

⁴⁴ *Healy v. Beer Inst., Inc.*, 491 U.S. 324, 325, 336–37 (1989) (opinion-part lineup; territorial language), <https://www.govinfo.gov/content/pkg/USREPORTS-491/pdf/USREPORTS-491-324.pdf>.

⁴⁵ *Id.* at 336–40 (practical effect and interaction with other States' rules). The application was to the price-affirmation mechanism; *Pork Producers*, 598 U.S. at 371–77, supplies the later contextual reading.

⁴⁶ *Baldwin v. G.A.F. Seelig, Inc.*, 294 U.S. 511, 519–24, 527–28 (1935).

⁴⁷ *Edgar v. MITE Corp.*, 457 U.S. 624, 641–43 (1982) (Part V-A, opinion of White, J., joined by Burger, C.J., and Stevens and O'Connor, JJ.).

⁴⁸ *Id.* at 643 (Part V-A); *see id.* at 626 n.† (opinion-part joins). This reports the plurality's analogy rather than treating it as an independent jurisdictional test for AI duties.

⁴⁹ *Id.* at 626 n.†, 643–46 (Part V-B, opinion of the Court, joined by Burger, C.J., and Powell, Stevens and O'Connor, JJ.), 646–47 (Powell, J., concurring in part). The four-Justice territorial rationale and five-Justice *Pike* rationale are distinct.

⁵⁰ *CTS Corp. v. Dynamics Corp. of Am.*, 481 U.S. 69, 88–94 (1987). Section II.A.5 explains the limit of the single-incorporation analogy.

⁵¹ *Brown-Forman Distillers Corp. v. N.Y. State Liquor Auth.*, 476 U.S. 573, 575–76, 579–84 (1986).

⁵² *Healy*, 491 U.S. at 337–40. The interacting-law analysis does not establish that any shared compliance system is a constitutionally prohibited price linkage.

also reserves territorial questions under due process, full faith and credit, sovereignty, and comity, and distinguishes the conditions described by the *MITE* plurality. A requirement addressed to an in-state actor needs a territorial account; a requirement addressed to a supplier outside the State is not invalid merely because compliance affects that supplier's work elsewhere.

The internet cases show why the factual and procedural setting matters. In *American Libraries Ass'n v. Pataki*, a district court preliminarily enjoined New York's criminal restriction on internet communications harmful to minors. Its Commerce Clause reasoning addressed extraterritorial application, the relation of benefits to burdens, and inconsistent state regulation on the record before it.⁵³ *American Booksellers Ass'n v. Hudnut* supplies a different constitutional example: the Seventh Circuit held Indianapolis's pornography ordinance unconstitutional under the First Amendment.⁵⁴ I do not derive a Dormant Commerce Clause rule for AI systems from that speech holding. Nor do the communications conditions described in a 1997 preliminary-injunction decision establish the technical constraints on an AI service today. *Pork Producers* requires that the practical-effects language in the earlier Commerce Clause cases be read in context.⁵⁵ The constitutional account must follow the actual command and record, including the role that distributing information across state lines plays in that record.

3. *Pork Producers*: The Effects-Alone Rule Rejected and the *Pike* Reasoning Divided

National Pork Producers Council v. Ross affirmed the dismissal, for failure to state a claim, of a Dormant Commerce Clause challenge to California's Proposition 12.⁵⁶ The proposition forbids the in-state sale of whole pork meat from breeding pigs, or their immediate offspring, when the breeding pig was confined in a manner the statute defines as cruel. The producers alleged that most of the resulting compliance costs would fall first on out-of-state firms, and they disclaimed any claim of discrimination.⁵⁷ They pressed two other theories: an "almost *per se*" rule against state laws with the "practical effect of controlling commerce outside the State," and a claim under *Pike v. Bruce Church, Inc.*⁵⁸ Five Justices rejected the first in an opinion of the Court.

⁵³ *Am. Libraries Ass'n v. Pataki*, 969 F. Supp. 160, 169–77, 181–84 (S.D.N.Y. 1997) (granting a preliminary injunction; Commerce Clause grounds and technological record), <https://law.justia.com/cases/federal/district-courts/FSupp/969/160/1808497/>. This is a historical district-court application, not a current AI holding or a survey of later treatment.

⁵⁴ *Am. Booksellers Ass'n v. Hudnut*, 771 F.2d 323 (7th Cir. 1985) (Easterbrook, J.) (First Amendment holding concerning viewpoint discrimination), <https://law.justia.com/cases/federal/appellate-courts/F2/771/323/379919/>.

⁵⁵ *Pork Producers*, 598 U.S. at 371–77 & n.1 (context of *Baldwin*, *Brown-Forman*, and *Healy*; rejection of proposed effects-alone rule and territorial reservation).

⁵⁶ *Nat'l Pork Producers Council v. Ross*, 598 U.S. 356, 361, 363–64, 391 (2023) (reporter's listing of opinions and joins; Gorsuch, J., announcing the judgment and delivering the opinion of the Court except as to Parts IV-B, IV-C, and IV-D; affirmance), https://www.supremecourt.gov/opinions/22pdf/598us2r18_g2bh.pdf (preliminary print). Later citations to *Pork Producers* are to the same print.

⁵⁷ *Pork Producers*, 598 U.S. at 364 (dismissal for failure to state a claim; affirmance), 365–66 (Proposition 12's sale prohibition and definition of cruel confinement), 367 (as alleged, compliance costs fall on California and out-of-state producers alike, with the majority initially borne by out-of-state firms), 370–71 (disavowal of any discrimination claim; concession that Proposition 12 imposes the same burdens on in-state and out-of-state producers).

⁵⁸ *Pike v. Bruce Church, Inc.*, 397 U.S. 137, 142 (1970), <https://www.govinfo.gov/content/pkg/USREPORTS-397/pdf/USREPORTS-397-137.pdf>; see *Pork Producers*, 598 U.S. at 371 (producers' "almost *per se*" theory and "practical effect of controlling commerce outside the State"), 377 (producers' *Pike* theory).

On the second, the Justices wrote in groups whose reasons overlap without adding up.

The opinion of the Court read the cases the producers invoked, *Baldwin v. G.A.F. Seelig, Inc.*, *Brown-Forman Distillers Corp. v. New York State Liquor Authority*, and *Healy v. Beer Institute*, as applications of the rule against purposeful discrimination in favor of in-state economic interests. *Baldwin* involved a New York law barring the sale of out-of-state milk unless its producers had been paid the New York minimum price; *Brown-Forman* and *Healy* involved laws requiring sellers to affirm that their in-state prices were no higher than their prices elsewhere. Language about practical effects outside the State did its work in that setting, and a rule reaching every such effect would sweep in tax, environmental, and many other state laws. The Chief Justice, for himself and Justices Alito, Kavanaugh, and Jackson, agreed that precedent does not support the proposed rule.⁵⁹ The Court did not say that a State's power has no territorial limit. It acknowledged that courts sometimes referee where one State's authority ends and another's begins, and it pointed to the Due Process Clause, the Full Faith and Credit Clause, and principles of sovereignty and comity. The Court distinguished *Edgar v. MITE Corp.*, in which a plurality addressed an Illinois law that directly regulated transactions wholly outside the State among persons with no connection to it. The producers did not allege those conditions here, and the Court did not settle the source of *Edgar's* territorial limit.⁶⁰ Part III therefore removes practical effect outside the State, standing alone, as a ground of invalidity. It leaves the reserved territorial limits, and it leaves *Pike*, where four Justices would count sweeping extraterritorial effects as pertinent.

On *Pike*, the opinion of the Court said comparatively little. *Pike* upholds an evenhanded law that serves a legitimate local interest and affects interstate commerce only incidentally unless the burden is "clearly excessive in relation to the putative local benefits." Part IV-A placed antidiscrimination at the core of that line, reading several *Pike* cases as turning in whole or in part on a law's discriminatory character. It also recognized that the Court has left the "courtroom door open" to nondiscriminatory burdens and that a small number of cases have invalidated laws that appear genuinely nondiscriminatory. It placed the producers' claim well outside *Pike's* heartland without holding the claim barred for that reason. Part V adds that, because the Court has inferred this limit, a court should act with extreme caution before preventing enforcement of a democratically adopted state law.⁶¹

Beyond that point the opinions separate. In Parts IV-B and IV-D, Justice Gorsuch, joined by Justices Thomas and Barrett, reasoned that no neutral legal rule guides a comparison of the

⁵⁹ *Pork Producers*, 598 U.S. at 371–74 (reading *Baldwin v. G.A.F. Seelig, Inc.*, 294 U.S. 511 (1935), *Brown-Forman Distillers Corp. v. N.Y. State Liquor Auth.*, 476 U.S. 573 (1986), and *Healy v. Beer Inst.*, 491 U.S. 324 (1989), as discrimination cases), 374 (earlier description of the rule as addressing price control or price affirmation statutes), 374–75 (many state laws affect conduct beyond the State; proposed rule would invite litigation); *id.* at 394 (Roberts, C.J., joined by Alito, Kavanaugh, and Jackson, JJ., concurring in part and dissenting in part) (agreeing that precedent does not support a per se rule against laws with extraterritorial effects).

⁶⁰ *Pork Producers*, 598 U.S. at 375–76 (territorial limits; Due Process and Full Faith and Credit Clauses; sovereignty and comity), 376–77 n.1 (distinguishing *Edgar v. MITE Corp.*, 457 U.S. 624 (1982), without settling the constitutional source of its territorial limit; conditions not alleged here); *id.* at 400 (Roberts, C.J., joined by Alito, Kavanaugh, and Jackson, JJ.) (sweeping extraterritorial effects "pertinent in applying *Pike*").

⁶¹ *Pike*, 397 U.S. at 142 (general rule; nature of the local interest; lesser impact on interstate activities); *Pork Producers*, 598 U.S. at 377–80 & n.2 (Part IV-A, opinion of the Court), 379 ("courtroom door open," quoting *Dept't of Revenue of Ky. v. Davis*, 553 U.S. 328, 353 (2008)), 379–80 (producers' claim outside *Pike's* heartland), 389–91 (Part V, opinion of the Court); *see id.* at 392 (Sotomayor, J., joined by Kagan, J., concurring in part) (Court does not shut the door on such claims); *id.* at 395 (Roberts, C.J., joined by Alito, Kavanaugh, and Jackson, JJ.) (same).

producers' higher production costs with California voters' moral and health interests, that the choice belongs to elected officials and to Congress, and that *Pike* counts burdens on commerce and not freestanding harms to out-of-state animal welfare.⁶² Part IV-C, joined by Justices Thomas, Sotomayor, and Kagan, rests on pleading. Relying on *Exxon Corp. v. Governor of Maryland*, it concluded that the complaint alleged harm to some producers' preferred methods of operation and a possible shift of California market share among out-of-state firms, and did not plausibly allege a substantial harm to the interstate market. It added that further experience might yield further facts.⁶³ The separate opinions show why neither rationale commands five votes. Justice Sotomayor, joined by Justice Kagan, joined Part IV-C but not Parts IV-B and IV-D, and wrote that courts can generally weigh disparate burdens and benefits and that a substantial burden is a threshold the plaintiff must allege first.⁶⁴ Justice Barrett joined Parts IV-B and IV-D and agreed that Proposition 12's benefits and burdens could not be compared, but she thought the complaint plausibly alleged a substantial burden and wrote that she would let the *Pike* claim proceed if balancing were possible. The Chief Justice, for himself and Justices Alito, Kavanaugh, and Jackson, would have found a substantial burden plausibly alleged, vacated the judgment, and remanded for the Ninth Circuit to decide whether the complaint plausibly alleged a burden clearly excessive in relation to the putative local benefits.⁶⁵

Counted proposition by proposition, six Justices (Justices Sotomayor and Kagan and the Chief Justice's four) accept that courts can generally weigh the burdens and benefits that *Pike* involves. Five (the Chief Justice's four and Justice Barrett) found a substantial burden plausibly alleged; four (Justices Gorsuch, Thomas, Sotomayor, and Kagan) did not. Only the Chief Justice's four would have let the claim go forward, because Justice Barrett's willingness depended on a condition she thought unmet for this statute.⁶⁶ These counts describe agreement on separate propositions. They do not supply an opinion of the Court explaining why this *Pike* claim fails, and I do not derive one from them. The narrower fact is that, as to the *Pike* claim, the five-vote judgment of affirmance rests on two reasons and neither commands five votes. A footnote to Part IV-D says the decision depends equally on both, and Justice Kavanaugh's separate opinion calls Part IV-C "controlling precedent for purposes of the Court's judgment as to the plaintiffs' *Pike* claim" while treating Parts IV-B and IV-D as not controlling, as he understands

⁶² *Pork Producers*, 598 U.S. at 380–83 (Part IV-B), 387–89 (Part IV-D); see *id.* at 361 (reporter's listing: Parts IV-B and IV-D, Gorsuch, J., joined by Thomas and Barrett, JJ.).

⁶³ *Pork Producers*, 598 U.S. at 383–87 (Part IV-C, applying *Exxon Corp. v. Governor of Md.*, 437 U.S. 117 (1978); Gorsuch, J., joined by Thomas, Sotomayor, and Kagan, JJ., see *id.* at 361), 385 (pleading standard: facts "plausibly" suggesting substantial harm), 386 (costs borne by in-state consumers; further facts), 387 (harm to some producers' methods of operation; substantial harm to interstate commerce a speculative possibility).

⁶⁴ *Pork Producers*, 598 U.S. at 391–93 (Sotomayor, J., joined by Kagan, J., concurring in part); *id.* at 396–97 (Roberts, C.J., joined by Alito, Kavanaugh, and Jackson, JJ.) (weighing possible); *id.* at 403 (Kavanaugh, J., concurring in part and dissenting in part) (six Justices retain *Pike* balancing).

⁶⁵ *Pork Producers*, 598 U.S. at 393–94 (Barrett, J., concurring in part); *id.* at 395, 397, 403 (Roberts, C.J., joined by Alito, Kavanaugh, and Jackson, JJ., concurring in part and dissenting in part) (substantial burden plausibly alleged; vacatur; remand to consider whether petitioners plausibly claimed a burden outweighing, or clearly excessive in relation to, the putative local benefits).

⁶⁶ *Pork Producers*, 598 U.S. at 361 (reporter's listing of opinions and joins), 391–93 (Sotomayor, J., joined by Kagan, J., concurring in part), 393–94 (Barrett, J., concurring in part), 394–403 (Roberts, C.J., joined by Alito, Kavanaugh, and Jackson, JJ., concurring in part and dissenting in part), 403 (Kavanaugh, J., concurring in part and dissenting in part). The groupings in the text are proposition-level counts drawn from these opinions: six on general weighability (Sotomayor, Kagan, Roberts, Alito, Kavanaugh, Jackson), five on a plausibly alleged substantial burden (Roberts, Alito, Kavanaugh, Jackson, Barrett), and four on the opposite pleading conclusion (Gorsuch, Thomas, Sotomayor, Kagan). They are not an independent *Marks* ruling.

it.⁶⁷ Each characterization belongs to the Justices who made it.

A local actor, a traditional state interest, or a regulatory label is no safe harbor, and *Pike* itself shows why. Arizona ordered that cantaloupes grown in the State be packed there, under a standardization statute whose stipulated primary purpose was to protect the reputation of Arizona growers. The Court treated reputation as a legitimate interest and assumed the statute valid as applied to growers who pack in Arizona. It held the order invalid as applied to a company that grew cantaloupes in Arizona and packed them in California. The State's tenuous interest could not justify requiring the company to build and operate an unneeded packing plant costing about \$200,000, and the Court regarded the nature of that burden, compelled in-state performance of operations that could be done more efficiently elsewhere, as more significant than its extent. Arizona argued that its order was regulatory and carried no protectionist purpose. The Court, relying on an earlier shrimp-packing case, indicated that such a burden can be unconstitutional even without that purpose, while allowing that a more compelling state interest might make the consequence tolerable; Arizona's interest was minimal at best. *Pork Producers* later read the order's practical effects as disclosing a discriminatory purpose. The case does not make local process duties suspect. It shows that the balance turns on the nature of the interest and of the burden, and a local framing fixes neither.⁶⁸

The strongest objection comes from the decision's own result. Proposition 12 conditioned sale in California, yet Part IV-C described the producers' choices as housing pigs to the required standard, segregating compliant product, or leaving the California market. The Chief Justice described allegations of an integrated national market in which farmers would be "forced to comply" whether or not they sell in California, and noted that the Ninth Circuit panel recognized allegations of "dramatic upstream effects."⁶⁹ The Court affirmed dismissal anyway. A duty confined to an in-state deployer's own decisions about the State's residents appears to reach less far from the regulated actor than that sales condition. If the sales condition survived, the objection concludes, caution about *Pike* is overdrawn.

The objection is right about Part III and overreaches on *Pike*. Part III establishes that upstream effects, standing alone, do not invalidate an in-state condition. It does not decide that such a condition satisfies *Pike*. The affirmance concerned the sufficiency of one complaint: the four-Justice pleading rationale turned on that complaint's allegations and noted that further experience might yield further facts, the three-Justice rationale turned on a comparison between production costs and moral and health interests, and five Justices thought a substantial burden plausibly alleged. The comparison is also imperfect. Proposition 12 regulated what could be

⁶⁷ *Pork Producers*, 598 U.S. at 389 n.4 (opinion of Gorsuch, J., joined by Thomas and Barrett, JJ.) (Part IV-D) (decision "depends equally" on the analysis in Parts IV-B and IV-C); *id.* at 403 (Kavanaugh, J., concurring in part and dissenting in part) (his statement that Part IV-C is "controlling precedent for purposes of the Court's judgment as to the plaintiffs' *Pike* claim" and that Parts IV-B and IV-D are not controlling, "as I understand it").

⁶⁸ *Pike*, 397 U.S. at 143 (stipulated primary purpose; reputation a legitimate interest; validity assumed as applied to growers who pack in Arizona), 144 (application to the appellee; capital expenditure of approximately \$200,000), 145 (tenuous interest; nature of burden more significant than extent; particular suspicion of in-state performance requirements; Arizona's argument that its order was regulatory; relying on *Toomer v. Witsell*), 145–46 (burden unconstitutional even without protectionist purpose; more compelling interest; interest minimal at best), 146 (judgment affirmed); *Pork Producers*, 598 U.S. at 377–78 (reading the order's practical effects as disclosing a discriminatory purpose).

⁶⁹ *Pork Producers*, 598 U.S. at 364–67 (case and alleged costs), 384 (Part IV-C: producers' choices to house, segregate, or withdraw), 386 (further facts); *id.* at 400 (Roberts, C.J., joined by Alito, Kavanaugh, and Jackson, JJ.) (complaint alleges an interconnected market in which producers will be "forced to comply" whether or not they sell in California), 402 (panel recognized allegations of "dramatic upstream effects").

sold in California, while a deployer duty regulates how an in-state actor makes AI-mediated decisions about residents, and that difference may change the burden and the benefit to be compared. The opinions cited here do not address it. The objection therefore supports a narrower conclusion than it claims. Extrastate practical effect is not itself the defect, and the survival of Proposition 12 at the pleading stage does not decide whether a differently pleaded or differently structured duty clears the threshold or the balance.

For a specified AI duty, the decision supports the following sequence of questions. The sequence is my inference, not a holding, and the questions have unequal support. Take a duty requiring a deployer operating in the State to complete and retain an impact assessment before making an AI-mediated decision about residents of the State. A reviewer first asks whether the duty discriminates against out-of-state interests in text, purpose, or practical effect; if it does, it lies in the antidiscrimination core that Part IV-A identified. If the challenge rests only on the duty's practical effect on how firms operate elsewhere, Part III answers it as stated. What remain are the territorial limits the Court reserved, including the concern about directly regulating transactions wholly outside the State among persons with no connection to it. The duty's definition of covered decisions bears on that question. The Chief Justice's four-Justice opinion also treated extrastate effects as pertinent to a *Pike* burden. Framing the duty around a local actor does not remove it from these questions. If the duty is nondiscriminatory, a *Pike* claim remains available in principle, and the Justices answered three further questions differently: whether a pleaded burden is a harm to the interstate market or only to particular firms' methods of operation; whether market-wide consequences beyond compliance cost count as burden; and whether the burdens and benefits can be compared at all.

A further inference is that a duty's stated benefit matters at the comparison step. The comparability rationale rested on a comparison of production costs with moral and health interests, and the cited opinions do not say how it applies to a duty whose stated benefit is documented review of decisions about the State's residents. The decision does not tell a court how to resolve the threshold or the balance for such a duty, and it does not decide the outcome for a different duty on a different record. That outcome depends on the burden pleaded and proved and on the local benefit the statute is shown to serve. If several States adopt differing duties over the same decision infrastructure, *Pork Producers*, which concerned one State's statute, does not decide whether or how their combined burden counts.⁷⁰

The Ninth Circuit's unpublished memorandum in *Iowa Pork Producers Ass'n v. Bonta* illustrates one subsequent application to Proposition 12.⁷¹ The panel rejected the discrimination claim and treated its own earlier *National Pork Producers Council* decision as controlling under the circuit's approach to fractured Supreme Court decisions. It then rejected the *Pike* claim under that circuit precedent.⁷² The memorandum is not precedent except as provided by Ninth Circuit Rule 36-3.⁷³ The Supreme Court denied certiorari on June 30, 2025; Justice Kavanaugh

⁷⁰ *Pork Producers*, 598 U.S. at 371–76, 376–77 n.1, 377–80, 381–82 (moral and health interests compared with production costs), 383–87, 397–402. The application to the duty described in the text is my inference; no cited opinion addresses an AI-mediated decision.

⁷¹ *Iowa Pork Producers Ass'n v. Bonta*, No. 22-55336 (9th Cir. June 25, 2024) (unpublished mem.), <https://cdn.ca9.uscourts.gov/datastore/memoranda/2024/06/25/22-55336.pdf>, reh'g denied (9th Cir. Aug. 6, 2024), cert. denied, No. 24-728 (U.S. June 30, 2025) (the Supreme Court's docket for No. 24-728 records the Ninth Circuit's decision date, the rehearing denial and the denial of certiorari).

⁷² *Iowa Pork Producers Ass'n v. Bonta*, No. 22-55336, mem. at 2–9 (9th Cir. June 25, 2024), <https://cdn.ca9.uscourts.gov/datastore/memoranda/2024/06/25/22-55336.pdf> (discrimination at 2–5; treatment of the fractured decision and circuit precedent at 6–8; *Pike* disposition at 8–9). This reports that panel's approach; it does not adopt a separate *Marks* analysis.

⁷³ 9th Cir. R. 36-3(a)–(b) (nonprecedential status and citation of post-2006 unpublished dispositions);

would have granted the petition.⁷⁴ That history concerns the pleaded challenge to one sales condition. It does not establish a comparative likelihood of success for challenges to AI duties.

The scholarly conversation about *Pork Producers* is ongoing and divided. Some commentators read the decision as a clarification of doctrine that was always more limited than its more aggressive readings suggested. Others read it as a transformation that materially shifts state regulatory capacity.⁷⁵ For purposes of this article, I take a moderate position consistent with the existing scholarly record: *Pork Producers* recalibrated the extraterritoriality strand and tied *Pike* more closely to antidiscrimination, but did not eliminate either strand or grant states unfettered authority to regulate technology in ways that fall on out-of-state actors. The architectural argument I develop in Parts III through V is designed to satisfy the doctrine as recalibrated, not to test its outer limits.

4. Model-Centric Duties Require Their Own Constitutional Account

A demand for training-data documentation and a demand to redesign a model impose different commands. The term model-centric identifies a regulatory focus; the legal analysis still needs the provision. California's 2024 AI Transparency Act illustrates the distinction. The original enactment imposed detection, audiovisual disclosure, and licensing duties on specified generative AI providers.⁷⁶ Its provider definition included a million-monthly-user threshold and public accessibility in California.⁷⁷ The definition is historical. The September 2026 SB 1000 amendment removes the size threshold, changes the verification and licensing provisions, and qualifies latent audiovisual disclosure by technical feasibility. The separate platform, hosting, and capture-device provisions have their own actors and later commencement dates.⁷⁸ A

Iowa Pork Producers, No. 22-55336, mem. at 1 n.* (publication designation).

⁷⁴ *Iowa Pork Producers Ass'n v. Bonta*, No. 24-728 (U.S. June 30, 2025) (certiorari denied; Justice Kavanaugh would grant), <https://www.supremecourt.gov/search.aspx?filename=/docket/docketfile/s/html/public/24-728.html>. The docket also records the June 25, 2024 decision and August 6, 2024 rehearing denial.

⁷⁵ For the emerging assessments, see Jack Goldsmith & Eugene Volokh, *The Relevance of Ross to Geolocation and the Dormant Commerce Clause*, 102 Tex. L. Rev. Online 30 (2023); Ruth Mason & Michael S. Knoll, *Bounded Extraterritoriality*, 122 Mich. L. Rev. 1623 (2024); Brannon P. Denning, *National Pork Producers Council v. Ross: Extraterritoriality Is Dead, Long Live the Dormant Commerce Clause*, 2022-2023 Cato Sup. Ct. Rev. 23; Bradley W. Joondeph, *The "Horizontal Separation of Powers" After National Pork Producers Council v. Ross*, 61 San Diego L. Rev. 45 (2024); Note, *The Dormant Commerce Clause and Moral Complicity in a National Marketplace*, 137 Harv. L. Rev. 980 (2024).

⁷⁶ California AI Transparency Act, 2024 Cal. Stat. ch. 291, § 1 (S.B. 942) (original Bus. & Prof. Code §§ 22757–22757.6), https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202320240SB942. This describes the original enactment, not the unchanged law on October 16, 2026.

⁷⁷ 2024 Cal. Stat. ch. 291, § 1 (original Bus. & Prof. Code § 22757.1(b), provider definition; §§ 22757.2–22757.3, detection, manifest and latent audiovisual disclosures, and licensee provisions). The original definition required over 1,000,000 monthly visitors or users. The text addresses image, video, and audio content; it is not a general all-text-output disclosure rule.

⁷⁸ 2025 Cal. Stat. ch. 674, §§ 1–6 (A.B. 853) (provider commencement August 2, 2026; platform and hosting provisions January 1, 2027; capture-device provision January 1, 2028), https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202520260AB853; 2026 Cal. Stat. ch. 861, §§ 2–4, 7–8 (S.B. 1000) (urgency amendment effective September 30, 2026; provider definition, verification tool, latent disclosures, licensing and exclusions), https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202520260SB1000; 2026 Cal. Stat. ch. 856, § 1 (A.B. 2713) (amending future platform provision, Bus. & Prof. Code § 22757.3.1; January 1, 2027 operation), https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202520260AB2713. Under S.B.

constitutional comparison needs those actual commands and dates before it can say what compliance reaches.

Another developer duty appears in California's AB 2013: public documentation of training data, including a high-level dataset summary for covered generative AI systems and services made available to Californians, with specified exclusions.⁷⁹ A documentation requirement is not, by its terms alone, a command to use a particular dataset or retrain a model. Part III.B discusses additional mechanisms, including Tennessee's ELVIS Act.⁸⁰ Their constitutional analysis requires their own operative provisions; the label model-centric supplies no common result. Likewise, *Hudnut's* First Amendment holding does not establish a Dormant Commerce Clause rule against state duties concerning expressive systems.⁸¹

New York City's Local Law 144 illustrates a decision-facing command. An employer or employment agency using a covered automated employment decision tool for city hiring or promotion must satisfy the law's prior bias-audit and public-summary conditions.⁸² That requirement and California's audiovisual provenance duties have different objects. A comparison between the approaches must account for those differences and identify which burdens follow from the actor regulated.

Consider a matched design question. A multistate employer buys a scoring system from a supplier based elsewhere and uses it to make hiring decisions for positions in the regulating State about that State's residents. The supplier serves other employers and the employer also hires in other States. Those are stipulated facts for the comparison, not a description of every AI service. One possible duty would require the supplier to provide the employer with validation methods, tested uses, and documented limitations for the system supplied for those hiring decisions. Another would require the employer to assess the proposed use before making those decisions, retaining the assessment for a designated reviewer. These are illustrative commands, not additions to the worked Appendix. The deployment stays the same while the duty bearer and required act change.

1000, § 22757.2(e) permits a qualifying third-party verification tool. Section 22757.3(b)(2) requires termination or specified notification within 72 hours after knowledge of an identifiable licensee's non-compliant modification; (b)(3) gives the notified licensee 96 hours to respond through compliance or cessation and a report; (b)(4) requires specified reporting to the Attorney General; (b)(5) creates no monitoring, investigation, or inquiry duty. Section 22757.5 excludes exclusively non-user-generated videogames and, before January 1, 2029, qualifying assistive technology. The comparison follows these enacted texts; implementation and later litigation are outside its scope.

⁷⁹ 2024 Cal. Stat. ch. 817, § 1 (A.B. 2013) (Cal. Civ. Code §§ 3110–3111), https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=202320240AB2013. Section 3111 covers specified releases and substantial modifications, beginning with a January 1, 2026 deadline and before later covered public availability; subsection (a) lists twelve dataset topics, and subsection (b) contains security/integrity, aircraft, and federal-only national-security, military, or defense exclusions. This is the enacted documentation requirement; no judicial validity conclusion follows.

⁸⁰ Tennessee Ensuring Likeness, Voice, and Image Security Act (ELVIS Act), 2024 Tenn. Pub. Acts ch. 588.

⁸¹ *Am. Booksellers Ass'n v. Hudnut*, 771 F.2d 323 (7th Cir. 1985) (First Amendment holding); see *supra* note 54. No Dormant Commerce Clause rule for AI duties is derived from that result. The separate Commerce Clause application in *Pataki* is discussed *supra* note 53.

⁸² N.Y.C. Local Law No. 144 (2021), § 1 (N.Y.C. Admin. Code §§ 20-870, -871(a)), <https://legistar.council.nyc.gov/View.ashx?GUID=F99584B7-57C8-469E-9637-46A0E780690E&ID=10399761&M=F> (certified enacted version). Section 20-871(a) requires a bias audit within the preceding year and a public summary of results and the tool's distribution date before covered use. Section 20-870 defines covered employment decisions by reference to screening for employment or promotion within the city; the separate resident-notice provision in § 20-871(b) is not treated as the whole coverage rule.

Both commands need a territorial connection and an account of discriminatory treatment. Both need evidence of compliance costs, effects on the interstate market, the benefit of the information or review, and less burdensome ways to obtain it. The supplier duty may require records the employer cannot otherwise obtain. The employer duty may depend on those same records. Neither the supplier's headquarters nor the location of a shared server answers whether the law legally commands changes to transactions elsewhere or whether compliance is operationally separable. The reserved territorial questions in *Pork Producers* remain distinct from its rejection of effects alone as a ground of invalidity.⁸³ Federal overlap is a further question under Section II.B; this comparison assumes no exemption from it. A record available to a reviewer also supplies no remedy by itself. The statute must specify who can require correction, who can challenge a decision, and what relief the responsible institution may provide.

Discrimination needs its own comparison. A size threshold does not become discriminatory merely because the firms it presently reaches are based outside the State. A local exemption or different burdens on similarly situated local and out-of-state firms would raise a different question. *Associated Industries of Missouri v. Lohman* rejects a de minimis escape from the determination that actual discrimination occurred, while preserving the narrow justification inquiry; *Camps Newfound/Owatonna* likewise distinguishes unequal treatment from a neutral rule whose burden happens to fall mainly on outsiders.⁸⁴ *Pork Producers* preserves the prohibition on purposeful discrimination. The prohibition applies to the employer duty as well as the supplier duty.⁸⁵

The interstate burden cannot be filled in from an assumed national code path. The present comparison establishes neither that a provider must rebuild its service everywhere nor that an employer can segregate every record, review, or remedy at low cost. *Kassel's* routing and trailer-separation facts were features of an evidentiary record, not findings about AI services.⁸⁶ A challenge to either hypothetical duty would need to identify its market burden and local benefit on the relevant facts. I prefer decision-linked accountability for the reasons that follow, while preserving the supplier obligations on which that accountability may depend.

5. Designing Deployer Duties Within the Same Constraints

I would attach accountability to the actor responsible for the covered AI-mediated decision. A decision-linked assessment can identify the use being reviewed, the people exposed to it, the evidence the reviewer needs, and the official or institution responsible for acting. Part III and the worked Appendix develop that design choice. It also exposes a dependency: where the necessary validation or limitations are known only to the supplier, a duty addressed to the

⁸³ *Pork Producers*, 598 U.S. at 371–77 & n.1 (effects-alone theory and territorial reservation), 377–80 (Part IV-A); *Healy*, 491 U.S. at 336–40 (price linkage and interacting laws); *Edgar v. MITE Corp.*, 457 U.S. 624, 641–43 (1982) (Part V-A, plurality opinion). The two duties in the text are the author's stipulated comparison; no cited case adjudicates either.

⁸⁴ *Associated Indus. of Mo. v. Lohman*, 511 U.S. 641, 647–50, 654–56 (1994) (actual discrimination, narrow justification and actual equal treatment); *Camps Newfound/Owatonna, Inc. v. Town of Harrison*, 520 U.S. 564, 575–81 & n.13 (1997) (Maine exemption and similarly situated actors); *CTS Corp.*, 481 U.S. at 87–88 (out-of-state incidence alone insufficient).

⁸⁵ *Pork Producers*, 598 U.S. at 364, 369–71 (opinion of the Court). The design comparison does not establish discrimination or its absence in any particular AI law.

⁸⁶ *Kassel v. Consol. Freightways Corp.*, 450 U.S. 662, 671–74 (1981) (plurality opinion). The analogy identifies an evidentiary question; it supplies no estimate of AI infrastructure costs or technical necessity.

deployer needs a lawful way to obtain that evidence. A state license, a process label, or the residence of the affected person supplies no automatic constitutional protection.

Colorado's legislative sequence shows why the duty bearers must be identified in the actual provisions. The 2024 enactment assigned reasonable-care and information duties to developers as well as reasonable-care, assessment, and notice duties to deployers.⁸⁷ Its enacted 2026 successor changes the definitions and consumer protections while retaining developer documentation, change-notice, and record duties. Its general operation and application to covered consequential decisions are scheduled for January 1, 2027, with specified provisions effective earlier. Its consumer definition also reaches some individuals through access, eligibility, or opportunity in Colorado, as well as residence.⁸⁸ These two enacted snapshots illustrate combinations of upstream and decision-facing obligations. They do not show that Colorado eliminated the developer side, prove why the legislature revised the law, or measure the relative constitutional burden of either arrangement.

CTS supplies a useful but limited analogy. The Court upheld Indiana's regulation of voting rights in corporations it had created.⁸⁹ It recognized the State's interest in creating corporations and defining the rights acquired through their shares.⁹⁰ Its answer to inconsistent state voting-rights regulation depended on each corporation being subject to one incorporating State's law.⁹¹ That legal allocation matters to the result. Operating under a state license or other state law gives a hospital, insurer, employer, or lender no comparable allocation.

The Appendix's resident-based State Nexus identifies covered decisions; it expressly leaves the deployer's jurisdictional position to applicable jurisdictional law. Coverage of a decision about a resident does not by itself place a deployer physically in the State. Nor does resident injury alone establish the defendant-created forum contacts required for personal jurisdiction. *Walden v. Fiore* supplies that adjudicative due-process qualification and expressly reserves questions about virtual contacts; it is not a Dormant Commerce Clause holding about the proposed duties.⁹² Coverage, adjudicative jurisdiction, and substantive territorial limits require separate accounts.

Public functions require another distinction. *United Haulers* upheld flow-control ordinances

⁸⁷ Colorado S.B. 24-205, § 1, Colo. Rev. Stat. §§ 6-1-1701(4), (6)–(7), -1702, -1703 (2024 signed-act snapshot), https://leg.colorado.gov/bill_files/47770/download. This comparison uses the enacted 2024 definitions and allocation of duties; it does not treat that act's original commencement dates as the current schedule.

⁸⁸ Colorado S.B. 26-189 (signed May 14, 2026), § 1, Colo. Rev. Stat. §§ 6-1-1701(4), (7)–(8), -1702 to -1705; bill § 5, https://leg.colorado.gov/bill_files/116489/download. Section 1702 retains bounded developer duties, including documentation, change notices, three-year records, marketed/contracted-use limits and a protected-information notice. Sections 1704–1705 address notices, adverse-outcome information, correction and human review/reconsideration; the human-review duty is qualified by commercial reasonableness. Bill § 5 states general January 1, 2027 effect and applicability to covered decisions on or after that date, with specified earlier provisions. The comparison addresses the enacted framework, not subsequent implementing rules or litigation.

⁸⁹ *CTS Corp. v. Dynamics Corp. of Am.*, 481 U.S. 69, 87–94 (1987).

⁹⁰ *Id.* at 89–93 (state-created corporate form, voting rights and shareholder interests).

⁹¹ *Id.* at 88–89 (equal treatment and single-incorporation answer to inconsistent regulation).

⁹² Appendix § 2.3(a) *infra* (decision coverage and jurisdiction distinguished); *Walden v. Fiore*, 571 U.S. 277, 283–91 & n.9 (2014), <https://www.govinfo.gov/content/pkg/USREPORTS-571/pdf/USREPORTS-571-277.pdf> (defendant-created forum contacts; residence/injury insufficient by themselves; virtual-contacts question reserved). This is an adjudicative-jurisdiction qualification, not a resolution of legislative jurisdiction or the Dormant Commerce Clause.

favoring a publicly owned waste facility.⁹³ The Court's discrimination analysis depended on the public beneficiary and equal treatment of private businesses; it did not establish that all government conditions on private suppliers receive that treatment.⁹⁴ *Department of Revenue of Kentucky v. Davis* applied the public-function reasoning to the municipal-bond tax arrangement before it, while leaving open the general applicability of *Pike* to a case of that kind and addressing the institutional comparison on that record.⁹⁵ Actual government purchasing is a distinct inquiry developed in Part IV.D. A condition on the State's own transaction and a rule governing the seller's wider business may raise different questions. These cases supply no constitutional ranking that places every procurement-anchored AI duty below every private-sector duty in exposure.

The strongest territorial objection to a deployer approach begins where the *CTS* analogy ends. A national employer, multistate hospital system, or insurer may face duties in several States over shared decision infrastructure. It lacks the singular governing jurisdiction on which the corporate voting-rights analysis depended.⁹⁶ A process requirement aimed at covered hiring decisions may induce the employer to use one reviewed pipeline for all hiring. The challenger may describe that consequence as state control of conduct elsewhere, invoking *Healy*'s practical-effects language.⁹⁷

Pork Producers answers the effects-alone version of that objection for both sides of the comparison. Large consequences for out-of-state production did not establish the proposed automatic ground of invalidity. Its judgment nevertheless concerned the sufficiency of a particular complaint, and its separate opinions gave different accounts of the *Pike* burden.⁹⁸ A legally delimited command and a firm's response to it need to be distinguished, without treating the response as constitutionally irrelevant. The price linkages in *Baldwin* and *Healy* made transactions elsewhere part of the legal constraint; a finding that an AI duty has the same feature requires the actual command and its operation.⁹⁹ A deployer's decision to standardize may instead reflect cost or compatibility. That possibility does not establish a winning *Pike* balance, just as a supplier's decision to standardize does not by itself establish invalidity.

The local benefit also needs evidence. *Kassel*'s plurality found a weak safety justification and substantial interference with interstate transport.¹⁰⁰ The fact that an AI duty concerns a resident's hiring, medical treatment, or access to credit identifies an interest the State seeks to serve. That interest alone does not establish the size of the benefit, the sufficiency of the re-

⁹³ *United Haulers Ass'n v. Oneida-Herkimer Solid Waste Mgmt. Auth.*, 550 U.S. 330, 334, 338–45 (2007).

⁹⁴ *Id.* at 342–45 (opinion of the Court; public beneficiary and equal treatment of private businesses), 346–47 (four-Justice *Pike* discussion); *see id.* at 332 (opinion-part lineup).

⁹⁵ *Dep't of Revenue of Ky. v. Davis*, 553 U.S. 328, 330, 339–43 & n.9, 353–56 (2008) (public-function reasoning; Part IV, joined by five Justices, leaves general *Pike* applicability to a case of that kind open and addresses the comparison on its record). This does not merge public-function analysis with the market-participant inquiry developed in Part IV.D.

⁹⁶ *CTS Corp.*, 481 U.S. at 88–94. Its corporate voting-rights holding does not establish a single governing State for the multistate actors described in the text.

⁹⁷ *Healy*, 491 U.S. at 336–40. The objection is stated against the later limits explained in *Pork Producers*, 598 U.S. at 371–77; *see supra* Section II.A.3.

⁹⁸ *Pork Producers*, 598 U.S. at 364, 371–80, 383–87 (four-Justice Part IV-C), 391–403 (separate opinions); *see supra* Section II.A.3.

⁹⁹ *Baldwin*, 294 U.S. at 519–24, 527–28; *Healy*, 491 U.S. at 336–40; *Pork Producers*, 598 U.S. at 371–77 & n.1.

¹⁰⁰ *Kassel*, 450 U.S. at 670–74 (plurality opinion); *see supra* note 86. The different regulatory context is material; neither the AI benefit nor its burden is measured here.

quired review, or the absence of less burdensome alternatives. Those are questions for the particular duty and record. I make no finding that documentation, monitoring, or recourse costs are always fixed, local, or small.

The cumulative-burden objection can coexist with national standardization. An employer might adopt a common review process and still face different state definitions, reports, retention periods, deadlines, and review routes. That is a possible configuration to investigate, not an empirical description of fifty current regimes. *Southern Pacific* itself considered both breaking trains at borders and conforming an entire route to the lowest limit. Conformity was an alternative burden, not proof that the burden disappeared.¹⁰¹ A challenger need not elect between all spillover and all patchwork as a condition of stating the problem.

Bibb v. Navajo Freight Lines shows a particularly concrete incompatibility. Illinois required contour mudguards inconsistent with Arkansas's rule, and the record included interchange work and disruption to interlining. The Court weighed that burden against the safety showing.¹⁰² Those facts explain a mechanism by which interacting rules impede a network. They do not make literal incompatibility, physical equipment, or a border alteration a prerequisite to a *Pike* challenge. An AI duty requires its own evidence of operational dependencies, costs, and effects on the interstate market.

I retain **local satisfiability** as a design question. In legal scope, does the duty require action only for the decisions it covers? In operation, can the actor comply for those decisions without changing other decisions or shared systems elsewhere? The first question can be examined in the statutory text. Operational separability requires facts about the actual records, systems, suppliers, and review work. A legally local duty may still have operational consequences beyond its coverage.

An assessment might be completed for one deployment if the required evidence is available and the review can be separated. An audit trail might be retained for covered decisions if the system can identify and retain the relevant records without conflicting duties elsewhere. A recourse pathway might address one person's decision, yet depend on shared data, supplier cooperation, or a system correction that affects other uses. These examples specify what must be checked; they do not establish that every proposed duty can be performed locally. A bounded supplier disclosure could also be separable. Local satisfiability is my analytical indicator, not a judicial test, an assertion of technical feasibility, or a substitute for the burden-benefit inquiry.

Costs can accumulate even when duties are separable. The three-Justice comparability reasoning in *Pork Producers* concerned production costs weighed against moral and health interests; six Justices accepted the general possibility of weighing disparate interests, as Section II.A.3 explains.¹⁰³ That division supplies no general answer to the combined burden of differing AI duties. The case leaves the question described here undecided; it does not establish that

¹⁰¹ *S. Pac. Co. v. Arizona ex rel. Sullivan*, 325 U.S. 761, 770–76, 779–83 (1945); see *supra* note 38. The border-adjustment and route-conformity alternatives appear at 772–74; the opinion considers operating costs, network effects, uniformity and the safety showing. This is not a holding that every aggregate compliance expense is unconstitutional.

¹⁰² *Bibb v. Navajo Freight Lines, Inc.*, 359 U.S. 520, 523–30 (1959) (Court opinion, with Harlan and Stewart, JJ., separately concurring in the judgment), <https://www.govinfo.gov/content/pkg/USREPO-RTS-359/pdf/USREPORTS-359-520.pdf>. The opinion addresses mudguard interchange, interlining and the safety evidence. It states no local-satisfiability test for AI.

¹⁰³ *Pork Producers*, 598 U.S. at 380–83 (Gorsuch, J., joined by Thomas and Barrett, JJ.), 391–403 (separate opinions); see *supra* note 62 and Section II.A.3. The discussion distinguishes the three-Justice rationale from proposition-level agreement on general weighability. The application to combined AI duties is an inference, not a case outcome. Pinpoints follow the preliminary print.

cumulative-burden doctrine is absent or that courts can never perform the comparison.

A legislature seeking to reduce duplication should examine the evidence fields, editions, review criteria, exceptions, clocks, and consequences it would share with other States. References to NIST AI RMF or ISO/IEC 42001 alone establish none of that compatibility. Reusable evidence and coordinated procedures are design objectives whose costs and benefits need to be tested. I do not infer savings or an effectively uniform regime from framework names or the corpus count.

Executive policy supplies a separate, narrower fact. Executive Order 14365 directs named presidential advisers to prepare a legislative recommendation and excludes otherwise lawful state-government procurement and AI-use laws, among other topics, from that recommendation's proposed preemption.¹⁰⁴ This is a choice about an executive recommendation. It neither binds Congress nor adjudicates the validity of procurement conditions, and it does not preserve every deployer or process duty. It matters to the policy environment without supplying the constitutional result.

The design I favor starts with a specified decision and follows responsibility, evidence, review, and relief back to the actors who can provide them. Legal scope should fit the State's connection to the activity; compliance should be examined for avoidable external demands and combined burdens. Upstream obligations may be necessary to make a deployer's accountability meaningful, and they require the same legal examination. The Dormant Commerce Clause constrains those choices without selecting an entire regulatory approach in advance. Distinct requirements follow from federal preemption (Section II.B) and state police powers (Section II.C).

B. Federal Preemption: Compare the Duty with the Federal Rule

1. Four Strands, Four Questions

Preemption begins with a comparison: what does the state duty require, and what does federal law require, permit or reserve in the same setting? Four familiar strands organize that inquiry. Express preemption turns on an enacted displacement clause. Field preemption asks whether Congress has left room for state supplementation.¹⁰⁵ Impossibility concerns an actor's ability to comply with both laws; obstacle preemption concerns interference with federal purposes established through the governing law.¹⁰⁶ These are different inquiries. Neither a deployer label nor an assessment requirement or traditional regulatory purpose answers all four.

An express clause requires close attention to its covered subject, actors and requirements. *Cipollone* illustrates the claim-by-claim approach, with the necessary distinction between the opinion of the Court and Justice Stevens's plurality application; *Altria* subsequently applied that approach to the asserted duty.¹⁰⁷ The medical-device cases show why the federal predicate

¹⁰⁴ Exec. Order No. 14365, §§ 4, 8(a)–(b), 9, 90 Fed. Reg. 58,499, 58,500–01 (Dec. 16, 2025) (Doc. 2025-23092), <https://www.govinfo.gov/content/pkg/FR-2025-12-16/pdf/2025-23092.pdf>. Section 8(b)(iii) concerns the legislative recommendation and “otherwise lawful State AI laws” relating to state-government procurement and use. Other listed topics include child safety and specified infrastructure. Section 4 also contemplates evaluating developer or deployer disclosure/reporting requirements.

¹⁰⁵ *Rice v. Santa Fe Elevator Corp.*, 331 U.S. 218, 230 (1947).

¹⁰⁶ *Hines v. Davidowitz*, 312 U.S. 52, 67 (1941).

¹⁰⁷ *Cipollone v. Liggett Grp., Inc.*, 505 U.S. 504, 516–18 (1992) (Part III for seven Justices; the claim applications in Parts V–VI spoke for a four-Justice plurality); *Altria Grp., Inc. v. Good*, 555 U.S. 70,

matters. In *Lohr*, the § 510(k) substantial-equivalence process did not preempt the challenged common-law duties.¹⁰⁸ In *Riegel*, device-specific premarket approval supplied federal requirements against which the state tort claims had to be compared.¹⁰⁹ Parallel state requirements may escape displacement by that clause, but similarity of purpose is not enough, and *Riegel* did not decide an unpreserved parallel-claim argument.¹¹⁰

The Court's 2026 decision in *Monsanto v. Durnell* reinforces the need for precision. It treated EPA-approved labeling as a federal requirement under FIFRA and held the additional state warning duty expressly preempted. Describing the state and federal standards at a high level of generality did not make them parallel. The statute separately preserves state authority over sale or use; the text establishes that preservation. A general product-versus-deployer rule cannot be inferred from it.¹¹¹ A saving clause also has to be read with the displacement clause. *Morales* refused to let a general remedies provision override the Airline Deregulation Act's specific substantive preemption provision.¹¹² Existing sectoral clauses can therefore reach an AI-related duty now. An AI-specific federal statute is not a prerequisite.

Field preemption asks a different question. *Rice* identifies both pervasive federal regulation leaving no room for supplementation and a dominant federal interest supporting exclusion of state law on the same subject.¹¹³ The relevant field must be identified, not assumed to be artificial intelligence as a whole. The absence of a comprehensive AI code would not establish that every use of AI lies outside an occupied field; nor would an executive preference for uniform AI policy establish that Congress has occupied one. For the state drafter, the task is to locate the proposed duty within the federal scheme that actually governs its subject.

Impossibility requires an equally concrete comparison. *PLIVA v. Mensing* asks whether the private party could independently do under federal law what state law requires.¹¹⁴ A generic-drug manufacturer unable independently to change the federally required label could not satisfy the state warning duty.¹¹⁵ *Bartlett* rejects an answer that would require the actor to leave the market to avoid the conflict.¹¹⁶ *Albrecht* addresses the different position of a brand-name

76-87 (2008); *Lorillard Tobacco Co. v. Reilly*, 533 U.S. 525, 548-51 (2001). An express clause does not eliminate ordinary conflict-preemption analysis. *Geier v. Am. Honda Motor Co.*, 529 U.S. 861, 869-74 (2000).

¹⁰⁸ *Medtronic, Inc. v. Lohr*, 518 U.S. 470, 492-502 (1996) (Part V is an opinion of the Court; Part VI resumes the plurality). The result concerns the federal and state requirements examined, rather than an exemption for every generally applicable duty. *Id.* at 500.

¹⁰⁹ *Riegel v. Medtronic, Inc.*, 552 U.S. 312, 322 (2008).

¹¹⁰ *Riegel*, 552 U.S. at 330 (distinguishing parallel requirements and declining to address the argument first raised in the Supreme Court). The clause's treatment of parallel duties does not foreclose an independent implied-preemption objection. See *Buckman Co. v. Plaintiffs' Legal Comm.*, 531 U.S. 341, 352-53 (2001); *infra* note 490.

¹¹¹ *Monsanto Co. v. Durnell*, 609 U.S. 371, 383-88 (2026) (Kavanaugh, J., for the Court; Thomas, J., concurring; Jackson, J., dissenting, joined by Gorsuch, J.); 7 U.S.C. § 136v(a)-(b). The preserved sale-or-use authority is discussed at 609 U.S. at 380 n.2; the Court reserves implied preemption at 388 n.7.

¹¹² *Morales v. Trans World Airlines, Inc.*, 504 U.S. 374, 384-86, 390-91 (1992) (Scalia, J., for a five-Justice Court; Souter, J., took no part) (specific preemption of state standards concerning rates, routes or services prevails over a general preservation of remedies; some state effects may be too remote or peripheral).

¹¹³ *Rice*, 331 U.S. at 230.

¹¹⁴ *PLIVA, Inc. v. Mensing*, 564 U.S. 604, 620 (2011).

¹¹⁵ *Id.* at 618.

¹¹⁶ *Mut. Pharm. Co. v. Bartlett*, 570 U.S. 472, 488 (2013).

manufacturer and treats the relevant clear-evidence question as one for the judge.¹¹⁷ A state requirement to document an AI-assisted decision may be compatible with the governing federal duties. That must be shown for the actor, information and action in question; it cannot be inferred from the word documentation. Access restrictions, federal reporting rules or device conditions can change the comparison.

Obstacle preemption examines the federal choice the state law would frustrate.¹¹⁸ *Hines* invalidated the state alien-registration scheme in the setting of the federal system Congress had enacted.¹¹⁹ *Crosby* found a conflict between Massachusetts's purchasing restriction and Congress's calibrated Burma sanctions regime.¹²⁰ In *Geier*, preserving a mix of passive-restraint options was a significant federal policy that the asserted airbag duty would defeat.¹²¹ *Williamson* shows the limit of that reasoning: a federal choice between options does not invariably establish an objective of protecting the choice from state law.¹²² Identifying the statutory and regulatory basis of the claimed objective is indispensable. A national preference for fewer AI rules does not by itself establish a conflict with each state assessment or remedy.

The doctrine also contains disagreement about how far courts may infer federal purposes. Justice Thomas's *Wyeth* concurrence and Justice Gorsuch's three-Justice lead opinion in *Virginia Uranium* criticize displacement based on purposes detached from enacted law. Those are identified opinions, not a holding abolishing obstacle preemption.¹²³ *Hencely v. Fluor* rejects preemption of the contractor claims before it and emphasizes the need for a constitutional or statutory source, while addressing the narrow federal-common-law route associated with *Boyle*.¹²⁴ It does not require Congress to enact an AI-specific statute before existing federal law can displace an AI-related state duty.

Traditional state authority remains relevant. *Wyeth* treats the presumption against preemption in an historically state-regulated field as an interpretive starting point, alongside congressional purpose.¹²⁵ That starting point is neither immunity nor a universal clear-statement requirement. The Court's approach to an express clause and the federal character of the regulated relationship can limit its use. The practical lesson is to identify the relevant presumption and then perform the statutory analysis needed to reach a result.¹²⁶

¹¹⁷ *Merck Sharp & Dohme Corp. v. Albrecht*, 587 U.S. 299, 303, 310-17 (2019) (clear evidence and the judge's legal inquiry, including the agency's legally effective action and the information provided to it). The decision is not a universal test for every form of impossibility preemption.

¹¹⁸ *Hines*, 312 U.S. at 67.

¹¹⁹ *Id.* at 74.

¹²⁰ *Crosby v. Nat'l Foreign Trade Council*, 530 U.S. 363, 373 n.7, 376-86 (2000).

¹²¹ *Geier v. Am. Honda Motor Co.*, 529 U.S. 861, 881 (2000).

¹²² *Williamson v. Mazda Motor of Am., Inc.*, 562 U.S. 323, 330-36 (2011) (distinguishing the policy significance of the choice preserved in *Geier*; a regulatory option alone does not establish that preserving it was a significant federal objective).

¹²³ *Wyeth v. Levine*, 555 U.S. 555, 583, 588 (2009) (Thomas, J., concurring in the judgment); *Virginia Uranium, Inc. v. Warren*, 587 U.S. 761, 778-79 (2019) (lead opinion of Gorsuch, J., joined by Thomas and Kavanaugh, JJ.). These identified opinions criticize purposes-and-objectives reasoning detached from enacted law; they do not overrule *Geier* or *Crosby*.

¹²⁴ *Hencely v. Fluor Corp.*, 608 U.S. 31, 38-39 (2026) (Thomas, J., for the Court; Alito, J., dissenting, joined by Roberts, C.J., and Kavanaugh, J.); *Boyle v. United Techs. Corp.*, 487 U.S. 500, 504-08, 512-13 (1988). *Hencely* addresses the asserted displacement of claims concerning contractor conduct the Government had not directed. Existing valid federal law may matter without an AI-specific enactment.

¹²⁵ *Wyeth v. Levine*, 555 U.S. 555, 565 (2009).

¹²⁶ *Wyeth*, 555 U.S. at 565 (quoting *Lohr*, 518 U.S. at 485, quoting *Rice*, 331 U.S. at 230); *Puerto Rico v.*

2. Traditional State Concern Does Not Supply an Exemption

The *Rice* presumption recognizes that Congress legislates against a background of state police powers.¹²⁷ Health and safety cases such as *Hillsborough County* and *Lohr* give that background real interpretive weight.¹²⁸ *Wyeth* does the same in pharmaceutical labeling.¹²⁹ For state AI law, this supports attention to the underlying decision environment: employment, housing, healthcare, education, public benefits, criminal justice and insurance. It does not establish a separate zone in which federal law can reach only lightly. Federal civil rights, benefits, device, banking and other laws operate within those same environments.

ERISA is an especially important test of the Article's recommendation. Section 514 displaces state laws with a prohibited reference to or connection with covered plans. *Travelers* and *Dillingham* reject an unlimited reading of relates to; *Egelhoff* shows that a generally applicable law may nevertheless intrude on plan administration. *Davila* separately addresses claims within ERISA's exclusive civil-enforcement scheme.¹³⁰ A state's consumer-protection purpose or choice of a general deployer category cannot supply a safe harbor. The proposed duty and remedy must be examined for their operation on the plan.

Gobeille makes that limit concrete. Vermont required health plans to report claims and member information to an all-payer database. The Court held the requirement preempted as applied to ERISA plans.¹³¹ Reporting, disclosure and recordkeeping were central components of the federal system of plan administration. Even parallel requirements could undermine that uniform system; the State's public-health purpose did not remove the problem.¹³² A state AI record-keeping law can encounter the same objection if it regulates those components. Changing the records' name from benefits report to impact assessment would not settle the question.

Rutledge supplies a different result on different facts. Arkansas regulated pharmacy-benefit-manager reimbursement rates and an associated pharmacy appeal process. The law neither referred impermissibly to ERISA nor governed a central matter of plan administration. The Court emphasized the distinction between affecting costs and dictating plan choices; it also recorded that the challenger had not argued that the enforcement mechanisms overlapped the fundamental administrative components at issue in *Gobeille*.¹³³ A state drafter can learn from

Franklin Cal. Tax-Free Tr., 579 U.S. 115, 125 (2016) (express clause); *Buckman*, 531 U.S. at 347-48 (federal-agency relationship). See *infra* Section II.B.2.

¹²⁷ *Rice*, 331 U.S. at 230.

¹²⁸ *Hillsborough County v. Automated Med. Labs., Inc.*, 471 U.S. 707, 718 (1985) (the quoted presumption formulation); *Lohr*, 518 U.S. at 485.

¹²⁹ *Wyeth*, 555 U.S. at 565.

¹³⁰ ERISA §§ 502, 514, 29 U.S.C. §§ 1132, 1144; *Shaw v. Delta Air Lines, Inc.*, 463 U.S. 85, 95-100, 102-09 (1983); *N.Y. State Conf. of Blue Cross & Blue Shield Plans v. Travelers Ins. Co.*, 514 U.S. 645, 655-62, 668 (1995); *Cal. Div. of Lab. Standards Enft v. Dillingham Constr., N.A., Inc.*, 519 U.S. 316, 324-32 (1997); *Egelhoff v. Egelhoff*, 532 U.S. 141, 146-52 (2001); *Aetna Health Inc. v. Davila*, 542 U.S. 200, 208-14 (2004). The § 514 reference-or-connection inquiry, § 502 complete preemption, and insurance/deemer provisions are distinct. See also *FMC Corp. v. Holliday*, 498 U.S. 52, 58-65 (1990); *infra* notes 455, 475.

¹³¹ *Gobeille v. Liberty Mut. Ins. Co.*, 577 U.S. 312 (2016). Kennedy, J., delivered the opinion of the Court, joined by Roberts, C.J., and Thomas, Breyer, Alito, and Kagan, JJ.; Thomas, J., and Breyer, J., filed concurring opinions; Ginsburg, J., filed a dissenting opinion, joined by Sotomayor, J. The Court sat with eight members.

¹³² *Gobeille*, 577 U.S. at 323-26 (central reporting, disclosure and recordkeeping; even parallel state requirements; rejecting reliance on a traditional public-health purpose). See *Egelhoff*, 532 U.S. at 148.

¹³³ *Rutledge v. Pharm. Care Mgmt. Ass'n*, 592 U.S. 80, 87-90 (2020) (Sotomayor, J., for the unanimous

that comparison. A duty addressing a deployer's own commercial conduct may leave plan administration untouched, but a general AI statute is not necessarily Arkansas's law in another form. Coverage, records, decision rules and remedies need separate analysis.

Express clauses require their own treatment of the presumption. In *Franklin California*, the Court declined to invoke it and focused on the clause's wording.¹³⁴ FCRA, the Airline Deregulation Act, Section 230 and the National Bank Act illustrate existing federal constraints that may matter to AI applications, although their texts and tests differ.¹³⁵ A state cannot establish compatibility by showing that it regulates a familiar field or uses a familiar institution.

Section 230 illustrates why the basis of liability matters. *Zeran* read the statute to bar claims treating an interactive service as publisher of third-party information, including claims pleaded as negligence. *Lemmon* allowed a design-duty claim independent of that publishing role; *Anderson* treated the particular algorithmic recommendation as the platform's own expression, while acknowledging circuit tension. *Moody* did not decide Section 230.¹³⁶ The Article's focus on an institution's decision about a person helps specify the asserted duty. It does not decide whether liability in a particular application would treat that institution as publisher or speaker of someone else's content. The same discipline applies to FCRA's enumerated subjects and the banking tests developed in Part IV.C.1: identify the duty and its statutory treatment. The design's name cannot decide the result.

3. Executive Pressure, Federal Law and Severability

Federal AI policy adds pressure without collapsing these distinctions. Executive Order 14365 directs challenges to state AI laws and preparation of a legislative framework recommending preemption.¹³⁷ The model-facing statutes discussed in Section II.A.4 are potential targets, but the existence of a target establishes neither a valid federal source nor a judicial result.¹³⁸ The order also directs a later Commerce policy notice concerning non-deployment BEAD funds and an FTC policy statement concerning state requirements that alter truthful model outputs. Those

participating Court; Thomas, J., concurring; Barrett, J., took no part). The reservation concerning the challenger's argument about fundamental plan-administration components is at 89 n.2.

¹³⁴ *Franklin California*, 579 U.S. at 125 (focusing on the plain wording of the express clause). This does not establish that every implied-preemption dispute receives an identical presumption. See *supra* note 126.

¹³⁵ 15 U.S.C. § 1681t; 49 U.S.C. § 41713; 47 U.S.C. § 230(c)(1), (e)(3); 12 U.S.C. § 25b. These are different statutory structures. In particular, § 25b(b)(1) states the applicable routes for state consumer financial laws, while § 25b(b)(4) disclaims field occupation. See *infra* Section IV.C.1.

¹³⁶ *Zeran v. Am. Online, Inc.*, 129 F.3d 327, 330, 332-34 (4th Cir. 1997) (unanimous panel; looking through negligence pleading to the asserted publisher duty); *Lemmon v. Snap, Inc.*, 995 F.3d 1085, 1091-93 (9th Cir. 2021) (independent design-duty theory); *Anderson v. TikTok, Inc.*, 116 F.4th 180, 184-88 (3d Cir. 2024) (specific algorithmic recommendation; circuit tension and reserved applications). *Moody v. NetChoice, LLC*, 603 U.S. 707 (2024), did not decide § 230; *id.* (Alito, J., concurring in the judgment), slip op. at 7 n.8. FCRA likewise requires attention to the enumerated subject: *Premium Mortgage Corp. v. Equifax, Inc.*, 583 F.3d 103, 106-07 (2d Cir. 2009) (per curiam); *Hinkle v. Midland Credit Mgmt., Inc.*, 827 F.3d 1295, 1301-05 (11th Cir. 2016) (furnisher investigation under the identified federal duty, not a general preemption holding).

¹³⁷ Exec. Order No. 14365, §§ 3, 8(a), 90 Fed. Reg. 58,499 (Dec. 16, 2025) (Doc. 2025-23092) (§ 3 establishing the AI Litigation Task Force; § 8(a) directing preparation of the federal legislative framework).

¹³⁸ See *supra* notes 76, 77, 78, 79, 80.

directions operate subject to applicable law.¹³⁹

The FTC's July 2026 publication is a proposal. It asserts deception and conflict-preemption concerns, including concerns about Colorado's revised law.¹⁴⁰ Its discussion of Colorado's developer-liability allocation states the Commission's position; it supplies no holding that either set of duties is invalid.¹⁴¹ The White House's March 2026 framework separately recommends preserving general state laws and requirements governing a state's own AI use. It recommends these boundaries to Congress; it does not establish constitutional reservations or enact saving clauses.¹⁴²

Litigation is another channel. The United States' April 24 complaint in intervention in *X.AI v. Weiser* alleges that Colorado's former law compels protected-category discrimination. The April 27 order adopts an agreed enforcement pause covering alleged violations through fourteen days after a future preliminary-injunction ruling, including specified replacement legislation. Neither document adjudicates the constitutional allegation or severability.¹⁴³ The FCC's June 2026 broadband-infrastructure notice is likewise a proposed rulemaking under its identified statutory authority, distinct from the AI-disclosure proceeding contemplated by the executive order. A proposed rule does not itself displace state law.¹⁴⁴

A duty on the deployer's own conduct can make the legal comparison more manageable. The drafter can specify who must retain a record, what information is available to that actor, and what action follows an identified risk. *Mensing* still asks whether that actor can comply independently with both laws.¹⁴⁵ Voluntary alignment with NIST AI RMF or ISO/IEC 42001 can

¹³⁹ Exec. Order No. 14365, §§ 5, 7, 9. Section 5(a) directs a later NTIA notice within statutory authority; it is not itself a finding that a particular State has lost eligibility. NTIA's BEAD Frequently Asked Questions, version 19 (Mar. 2026), questions 2.21, 2.30, discusses non-deployment approvals and further guidance. Neither document establishes the operative AI-specific eligibility notice, award condition or denial in a particular State.

¹⁴⁰ Policy Statement Concerning the Suppression of Accuracy in Artificial Intelligence Systems, 91 Fed. Reg. 41,638, 41,639 & n.12, 41,641 & n.43 (proposed July 7, 2026); Press Release, Fed. Trade Comm'n, *FTC Seeks Public Comment on Policy Statement Addressing AI Accuracy* (July 1, 2026) (2-0 publication vote; July 31 comment deadline). The FTC materials identify a proposed statement as of October 16, 2026. The statement's preemption position is an agency argument, not a judicial holding.

¹⁴¹ 91 Fed. Reg. at 41,641 & n.43; Colorado SB 26-189, Colo. Rev. Stat. § 6-1-1707(5) (2026). The provision specifies conditions concerning qualifying unlawful consequential decisions and the developer's intended, documented, marketed or otherwise specified uses. Its conditions, not merely the developer label, govern the allocation. *See supra* note 88.

¹⁴² White House, *A National Policy Framework for Artificial Intelligence: Legislative Recommendations* 3-4 (Mar. 20, 2026), <https://www.whitehouse.gov/wp-content/uploads/2026/03/03.20.26-National-Policy-Framework-for-Artificial-Intelligence-Legislative-Recommendations.pdf> (recommendations concerning sector-specific institutions, laws of general applicability and state procurement and use); Exec. Order No. 14365, § 8. The general-law recommendation expressly includes developers as well as users; it does not establish a constitutional ranking between them.

¹⁴³ Complaint in Intervention, *X.AI LLC v. Weiser*, No. 1:26-cv-01515, ECF No. 17, ¶¶ 46-59 (D. Colo. Apr. 24, 2026); Minute Order, *id.*, ECF No. 24, at 1-2 (Apr. 27, 2026) (agreed enforcement and investigation pause covering SB 24-205 and specified replacement or amending legislation). The complaint is a litigant's Equal Protection theory; the order is procedural. *See infra* notes 182-183 (substantive constitutional scope).

¹⁴⁴ FCC, Notice of Proposed Rulemaking, WC Docket No. 25-253, FCC 26-40, ¶¶ 1-11, 74-75 (adopted June 25, released June 26, 2026) (47 U.S.C. § 253 wireline-infrastructure proceeding). EO 14365 § 6 concerns a different AI-disclosure inquiry and measures its deadline from publication of the § 4 identification.

¹⁴⁵ *PLIVA*, 564 U.S. at 620.

help organize the duty; it does not make the duty a parallel federal requirement under *Riegel*. A management framework's endorsement by government also does not prove that every state mandate built from it avoids conflict.¹⁴⁶ The argument for a deployer-centered design is that responsibility can be specified around a decision and the institution making it. Legal compatibility remains a condition of that design.

Colorado's successive statutes demonstrate why actor labels are insufficient. SB 24-205 combined developer documentation and disclosure with deployer assessments and consumer protections.¹⁴⁷ SB 26-189 reallocates duties within a revised automated-decision framework, but retains developer information, documentation and related responsibilities alongside deployer obligations. It also addresses responsibility for qualifying unlawful consequential decisions. The revision is legislative action, not judicial severance, and the sequence does not establish that preemption pressure caused the legislature's choices.¹⁴⁸

Severability nevertheless belongs in the design. When part of a state law is invalid, state law governs whether the rest can remain.¹⁴⁹ The familiar federal formulation asks about an operative remainder and legislative intent, but it does not replace the relevant state's rule.¹⁵⁰ Colorado's default rule, for example, addresses provisions inseparably dependent on the invalid part and a remainder incapable of execution in accordance with legislative intent.¹⁵¹ A drafter should examine both dependencies and the legislature's stated choice.

Leavitt v. Jane L. shows the value of a specific legislative statement that provisions would have been enacted independently. The per curiam Court rejected the contrary account of legislative intent; four Justices dissented from taking the state-law issue for review.¹⁵² Such a statement can assist the intent inquiry. It cannot supply information or authority the surviving duty lacks. An assessment dependent on a developer's invalidated furnishing obligation may need a lawful alternative source of information. An assessment based on records the deployer holds may avoid that dependency, but only to the extent those records actually suffice. Appendix § 3.4(a)'s treatment of reasonably obtainable lineage information illustrates a drafting response

¹⁴⁶ *Riegel*, 552 U.S. at 330; NIST, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, at 2, 22 (Jan. 2023) (voluntary framework; applicable legal requirements); ISO/IEC 42001:2023, *Information Technology — Artificial Intelligence — Management System* (public scope description). ISO supplies an organizational analogy; binding incorporation would require a separate legal source.

¹⁴⁷ Colorado SB 24-205, Colo. Rev. Stat. §§ 6-1-1702 to -1703 (2024) (historical developer and deployer duties); see *supra* note 87. The historical text should be distinguished from the later reenactment.

¹⁴⁸ Colorado SB 26-189, §§ 6-1-1702, -1706 to -1707, § 5 (2026) (developer and deployer duties, allocation and effective provisions); see *supra* note 88. Principal duties apply beginning January 1, 2027, with the Act's specified exceptions. Legislative replacement is not a ruling on the earlier Act or evidence of a particular legislative motive.

¹⁴⁹ *Leavitt v. Jane L.*, 518 U.S. 137, 139 (1996) (per curiam) ("Severability is of course a matter of state law."). The Court summarily reversed over four dissents. The dissent did not contest the proposition, describing the severability issue as "purely a question of Utah law" and objecting only to the certiorari grant. *Id.* at 146 (Stevens, J., dissenting, joined by Souter, Ginsburg & Breyer, JJ.).

¹⁵⁰ *Alaska Airlines, Inc. v. Brock*, 480 U.S. 678, 684 (1987) (quoting *Buckley v. Valeo*, 424 U.S. 1, 108 (1976) (per curiam), quoting *Champlin Refining Co. v. Corporation Comm'n of Oklahoma*, 286 U.S. 210, 234 (1932)).

¹⁵¹ Colo. Rev. Stat. § 2-4-204.

¹⁵² *Leavitt*, 518 U.S. at 138-40, 144-48 (per curiam summary reversal, relying on Utah's specific independent-enactment statement; Stevens, J., dissenting, joined by Souter, Ginsburg and Breyer, JJ., objecting to review of the state-law question). The fact that the dissent did not decide the severability merits does not make the decision unanimous.

whose operability still needs to be examined.

Nor is a severability clause an invitation to judicial reconstruction. *Ayotte* emphasizes legislative intent and restraint in rewriting a state statute.¹⁵³ The April 27 *Weiser* order adopts a procedural pause; it does not decide severability.¹⁵⁴ The recommendation here is prospective: specify which duties should operate independently, inspect their information and enforcement dependencies, and supply the legislature's intended fallback without assuming that a court can invent one.

Figure 1 collects the four duty-specific questions. It assigns no risk grade to a design category.¹⁵⁵ Procurement belongs in the same analysis: *Crosby* rejects the premise that a purchasing restriction escapes preemption merely because the State spends its own money.¹⁵⁶ The design recommendation is therefore concrete. Locate responsibility at the regulated decision, identify the applicable federal law, and make any surviving duty workable on its own terms.

Express	What does the enacted clause displace? Read its actors, subject matter, requirements and saving provisions.
Field	Has federal law left room for this state duty? Identify the field and the scope of federal occupation.
Impossibility	Can this actor independently comply with both laws? Compare the required action with the actor's federal authority.
Obstacle	Would this duty frustrate a binding federal choice? Identify the legal source and the federal objective actually at issue.

Compare the same duty, actor and conduct with the applicable federal law.

General applicability, guidance alignment and purchasing form do not decide the result.

Figure 1. Four preemption questions for a specified state duty. Compare the same actor, conduct and federal rule. Source: Section II.B.1 and the authorities cited there.

¹⁵³ *Ayotte v. Planned Parenthood of N. New Eng.*, 546 U.S. 320, 329-30 (2006) (quoting *Virginia v. American Booksellers Ass'n*, 484 U.S. 383, 397 (1988)).

¹⁵⁴ *Weiser*, ECF No. 24, *supra* note 143. The April 27 order is a procedural pause, not a severability ruling.

¹⁵⁵ See *supra* Section II.B.1. The figure tracks express and field preemption, independent compliance, and obstacle preemption based on a valid federal source. For the corresponding authorities, see *supra* notes 107-113, 114-117, 118-124. Figure 1 is an analytic aid for a specified duty, not an empirical score or comparative ranking of architectures.

¹⁵⁶ *Crosby*, 530 U.S. at 386-87 (holding Massachusetts's Burma sanctions law preempted as an obstacle to the federal scheme); *id.* at 373 n.7 (rejecting the State's argument that the law escaped preemption as an exercise of its spending rather than regulatory power (citing *Wisconsin Dep't of Indus., Labor & Human Relations v. Gould Inc.*, 475 U.S. 282, 287 (1986))).

4. Four Authorities, with Different Legal Supports

Consumer protection, civil rights, sectoral licensing and public purchasing give the state several ways to address an AI-mediated decision. Their value is institutional as well as legal: each identifies actors who already have responsibilities and institutions that may be able to investigate. Each remains subject to federal law.

State unfair-and-deceptive-acts-and-practices statutes supply the first route. The NCLC survey dated March 2018 found a statute in every state and the District of Columbia; California, New York and Massachusetts illustrate different texts and enforcement arrangements.¹⁵⁷ Deceptive claims about an AI system's capabilities or promised review process may be actionable under the applicable statute. The FTC's AI enforcement materials supply a federal analogue, with the Rytr reversal also showing why a complaint, consent order and subsequent agency action must be kept distinct.¹⁵⁸ State authority does not expand automatically when the FTC retreats. Incorporation of federal interpretations, statutory defenses, exemptions and other federal laws can affect its operation. The conduct and remedial limits are mapped in Section IV.A.

Civil rights law supplies substantive standards as well as enforcement institutions. Title VII expressly accommodates qualifying state proceedings and preserves state duties within its terms.¹⁵⁹ *Griggs* established that a facially neutral employment practice can be unlawful without discriminatory intent.¹⁶⁰ *Albemarle* examines validation and job relatedness in that setting.¹⁶¹ Congress's 1991 legislation responded to *Wards Cove* and codified the burden framework, including its particular-practice and alternative-practice provisions.¹⁶² That history supports an argument about an institutional and statutory commitment. It does not establish that every state AI bias duty is lawful, every disparity proves liability, or every affected person can obtain a remedy. The statutes' limits remain part of the proposed design.

Sectoral licensing supplies authority tailored to an institution's activity. McCarran-Ferguson gives state insurance law a protection whose application depends on its statutory predicates and the federal law involved.¹⁶³ Medical and hospital law distinguish the professional from the institutional deployer; California's rule that only natural persons receive medical licenses

¹⁵⁷ Cal. Bus. & Prof. Code § 17200; N.Y. Gen. Bus. Law § 349; Mass. Gen. Laws ch. 93A, § 2(a)-(c). Massachusetts expressly addresses FTC and federal-court interpretations; New York § 349(d) provides the specified federal-compliance defense. For the dated nationwide survey, Carolyn Carter, Nat'l Consumer Law Ctr., *Consumer Protection in the States: A 50-State Evaluation of Unfair and Deceptive Practices Laws* 2, 9, 28 (Mar. 2018). These sources do not establish current uniform coverage or remedies across jurisdictions.

¹⁵⁸ Press Release, Fed. Trade Comm'n, *FTC Announces Crackdown on Deceptive AI Claims and Schemes* (Sept. 25, 2024); *FTC v. Wyndham Worldwide Corp.*, 799 F.3d 236, 247, 255, 259 (3d Cir. 2015). See *infra* notes 416-424 (pleadings, orders, reversal and the limited state-law analogy).

¹⁵⁹ Title VII, 42 U.S.C. §§ 2000e to 2000e-17; concurrent state enforcement at 42 U.S.C. §§ 2000e-7, 2000e-5(c).

¹⁶⁰ *Griggs v. Duke Power Co.*, 401 U.S. 424, 431 (1971).

¹⁶¹ *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 425, 431 (1975).

¹⁶² Civil Rights Act of 1991, Pub. L. No. 102-166, § 105(a), 105 Stat. 1071 (codified at 42 U.S.C. § 2000e-2(k)(1)) (statutory restoration of pre-*Wards Cove* burden structure); see also *Wards Cove Packing Co. v. Atonio*, 490 U.S. 642 (1989) (decision statutorily overridden in critical part by CRA 1991 § 105(a)).

¹⁶³ McCarran-Ferguson Act, 15 U.S.C. §§ 1011-1012, especially § 1012(b); *U.S. Dep't of Treasury v. Fabe*, 508 U.S. 491, 500-10 (1993); *Humana Inc. v. Forsyth*, 525 U.S. 299, 309-14 (1999). The statutory protection concerns qualifying state law and the operation of the federal law, not all conduct of an insurer. See *infra* Section IV.C.2.

is one reason to keep that distinction explicit.¹⁶⁴ Banking's model-risk guidance supplies a process analogue, while state banking codes and the federal banking statutes determine actual supervisory and enforcement powers.¹⁶⁵ These are different legal routes. Their usefulness comes from matching a proposed duty to a regulator's jurisdiction. The license itself proves nothing about displacement.

The State can also act through its purchases. *United Haulers* addresses nondiscrimination where government performs the public function at issue and treats private competitors alike.¹⁶⁶ The market-participant doctrine addresses a related but distinct question about government participation in a market. Executive Order 14365's procurement and own-use exception addresses something else again: what its legislative recommendation should preserve.¹⁶⁷ Neither an affirmative reading of that political choice nor a negative inference from the exception determines the reach of existing federal law.¹⁶⁸ *Crosby* requires the preemption inquiry even when the state measure takes purchasing form.¹⁶⁹ A genuine connection between the state's transaction and the vendor's duties is central to the procurement recommendation in Section IV.D, alongside market scope, downstream conditions and federal overlap.

Anti-commandeering limits a different kind of federal action. *Printz* forbids commands that state officers administer a federal program.¹⁷⁰ *Murphy* distinguishes such commands to state legislatures from valid federal regulation of private conduct; it also recognizes evenhanded federal regulation of activities undertaken by public and private actors.¹⁷¹ The doctrine cannot be turned into immunity for all state purchasing, agency use or licensing decisions. It identifies a constitutional limit on the means Congress may use, while valid substantive federal rules and lawful funding conditions remain separate inquiries.

The portfolio's resilience is thus a design argument, not a legal conclusion produced by addition. More than one institution may be able to address harmful conduct, but a displaced licensing duty cannot necessarily be recreated through consumer protection or a contract. Each alternative needs its own authority, elements, access and remedy. The case for this approach is that it makes those choices explicit and can preserve useful, lawful functions when another route

¹⁶⁴ Cal. Bus. & Prof. Code §§ 2032, 2234(b)-(d), 2400; Cal. Health & Safety Code §§ 1250(a), 1275; *infra* notes 485-486 (professional and institutional powers distinguished). Section 2032 defines persons for the Medical Practice Act and permits licensing only natural persons; the discipline grounds are separately specified. These California provisions do not establish a national rule.

¹⁶⁵ See FDIC, FIL-22-2017, *Adoption of Supervisory Guidance on Model Risk Management* (June 7, 2017), and accompanying guidance; Bd. of Governors of the Fed. Rsrv. Sys., FDIC & OCC, SR 26-2, *Revised Guidance on Model Risk Management* (Apr. 17, 2026); OCC Bulletin 2026-13, *Model Risk Management: Revised Guidance* (Apr. 17, 2026). The OCC's 2026 background identifies the original 2011 joint guidance; the FDIC's 2017 version is the historical operative text used here. See *infra* Section IV.C.1.

¹⁶⁶ *United Haulers Ass'n v. Oneida-Herkimer Solid Waste Mgmt. Auth.*, 550 U.S. 330, 342-47 (2007). The nondiscrimination holding is distinct from Part II-D's four-Justice *Pike* analysis. See *supra* Section II.A.5; *infra* note 494.

¹⁶⁷ Exec. Order No. 14365, § 8(b)(iii).

¹⁶⁸ The order's exception governs the requested legislative recommendation. An inference from its presence or specificity does not determine the validity of a particular state duty under existing law. See Exec. Order No. 14365, § 8(b); *Crosby*, 530 U.S. at 373 n.7.

¹⁶⁹ *Crosby*, 530 U.S. at 373 n.7, 386-87. A relevant federal scheme can constrain purchasing without comprehensive AI legislation. See also *Wis. Dep't of Indus., Lab. & Hum. Rels. v. Gould Inc.*, 475 U.S. 282, 287-91 (1986).

¹⁷⁰ *Printz v. United States*, 521 U.S. 898, 935 (1997).

¹⁷¹ *Murphy v. NCAA*, 584 U.S. 453, 475-80 (2018) (distinguishing legislative commands, private-actor preemption and evenhanded regulation of public and private activity). See *supra* note 170 (*Printz*).

narrows. The institutional mapping in Part IV informs the proposed duties in Part V.

C. State Police Powers and the Regulatory Heartland

1. Public Health, Safety, Morals, and Welfare Authority Reaches AI-Mediated Decisions

State regulatory authority over the deployer-centered domains at issue, including employment, housing, insurance, healthcare, education, public benefits, and criminal justice, rests on the classical state police power. The doctrine’s canonical formulation reaches public health, safety, morals, and welfare, with these categories illustrative rather than exhaustive of the power’s scope. The classical formulation operates against the federal supremacy limits Section B developed and against individual constitutional rights addressed at Section C.2 below. Chief Justice Marshall recognized state regulatory authority over “[i]nspection laws, quarantine laws, [and] health laws of every description” in *Gibbons v. Ogden*, framing it as a residuum of state sovereignty preserved by the Constitution rather than granted by it.¹⁷² *Lawton v. Steele* observed that the police power “is universally conceded to include everything essential to the public safety, health, and morals.”¹⁷³ *Berman v. Parker* placed “[p]ublic safety, public health, morality, peace and quiet, law and order” among the “more conspicuous examples of the traditional application of the police power,” explicitly noting that the listed categories “merely illustrate the scope of the power and do not delimit it.”¹⁷⁴ Modern federalism scholarship reframes the structural significance of these state authorities, treating them less as a sovereignty account and more as a structural allocation of regulatory power across institutions, but the underlying authority itself remains undisturbed.¹⁷⁵ The “regulatory heartland” framing this Section uses draws on the New Nationalism federalism literature that locates state regulatory authority not in pre-constitutional residuum but in the structural allocation the dual-sovereignty architecture continuously reproduces; the framing used here emphasizes the institutional allocation of regulatory power. That allocation remains subject to valid federal displacement of particular state duties.¹⁷⁶ State AI regulation in regulated decision environments is the application of pre-existing state police-power authority to a new technological substrate within which consequential decisions are made.

States retain quasi-sovereign authority to protect the welfare of their populations against harms operating across state borders. *Massachusetts v. EPA* held that “States are not normal litigants for the purposes of invoking federal jurisdiction” and that Massachusetts could press its quasi-sovereign interest in protecting state territory and citizens against greenhouse-

¹⁷² *Gibbons v. Ogden*, 22 U.S. (9 Wheat.) 1, 203 (1824) (Marshall, C.J.).

¹⁷³ *Lawton v. Steele*, 152 U.S. 133, 136 (1894).

¹⁷⁴ *Berman v. Parker*, 348 U.S. 26, 32-33 (1954).

¹⁷⁵ See Heather K. Gerken, *The Supreme Court, 2009 Term—Foreword: Federalism All the Way Down*, 124 Harv. L. Rev. 4, 8 (2010).

¹⁷⁶ The “regulatory heartland” framing draws on the New Nationalism federalism literature. See Jessica Bulman-Pozen, *Partisan Federalism*, 127 Harv. L. Rev. 1077 (2014); Jessica Bulman-Pozen & Heather K. Gerken, *Uncooperative Federalism*, 118 Yale L.J. 1256 (2009); Cristina M. Rodríguez, *The Significance of the Local in Immigration Regulation*, 106 Mich. L. Rev. 567 (2008); Abbe R. Gluck, *Our [National] Federalism*, 123 Yale L.J. 1996 (2014). The framing’s load-bearing content for this paper is the structural-allocation account of state regulatory authority in police-power domains, which the New Nationalism scholarship developed against the older sovereignty-residuum account associated with the Rehnquist Court federalism revival.

gas emissions originating outside the state.¹⁷⁷ The Court drew on *Alfred L. Snapp & Son, Inc. v. Puerto Rico*, which had articulated the modern quasi-sovereign-interest doctrine: a state has cognizable interests in the physical and economic well-being of its residents, and in not being discriminatorily denied the benefits of the federal system.¹⁷⁸ The translation from quasi-sovereign-interest standing in federal court to prescriptive state regulatory authority is by analogy rather than by direct doctrinal extension: *Snapp* articulated the quasi-sovereign interest as a *justiciability* anchor for *parens patriae* federal-court suits, and the prescriptive question, whether a state may legislatively regulate out-of-state conduct producing in-state effects, is governed by the extraterritoriality doctrine treated at Section A.2 above and by substantive Due Process limits on punitive state regulation of out-of-state conduct.¹⁷⁹ The analogy is doctrinally available because the state interest the *Snapp* line recognizes (protection of in-state residents in regulated decision environments) is the same state interest the police power affirmatively asserts. AI-mediated decisions made by systems developed, trained, and operated by entities outside the regulating state produce in-state effects on in-state residents in regulated decision environments. The state's interest in regulating those in-state effects is one the Court has long recognized in the standing context and one that pre-AI sectoral state regulation already operationalizes through legislative means.

Anti-commandeering limits federal commands to state legislatures and officers. *Printz v. United States* and *Murphy v. NCAA* distinguish those commands from valid preemption and other permissible federal regulation.¹⁸⁰ Police powers provide an affirmative state source, while the relevant federal law determines whether a particular state duty is displaced. The presumption discussed in Section B.2 operates within that inquiry; it is not a categorical floor protecting all four authority channels.

Conditional spending requires a separate inquiry into the funding statute, the condition, notice and the program relationship. *South Dakota v. Dole* states limits on the spending power, and *NFIB v. Sebelius* addresses coercion in the Medicaid expansion.¹⁸¹ A state may avoid a particular funding condition by declining the funds to which it lawfully attaches. That choice does not eliminate otherwise applicable federal requirements, and an apparently independent licensing or procurement channel is not automatically insulated from every funding consequence. The enabling law and actual program facts must support the proposed design.

¹⁷⁷ *Massachusetts v. EPA*, 549 U.S. 497, 518 (2007).

¹⁷⁸ *Alfred L. Snapp & Son, Inc. v. Puerto Rico ex rel. Barez*, 458 U.S. 592, 600-08 (1982) (articulating modern quasi-sovereign-interest doctrine; canonical formulation at 607-08).

¹⁷⁹ The translation from quasi-sovereign-interest standing to prescriptive state regulatory authority operates by analogy rather than by direct doctrinal extension; the prescriptive limits on state regulation of out-of-state conduct producing in-state effects are governed by extraterritoriality doctrine treated *supra* Section A.2 and by substantive Due Process limits on punitive damages and regulation of out-of-state conduct. See *BMW of N. Am., Inc. v. Gore*, 517 U.S. 559 (1996); *State Farm Mut. Auto. Ins. Co. v. Campbell*, 538 U.S. 408 (2003); see also Margaret H. Lemos, *Aggregate Litigation Goes Public: Representative Suits by State Attorneys General*, 126 Harv. L. Rev. 486 (2012) (theorizing state quasi-sovereign authority across justiciability and regulatory contexts).

¹⁸⁰ See *supra* notes 170-171 (citing *Printz*, 521 U.S. at 935, and *Murphy v. NCAA*, 584 U.S. at 479).

¹⁸¹ *South Dakota v. Dole*, 483 U.S. 203, 207-08, 211-12 (1987) (conditional-spending requirements and the separate coercion limit); *NFIB v. Sebelius*, 567 U.S. 519, 575-86 (2012) (Roberts, C.J., announcing the judgment; Part IV joined by Breyer and Kagan, JJ.) (threat to existing Medicaid funds distinguished from conditions on newly offered expansion funds). EO 14365 § 5(a) directs a later BEAD eligibility notice subject to law; the order directs later implementation rather than imposing a State-specific AI eligibility denial. See *supra* note 139. Declining funds avoids an attached condition, not every other federal constraint. Broader program-specific applications require the enabling statute, condition, notice, funding facts and constitutional analysis; they are reserved for further work.

These inquiries remain distinct when a state designs AI duties. Valid preemption can displace a requirement even where the state declines federal funds. A funding condition requires statutory authority and constitutional analysis; an executive recommendation alone does not establish either a denial of funds or a preemption rule. Anti-commandeering protects the state decision-making role within its own limits. Consumer protection, civil rights, licensing and procurement must each be examined for the duty, federal source, institution and remedy actually involved.

2. AI Decision Environments Sit Within Traditional State Police-Power Domains

The seven high-impact domains the deployer-centered architecture targets are state-regulated subjects in which AI changes the substrate of consequential decisions but does not change the regulatory subject. Employment decisions are subject to state labor and employment law before AI mediates them; AI-mediated employment decisions are subject to the same body of law. Housing decisions are subject to state landlord-tenant law and concurrent Fair Housing Act enforcement; AI-mediated housing decisions are subject to the same framework. Insurance underwriting is subject to state insurance commissioner regulation and McCarran-Ferguson Act protection; AI-mediated underwriting is subject to the same regime. Healthcare decisions are subject to state medical licensing, state hospital licensing, and FDA regulation; AI-mediated clinical decisions sit within those licensing regimes. Education decisions are subject to state education authority and FERPA's concurrent federal layer; AI-mediated educational decisions sit within state and federal authority alike. Public benefits administration is state-administered federal-program territory; AI-mediated benefits decisions remain within state administrative law and federal due-process review. Criminal justice decisions are subject to state criminal procedure; AI-mediated risk assessment in criminal justice sits within state criminal-justice authority. The convergence is doctrinal rather than novel: AI does not create new regulatory subjects; it deploys a new technological substrate within which existing regulatory subjects produce consequential decisions affecting persons.

The Article focuses on duties governing consequential decisions about people. Some such duties may still implicate speech, due process or equal protection; the focus is not an exemption from those constraints. *Moody v. NetChoice* requires an application-specific analysis of the challenged laws, and *X Corp. v. Bonta* concerns preliminary relief against identified content-reporting provisions. DOJ's *Weiser* pleading advances an Equal Protection theory; it is not a judicial determination about AI audits. These issues can matter within the decision environments treated here as well as in deepfake, provenance or content-moderation regulation. This Article maps the authority and preemption questions without supplying a complete substantive constitutional analysis of each proposed duty.¹⁸²

A note on what this Part does not address. The architecture engages structural constitutional constraints: Dormant Commerce Clause scrutiny, federal preemption pressure, and the police-power and anti-commandeering boundaries of state regulatory authority. State AI

¹⁸² See *Moody v. NetChoice, LLC*, 603 U.S. 707 (2024) (slip op.) (bound volume 603 had not issued as of this writing) (First Amendment review of state social-media regulation; methodological discomfort with breadth in state regulation of algorithmic curation); *X Corp. v. Bonta*, 116 F.4th 888 (9th Cir. 2024) (reversing the denial of a preliminary injunction as to Cal. Bus. & Prof. Code sections 22677(a)(3), (a)(4)(A) and (a)(5) on First Amendment grounds, and remanding for severability; the Dormant Commerce Clause count was pleaded but not pressed on the injunction motion); *X.AI LLC v. Weiser*, No. 1:26-cv-01515 (D. Colo.) (DOJ complaint in intervention alleging Equal Protection violations from Colorado SB 24-205 on a *Students for Fair Admissions*-derived theory); *Students for Fair Admissions, Inc. v. President & Fellows of Harv. Coll.*, 600 U.S. 181 (2023); see also *infra* note 183 (substantive constitutional review routed to successor scholarship).

law also faces substantive constitutional review under the First Amendment (content-based scrutiny for state AI laws affecting speech), federal Due Process (vagueness in consequential-decision definitions, and procedural due process for AI-mediated benefits and adverse decisions), and federal Equal Protection (state AI laws that produce disparate impact may face constitutional review on top of the statutory disparate-impact framework discussed in Section B.4). These substantive constitutional questions are reserved for separate treatment in future work.¹⁸³ Process-based deployer-centered regulation in regulated domains presents a different substantive-constitutional profile than model-centric content-based alternatives; full comparative treatment is reserved.

The architecture this Part develops reaches AI applications that fit within settled state regulatory domains. AI compute, AI in critical infrastructure environments, and emergent AI categories raise questions the architecture identifies but does not fully answer. I return to these limits in Part VI.C.

The constitutional analysis supports a design choice with conditions. The Dormant Commerce Clause calls for attention to the conduct regulated and the burden imposed. Preemption requires comparison with the applicable federal law. State police powers supply affirmative authority subject to those and other constitutional limits. I recommend locating responsibility at the consequential decision and matching each duty to an institution able and authorized to enforce it. The institutional mapping and proposed implementation include worked applications in healthcare, employment and insurance underwriting.

Part III: The Deployer-Centered Model: Accountability for AI-Mediated Decisions

A. Accountability Begins with the Decision and Its Author

1. The Decision Supplies a More Useful Starting Point than the Model Category

A state seeking accountability for an AI-assisted hiring rejection should begin with the rejection: who made it, what the system contributed, what information could explain it, and who can correct it. Starting there does not dispose of the constitutional questions in Part II. A supplier's documentation duty may be both necessary and legally defensible; an employer's assessment duty may impose interstate burdens or conflict with federal law. The advantage of beginning with the decision is that the questions acquire an actor, a use and a consequence. They can be answered against an actual command rather than against a label such as "model-centric" or "deployer-centered."

The same starting point helps a drafter accommodate technical change. A statute can specify the information needed to review a hiring decision without committing itself to the screening model's architecture. That is a design choice, not a promise that definitions never need revision.

¹⁸³ State AI law's substantive constitutional review under the First Amendment, federal Due Process, and federal Equal Protection is the subject of related work. *See, e.g.,* X Corp. v. Bonta, 116 F.4th 888 (9th Cir. 2024) (First Amendment review of California AB 587 content-moderation reporting statute); *Moody v. NetChoice*, 603 U.S. 707 (2024) (First Amendment review of state social-media regulation); cf. *Grayned v. City of Rockford*, 408 U.S. 104, 108-09 (1972) (vagueness standard). These cases illustrate the substantive-constitutional layer that operates concurrently with the structural-constitutional architecture this paper develops; full treatment is reserved for a successor project.

The banking agencies' April 2026 revision of their model-risk guidance illustrates the need to state a framework's limits. SR 26-2 excludes generative and agentic AI models as novel and rapidly evolving.¹⁸⁴ The guidance directs banking organizations to use their own risk management and governance for tools outside its coverage. Its scope illustrates the need to connect a framework to the risks it addresses; the choice says little about the capacity of state legislatures or how long a particular statute will remain useful.

Hiring, tenant screening and insurance underwriting give the inquiry a continuing subject even as their tools change. A large language model may produce a narrative rather than a score; either output can influence the offer, lease or premium. The decision remains a useful place to require reasons, records and review. The system's design still matters when it determines what information exists, which errors can be detected or whether an intervention will work. I recommend decision-linked accountability because it makes those dependencies explicit and assigns responsibility for acting on them.

2. The Deployer Bears Responsibility for Using AI to Make a Consequential Decision

The *deployer* is the entity that uses an AI system to make or substantially inform a consequential decision affecting a person. The definition follows the act of decision-making. It does not depend on the model's training method, capability class or ownership. A person who supplies a system and also uses it to decide may perform both roles; assigning responsibility to the deployer does not extinguish the supplier's duties.

Colorado's 2026 successor act supplies a statutory comparison. It reaches the use of automated decision-making technology to materially influence consequential decisions in specified domains, while retaining developer documentation and other duties. Its general commencement is January 2027, with specified provisions taking effect earlier.¹⁸⁵ The act is an enacted design from which to learn, rather than evidence that a wholly downstream regime is already operating successfully. The definition used here is the paper's own organizing term; each existing statute retains its actual coverage and exceptions.

Many deployers already act within established legal relationships. An insurer writes a policy, an employer selects an applicant, a hospital delivers care and a landlord screens a tenant. Those relationships give a state possible institutions and legal predicates for imposing a duty. They do not confer unlimited authority over every actor in the chain. Licensing coverage, federal overlap, public and private roles, and available remedies differ across sectors. Part IV examines those differences rather than treating the existence of a regulator as sufficient authority for an AI obligation.

The corporate-law analogy has a narrower use. *CTS Corp. v. Dynamics Corp. of America* upheld Indiana's control-share statute in its setting of state-chartered corporations and voting rights. Its single-incorporation rationale helps explain why conflicting commands were less likely there; it does not establish a single licensing state for national employers, hospitals or

¹⁸⁴ Federal Reserve, FDIC & OCC, *Revised Guidance on Model Risk Management*, SR 26-2 (Apr. 17, 2026), attachment (scope exclusion for generative and agentic AI models as novel and rapidly evolving). The exclusion states the guidance's scope; it does not establish regulatory incapacity or a measured rate of statutory obsolescence.

¹⁸⁵ S.B. 26-189, 2026 Reg. Sess. (Colo.), §§ 1, 5 (repealing and reenacting Colo. Rev. Stat. §§ 6-1-1701 to -1707; definitions, developer duties and commencement); see *supra* note 88. The definition of covered consumer is not confined to residence; the statutory categories and exceptions govern.

landlords.¹⁸⁶ Decision-linked regulation must confront the possibility that several states will govern the same enterprise. The territorial and cumulative-burden questions developed in Part II remain questions about the particular duty.

The European Union's AI Act offers a comparison of legal design. Article 26 directs deployers of high-risk systems to follow instructions, provide competent human oversight, address input data within their control, monitor operation, preserve logs subject to the applicable qualifications, and give specified workplace notices. Article 27 adds a fundamental-rights impact assessment for a narrower class.¹⁸⁷ These duties sit alongside provider, system and general-purpose-model obligations. The Act therefore supplies an example of responsibility distributed across the chain, not a demonstration that regulating the deployer requires abandoning upstream regulation.

The difference between Articles 26 and 27 matters for borrowing. Article 27(1) reaches deployers that are public-law bodies or private entities providing public services, together with deployers of the creditworthiness and life-and-health-insurance systems specified in Annex III. It concerns high-risk systems under Article 6(2), with the critical-infrastructure category excepted.¹⁸⁸ Annex III's creditworthiness provision excludes financial-fraud detection; the insurance provision concerns risk assessment and pricing for natural persons in life and health insurance.¹⁸⁹ A private employer or school does not owe the assessment merely because employment or admissions appears in Annex III. Public-body or public-service status may bring it within Article 27. Nor does Annex III listing alone settle high-risk status: Article 6(3) supplies a qualified exception, including the proviso for profiling natural persons.

The worked state instrument takes a different approach. Its assessment coverage follows the specified domains and tiers, with public and private coverage governed by its actual provisions. That choice seeks accountability for consequential decisions by private actors as well as public ones. It is a recommendation to be justified against state authority, information access and cost. The EU's different boundary is not proof that private employment or housing requires no assessment, and the wider proposed boundary is not proved by describing it as deployer-centered. Borrowing a procedure requires a separate decision about whom it binds.

The filing mechanisms also need to be kept distinct. Article 27(3) requires notification of assessment results to the market-surveillance authority through the prescribed template. The July 2026 amendment permits overlapping assessment elements to be met through cross-references to, or incorporation of relevant parts of, a data-protection impact assessment.¹⁹⁰ Reuse can reduce duplicative preparation; this amendment alone does not establish an unread-document crisis or its cause. Notification to an authority is also different from making a full assessment

¹⁸⁶ *CTS Corp. v. Dynamics Corp. of Am.*, 481 U.S. 69, 88–94 (1987); see *supra* notes 89, 90, 91. The incorporation and voting-rights setting bounds the analogy; sectoral licensing does not reproduce its single-state allocation.

¹⁸⁷ Regulation (EU) 2024/1689, arts. 26–27, 2024 O.J. (L 2024/1689). Article 26 contains qualifications governing control, records, notice and sectoral duties. The comparison concerns legal design under different constitutional arrangements.

¹⁸⁸ Regulation (EU) 2024/1689, arts. 6(2)–(3), 27(1), Annex III pts. 2–5, consolidated July 27, 2026. Article 27's institutional and specified-use predicates differ from Article 26's high-risk-deployer scope.

¹⁸⁹ Regulation (EU) 2024/1689, Annex III pts. 5(b)–(c) (creditworthiness and life-and-health-insurance risk assessment and pricing).

¹⁹⁰ Regulation (EU) 2024/1689, arts. 27(3)–(5), 49, 71, Annex VIII sec. C; Regulation (EU) 2026/1744, art. 1(13), 2026 O.J. (L 2026/1744)(July 24, 2026). Assessment notification, reuse of overlapping DPIA elements and registration are distinct requirements; Article 27 does not itself require universal public release of the full assessment.

public. Articles 49 and 71 separately require qualifying public-authority and acting-on-behalf deployers to register specified uses, with public assessment summaries under the applicable fields and nonpublic treatment for specified sensitive registrations. That scope differs from Article 27’s public-service and specified financial-use triggers. A state choosing public summaries, confidential regulator submissions or both must specify the recipient, content and lawful withholding rules rather than borrow the word “filing.”

Commencement limits what the comparison can establish. The amended Article 113 generally places the Chapter III Sections 1–3 requirements for Annex III high-risk systems at December 2, 2027, and those for the Annex I route at August 2, 2028, subject to its specified exceptions and the Act’s transitional rules.¹⁹¹ These are future design commitments at the date of writing. They supply neither measured costs nor demonstrated improvements in decisions. Other provisions of the Act have different commencement rules; the dates do not mean that the entire Act is awaiting operation.

Two narrower features are useful for state design. Article 26(5) treats a specified monitoring obligation as fulfilled for financial institutions complying with the relevant internal-governance rules under Union financial-services law.¹⁹² That is a defined instance of using existing supervision, rather than a general exemption for regulated entities. Article 26(4)’s input-data obligation is limited by the deployer’s control over those data. The proposed audit trail likewise distinguishes information the deployer holds or can reasonably obtain from information it cannot secure. The tasks differ, but both require a drafter to specify access instead of assuming that the decision-maker possesses every relevant record.

The comparison supports a practical recommendation: assign the deployer responsibility for its use, identify the supplier information that responsibility requires, and fit each duty to the institution that can administer it. Neither a shared design feature nor a difference in scope establishes independent convergence, constitutional necessity or effectiveness. Those are further propositions with further evidence to supply.

3. AI-Mediated Decisions Identify the Use for Which Accountability Is Owed

An *AI-mediated decision* is a decision affecting a person’s rights or interests in which an AI system contributes to or determines the outcome. The phrase connects the technology to the act for which someone must answer. It includes a generated recommendation or narrative that materially informs a decision as well as a score that directly determines it.

Existing statutory terms can perform that work too. New York City’s “automated employment decision tool” and California’s employment “automated-decision system” definitions connect systems to employment decisions; their use of “tool” or “system” does not itself make their duties regulation of model development.¹⁹³ “Algorithmic decision-making” need not exclude generative uses merely because the term predates them. Colorado’s “consequential decision” identifies the significance and domains of the outcome; the surrounding definitions and commands

¹⁹¹ Regulation (EU) 2024/1689, art. 113(c), as amended by Regulation (EU) 2026/1744, art. 1(40)(b); see *supra* note 190. The dates are subject to the specified Article 6(5) exception and applicable transitional provisions; they are not a single commencement date for the Act.

¹⁹² Regulation (EU) 2024/1689, art. 26(5), second subparagraph (specified monitoring treatment for financial institutions subject to relevant Union financial-services internal-governance requirements); see also *id.* art. 26(4)(input data within deployer control).

¹⁹³ N.Y.C. Admin. Code §§ 20-870 to -874; 6 R.C.N.Y. § 5-300; *infra* notes 196,233. Cal. Code Regs. tit. 2, § 11008.1(a); see *infra* note 218. The definitions connect computational systems or outputs to employment decisions; terminology alone does not classify the entire regime’s architecture.

establish the technology's role.¹⁹⁴ The choice of vocabulary helps exposition. The actor, trigger and duty determine the legal obligation.

A **decision environment** is a recurring context in which an identified actor exercises authority over an outcome that materially affects a person's rights or interests. The decision and the actor's role fix the inquiry: selecting an applicant, underwriting a policy or deciding eligibility for a benefit. An *AI-mediated decision environment* is such a context in which AI contributes to the decisions. It is not the model or the technical infrastructure standing alone.

A decision environment is **regulated** in the sense used here when the state already governs the decision-maker or the decision through consumer protection, civil rights, sectoral licensing or public procurement. The seven domains in Section C organize consequential uses within those relationships. They are the proposed instrument's selected domains, not a claim that other domains lack state authority or that each of the seven has identical legal coverage. Part IV identifies the institution, predicate and remedy available for the particular obligation.

These definitions do not locate every decision inside the regulating state. A national employer may use a shared system, draw on records held elsewhere and make decisions involving people with different connections to the forum. Resident coverage, personal jurisdiction, substantive territorial reach and federal overlap remain separate inquiries. A law does not satisfy them by calling its subject an in-state decision. The same is true of interstate burden: a locally framed assessment can require national changes, while a supplier information duty may concern a specified covered use.

The contribution is the connection between a decision, a justified duty and an available institution. Existing assessment, audit, civil-rights and sectoral regimes already supply important components; Section B identifies several. The proposed architecture asks how to arrange those components so that the person making the decision can explain it, the reviewer can obtain relevant evidence, and an institution has lawful power to act. Beginning with the decision makes that design test concrete. It does not decide in advance which duties will withstand constitutional or federal-law objections.

B. Existing Laws Supply Components, Not a Single Architectural Verdict

The historical inventory records 172 instruments from 46 jurisdictions among 249 research records. The count covers qualifying state-level enactments and adopted regulations from January 1, 2021 through May 31, 2026; it counts instruments, including qualifying amendments to in-window base laws, rather than distinct regimes or obligations. A future operative date does not exclude an enacted instrument, and a superseded enactment may remain in this historical population. New York City's law, model guidance and selected earlier or later enactments appear below as comparisons, with their different status identified.¹⁹⁵

¹⁹⁴ S.B. 26-189, 2026 Reg. Sess. (Colo.), § 1 (Colo. Rev. Stat. § 6-1-1701); see *supra* note 88.

¹⁹⁵ The historical inclusion rule admits a state or District of Columbia statute or finally adopted regulation with a qualifying operative provision concerning the development, provision, deployment, procurement, use or output of an AI system, including specified automated decision and identification systems. A technology-specific definition or exemption can qualify without imposing a new duty. Amendments to a pre-window base law remain outside the numerator under the adopted convention; implementing regulations are folded into an already included base instrument rather than counted again. The rule distinguishes obligations governing systems or their use from content-only measures; it excludes general privacy statutes' profiling provisions but admits ADMT-specific regulations. It also

The expanded inventory supports a wider account of legal mechanisms. It does not establish which architecture predominates, which domain receives adequate regulation or why a legislature chose a design. An act may combine developer disclosures, deployer assessments, government procurement, conduct prohibitions and civil remedies. The descriptions below follow those commands and their bearers. The New York/California comparison in Section B.4 then places selected mechanisms alongside one another without treating their labels as constitutional outcomes.

1. Duties Run to Users, Suppliers and Public Bodies, Often in the Same Instrument

Several established measures attach duties to an entity using AI to act on another person. New York City's Local Law 144 requires employers and employment agencies to satisfy an independent bias-audit and publication condition before using a covered automated employment decision tool. Its separate notice rule concerns city-resident candidates and employees; that residency condition is not the whole law's coverage test.¹⁹⁶ Illinois's 2024 human-rights amendment addresses discriminatory AI use in employment and notice of such use.¹⁹⁷ California's privacy regulations combine risk-assessment duties with significant-decision ADMT notice, opt-out and access rights, subject to their definitions, exceptions and phased dates.¹⁹⁸ Utah regulates specified generative-AI interactions in consumer transactions, with separate treatment of regulated occupations.¹⁹⁹ Colorado's insurance law and regulation address insurers' use of external consumer data, algorithms and predictive models; the NAIC model bulletin offers a different, subregulatory form of insurer governance.²⁰⁰ Illinois's Video Interview Act requires notice, explanation and consent for covered AI-analyzed interviews, alongside sharing

excludes, absent a separate qualifying limb, physical-autonomy regulation, transaction-automation software, fiscal incentives, instruction about AI, task-force mandates, court rules, guidance and executive policy. These choices bound the count. The 249 records comprise 172 counted instruments, 46 pre-window records, 21 post-window records, 2 enactment-date holds, 6 near misses, 1 multistate model record and 1 negative record. All 51 jurisdictions have search records; 26 disclosed search-limb gaps span 22 jurisdictions. A bounded record audit covered 17 rows and 425 fields, with substantive primary reads for 15 rows. The vetoed Virginia bill retains an operative-text access limit. Indiana's administrative-filing generative-AI statement rule, 41 Ind. Admin. Code § 2-2-5(c), was subsequently opened in its 2026 final-rule text; that later access does not enlarge the historical audit. These are recorded historical quantities and disclosed limitations, not a certification of exhaustive recall, every field or law in force on October 16, 2026. Municipal Local Law 144, the NAIC model, the original 2019 Illinois Video Interview Act, vetoed Virginia H.B. 2094 and unenacted New Jersey A. 3854 do not enter the historical numerator. Selected post-May 2026 enactments discussed below are not silently added to it.

¹⁹⁶ N.Y.C. Admin. Code §§ 20-870 to -874; 6 R.C.N.Y. §§ 5-300 to -304. The audit/publication condition and resident-notice requirement have distinct triggers; no passing audit score or unconditional alternative-selection right is supplied by these provisions.

¹⁹⁷ 775 Ill. Comp. Stat. 5/2-101, 5/2-102, as amended by Pub. Act 103-0804(2024)(effective Jan. 1, 2026).

¹⁹⁸ Cal. Code Regs. tit. 11, §§ 7001, 7150–7157, 7200, 7220–7222 (approved 2025; effective Jan. 1, 2026, with separate compliance dates). Section 7153's provider-business information duty and § 7222's access/assistance provisions qualify an exclusively user-facing description.

¹⁹⁹ Utah Code ch. 13-77 (generative-AI consumer-transaction disclosure provisions, including regulated occupations). Chapter 13-75 is not the current generative-AI chapter; the enacted-bill and codified locators differ.

²⁰⁰ Colo. Rev. Stat. § 10-3-1104.9; 3 Colo. Code Regs. § 702-10:10-1-1; NAIC, *Model Bulletin: Use of Artificial Intelligence Systems by Insurers* (Dec. 4, 2023). The Colorado statute/regulation and model bulletin have different legal forms. A model's availability does not establish adoption or identical duties in every jurisdiction.

and deletion restrictions.²⁰¹

These examples establish available mechanisms, not an exclusively downstream family. California's assessment rules include a limited information duty for a business providing specified ADMT to another business. An employer may need supplier records to assess its use; a vendor may perform functions that bring it within an existing employment-law definition. The location of a duty in a consumer, employment or insurance regime tells a drafter where to begin the authority inquiry. It does not establish that all compliance occurs in the forum, that every supplier is excluded, or that the duty will survive every federal overlap.

Other provisions address development, information about a system or the means of producing harmful outputs. California's AI Transparency Act is one evolving regime: the 2024 enactment, the 2025 expansion and the September 2026 amendments must be read together. The latest amendments remove the original provider-size threshold and revise the licensing response; later platform and device duties have their own dates.²⁰² California's training-data law separately requires high-level public documentation of specified datasets used in developing covered generative-AI systems. That is a documentation command, not by its terms an obligation to publish the entire dataset or retrain the model.²⁰³

California's frontier act also differentiates its bearers. It places framework duties on large frontier developers while imposing specified transparency and incident duties on frontier developers more generally. Its civil-penalty provision expressly concerns large frontier developers.²⁰⁴ Tennessee's ELVIS Act reaches both specified uses of an individual's identity and the knowing provision of technology whose primary purpose or function is producing a particular identifiable person's photograph, voice or likeness, subject to the statute's conditions and exceptions.²⁰⁵ New York's RAISE enactment was materially replaced in March 2026; the amended framework, generally commencing January 2027, likewise distinguishes frontier developers from large frontier developers.²⁰⁶ A threshold used for one obligation cannot be assumed to govern every duty

²⁰¹ 820 Ill. Comp. Stat. 42/1–20 (Artificial Intelligence Video Interview Act, original 2019 act and 2021 amendment). The historical count excludes the pre-window base act and, under its stated convention, its 2021 amendment; it does not treat that amendment as a counted instrument.

²⁰² Cal. Bus. & Prof. Code §§ 22757–22757.6; 2024 Cal. Stat. ch. 291 (S.B. 942), as amended by 2025 Cal. Stat. ch. 674 (A.B. 853), 2026 Cal. Stat. ch. 861 (S.B. 1000) and ch. 856 (A.B. 2713); see *supra* notes 76, 78. Original and amended provisions form a regime history; the original million-user threshold and 96-hour automatic provider-revocation description do not state the September 2026 licensing mechanism.

²⁰³ Cal. Civ. Code §§ 3110–3111; 2024 Cal. Stat. ch. 817 (A.B. 2013); see *supra* note 79. The public documentation duty, release/modification timing and specified sole-purpose exceptions govern its scope.

²⁰⁴ Cal. Bus. & Prof. Code §§ 22757.10–22757.16, especially §§ 22757.11(i)–(j), 22757.12–22757.15; Cal. Gov't Code § 11546.8; Cal. Lab. Code § 1107; 2025 Cal. Stat. ch. 138 (S.B. 53). The training threshold is greater than 10²⁶ operations; the large-developer revenue threshold exceeds \$500 million with affiliates in the preceding calendar year. Transparency/incident duties and framework/penalty coverage must be distinguished.

²⁰⁵ Tenn. Code Ann. §§ 47-25-1101 to -1108, especially §§ 1105–1107 (2025 code reproduction; reflecting 2024 Tenn. Pub. Acts ch. 588); see *supra* note 80. Liability for the technology limb depends on the enacted knowledge and purpose/function language and surrounding exceptions, not on the mere existence of a generative model.

²⁰⁶ 2025 N.Y. Laws ch. 699 (S. 6953-B), as amended by 2026 N.Y. Laws ch. 96; N.Y. Gen. Bus. Law art. 44-B (substantive amended framework commencing Jan. 1, 2027). The original and replacing acts are one regime history, not two independent current frameworks. The original ninety-day commencement and the March 27, 2026 replacement require a separate interval qualification; the amendment alone does not prove that the original never operated.

or remedy in the act.

Illinois's July 2026 frontier enactment supplies a later comparison outside the historical window. It combines transparency before or concurrently with deployment and incident reporting, including a shorter imminent-harm clock, with additional framework and independent-audit duties for large developers phased into 2028. Its filing provisions also permit agency sanctions alongside the Attorney General civil-penalty route. Its express in-state language for a disclosure-statement filing should not be silently carried into every other provision.²⁰⁷ These upstream requirements may demand substantial information or changes in practice. Their constitutional position depends on the command and relevant facts, not on their inclusion in a family called model-centric.

Colorado shows why the categories overlap. The 2024 act assigned duties to both developers and deployers.²⁰⁸ The 2026 replacement changes those duties and the consequential-decision framework, but retains developer documentation, defined uses and limitations, records and allocation of responsibility.²⁰⁹ Its consumer-facing provisions distinguish pre-use notice, post-adverse-outcome explanation, correction of factually inaccurate personal data, and requested meaningful review subject to the applicable conditions. Section 1707 also preserves specified antidiscrimination liability and allocates responsibility between developers and deployers.²¹⁰ The replacement changes the design; it does not demonstrate that mixed responsibility became constitutionally untenable or that the legislature proved a downstream alternative superior. The text alone cannot decide among constitutional concern, implementation experience, lobbying or other explanations for the change. Virginia's vetoed 2025 bill is a failed proposal, not another operative regime or a judicial verdict on a shared architecture.²¹¹

Texas uses another combination. TRAIGA prohibits specified conduct by governmental and private actors, includes healthcare disclosure, and creates a sandbox alongside enforcement and governance provisions. The reach and exceptions differ by command.²¹² The healthcare statutes discussed below are also distinct: California governs specified utilization-review technology and professional determinations; Maryland imposes its own conditions on AI-supported review; Indiana's mandatory human review concerns a narrower downcoding use; Arizona's medical-director review provision does not name AI.²¹³ They are neighboring responses to med-

²⁰⁷ Artificial Intelligence Safety Measures Act, Pub. Act 104-0538(2026)(S.B. 315), §§ 5, 10, 15, 18, 25, 99. The act postdates May 31, 2026 and is outside the 172-instrument numerator. The enacted text, rather than the introduced technical-amendment synopsis, controls. General commencement is Jan. 1, 2027; the large-developer framework and independent-audit provisions have specified 2028 timing. The disclosure-statement filing's territorial language does not itself resolve every other provision's reach.

²⁰⁸ S.B. 24-205, 2024 Reg. Sess. (Colo.), § 1 (Colo. Rev. Stat. §§ 6-1-1701 to -1707); see *supra* note 87. Developer and deployer duties belong to the historical act, whose later replacement and procedural status are separate questions.

²⁰⁹ S.B. 26-189, 2026 Reg. Sess. (Colo.), §§ 1, 5 (signed May 14, 2026; general commencement Jan. 1, 2027, with specified earlier provisions); see *supra* note 88.

²¹⁰ *Id.* § 1 (Colo. Rev. Stat. §§ 6-1-1702 to -1707). The notice, explanation, correction and review provisions have distinct triggers and exceptions; the act's no-new-private-action provision does not eliminate the specified liability under other law addressed in § 1707.

²¹¹ H.B. 2094, 2025 Reg. Sess. (Va.), bill-status history reproduced by LegiScan (veto Mar. 24, 2025; sustained Apr. 2, 2025), <https://legiscan.com/VA/bill/HB2094/2025>. The vetoed proposal supplies no operative duty or penalty.

²¹² Tex. Bus. & Com. Code chs. 551–554; H.B. 149, 89th Leg., R.S.(2025)(TRAIGA). Territory, prohibited conduct, healthcare disclosure, sandbox and enforcement provisions must be examined separately.

²¹³ 2024 Cal. Stat. ch. 879 (S.B. 1120)(Cal. Health & Safety Code § 1367.01(k); Cal. Ins. Code §

ical decisions, not interchangeable AI prohibitions.

Public use supplies further combinations. Connecticut's 2023 act assigns inventory and assessment responsibilities to executive institutions and the Judicial Department. Its 2026 act adds employment-related AI duties with later deployment dates, developer information provisions and a separate frontier-whistleblower mechanism.²¹⁴ Treating Connecticut as procurement-only would miss both its earlier assessment structure and its enacted expansion. Public and private coverage must be read from the provision being compared.

Litigation gives another reason to resist an architectural scorecard. The April 27, 2026 stay in *X.AI LLC v. Weiser* has specified procedural triggers and reaches same-session replacement legislation on its stated terms. A stay is not a merits determination that developer duties violate the Dormant Commerce Clause. Nor do the complaint's First Amendment claims, the government's Equal Protection allegations or the replacement statute establish why Colorado legislated.²¹⁵ The legal design inquiry in Part II remains necessary after the inventory is complete.

2. Sectoral Examples Show Different Ways to Govern the Use

The laws can also be compared by the decisions and relationships they address. Cross-cutting acts such as Colorado's, Texas's and Utah's differ from California's privacy rules in coverage, prohibited conduct, exceptions and procedural rights. None supplies a single template that can be carried unchanged into every domain.²¹⁶ Sectoral examples make the design choices easier to see because the decision, affected person and existing regulator are more readily identified. The selection below illustrates mechanisms; it is not a ranked measure of regulatory intensity across the 172 instruments.

Employment supplies several distinct mechanisms. New York City's audit condition, Illinois's human-rights amendment and the Video Interview Act allocate responsibility differently. California's Civil Rights Council addresses automated systems through employment-discrimination regulations rather than by imposing a universal annual audit. New Jersey's 2024 Assembly Bill 3854 illustrates a proposed audit-and-notice design; its introduced text is used here as

10123.135(j)); 2025 Ariz. Sess. Laws ch. 165 (H.B. 2175)(Ariz. Rev. Stat. §§ 20-3103, 20-3407); 2025 Md. Laws ch. 747 (H.B. 820)(Md. Code Ann., Ins. §§ 15-10B-05.1, 15-10B-07); Ind. Code § 27-1-52-9 (Pub. L. No. 88-2026); *see also* Tex. Bus. & Com. Code § 552.051. Arizona's provision is an adjacent technology-neutral review requirement, not an AI-specific law in the counted population; *see infra* note 298.

²¹⁴ 2023 Conn. Pub. Acts 23-16, §§ 1–3; 2026 Conn. Pub. Acts 26-15, especially §§ 2, 8–12. The 2026 employment provisions' effectiveness and later deployment-applicability dates differ; the frontier-whistleblower mechanism has a separate predicate and remedy. This enactment was already discussed elsewhere in the paper.

²¹⁵ *X.AI LLC v. Weiser*, No. 1:26-cv-01515 (D. Colo.), ECF No. 24, ¶¶1–2 (Apr. 27, 2026); *see supra* note 143. Paragraph 1 prohibits initiating enforcement, including investigations, for alleged violations of SB 24-205 or specified same-session replacement or amending legislation occurring on or before fourteen days after the preliminary-injunction ruling. Paragraph 2 requires the motion, and any necessary amended complaint, within twenty-eight days after final implementing rulemaking. The dated order supplies procedure, not a constitutional holding, legislative motive, automatic mootness determination or certification of later docket status. First Amendment merits remain outside this paper's structural-federalism scope.

²¹⁶ *See supra* notes 185,198–199,208–212. Cross-cutting coverage is not identical across the named instruments; sectoral exceptions and information duties remain relevant.

proposal history, outside the operative-law denominator.²¹⁷ The California rules, effective October 2025, define automated-decision systems in relation to employment benefits, address discriminatory use and agents performing covered employer functions, and require specified employment records to be preserved for four years. Anti-bias testing may be relevant evidence under the rules without being a general annual-testing mandate.²¹⁸ New Jersey's December 2025 disparate-impact regulations illustrate a further legal form: an adopted, generally applicable discrimination rule with automated-tool applications. Its commands and available routes should be read from the rule rather than inferred from the existence of agency guidance.²¹⁹

Insurance illustrates supervision of an existing licensed relationship. Colorado's 2021 statute and implementing regulation address the use of external consumer data, algorithms and predictive models, with governance and testing requirements defined by the relevant provisions.²²⁰ The Colorado rule keeps responsibility with the insurer when it uses third-party resources and permits the vendor to provide requested records directly to the Division. New York's July 2024 insurance circular addresses governance, fairness analysis, documentation and information about third-party systems used in underwriting and pricing.²²¹ The NAIC model bulletin offers a common supervisory vocabulary, but each jurisdiction's adopted instrument and powers still matter. A shared model does not itself prove uniform duties, equivalent enforcement or lower compliance costs.

The NAIC's 2026 evaluation-tool pilot is a further institutional experiment. Its documents describe use in different supervisory work and a process for revising the tool; the August meeting materials record the name "AI Risk Evaluation Supplement" and a goal of later adoption.²²² A pilot plan and draft minutes do not establish that the instrument was finally adopted, that every participating regulator used it identically, or that it improved insurer decisions. The useful comparison is the effort to fit requests for evidence into existing supervision, together with the question of what information the insurer can obtain from a vendor.

Healthcare makes precision especially important. California's 2024 act constrains specified utilization-review technology and reserves covered medical-necessity determinations to qualified professionals. Maryland's 2025 law governs AI-supported review under its own conditions

²¹⁷ A. 3854, 221st Leg. (N.J. introduced Feb. 22, 2024), §§ 1–2. The introduced text proposes audit, publication and notice conditions; it supplies no enacted duty here. Later substitutes and final session status are not established by the introduced text.

²¹⁸ Cal. Code Regs. tit. 2, §§ 11008.1, 11009, 11013(c), 11017.1, 11021, 11028, 11031, 11051–11052 (Civil Rights Council automated-decision-system regulations, effective Oct. 1, 2025). The definition and four-year record rule differ from the CPPA's ADMT regime; testing is not a universal annual audit command.

²¹⁹ N.J. Admin. Code ch. 13:16, especially §§ 13:16-1.3, 13:16-2.1, 13:16-3.2(c) (effective Dec. 15, 2025). The DCR courtesy copy identifies its rulemaking history and states that the official published code controls. These are adopted disparate-impact rules, not an AI-only statute or the separate agency guidance.

²²⁰ Colo. Rev. Stat. § 10-3-1104.9; 3 Colo. Code Regs. § 702-10:10-1-1; see *supra* note 200. The applicable lines, governance provisions and implementation dates govern; no first-in-the-nation or statewide-outcome inference is claimed.

²²¹ N.Y. Dep't of Fin. Servs., Insurance Circular Letter No. 7(2024)(July 11, 2024), especially §§ 35–45 (third-party governance and information considerations). A circular's supervisory expectations and underlying insurance authority are distinct from an independently enacted general AI statute.

²²² NAIC, *AI Systems Evaluation Tool Pilot: Pilot Project Background*(2026); NAIC Big Data & Artificial Intelligence (H) Working Group, *Summer National Meeting Materials*, Attachment A (Aug. 13, 2026)(draft July 22, 2026 minutes, renaming and later-adoption objective). The described pilot and meeting plans are dated records, not proof of subsequent adoption or practical effectiveness.

and decision-maker requirements. Indiana's 2026 provision requires human review for specified AI-based medical-necessity downcoding; its prior-authorization disclosure has a different function. Arizona's neighboring requirement concerns individual review by a medical director without naming AI. Texas adds disclosure when AI is used in relation to healthcare services or treatment.²²³ These are different responses to different decisions. Similar concerns about human responsibility do not establish identical commands, causal borrowing among states or a measured rate of legislative diffusion.

Generative and frontier measures address additional objects. California's provenance regime concerns the detection and identification of specified generated audiovisual content; its dataset-documentation law concerns development information; its frontier act and New York's amended RAISE framework concern defined frontier risks and reporting. Tennessee's identity protections can reach specified enabling technology as well as harmful uses.²²⁴ A law can require information from a supplier without treating every later decision as the supplier's act. Conversely, naming the user does not guarantee access to the information necessary for meaningful review.

Relational chatbots provide another comparison. California's 2025 companion-chatbot act imposes specified disclosure and safety duties and supplies a civil action. Utah's mental-health-chatbot provisions use a different definition and make enumerated safeguards conditions of an affirmative defense; that structure should not be recast as an unconditional crisis-response mandate. Washington and Oregon enacted companion-chatbot measures with 2027 commencement, while Nebraska's conversational-AI provisions have July 2027 operation and an Attorney General route without a new private action.²²⁵ Definitions, minor protections, reminders, crisis provisions and remedies differ. "Companion-chatbot cluster" is a descriptive grouping used here, not a shared statutory category or evidence that the regimes already have comparable operating records.

Maryland's pricing law shows how a narrow decision can be governed through a conduct prohibition. The enacted 2026 act targets covered food retailers and third-party food-delivery providers: subject to enumerated exceptions, it prohibits specified personalized upward pricing of tax-exempt food using personal data and separately addresses protected-class data used to withhold advantages. It requires a notice-and-cure opportunity before the specified enforcement action and disclaims a private right of action under that section.²²⁶ Its scope does not

²²³ See *supra* note 213. Cal. Health & Safety Code § 1367.01(k); Cal. Ins. Code § 10123.135(j); Ariz. Rev. Stat. §§ 20-3103, 20-3407; Md. Code Ann., Ins. §§ 15-10B-05.1, 15-10B-07; Ind. Code § 27-1-52-9; Tex. Bus. & Com. Code § 552.051(d), (f). The statutory decision-makers, triggers and exceptions differ; Arizona's review provision does not itself name AI.

²²⁴ See *supra* notes 202–206. The current California provider-provenance duties must be distinguished from the later platform/device provisions and from dataset or frontier-risk documentation.

²²⁵ 2025 Cal. Stat. ch. 677 (S.B. 243)(Cal. Bus. & Prof. Code §§ 22601–22606); Utah H.B. 452, 2025 Gen. Sess. (mental-health-chatbot affirmative-defense provisions); 2026 Wash. Sess. Laws ch. 168 (ESHB 2225); 2026 Or. Laws ch. 85 (S.B. 1546); L.B. 525, 109th Leg., 2d Sess., §§ 12–18 (Neb. 2026). Washington/Oregon general commencement is Jan. 1, 2027; Nebraska's specified provisions operate July 1, 2027. The Utah safeguards condition a defense; Nebraska § 18 creates no private action. Washington's consumer-protection route must be read separately from an Attorney General-only model. The group label is the paper's; Utah and Nebraska use different statutory terms.

²²⁶ 2026 Md. Laws ch. 154 (H.B. 895), § 1 (Md. Code Ann., Com. Law §§ 13-321(a)–(e), 13-408(a), 13-411(a)), § 2 (approved Apr. 28, 2026; effective Oct. 1, 2026). The enacted marked text deletes proposed § 13-322 and its other-merchants disclosure. Section 13-321(a) defines the covered 15,000-square-foot food retailer and third-party delivery provider; subsection (b) states nine exceptions and the upward-food-pricing prohibition; (c) addresses protected-class data; (d) provides 45 days to cure; (e) disclaims a private action under this section or other law. The labor amendment was deleted; the enacted

establish a general prohibition on changing prices, and the proposal for an other-merchants disclosure rule was deleted before enactment. That difference changes both the duty and who can seek relief.

Synthetic media and identity laws form an overlapping comparison rather than a uniform AI-governance regime. Kentucky's broader state-use enactment also provides an electioneering-disclosure remedy; Arkansas amended publicity-rights definitions to include AI-generated or simulated likeness and voice; Montana addresses specified unauthorized commercial digital depictions and enabling technology, with knowledge requirements and exceptions. California's deepfake-pornography remedies and its restriction on implying that an AI health service comes from a licensed natural person address still different harms.²²⁷ Some content-only provisions are outside the historical count under its stated inclusion rule even though they bear on the paper's design choices.

These examples explain why the seven domains in Section C need separate applications. A prohibition on a use, an assessment before a decision, a disclosure to an affected person and an action for a harmful depiction do different work. The inventory does not by itself establish that employment is adequately governed or that education, public benefits and criminal justice receive less protection. Such comparisons would require a defined measure of coverage, obligation, enforcement and outcome, with the search gaps carried through. The proposal can justify attention to those domains by the importance of their decisions without turning incomplete descriptive evidence into a nationwide absence claim.

3. Enforcement Allocation Determines Who Can Act and What Relief Is Available

The same substantive duty can operate differently depending on who receives a complaint, obtains the records and has power to respond. The examples below show Attorney General actions, specialist supervision, civil-rights procedures and specific private remedies. Those routes may coexist within an instrument. The allocation is part of the design; its existence does not establish enforcement capacity or practical effectiveness.

Colorado's 2026 framework gives the Attorney General the specified enforcement role and creates no new private cause of action under that part. That limitation must be read with its express treatment of antidiscrimination liability and responsibility under other law.²²⁸ Texas likewise supplies Attorney General enforcement of Chapter 552, a sixty-day notice-and-cure procedure and graduated civil penalties. Its separate provision permits specified licensing-agency

amendments to §§ 13-408(a) and 13-411(a) exclude § 13-321 violations from those general private-action and criminal provisions.

²²⁷ See *supra* note 205 (Tennessee); 2025 Ky. Acts ch. 66 (S.B. 4), §§ 3–6 (state-use governance and separate synthetic-electioneering provisions); 2025 Ark. Acts 159 (H.B. 1071), §§ 1–3 (publicity definitions, not a statute using the term “digital replica”); 2025 Mont. Laws ch. 685 (H.B. 513), §§ 1, 3 (specified commercial uses/enabling technology, conditions and exceptions; effective Jan. 1, 2026); 2025 Cal. Stat. ch. 673 (A.B. 621)(Cal. Civ. Code § 1708.86); 2025 Cal. Stat. ch. 615 (A.B. 489)(Cal. Bus. & Prof. Code §§ 4999.8–4999.9). A publicity definition, political disclosure remedy, depiction action and licensed-person representation rule are distinct mechanisms; this note does not place every content-only provision in the historical numerator.

²²⁸ S.B. 26-189, 2026 Reg. Sess. (Colo.), § 1 (Colo. Rev. Stat. §§ 6-1-1706 to -1707); see also *supra* notes 208–210 for the historical and successor frameworks. The express enforcement/no-new-action provision and other-law liability allocation are distinct; an older general consumer-protection penalty is not presented as a measured AI enforcement result.

sanctions after an Attorney General recommendation and a qualifying violation.²²⁹ Calling either framework “public enforcement” is useful only if the reader retains the boundaries of the no-private-action clause and the other institutions’ powers.

California’s frontier act permits the Attorney General to recover the specified civil penalties from large frontier developers. The penalty provision does not define the scope of every substantive duty, and the act’s labor provisions supply a distinct employee route.²³⁰ New York’s amended RAISE framework assigns functions to DFS and the Attorney General under its specified provisions, with the civil-penalty language expressly naming large frontier developers.²³¹ A duty imposed on frontier developers generally and a penalty clause naming large developers require separate readings. The vetoed Virginia proposal supplies no operative public-enforcement mechanism.²³²

Specialist administration does not necessarily mean exclusive agency enforcement. New York City’s Local Law 144 provides civil penalties through the designated city tribunal and authorizes Corporation Counsel or designees to seek correction or compliance in court. It also preserves candidate and employee civil actions and the city human-rights commission’s authority. That savings language should not be mistaken for a newly created, freestanding private damages action for every Local Law 144 violation.²³³ The law’s audit, notice, penalty and preserved-rights provisions perform different functions.

California’s privacy regime gives the CCPA and Attorney General specified public powers; the CCPA’s limited private action for certain security incidents does not become a general ADMT claim merely because the ADMT rules implement the same statute.²³⁴ Utah’s generative-AI

²²⁹ Tex. Bus. & Com. Code §§ 552.101–552.106. Section 552.104 bars an action before the sixtieth day and after a qualifying cure with the required written assurance. Section 552.105 distinguishes curable violations or assurance breaches (\$10,000–\$12,000), incurable violations (\$80,000–\$200,000) and continued violations (\$2,000–\$40,000 per day), with additional qualifications. Section 552.106 conditions licensing-agency sanctions on a qualifying violation and Attorney General recommendation; its separate monetary cap is \$100,000. These are statutory authorizations, not fines shown to have been imposed.

²³⁰ Cal. Bus. & Prof. Code § 22757.15; Cal. Lab. Code § 1107.1; *see supra* note 204. The former permits Attorney General civil penalties against a large frontier developer up to \$1 million per violation, taking severity into account; the labor remedy prevents treating that clause as the entire act’s exclusive enforcement structure.

²³¹ N.Y. Gen. Bus. Law §§ 1422–1428, as enacted by 2026 N.Y. Laws ch. 96; *see supra* note 206. Section 1427’s civil penalties expressly concern a large frontier developer, with the specified first/subsequent limits; neither that subject nor the administering office substitutes for reading each duty’s actual bearer.

²³² H.B. 2094, 2025 Reg. Sess. (Va.); *see supra* note 211. Vetoed proposal history supplies no presently operative penalty or cure period.

²³³ N.Y.C. Admin. Code §§ 20-872 to -874; *see supra* note 196. Section 20-872 provides civil penalties and tribunal procedure; § 20-873 authorizes correction/compliance actions by Corporation Counsel or designees; § 20-874 preserves civil actions and Human Rights Commission authority. DCWP administration is not the same proposition as exclusive enforcement or a new universal private damages action.

²³⁴ Cal. Civ. Code §§ 1798.150, 1798.155, 1798.199.45, 1798.199.90(a); Cal. Privacy Prot. Agency, *Adjusted Monetary Thresholds and Fines Under the CCPA* (Dec. 17, 2024) (effective Jan. 1, 2025). The adjusted public-penalty ceilings are \$2,663 and \$7,988 under the respective statutory conditions. The agency’s higher-fine clause addresses intentional violations or actual knowledge of consumers under 16; the Attorney General clause uses intentional violations or violations involving minor consumers. Section 1798.150’s private action has its own security-incident predicates; it is not a general private ADMT cause of action.

chapter permits Division action, provides for court enforcement and preserves other remedies. Its disclosure safe harbor and underlying consumer-protection relationship also matter.²³⁵ Insurance supervision similarly depends on the actual state statute, rule or adopted bulletin and the commissioner's powers. The NAIC model does not itself confer a uniform national license-sanction or private-action regime.²³⁶

Civil-rights instruments use another combination. California's employment rules operate within FEHA's public and private enforcement structure, with the applicable charge and right-to-sue requirements. New Jersey's disparate-impact regulations expressly contemplate the Division on Civil Rights and Superior Court routes.²³⁷ Connecticut's 2023 inventory and assessment provisions instead assign duties within government; the 2026 employment and frontier-whistleblower provisions add distinct mechanisms and dates.²³⁸ These examples cannot be sorted reliably into an "agency-only" category by naming the administering institution.

Illinois makes the procedural choice visible. Its Human Rights Act uses a Department charge process with specified routes to the Commission or a civil action, including an opt-out procedure and time limits.²³⁹ A private route can supplement administrative work, but the existence of both routes does not establish that duplication disappears, that complainants can afford litigation or that either institution has adequate resources. Those are questions for the capacity account in Part V.F.

Specific private remedies address still different wrongs. Tennessee provides a civil action within its publicity-rights framework and also criminal liability under the specified provision; its identity protections are not a purely private-enforcement regime.²⁴⁰ California's companion-chatbot act permits an injured person to seek its specified damages and other relief,²⁴¹ while Oregon's 2026 act provides a civil action with its own predicates and minimum recovery.²⁴² These remedies do not automatically travel to every assessment or disclosure duty discussed here. The

²³⁵ Utah Code §§ 13-77-103 to -106, especially § 13-77-105(4)(a) (administrative fine of up to \$2,500 per violation) and § 13-77-106 (preservation of other remedies); *see supra* note 199. Inquiry-based consumer disclosure, defined high-risk regulated-occupation interactions and the at-outset/throughout disclosure safe harbor are distinct provisions.

²³⁶ Colo. Rev. Stat. § 10-3-1104.9; 3 Colo. Code Regs. § 702-10:10-1-1; NAIC, *Model Bulletin: Use of Artificial Intelligence Systems by Insurers* (Dec. 4, 2023); *see supra* notes 200, 220–222. A particular jurisdiction's adoption, delegated authority and procedural law govern its available measures.

²³⁷ Cal. Gov't Code § 12965(a), (c)(1)(D); Cal. Code Regs. tit. 2, § 11009; N.J. Admin. Code § 13:16-2.1(c); *see supra* notes 218–219. Public investigation/civil enforcement and private court access retain their statutory conditions.

²³⁸ 2023 Conn. Pub. Acts 23-16, §§ 1–3; 2026 Conn. Pub. Acts 26-15, §§ 2, 8–12; *see supra* note 214. The frontier-whistleblower provisions and employment-related disclosure/developer provisions do not share one undifferentiated remedy or commencement date.

²³⁹ 775 Ill. Comp. Stat. 5/7A-102, especially (A), (C-1), (D), (G) (charge, investigation opt-out and specified Commission/court routes). A private court route does not eliminate the governing charge, notice, election or filing requirements.

²⁴⁰ Tenn. Code Ann. §§ 47-25-1105(b), 47-25-1106 to -1107 (2025 code reproduction); *see supra* note 205. The civil cause, criminal provision and exceptions defeat an exclusive-private-enforcement description.

²⁴¹ Cal. Bus. & Prof. Code § 22605 (specified injury, actual damages or \$1,000 per violation, whichever greater, and other relief); *see supra* note 225. Later amendments have their own dates and do not supply evidence that the original remedy has produced successful claims.

²⁴² 2026 Or. Laws ch. 85 (S.B. 1546), § 2 (individual suffering ascertainable monetary/property loss or other injury in fact as a result of a violation; greater of actual damages or \$1,000 for each violation, injunction and specified fees); *see supra* note 225. The official chapter records approval and filing March 31, 2026 and effectiveness Jan. 1, 2027. The prospective remedy is not an operating record.

drafter must identify injury, standing, proof, relief and any interaction with existing causes of action.

New York City's record shows why enacted procedure and administered accountability must be examined separately. The state Comptroller's December 2025 audit reported problems with complaint intake and enforcement of Local Law 144. In twelve test calls during August–September 2024, nine were not routed to DCWP: eight went to the Department of Labor and one to the employer. For a separate comparison, auditors reviewed DCWP's selection of 32 company websites identified in Cornell and ACLU publications. The auditors classified them as 22 employers and 10 vendors and compared publicly available information for the 22 employers. Across three categories, they identified 17 potential instances where DCWP identified one. The 17 is a sum of issues, which may overlap; it is not 17 adjudicated violations, 17 distinct employers or a population estimate. The report excluded posting and annual-audit comparisons from that table, and DCWP separately identified posting issues and sent letters. DCWP disputed the audit's legal interpretations and explained the limits of its reviews and resources.²⁴³

The audit supports a bounded concern: a written audit condition can fail to produce reliable intake and consistent examination of the evidence available in particular cases. It does not establish that the law lacked an independence standard, that every identified issue was a violation, or that additional funding alone would have changed the result. The institutional question is concrete. Who receives the complaint? Which records must the employer or auditor produce? Who resolves a disputed calculation or independence question, under what standard, and with what consequence? The proposed instrument must answer those questions rather than rely on the word “audit.”

Texas's online-complaint provision illustrates another intake choice.²⁴⁴ The mechanism's statutory existence supplies no complete account of announced or unannounced enforcement. Nor can the Colorado stay be counted as a constitutional defeat or ordinary nonenforcement: the April 2026 order imposes a defined procedural restraint, and subsequent posture requires its own dated evidence.²⁴⁵ California's ADMT rules present a different temporal problem. Their January 2026 effectiveness, January 2027 ADMT compliance date and later legacy-assessment and reporting schedules are distinct. A future compliance deadline for a particular duty does not render the whole privacy regime unenforceable in the meantime.²⁴⁶

²⁴³ N.Y. Off. of the State Comptroller, *Enforcement of Local Law 144: Automated Employment Decision Tools*, Report 2024-N-6 (Dec. 2025), especially the enforcement findings, website-comparison methodology and appended agency response, <https://www.osc.ny.gov/state-agencies/audits/2025/12/02/enforcement-local-law-144-automated-employment-decision-tools>. The website selection was not a probability sample; the 17 potential instances sum independence, calculation and historical-data issues. Posting and annual-audit comparisons were excluded from that table because sites could have changed. The agency response disputes the legal significance of identified issues and explains its limited review; the Comptroller separately cautions that reviewing published bias audits cannot measure broader compliance; the report is not a finding of 17 adjudicated violations or a national enforcement rate.

²⁴⁴ Tex. Bus. & Com. Code § 552.102 (online consumer-complaint mechanism); H.B. 149, 89th Leg., R.S. (2025), §§ 8–9 (posting deadline and agency-appropriation qualification). The statute is evidence of an intake requirement, not a complete public or nonpublic enforcement history.

²⁴⁵ *XAI LLC v. Weiser*, No. 1:26-cv-01515 (D. Colo.), ECF No. 24 (Apr. 27, 2026); see *supra* note 215. The dated procedural order is not a merits judgment or a certification that no later event occurred.

²⁴⁶ Cal. Code Regs. tit. 11, §§ 7155, 7157, 7200; Cal. Privacy Prot. Agency, approved 2025 regulatory text; see *supra* note 198. Specific ADMT compliance is Jan. 1, 2027; legacy-assessment and 2026–2027 reporting schedules differ, including the April 1, 2028 submission date. Those dates cannot be collapsed into a claim that the CCPA or every rule is awaiting enforcement.

The lesson is an assignment of work, not a prediction that one enforcement family succeeds. For each duty, a state should specify the institution that can compel the relevant information, the person entitled to invoke review, the legal standard and the available relief. Part IV asks which existing powers can support that assignment. Part V asks what records, staff and incentives it requires.

4. New York and California Show Why the Comparison Must Follow the Duty

New York and California are useful comparisons because each combines upstream and use-related mechanisms. The comparison here concerns selected instruments and their source-specific dates through October 16, 2026. It does not treat either state as a single architecture, count original and amending acts as separate current regimes, or certify all law and litigation in either jurisdiction.

Their frontier statutes provide the closest comparison of like commands. California's SB 53 and New York's amended RAISE framework use a training-compute threshold above 10^{26} operations and a separate large-developer revenue threshold above \$500 million with affiliates. Both require specified model transparency and incident reporting from frontier developers while reserving additional framework duties to large developers. Both also use penalty provisions that expressly name large frontier developers. The common structure does not erase differences in filing, report content, incident triggers or administering institutions. California assigns specified incident functions to its emergency-services office and generally requires reporting within fifteen days after discovery; New York's amended framework uses DFS and a seventy-two-hour determination or reasonable-belief clock. Both include a shorter twenty-four-hour notice for specified imminent threats, with their actual conditions governing.²⁴⁷

Timing changes the comparison further. California's 2025 frontier act generally operates from January 2026. New York's December 2025 RAISE text was replaced by the March 2026 amendment, whose substantive framework generally commences January 2027. The original ninety-day commencement falls before the replacing act's signature, so the history cannot be compressed into an assertion that the original never operated. The comparison identifies enacted commitments, not equivalent periods of enforcement or demonstrated risk reduction. The relevant implementation and legal challenges require their own evidence.

Decision-related duties show a different pairing. California's privacy rules govern covered businesses processing consumers' personal information, with specific significant-decision ADMT duties, exceptions and phased compliance. They also impose a limited information duty on a business providing specified ADMT to another business, and assistance requirements relevant to consumer access. New York City's employment law is a municipal comparator: it conditions covered hiring and promotion uses on an independent bias audit and public summary, with separate resident notice. California's employment-discrimination rules address a broader automated-system definition and discriminatory employment practices through FEHA. A privacy opt-out, a bias-audit condition and a discrimination claim are different protections; none can be substituted for the others merely by calling all three deployer accountability.²⁴⁸

²⁴⁷ Cal. Bus. & Prof. Code §§ 22757.11–22757.15; N.Y. Gen. Bus. Law §§ 1420–1427, as enacted by 2026 N.Y. Laws ch. 96; *see supra* notes 204,206,230,231. Compare each actor's framework, transparency, incident and filing obligations, including California's fifteen-day and New York's seventy-two-hour ordinary incident clocks and their specified twenty-four-hour urgent notices. *See also* Cal. Gov't Code § 9600 (ordinary enactment-effectiveness rule). The large-developer civil-penalty language does not imply that every substantive obligation is large-developer-only.

²⁴⁸ Cal. Code Regs. tit. 11, §§ 7153, 7200, 7220–7222; Cal. Civ. Code §§ 1798.140, 1798.145, 1798.150,

The territorial predicates and remedies likewise remain separate. The California privacy provisions inherit the CCPA's business, consumer and exemption structure. The NYC audit command concerns covered employment decision tools and employment uses; its city-resident notice rule does not by itself define the entire law's reach. New York's amended frontier framework expressly reaches development, deployment or operation wholly or partly in the state. California's chapter does not contain the same territorial clause. That textual difference calls for analysis of the particular command and contacts; it supplies no inference of unrestricted California jurisdiction. Public reporting, confidential regulator access, an administrative penalty and an affected person's civil action also perform different work. Neither the use of an existing regulator nor a statutory savings clause resolves jurisdiction, federal conflict or constitutional burden for every duty.

Public-sector laws show why the version history matters as much as the label. New York's original LOADinG framework contained broad assessment and human-review provisions. Its 2025 amendment replaced that structure with employment-tool lists alongside a separate state-agency public-impact AI inventory; the 2026 amendment repealed another general government-ADM article and instead added specified local and education employment-tool lists, with feasibility and sunset qualifications. California's corresponding examples require an inventory of high-risk automated decision systems for specified agencies and legislative reporting, while a separate provision requires AI disclosure and human-contact information in direct government-service communications. These are inventory, communication and employment protections with different coverage and recipients, not interchangeable impact-assessment mandates. None of those descriptions establishes that agencies completed the required work or that an affected person obtained a remedy.²⁴⁹

California's provenance and training-data laws add upstream information requirements that a decision-accountability proposal should confront directly. The provenance regime's 2026 provider amendments and its later platform and device provisions have different bearers and dates. Dataset documentation concerns development information; it does not automatically provide validation of a particular employer's or insurer's use. New York's frontier transparency and incident reports likewise concern specified model risks rather than an individual's explanation or appeal. These measures show why upstream and downstream duties may be complements. The design task is to identify the information a covered decision actually needs, which actor can lawfully provide it, and what the reviewer may do when adequate information remains unavailable.

This comparison supports borrowing with specified changes. An assessment is already an existing component; so are notices, audit summaries, professional oversight and public reporting. The proposed added step is their arrangement into a duty and evidence chain linked to a consequential decision, an institution and an available remedy. Part V and the worked text must

1798.155; Cal. Code Regs. tit. 2, §§ 11008.1, 11009, 11013(c); N.Y.C. Admin. Code §§ 20-870 to -874; *see supra* notes 196–198, 218, 233–234, 246. The municipal audit comparator is not New York State's general employment or privacy law. CCPA exemptions and ADMT exceptions govern the privacy comparison; neither a notice nor access right is an unconditional right to a different selection process or a passing audit score.

²⁴⁹ 2024 N.Y. Laws ch. 674, as amended by 2025 N.Y. Laws ch. 96; 2025 N.Y. Laws ch. 711, as amended by 2026 N.Y. Laws ch. 86; N.Y. State Tech. Law §§ 103-e, 401–402-a. The original and amended acts have mixed commencement and provision-specific July 1, 2028 sunsets; the replaced general assessment provisions are not current duties merely because the original act appears in the inventory. Cal. Gov't Code §§ 11546.45.5, 11549.63–11549.66; 2023 Cal. Stat. ch. 800(A.B. 302); 2024 Cal. Stat. chs. 843(A.B. 2885), 928(S.B. 896). The California inventory includes specified proposed, developed, procured or used systems and risk measures “if any”; that language does not mandate every listed mitigation. Its annual legislative-reporting requirement has a January 1, 2029 sunset.

show that arrangement: coverage, lawful access, withholding, timing, review and consequence. They must also identify where a proposed duty asks more than the comparator supplies. Neither state's enactments establish that the arrangement will cost less or improve outcomes without further evidence.

Colorado reinforces the methodological point from a different history. Its 2024 and 2026 acts contain both developer and deployer responsibility, and their changes do not by themselves identify legislative motive or comparative constitutional safety.²⁵⁰ The same caution applies to claims that states are converging on an institution, that a particular subject is neglected, or that sectoral laws spread more quickly. The historical inventory lacks an adjudicated, mutually exclusive coding scheme and a measure of regulatory adequacy capable of supporting those conclusions. The named laws can support provision-specific comparisons without carrying that additional inference.

Washington's task force supplies a more revealing legislative-process record than a success story for one architecture. Its July 2026 report tracks eleven adopted recommendations: four enacted in whole or part, four included in bills that did not pass, and three with no legislative action reported. Those are the report's categories, not eleven independent tests of a constitutional thesis. The task force's statutory remit expressly included how federal, state and local AI laws align, differ, conflict and interact; the record does not establish that questions about legal conflict were never put to it.²⁵¹

The four enacted-in-whole-or-part entries concern companion chatbots, healthcare prior authorization, law-enforcement AI disclosure and child sexual-abuse material. Their subjects and legal forms differ. The law-enforcement entry records a partial ALPR enactment rather than the full recommended disclosure program. The chatbot law was signed in March 2026, before the final task-force recommendation vote in April; the chronology does not show that the vote caused enactment. A recommendation and a later table entry cannot establish that all of its proposed provisions became law.²⁵²

The unsuccessful-bill entries concern high-risk-AI governance, development transparency, an innovation grant program and workplace guidelines. The workplace recommendation expressly names employers and employment decisions and proposes disclosure, human decision-makers and appeals. It did not become law through the reported vehicle. That is a material counterexample to an account in which named actors in established domains succeed while technology-wide measures fail. The grant proposal passed the House by 96–1 before failing to advance in the Senate committee; the record does not support an explanation that it lacked sponsors, constituencies or bipartisan support.²⁵³

²⁵⁰ See *supra* notes 208–210, 215, 228, 245. The successor preserves both developer functions and specified other-law liability. Enactment, replacement, procedural restraint and constitutional merits remain distinct propositions.

²⁵¹ Wash. Att'y Gen., *Artificial Intelligence Task Force: Final Report* (July 1, 2026), at 4–8; 2024 Wash. Sess. Laws ch. 163 (E2SSB 5838), § 2(4)(b). The report records eleven adopted recommendations and its 4/4/3 legislative-action grouping. The separate unapproved packages are outside that eleven-row denominator. The enabling act expressly calls for analysis of interaction and conflict among federal, state and local laws.

²⁵² *Id.* at 9–15, 41 (four enacted-in-whole-or-part entries: ESHB 2225, E2SSB 5395, ESSB 6002 and 2ESSB 5105). The chatbot recommendation's Apr. 24, 2026 final vote follows the Mar. 24 statutory signature; law-enforcement disclosure was partially implemented through an ALPR measure. These are report-based descriptions, not a whole-bill audit of the recommended and enacted text.

²⁵³ *Id.* at 16–24 (HB 2157/SB 6284; HB 1168; HB 1833; workplace-guidelines striker to 2SHB 1833). The workplace recommendation's named employment actors, decisions and review mechanisms are contrary evidence to a domain-versus-technology success rule. The grant bill's 96–1 House vote and

The three no-action entries concern an emerging-technology advisory body, education investment and ethical AI principles. The principles recommendation drew on NIST and an added public-purpose principle; its 13–1 vote, with an abstention and absences, did not itself create an operative duty. The table establishes the reported action, not why none followed or whether a framework could be translated into enforceable obligations.²⁵⁴

The report separately preserves three packages the task force did not approve. One would have required adoption of specified state generative-AI guidance and included six months' notice to labor organizations before certain consequential workplace changes; it tied 8–8. Data-center guardrails failed 6–10, and expanded public-sector bargaining rights failed 3–12, with abstentions and absences in each tally. Those packages concern procurement, public employment and other established state functions. They are relevant contrary examples even though they differ from the worked statute's terms. The votes do not identify which feature of a package produced its result, and task-force rejection is not legislative rejection.²⁵⁵

The Washington record leaves room for the institutional hypothesis developed here: a proposal may be easier to administer or enact when its duty, bearer and institution are clear. It does not validate that hypothesis, establish a ranking of architectures or show that constitutional pressure caused the observed choices. A serious design recommendation must retain the failed domain-specific proposals and the enacted technology-facing measures alongside the examples that fit it.

Federal policy supplies another comparator with a different legal form. The March 2026 White House framework and EO 14365's legislative-recommendation provisions are executive proposals, including qualified treatment of specified state activities. They are not congressional limits on preemption. H.R. 8516's introduced text combines capacity-building with targeted civil-remedy, fraud and whistleblower provisions; a section's disclaimer of new agency regulatory power cannot be expanded into a description of the whole bill as nonregulatory.²⁵⁶ These selected dated documents do not prove the absence of other federal law or later administrative action.

The inventory thus gives the proposal both material to use and limits to respect. Section C identifies the seven decision domains; Section D specifies six process functions. Part IV then asks which institution can impose each duty, against which actor, under what federal and constitutional constraints. The practical test is whether the arrangement supplies usable information and lawful review for the decision in question. Continued existence of an institution is the beginning of that inquiry, not proof that its AI duty or an equivalent remedy will remain available.

Senate committee outcome do not establish absence of constituency or a reason for failure.

²⁵⁴ *Id.* at 25–32, 48 (advisory body, education investment and ethical principles). The ethical-principles vote was 13 in favor, 1 opposed, 1 abstaining and 4 absent; no-action status is the table's classification. The report does not establish a causal explanation for that outcome.

²⁵⁵ *Id.* at 33–37 (state-use guidance/labor-notice package 8–8–2–1; data centers 6–10–1–2; public-sector bargaining 3–12–3–1, each listed as yes/no/abstain/absent). Different features within each package and the distinction between a task-force vote and legislative action preclude treating these as controlled tests of procurement or domain-based regulation.

²⁵⁶ White House, *National Policy Framework for Artificial Intelligence* (Mar. 20, 2026), pillar VII; Exec. Order No. 14365, §§ 8–9; H.R. 8516, 119th Cong. (introduced Apr. 27, 2026), especially § 101 (proposed institute and section-specific limitation), § 501 (digital-forgery remedy and subtitle-specific nonpreemption), §§ 511–512 (fraud/impersonation) and § 522 (whistleblower provisions); see *supra* note 142. The introduced bill is used as a dated proposal, not enacted law or a verified account of its later status. Section-specific limitations do not establish that the whole bill contains no regulatory provisions or every possible preemption effect.

C. High-Impact Domains Define Different Accountability Tasks

1. Employment Duties Must Reach the Practice That Causes the Harm

An employer using AI to screen, rank or evaluate people remains subject to the employment laws that cover its conduct. Title VII's federal enforcement and deferral arrangements operate alongside state laws, whose actors, thresholds and remedies must be identified separately.²⁵⁷ The disparate-impact framework addresses a challenged employment practice, its causal contribution to a protected group's disadvantage, the applicable justification and available alternatives. A statistical disparity alone does not complete that inquiry.²⁵⁸

The decision may occur at screening, interview selection, hiring, evaluation, promotion, discipline or termination. The data used can carry prior patterns of disadvantage into each setting; their legal significance depends on how they affect the actual practice.²⁵⁹ Scholarship on algorithmic auditing and classification supplies reasons to demand evidence and examine institutional choices. Its proposed audit arrangements and critiques should be distinguished from duties already imposed by employment law.²⁶⁰ The proposal attaches assessment and records to the employer's use while specifying what information must be sought from vendors. It does not assume the employer can inspect or alter everything a supplier controls.

Existing instruments illustrate different interventions. New York City's Local Law 144 restricts covered employers' and employment agencies' use of covered automated employment decision tools through audit, publication and notice conditions.²⁶¹ Illinois HB 3773 amends employment-discrimination law to address AI use under its statutory terms.²⁶² The state's Video Interview Act addresses a narrower setting, with its own consent, disclosure, deletion and specified reporting provisions.²⁶³ California's employment regulations address automated systems within discrimination law and require preservation of covered records; their reach is not confined to an employer's direct operation of a model.²⁶⁴ New Jersey's disparate-impact rules likewise operate

²⁵⁷ 42 U.S.C. §§ 2000e-5(c), 2000e-7. The state-law saving provision excludes state laws requiring or permitting acts that would be unlawful under Title VII. *See supra* note 159.

²⁵⁸ 42 U.S.C. § 2000e-2(k); *Griggs v. Duke Power Co.*, 401 U.S. 424, 431–32 (1971); *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 425–26 (1975); *Connecticut v. Teal*, 457 U.S. 440, 452–55 (1982). The nondiscriminatory bottom line does not excuse an unlawful intermediate barrier. *See supra* notes 160–162.

²⁵⁹ Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 Calif. L. Rev. 671 (2016), author abstract at <https://doi.org/10.2139/ssrn.2477899> (explaining how imperfect data and existing patterns of exclusion can produce unintentional discrimination).

²⁶⁰ Ifeoma Ajunwa, *An Auditing Imperative for Automated Hiring Systems*, 34 Harv. J.L. & Tech. 621, 667–68 (2021) (proposing certification and periodic audits with institutional safeguards); Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 Wm. & Mary L. Rev. 857, 916–17 (2017) (proposing a framework that departs from aspects of traditional disparate-impact analysis). These are scholarly proposals and critiques.

²⁶¹ N.Y.C. Admin. Code §§ 20-870 to 20-874; 6 R.C.N.Y. §§ 5-300 to 5-304. Coverage, audit independence, publication and notice are separately specified. *See supra* notes 196, 233.

²⁶² 775 Ill. Comp. Stat. 5/2-102(L), as amended by Pub. Act 103-804 (eff. Jan. 1, 2026). The provision addresses specified AI use, protected-class discrimination, ZIP-code proxies and notice. *See supra* notes 197, 239.

²⁶³ 820 Ill. Comp. Stat. 42/5, /10, /15, /20 (AI Video Interview Act). Notice, explanation, consent, limited sharing, requested deletion and specified reporting have separate predicates. *See supra* note 201.

²⁶⁴ Cal. Code Regs. tit. 2, §§ 11008–11009, 11013(c), 11017 (effective Oct. 1, 2025). The provisions apply FEHA rules to automated decision systems and specify relevant recordkeeping. *See supra* notes 218, 237.

through their own employment-law elements and burdens.²⁶⁵

Proposals for additional audit duties require the same attention to scope. New Jersey's A3854 proposed additional audit and disclosure duties with its own coverage choices.²⁶⁶ California's AB 1018 illustrates another proposed allocation between developers and deployers in the dated version discussed in the accompanying note.²⁶⁷ A supplier-facing duty and an employer-facing duty should be compared for the information each can obtain and the burden each creates.

Federal statutes, regulations and technical assistance also perform different work. The EEOC's 2022 disability assistance discusses how software use can exclude a person with a disability and the need to consider accommodation; its legal role is technical assistance under the existing disability statute.²⁶⁸ The Uniform Guidelines' four-fifths comparison is a qualified screen. Smaller differences may matter, and sample size and other evidence can alter the inference.²⁶⁹ Changes in federal contractor rules and executive enforcement policy require their own dates and legal bases; they do not by themselves amend Title VII or establish the posture of every federal employment program.²⁷⁰

The hiring application in V.E.2 connects an assessment, retained records and access to review. *Ricci* addresses an employer's race-conscious response in its particular setting; it does not mandate this Article's proposed state presumption or immunize an employer from federal law.²⁷¹ The Comptroller's selected New York City audit identified problems with notices, audits and oversight, while the agency disputed aspects of the methodology and described resource constraints. A statutory audit duty therefore needs standards for the auditor, usable evidence and an institution able to examine the result.²⁷²

²⁶⁵ N.J. Admin. Code §§ 13:16-2.2, -4.1 (effective Dec. 15, 2025). The employment rule allocates the impact, necessity and less-discriminatory-alternative showings in specified stages. *See supra* notes 219, 237.

²⁶⁶ New Jersey A. 3854, 221st Leg., Reg. Sess., § 1(b)–(h) (introduced Feb. 22, 2024) (proposed audit, notice and enforcement provisions). The cited version illustrates a proposal, not an operative duty. *See supra* note 217.

²⁶⁷ California AB 1018 (2025–2026 Reg. Sess.), Senate amendments of Aug. 21, 2026, reproduced bill text (developer instructions and deployer adverse-outcome disclosures and appeal). The proposed allocation differs from the enacted FEHA regulations; the dated version does not establish later enactment or operative duties.

²⁶⁸ EEOC, *The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees* 1–4 (May 12, 2022). The technical assistance applies existing disability-law concepts to software-assisted assessment; it is distinct from the statute and from a judicial determination.

²⁶⁹ 29 C.F.R. § 1607.4(C)–(D); *Teal*, 457 U.S. at 453 n.12; *Ricci v. DeStefano*, 557 U.S. 557, 586–87 (2009). The four-fifths measure is a qualified enforcement screen: small samples, statistical significance and other evidence matter. *See also* 42 U.S.C. § 2000e-2(k), (l).

²⁷⁰ 42 U.S.C. § 2000e-5 (Title VII enforcement powers and procedures). The statutory authority and the intensity or priorities of its administration are different questions. *See supra* note 268.

²⁷¹ *Ricci v. DeStefano*, 557 U.S. 557, 585 (2009); 42 U.S.C. § 2000e-2(l). *Ricci* requires a strong basis in evidence for the particular race-conscious action taken to avoid disparate-impact liability; it does not make every response to an audit lawful.

²⁷² N.Y. State Comptroller, *Enforcement of Local Law 144 – Automated Employment Decision Tools*, Report 2024-N-6, at 16–17, 21–24, 26–30, 34 (Dec. 2025). The Comptroller examined thirty-two companies, classified as twenty-two employers and ten vendors, and identified seventeen potential instances of noncompliance in overlapping categories. DCWP disputed aspects of the method, company classification, inferences and resource demands. These were not seventeen adjudicated violations or a citywide prevalence estimate. *See supra* note 243; *infra* Part V.E.2.

2. Housing Accountability Follows Different Decisions and Legal Routes

Housing decisions include tenant screening, advertising, rent setting and mortgage credit. Each identifies a different actor, use of information and potential harm. The Fair Housing Act prohibits specified discrimination in covered housing activity; state and local laws may add protections and routes of enforcement. Coverage, exemptions, the theory of liability and the available public or private remedy still have to be established.²⁷³ A landlord, screening company, platform and lender cannot be treated as interchangeable deployers for every duty.

The distinction is especially important for pricing. The RealPage proceedings concern alleged anticompetitive conduct, whose claims and procedural posture must be read on their own terms. An allegation in a complaint is not a finding that all algorithmic pricing is unlawful.²⁷⁴ Screening and advertising raise different questions about information, classification and access to housing. The data-mining literature helps identify possible pathways to disadvantage; whether a practice violates a housing statute requires that statute's own elements and causal inquiry.²⁷⁵

Cross-cutting AI and privacy instruments add other duties. Colorado includes housing among its covered consequential decisions, subject to the framework's definitions, exceptions and operation dates.²⁷⁶ California's ADMT rules address specified processing by covered businesses, with coverage and exceptions that matter to housing uses.²⁷⁷ New Jersey's civil-rights rules apply their own protected-class and disparate-impact framework.²⁷⁸ Virginia's vetoed HB 2094 provides a separate historical proposal concerning housing decisions.²⁷⁹

New York's General Business Law § 340-b and New Jersey's FAIR Act supply concrete rent-setting examples through antitrust legislation.²⁸⁰ The two measures reach specified forms of software-assisted coordination and pricing. A fair-housing complaint, an antitrust action and a duty to provide a decision record address different conduct and require different evidence.

HUD's 2024 screening and advertising materials illustrate an agency's application of fair-

²⁷³ 42 U.S.C. §§ 3604–3605 (covered housing and residential real-estate-related transactions). State analogues require their own coverage and liability analysis. *See infra* Part IV.B.2.

²⁷⁴ Amended Complaint, *United States v. RealPage, Inc.*, No. 1:24-cv-00710, ECF No. 47, ¶¶ 5–6 (M.D.N.C. Jan. 7, 2025) (United States and ten named states alleging coordination involving nonpublic rental data, software and landlord defendants). These are allegations in that dated pleading, not findings of liability or a statement of later docket status.

²⁷⁵ Barocas & Selbst, *supra* note 259 (author abstract). Applying the data-discrimination analysis to housing is the Article's inference; FHA and employment liability rules remain distinct.

²⁷⁶ Colorado SB 26-189 (2026 Reg. Sess.), reenacted Colo. Rev. Stat. §§ 6-1-1701(4)–(8), -1702, -1704. The covered residential decisions and developer and deployer duties are defined separately. *See supra* note 88.

²⁷⁷ Cal. Code Regs. tit. 11, §§ 7001, 7150(b)(3), 7153, 7155(b), 7200–7222. The covered business, consumer, significant decision, provider assistance and phased obligations control the application. *See supra* note 198.

²⁷⁸ N.J. Admin. Code § 13:16-2.3. The housing respondent bears the necessity and no-less-discriminatory-alternative showings after the complainant establishes the specified impact. This differs from the employment sequence in § 13:16-2.2. *See supra* note 265.

²⁷⁹ Virginia HB 2094 (2025 Reg. Sess.), enrolled proposed §§ 59.1-608 to -609 (including housing consequential decisions). The cited text is a legislative proposal. *See supra* note 211.

²⁸⁰ N.Y. Gen. Bus. Law § 340-b(1)–(4), added by 2025 N.Y. Laws ch. 437 (knowing or reckless software facilitation of specified anticompetitive agreements and landlord reliance on a defined coordinating function); 2026 N.J. Laws ch. 43, §§ 3–4, 9 (FAIR Act, approved July 20, 2026) (specified coordinator arrangements and functions, with exclusions; effective on the first day of the twelfth month following enactment). The commands reach different actors through competition law. *See supra* note 274.

housing concerns to automated tools.²⁸¹ Their later archival status is separate from the underlying statute, causes of action and applicable legal tests.²⁸² A private claimant relying on discriminatory effects must satisfy the governing housing doctrine; an agency's discussion of a risk does not itself establish liability. Credit decisions may also engage ECOA, with its own covered credit transaction, applicant and creditor requirements. A score connects to these legal routes through the activity and actors the statute covers.

The proposed duties would require a covered decision maker to identify the information used, preserve a usable record and provide a route to challenge a consequential error. That recommendation must be specified for each actor. A screening company may hold a disputed record while the landlord makes the rental decision, and a lender may be subject to a separate federal regime. Meaningful recourse requires an allocation of requests, review and relief across those relationships, without promising that a deployer can correct a vendor's entire dataset or compel a different credit decision.

3. Insurance Supervision Supplies Authority with Defined Statutory Limits

Insurance supplies established institutions for supervising consequential decisions, but no sector-wide exemption from federal law. McCarran-Ferguson's reverse-preemption rule depends on the state law's purpose, the federal statute and the effect on the state insurance regime. Its regulatory and fee-or-tax routes, and the different treatment of particular antitrust questions, must be kept distinct.²⁸³ An insurer's license or the description of a duty as procedural does not settle those predicates.

State insurance supervision supplies powers over specified licensing, solvency, forms, rates and market conduct. Their availability depends on the jurisdiction, product and activity.²⁸⁴ A duty to document underwriting, for example, poses a different question from a claim about an employee health plan's administration. Independent vendors may also fall outside powers directed at insurers, leaving the insurer's contractual access to evidence important.

Underwriting, pricing, claims and fraud investigation should be examined separately. Non-traditional data can affect a person's eligibility or price without making every input an external consumer data source under a particular regulation.²⁸⁵ Health-claim litigation supplies allega-

²⁸¹ HUD, *Guidance on Application of the Fair Housing Act to the Screening of Applicants for Rental Housing* 1–4 (Apr. 29, 2024); HUD, *Guidance on Application of the Fair Housing Act to the Advertising of Housing, Credit, and Other Real Estate-Related Transactions through Digital Platforms* 1–4 (Apr. 29, 2024). These dated guidance documents are available in HUD's archives; they are not independent statutes or findings of liability.

²⁸² 42 U.S.C. §§ 3604–3605, 3613. Section 3613 supplies a conditional aggrieved-person civil action, subject to its timing, prior-proceeding and remedial provisions. *See supra* note 273.

²⁸³ 15 U.S.C. § 1012(b); *U.S. Dep't of Treasury v. Fabe*, 508 U.S. 491, 501–05, 508–09 (1993); *Humana Inc. v. Forsyth*, 525 U.S. 299, 307–11 (1999). The reverse-preemption inquiry depends on the state law's insurance-regulating purpose, the federal statute and impairment; *Fabe* does not reduce the first clause to the narrower antitrust test. ERISA separately distinguishes insured and self-funded plans: *FMC Corp. v. Holliday*, 498 U.S. 52, 61 (1990); *Pilot Life Ins. Co. v. Dedeaux*, 481 U.S. 41, 52–54 (1987). A state licensing label does not settle the command or remedy. *See infra* Parts IV.C.2, V.E.3.

²⁸⁴ NAIC, *About the NAIC*, <https://content.naic.org/about>; NAIC, *Model Bulletin: Use of Artificial Intelligence Systems by Insurers* § 1 (Dec. 4, 2023). The association coordinates regulators and model instruments; the adopted state instrument and underlying authority determine legal force.

²⁸⁵ Anya E.R. Prince & Daniel Schwarcz, *Proxy Discrimination in the Age of Artificial Intelligence and Big Data*, 105 Iowa L. Rev. 1257, 1257–58 (2020). The analysis identifies a discriminatory mechanism,

tions and rulings concerning particular systems and plans, including the limits of the claims allowed to proceed.²⁸⁶ A proposed record or review duty needs to identify the relevant line, data, decision and responsible organization before selecting an enforcement channel.

Colorado SB 21-169 and its implementing rules illustrate a defined governance regime for covered insurance practices and data uses.²⁸⁷ Regulation 10-1-1 has distinct line-specific scope and implementation dates; its treatment of life, private passenger auto and health-benefit insurance must be read with its exclusions.²⁸⁸ The NAIC Model Bulletin sets out expectations for a risk-appropriate written AI Systems Program and regulator review under existing authority. A model bulletin is not itself a law binding every insurer.²⁸⁹ The NAIC's October 8, 2026 roster groups twenty-nine jurisdictions under its model-bulletin category, separately identifying other arrangements in four states. The listed instruments still require jurisdiction-specific examination; the map supplies no uniform rule or outcome measure.²⁹⁰

New York's Circular Letter No. 7 addresses AI systems and external consumer data in covered underwriting and pricing, with specified exclusions and responsibilities for governance, testing, controls and third-party systems.²⁹¹ The NAIC evaluation pilot and later draft supplement concern tools for regulatory review. Planned work and an exposure draft do not establish completed adoption or successful examinations.²⁹²

The worked application in V.E.3 uses these antecedents to specify proposed evidence, access and recourse. It distinguishes insurer underwriting from plan administration, and a state in-

not a finding about the legality of a particular insurer's model.

²⁸⁶ *Estate of Lokken v. UnitedHealth Group, Inc.*, No. 0:23-cv-03514, ECF No. 91, at 20–24 (D. Minn. Feb. 13, 2025) (dismissing most claims as Medicare Act-preempted while permitting contract and implied-covenant claims after the exhaustion analysis); *Kisting-Leung v. Cigna Corp.*, No. 2:23-cv-01477, ECF No. 55, at 25–27 (E.D. Cal. filed Mar. 31, 2025) (dismissing the § 1132(a)(1)(B) benefits claim with leave to amend; retaining specified § 1132(a)(3) and UCL claims for some plaintiffs). These are pleading-stage dispositions of allegations about automated review, not adjudications that the algorithms were unlawful.

²⁸⁷ Colo. Rev. Stat. § 10-3-1104.9(1)–(4), (6)–(7), enacted by Colorado SB 21-169 (2021 Reg. Sess.). Coverage, exclusions and rulemaking instructions qualify the external-data and algorithmic unfair-discrimination provisions.

²⁸⁸ 3 Colo. Code Regs. § 702-10, Regulation 10-1-1, §§ 2–6, 10–11 (effective Nov. 14, 2023; amended effective Oct. 15, 2025). The amended governance rule covers specified life, private-passenger automobile and health-benefit-plan uses, with compliance periods and exemptions defined in the rule.

²⁸⁹ NAIC, *Model Bulletin: Use of Artificial Intelligence Systems by Insurers* §§ 1, 3–4 (Dec. 4, 2023). Written governance, testing, documentation and supervisory information requests operate within the model's risk-based structure and state-law predicates.

²⁹⁰ NAIC, *Implementation of NAIC Model Bulletin: Use of Artificial Intelligence Systems by Insurers* 1–2 (status Oct. 8, 2026). Its model-bulletin roster lists twenty-seven states, the District of Columbia and Puerto Rico; four other states are identified separately. The earlier map dated April 1, 2026 lists twenty-four states and the District of Columbia in that category, with California, Colorado, New York and Texas identified separately. Neither map determines that a jurisdiction adopted all model elements or meets accreditation standards; the rosters do not establish uniform legal force or efficacy.

²⁹¹ N.Y. Dep't of Fin. Servs., *Insurance Circular Letter No. 7: Use of Artificial Intelligence Systems and External Consumer Data and Information Sources in Insurance Underwriting and Pricing* ¶¶ 1–8, 12–17, 35–45 & n.3 (July 11, 2024). Its scope and expectations concern specified underwriting and pricing uses; the letter supplies its own treatment of traditional underwriting or medical information.

²⁹² NAIC, *AI Systems Evaluation Tool Pilot Project* 1–3 (2026); Big Data and Artificial Intelligence (H) Working Group, *AI Systems Evaluation Tool*, version 5 (exposure). The pilot develops supervisory questions and permits jurisdictional adaptation; it is distinct from an adopted uniform legal requirement. See *supra* note 222.

surance remedy from an available federal civil-rights claim. McCarran-Ferguson, ERISA and the particular substantive statute each require their own inquiry. The institutional advantage is concrete: an existing supervisor may have staff, examination powers and relationships capable of supporting a defined duty. Whether those resources are sufficient, and whether an affected person can obtain relief, remains an implementation question.

4. Healthcare Duties Depend on the Professional, Facility, Payer and Device

A clinical recommendation and an insurer's coverage decision can concern the same patient while engaging different legal responsibilities. Professional licensing, facility oversight and payer supervision supply distinct state-law starting points. *Dent v. West Virginia* illustrates the licensing tradition; it does not decide the scope or validity of a contemporary AI duty.²⁹³ Corporate-practice and professional-delegation rules vary by state. Whether a particular use improperly substitutes for clinical judgment requires the relevant rule and facts, without equating every software recommendation with the practice of medicine.²⁹⁴

For a hospital considering a diagnostic tool, the first question is the function the system will perform. The FDA's device materials identify regulated products and pathways, with the function and authorization defining the relevant inquiry.²⁹⁵ A payer's use of an output to assess medical necessity raises separate coverage and plan-administration questions, illustrated by the allegations and rulings in the health-claim litigation.²⁹⁶ A documentation tool, a diagnostic function and a utilization-review system need not share either a regulator or the same access and review duties.

The selected state statutes show why their commands must be compared directly. California SB 1120 governs specified health plans' and insurers' utilization review and reserves medical-necessity determinations to a licensed physician or other competent licensed healthcare professional under its conditions.²⁹⁷ Arizona HB 2175 addresses individual medical-director review of specified medical-necessity denials without naming AI.²⁹⁸ Maryland's measure addresses

²⁹³ *Dent v. West Virginia*, 129 U.S. 114, 121–25 (1889) (upholding the medical-licensing requirements at issue). State authority over professional qualifications does not dispose of later federal preemption questions. See *supra* notes 173–174.

²⁹⁴ Cal. Bus. & Prof. Code § 2400 (professional rights and its specified charitable-employment exception); N.Y. Educ. Law § 6522 (licensed or otherwise authorized persons). These illustrate distinct state constraints on the practice of medicine, not an exception-free national rule barring all corporate employment of physicians.

²⁹⁵ FDA, *Artificial Intelligence-Enabled Medical Devices*, <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-device> s. FDA describes a noncomprehensive list of devices authorized for marketing; a list entry is not a general approval of every later clinical use.

²⁹⁶ See *supra* note 286 (the dated *Lokken* and *Kisting-Leung* dispositions concerning alleged automated claims review).

²⁹⁷ 2024 Cal. Stat. ch. 879 (SB 1120); Cal. Health & Safety Code § 1367.01(k); Cal. Ins. Code § 10123.135(j). The covered plan or insurer must meet the specified tool-use conditions; the tool may not deny, delay or modify services based wholly or partly on medical necessity. The medical-necessity determination is reserved to a competent licensed physician or healthcare professional. The statute does not use a general finalization-or-override formula.

²⁹⁸ 2025 Ariz. Sess. Laws ch. 165 (HB 2175), adding Ariz. Rev. Stat. §§ 20-3103, 20-3407 (effective from and after June 30, 2026). The medical director must individually review the denial, exercise independent medical judgment and not rely solely on recommendations from another source. The enacted command does not name AI; the House-engrossed technology-specific prohibition was not

covered utilization review by carriers and specified contracted entities, including pharmacy benefits managers and private review agents.²⁹⁹ Indiana HB 1271 combines provider-facing requirements with restrictions on sole reliance on AI for downcoding; its required employee or contractor review is not a general physician-review mandate.³⁰⁰ These related interventions attach to different decisions and actors. Their review requirements must therefore remain distinct when a drafter adapts one to another setting.

Federal law adds function-specific constraints. The Cures Act's software exclusions and FDA clinical-decision-support guidance help determine whether a function is a device; marketing pathways and requirements then require their own analysis.³⁰¹ Predetermined Change Control Plan guidance addresses defined modifications under the relevant authorization; changes outside that plan require their own regulatory assessment.³⁰² Information-blocking rules concern specified actors, electronic health information and regulated practices, with exceptions; they do not guarantee access to a supplier's training corpus or proprietary model.³⁰³ CMS's prior-authorization rule imposes defined requirements on impacted payers on its implementation schedule.³⁰⁴ EMTALA supplies emergency-screening and stabilization obligations for covered hospitals under its own statutory conditions.³⁰⁵

The state duty must then be compared with the actual federal requirement. A requirement concerning a PMA-approved device may encounter *Riegel* even if it is generally applicable or addressed to a licensed user. A parallel-duty theory and a theory that a requirement falls outside the device provision are different arguments. *Buckman* also requires attention to a claim's dependence on the federal scheme; an independent state duty cannot be replaced by private enforcement of the FDCA. *Lohr* concerns a different pathway and record, and its reasoning must

enacted. The textual sequence does not establish its cause.

²⁹⁹ 2025 Md. Laws ch. 747 (HB 820); Md. Code Ann., Ins. §§ 15-10A-06, 15-10B-05.1(b)–(d), -07. The provisions address the covered carrier, contracted utilization review, and specified pharmacy benefit managers and private review agents. Their decision-maker and tool-use restrictions differ from a generic human-signoff duty.

³⁰⁰ Enrolled Indiana HB 1271, Pub. L. No. 88-2026, at 5 (adding Ind. Code § 27-1-52-9, eff. July 1, 2026). The downcoding restriction requires medical-record review by an insurer employee or contractor, not necessarily a physician; provider claims and AI-use disclosures have separate requirements. Section 27-1-52-0.3 excludes the Medicaid program and managed care organizations providing services to Medicaid recipients.

³⁰¹ 21 U.S.C. § 360j(o)(1)(E); FDA, *Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff* (Jan. 29, 2026). The clinical-decision-support exclusion requires all four statutory criteria, including independent review of the recommendation's basis; the guidance is nonbinding.

³⁰² FDA, *Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions* (Aug. 18, 2025), <https://www.fda.gov/media/166704/download>. The authorized plan and modification protocol bound the changes addressed; the guidance supplies no permission for unbounded postmarket learning.

³⁰³ 42 U.S.C. § 300jj-52; 45 C.F.R. §§ 171.102–171.103 and pt. 171, subpts. B–D. Information blocking depends on the covered actor, practice, electronic health information, applicable knowledge standard and exceptions. Application to an AI-mediated practice is the Article's inference under those elements.

³⁰⁴ CMS, *CMS Interoperability and Prior Authorization Final Rule CMS-0057-F*, 89 Fed. Reg. 8758, 8873 (Feb. 8, 2024). Prior-authorization process requirements and the substantive legality of a decision made with AI are separate inquiries; compliance with a process rule alone does not resolve the latter.

³⁰⁵ 42 U.S.C. § 1395dd(a)–(c), (f)–(h) (screening, stabilization, transfer conditions, limited express pre-emption and related provisions). EMTALA applies through its own hospital and emergency-condition predicates.

be applied on those terms.³⁰⁶ General professional or facility authority cannot resolve these questions in advance.

The hospital application in V.E.1 specifies the proposed allocation of records, requests to external sources, review and conditional re-decision. That allocation must distinguish the provider's clinical decision from the payer's coverage decision. A patient may be able to seek amendment of a held record without controlling every input used by a vendor. An implementation plan must identify who can correct the relevant error, who can reconsider the result, and which existing legal route can obtain relief.

5. Education Requires Decisions About Records, Individualization and Institutional Responsibility

A school using AI to recommend placement, flag a student or assess work makes decisions within a system of public and institutional responsibility. State constitutional and statutory provisions allocate authority over education, with different arrangements for districts, public higher education and private institutions.³⁰⁷ The constitutional education cases illustrate constraints on those arrangements as well as the significance of the state role; a familiar field of state concern supplies no general answer to an AI-related rights claim.³⁰⁸

The federal obligations attach to particular records, institutions and students. FERPA addresses covered education records, access and amendment under its statutory conditions. Inspection of a model or supplier data requires a separate basis.³⁰⁹ IDEA's individualized educational planning and procedural protections require attention to the child and the educational decision, even where an institution uses a standardized tool.³¹⁰ Section 504 and ADA Title II impose their own disability-discrimination requirements on covered entities.³¹¹ Title VI and Title IX have distinct statutory scope, enforcement routes and regulatory requirements. A private damages claim, an administrative complaint and an effects-based regulation cannot be

³⁰⁶ *Riegel v. Medtronic, Inc.*, 552 U.S. 312, 321–30 (2008); *Buckman Co. v. Plaintiffs' Legal Comm.*, 531 U.S. 341, 348–53 (2001); *Medtronic, Inc. v. Lohr*, 518 U.S. 470 (1996). *Riegel* treats the actual device requirement and federal approval as material, including generally applicable duties in operational effect; its parallel-claim discussion did not decide the unpreserved claim there. *Buckman* addresses fraud-on-the-FDA claims. In *Lohr*, Parts I, II, III, V and VII are the Court's opinion; Parts IV and VI are plurality portions. A deployer duty must be tested against the relevant federal pathway and the requirement or remedy it would add. *See supra* notes 108–109; *infra* Part V.E.1.

³⁰⁷ Cal. Const. art. IX, § 5; N.Y. Const. art. XI, § 1. These clauses illustrate affirmative public-school responsibilities; their wording does not establish identical governance or remedies.

³⁰⁸ *Plyler v. Doe*, 457 U.S. 202, 221–24 (1982) (rejecting the exclusion at issue while distinguishing education's importance from a fundamental federal right); *San Antonio Independent School District v. Rodriguez*, 411 U.S. 1, 35, 49–51 (1973) (federal equal-protection review of the financing system and local-control interests). Neither resolves the validity of a particular educational AI duty. *See supra* notes 173–174.

³⁰⁹ 20 U.S.C. § 1232g(a)–(b); *Gonzaga University v. Doe*, 536 U.S. 273, 287, 290 (2002). The funding, record-access and disclosure provisions have specific institutional predicates and exceptions; the nondisclosure provisions do not supply § 1983-enforceable rights under *Gonzaga*.

³¹⁰ 20 U.S.C. § 1414(b), (d) (evaluation procedures and IEP teams); 34 C.F.R. pt. 300. The evaluation uses multiple assessment tools, and the IEP process has defined participants and duties. AI use must be assessed within those requirements.

³¹¹ 29 U.S.C. § 794; 42 U.S.C. §§ 12131–12132. Funding, covered public-entity status, qualification and the relevant disability-discrimination claim determine the applicable duty.

treated as interchangeable protections.³¹²

Admissions screening, automated scoring, tutoring and intervention tools also perform different work. An assessment of an automated score needs a defined task and evidence suited to that task.³¹³ For the proposed design, the decision must remain identifiable: which placement, score or intervention used the output, what evidence the institution retained, and how a student or parent can question the relevant information. A general institutional AI policy need not answer those questions.

Selected legislative texts supply different levels of specification. Idaho's SB 1227 provides a framework, local-policy and procurement example in the bill text discussed in the note.³¹⁴ Colorado's consequential-decision framework includes education under its own coverage and exception provisions.³¹⁵ California's ADMT regulations concern specified processing by covered businesses, requiring separate attention to whether the educational institution or supplier falls within the regime.³¹⁶ Education-agency guidance provides a practical model for local policies and use; guidance, a binding rule and a school-board policy have different legal force.³¹⁷ Selected legislative proposals illustrate additional choices about classroom uses and statewide guidance in the dated versions discussed in the note.³¹⁸

A deployer-centered duty should preserve institutional responsibility without treating every educational AI use as the same decision. A district assessing a placement tool may need validation evidence from a vendor and a record of its own use; an individual review may require correction of an input or reconsideration through an existing education procedure. The responsible institution, authorized reviewer and available remedy must be specified. Neither a framework document nor a generic explanation replaces IDEA procedures or creates a private disparate-impact action where the governing law supplies none.

³¹² 42 U.S.C. §§ 2000d to 2000d-7; 20 U.S.C. § 1681; *Alexander v. Sandoval*, 532 U.S. 275, 285–86 (2001) (the private Title VI action does not extend to § 602 disparate-impact regulations); 34 C.F.R. § 106.21(b)(2) (conditional admissions criterion rule under Title IX). Agency administrative enforcement and private relief require separate authority. See *infra* note 437 for the dated Education Department Title VI regulatory change.

³¹³ The task-specific distinction matters to the proposed assessment: evaluating an essay score, placement recommendation or tutoring response requires evidence suited to that output and its educational use. See *supra* Part III.A; *infra* Part V.C. A generic accuracy statement cannot identify all of those tasks.

³¹⁴ Idaho SB 1227 (2026 Reg. Sess.), bill text adding Idaho Code §§ 33-7003 to -7008, bill-text reproduction at Tally, *Idaho SB 1227*. The reproduced text directs a statewide framework subject to board approval, aligned district and charter policies, and specified vendor disclosures and assurances in procurement. The text is used for those mechanisms; the comparison does not depend on a corpus-absence claim or the chapter's approval and commencement dates.

³¹⁵ Colorado SB 26-189 (2026 Reg. Sess.), reenacted Colo. Rev. Stat. §§ 6-1-1701(3)–(8), -1704. Coverage of education remains subject to the consequential-decision definition, obligated actor and exclusions. See *supra* note 88.

³¹⁶ Cal. Code Regs. tit. 11, §§ 7001, 7200–7222. The defined significant decisions do not dispense with the CCPA's covered-business and processing predicates. See *supra* note 198.

³¹⁷ Cal. Dep't of Educ., *Learning With AI, Learning About AI* (2023), <https://www.cde.ca.gov/ci/pl/ai/incaifornia2023.asp> (including evaluation of terms of use and data-collection guidelines). This is agency guidance, not a census of binding state duties.

³¹⁸ N.Y. A. 9190 (introduced Nov. 3, 2025), proposed Educ. Law § 115 (classroom restrictions below ninth grade with specified exceptions); N.Y. A. 6972 (introduced Mar. 18, 2025), § 2 (working group, guidance and model policy). These particular introduced versions illustrate legislative choices; they do not state later committee, enactment or implementation status.

The resulting duty should connect the educational decision to existing student rights, information that can lawfully be obtained and an institution able to provide meaningful review. That inquiry gives each policy, procurement term and statutory protection a concrete role.

6. Benefits Decisions Require Program-Specific Process and Usable Reasons

A state agency using a system to determine eligibility or calculate assistance administers a program with its own legal standards. Medicaid's federal-state arrangement supplies one set of statutory responsibilities.³¹⁹ SNAP assigns eligibility work to state agencies under its own federal requirements.³²⁰ TANF's block-grant design leaves different choices to states and has different entitlement consequences.³²¹ Unemployment insurance adds another federal-state framework.³²² Program differences matter for the responsible actor, applicable standards, hearings and relief; they cannot be replaced by a general description of benefits AI.

An automated flag can initiate investigation, a rules engine can calculate an amount, and a recommendation can affect caseworker attention. The evidence needed to explain each action differs. Fraud-detection records illustrate why an agency must identify what a flag means before treating it as a determination of wrongdoing.³²³ A useful record connects the governing rule, relevant inputs, system output and action taken. Preserving only the output may leave a claimant unable to identify the error.

Where a protected interest is implicated, due process requires the procedure appropriate to the context. *Goldberg* addresses welfare termination; *Mathews* supplies a balancing inquiry in a different benefits setting. They do not require the same hearing at the same time for every automated determination.³²⁴ Citron's technological-due-process account makes examination and challenge central to computerized administration.³²⁵

The litigation examples likewise require their particular posture and theory. *Cahoo* concerns procedural claims arising from Michigan's MiDAS system; a pleading or immunity ruling does not establish liability of every official or contractor.³²⁶ *Barry* addresses the adequacy of no-

³¹⁹ 42 U.S.C. § 1396a(a)(3), (5) (state-plan hearing and administration provisions). The federal-state program's particular eligibility and procedural requirements govern the decision at issue.

³²⁰ 7 U.S.C. § 2020(e)(10) (state-plan fair-hearing requirements for households aggrieved by specified agency action).

³²¹ 42 U.S.C. §§ 601–604. Section 601(b) expressly disclaims an individual or family entitlement to assistance under that part; the block-grant structure differs from other benefits programs.

³²² 42 U.S.C. § 503(a)(3) (opportunity for a fair hearing before an impartial tribunal for persons whose unemployment-compensation claims are denied).

³²³ *Cahoo v. SAS Analytics Inc.*, 912 F.3d 887, 892–96, 900–04 (6th Cir. 2019) (describing the alleged MiDAS process and reviewing claims at the pleading stage). The example supports inquiry into the relationship between a fraud flag and the agency's determination. *See infra* note 326.

³²⁴ *Goldberg v. Kelly*, 397 U.S. 254, 263–66 (1970); *Mathews v. Eldridge*, 424 U.S. 319, 335 (1976). The protected interest, risk of error, value of added safeguards and governmental interests determine the appropriate procedure; the cases concern different benefit settings.

³²⁵ Danielle Keats Citron, *Technological Due Process*, 85 Wash. U. L. Rev. 1249, 1281–84, 1305–09 (2008) (examining notice, hearings and procedures adapted to automated administration). The Article's proposed institutional safeguards must be distinguished from what existing doctrine already compels.

³²⁶ *Cahoo v. SAS Analytics Inc.*, 912 F.3d 887, 900–04 (6th Cir. 2019). The court accepted the pertinent allegations at the dismissal stage and denied qualified immunity to specified individual defendants on the due-process theory, while treating other claims and defendants separately. It did not adjudicate that every contractor or official was liable.

tice and the opportunity to contest the particular benefits disqualification.³²⁷ The Arkansas assessment litigation illustrates a separate state-administrative-law route; an interim order or remand cannot be reported as a general final constitutional holding.³²⁸

Cross-cutting legislation may provide additional duties, but the covered domain and obligated actor remain separate questions. Colorado's reference to essential government services must be read with its deployer definition and exclusions before treating a state agency or contractor as covered.³²⁹ State administrative procedure acts govern proceedings within their own definitions and exceptions; the federal APA provides a useful analogy with its express on-the-record hearing predicate.³³⁰ Section 1983 likewise requires an enforceable right, action under color of law, a proper defendant and an available remedy. A contract with the state does not automatically make every vendor amenable to the same claim.

Benefits-specific duties should connect records and review to the agency that has power to determine the benefit. That may require procurement terms for evidence access, reasons a claimant can use, a correction route and a reviewer able to address the disputed determination. Public enforcement, contractor accountability and private litigation have different limits, including immunity and the relief available. Their coordination requires legal specification; the six process obligations are not all already mandated by due-process doctrine in the same form.

Statutory design can seek to prevent the failures illustrated by litigation without claiming that the proposed paperwork would necessarily have prevented them. The relevant test is whether a claimant can discover and contest an error before the consequence becomes irreversible, and whether the agency has the information and authority needed to respond.

7. Criminal Justice Requires Evidence Access and Decision-Specific Safeguards

State and local police, prosecutors, courts and corrections bodies make consequential decisions under public authority, alongside a substantial federal criminal system. Both levels require attention to the source and limits of the decision maker's authority.³³¹ Regulating a public deployer requires more than transferring commercial-licensing duties to an agency. The drafter must identify the institution that can supervise the use, the relevant judicial or administrative procedure, and the rights of the person affected.

³²⁷ *Barry v. Lyon*, 834 F.3d 706, 718–23 (6th Cir. 2016) (due-process notice and opportunity to contest the particular Michigan food-assistance disqualification). The decision turns on the notices, legal predicates and hearing opportunity at issue, not a rule requiring explanation of every model in every benefits program.

³²⁸ *Arkansas Department of Human Services v. Ledgerwood*, 2017 Ark. 308, 530 S.W.3d 336 (interlocutory review affirming temporary restraint on the challenged assessment methodology under the state APA). The temporary relief and state-law theory do not establish a final general constitutional holding about algorithms.

³²⁹ Colorado SB 26-189 (2026 Reg. Sess.), reenacted Colo. Rev. Stat. §§ 6-1-1701(6)(g), (7), -1705(1)(b). Listing essential government services and benefits does not by itself answer whether a particular agency or contractor meets the obligated-actor definition. *See supra* note 88.

³³⁰ 5 U.S.C. § 554(a) (federal adjudications required by statute to be determined on the record after an opportunity for an agency hearing). State APA coverage requires the particular state's definitions and exceptions; the federal provision is an analogy, not a census of state hearing rights. Program statutes and constitutional requirements supply separate predicates. *See supra* notes 319–324.

³³¹ *United States v. Lopez*, 514 U.S. 549, 561–63 (1995) (the federal commerce power and the criminal provision at issue). The state role does not exclude federal criminal jurisdiction.

Pretrial instruments estimate defined risks relevant to release, while their use by a decision maker requires separate evaluation.³³² Predictive policing raises questions about how attention and intervention affect the data used for later predictions.³³³ Facial identification presents another task: an investigative lead must not be treated as proof identifying a suspect. Wrongful-arrest claims and their distinct dispositions illustrate why the system output, additional investigation and officer conduct need separate examination.³³⁴ Sentencing, parole and supervision involve still other decisions. A tool's presence in several settings does not validate the same safeguards or legal use in each.

State v. Loomis upheld the particular sentencing use of COMPAS under limitations and cautions; it did not approve reliance on a risk score for every sentencing purpose or establish the legality of all proprietary tools.³³⁵ The later denial of certiorari supplies no endorsement of that reasoning.³³⁶ Mayson's analysis of risk and inequality challenges an assumption that a prediction's accuracy settles its fairness.³³⁷ Huq's and Ferguson's work raises further questions about the institutional choices surrounding prediction and policing.³³⁸ Wexler's trade-secrecy analysis identifies the conflict between proprietary claims and a defendant's ability to examine evidence.³³⁹ These criticisms support evidence access that can be tested in the relevant proceeding. Disclosure of a score leaves open whether the person has enough information to contest its use.

Existing measures illustrate distinct interventions. Idaho's pretrial-tool provision addresses disclosure under its particular terms.³⁴⁰ California's AB 1215 imposed a time-limited restric-

³³² Megan T. Stevenson, *Assessing Risk Assessment in Action*, 103 Minn. L. Rev. 303, 305–07 (2018) (distinguishing predictive assessment from its implementation in pretrial decision making). The passage does not establish performance in every jurisdiction or every later version of a tool.

³³³ Andrew Guthrie Ferguson, *Policing Predictive Policing*, 94 Wash. U. L. Rev. 1109, 1148–49 (2017) (data collection, targeting and feedback concerns). This analytical account does not measure the same effect in every policing system.

³³⁴ *Williams v. City of Detroit*, No. 2:21-cv-10827 (E.D. Mich.), filed settlement agreement and order (June 2024) (specified constraints on facial-recognition-derived investigation and arrests); *Woodruff v. City of Detroit*, No. 5:23-cv-11886 (E.D. Mich. Aug. 5, 2025) (granting the officer summary judgment on the claims presented). *Woodruff* left an unadopted facial-recognition-tainted-lineup theory undecided. A settlement and the adverse summary-judgment disposition have distinct evidentiary and precedential force.

³³⁵ *State v. Loomis*, 2016 WI 68, ¶¶ 50–58, 98–105, 881 N.W.2d 749. The permitted use depended on the sentencing court's independent reasons and specified cautions; COMPAS scores could not determine incarceration or sentence severity. The decision does not validate every proprietary assessment or use.

³³⁶ *Loomis v. Wisconsin*, 582 U.S. 933 (2017) (mem.) (denying certiorari, No. 16-6387, June 26, 2017). The denial supplies no decision on the merits of the constitutional question.

³³⁷ Sandra G. Mayson, *Bias In, Bias Out*, 128 Yale L.J. 2218, 2223–24 (2019) (prediction's relationship to existing racial inequality).

³³⁸ Aziz Z. Huq, *Racial Equity in Algorithmic Criminal Justice*, 68 Duke L.J. 1043 (2019) (proposing evaluation of dynamic racial effects); Ferguson, *supra* note 333. These are scholarly analyses of institutional and distributive questions, not statements that constitutional doctrine supplies every proposed safeguard.

³³⁹ Rebecca Wexler, *Life, Liberty, and Trade Secrets: Intellectual Property in the Criminal Justice System*, 70 Stan. L. Rev. 1343, 1369–71 (2018) (trade-secret claims and access to criminal evidence).

³⁴⁰ Idaho Code § 19-1910(1)–(2) (2019 enactment, reproduced in the 2025 code). The defined pretrial tools are subject to public inspection, auditing and testing; a party may review the data and calculations used for that defendant's score, and specified intellectual-property claims cannot quash discovery of the enumerated materials. The definition and the material sought control the duty.

tion on facial recognition associated with body cameras; the historical restriction must be distinguished from the law governing a later use.³⁴¹ San Francisco’s surveillance controls operate at a municipal level, with specified exceptions and qualifications.³⁴² Washington’s facial-recognition statute specifies reporting, oversight, testing and use conditions, with exceptions that matter to any account of its warrant requirements.³⁴³ Massachusetts’s reform legislation provides another defined allocation of facial-recognition use and review.³⁴⁴ These measures intervene at different stages: acquiring a tool, identifying a person, disclosing evidence or authorizing a particular use.

The design should follow the challenged use. A pretrial recommendation calls for an account of the data and factors used, their relevance to the governing release decision, and the opportunity to correct or contest them. A face-identification lead calls for controls on confirmation and further investigation. The institution must decide what evidence can be examined, under which confidentiality safeguards and in which proceeding, without assuming that a vendor’s secrecy claim or a generic audit settles the person’s rights.

Public oversight and individual remedies also require separate specification. Judicial review, legislative supervision, an authorized regulator and civil-rights litigation are not interchangeable routes. State action, standing, immunity, procedural rules and the relief sought constrain their availability. The Appendix’s general recourse template does not replace criminal procedure or resolve confrontation, disclosure and sentencing questions. The six process obligations developed next provide a design vocabulary; each must be fitted to the legal power and safeguards governing the actual decision.

D. Process-Based Obligations as the Regulatory Mode

1. Documentation as the Foundational Obligation

The deployer-centered model regulates how an institution uses AI to make consequential decisions. Its six duties are documentation, risk assessment, auditability, monitoring, incident response, and recourse.³⁴⁵ Together they make a deployment open to scrutiny: what was used, what risks were considered, what happened, and what can be done when a decision goes wrong. They also give the substantive laws developed in Part IV something to work with. A record can expose discrimination or deception; producing the record does not excuse either.

³⁴¹ California AB 1215 (2019–2020 Reg. Sess.), chaptered text, § 2, former Cal. Penal Code § 832.19(b)–(e) (officer-camera biometric-surveillance restriction, specified equitable or declaratory action, and Jan. 1, 2023 sunset). The enacted text supplies a historical intervention with an express end date, not a later statewide absence certificate.

³⁴² San Francisco Administrative Code ch. 19B (acquisition and use of surveillance technology); *see also* San Francisco Police Department, *19B Surveillance Technology Policies* (describing the 2019 ordinance and its facial-recognition restrictions with limited exceptions). The municipal level supplies a distinct institutional example, not a nationwide count or a claim of unchanged local law.

³⁴³ Wash. Rev. Code ch. 43.386, including §§ 43.386.020–.060, .080–.090. Reporting, testing, meaningful human review and law-enforcement-use restrictions have separate conditions and exceptions; facial recognition cannot be the sole basis for probable cause.

³⁴⁴ 2020 Mass. Acts ch. 253, § 26 (adding Mass. Gen. Laws ch. 6, § 220). The defined searches, court-order or emergency routes, recordkeeping and exceptions delimit the restriction.

³⁴⁵ The six duties are this Article’s proposed regulatory design. Parts II and IV address their legal authority and limits; process form does not itself establish compatibility with federal law.

The constitutional distinction comes first. Procedural due process applies when conduct attributable to government deprives a person of a protected interest. Government use of AI does not, by itself, make all six duties constitutionally mandatory. The required procedure depends on the interest, the risk of error, the value of additional safeguards, and the government's burdens. Nor does a private deployer's label settle the state-action inquiry. The constitutional cases supply a design source for rights against private actors; the authority to enact those rights must come from the state's consumer-protection, civil-rights, licensing, or procurement powers.³⁴⁶

Documentation supports the other duties. An assessment needs information about the system and its use; an auditor needs a record to examine; monitoring needs a baseline; an incident investigation needs an account of what failed; and an affected person needs enough information to challenge the decision. Each needs different information. A regulator's confidential examination file, a public assessment summary, and a person's explanation of an adverse decision should not be treated as interchangeable disclosure obligations.

Existing law supplies useful, bounded examples. *Goldberg* and *Mathews* address fair procedures for government deprivations, without establishing a universal right to an algorithm's source code.³⁴⁷ California requires written reasons and a factual basis in covered administrative adjudications; New York's Article 78 review distinguishes arbitrary-and-capricious review from substantial-evidence review of determinations after legally required hearings.³⁴⁸ Employment-selection rules impose recordkeeping and, where their conditions are met, validation requirements. Those rules do not make every deployer subject to one uniform documentation schedule across employment, housing, and other domains.³⁴⁹

Management frameworks help specify the work. NIST AI RMF's GOVERN and MAP functions address organizational accountability, deployment context, system characteristics, and impacts.³⁵⁰ ISO/IEC 42001 supplies an organizational AI management-system approach, including policy and risk-assessment processes.³⁵¹ Neither framework alone establishes an affected person's legal entitlement to records. The legislative step is to select the information, duty-bearer, access route, and consequence that a particular deployment requires. Framework conformity is evidence about a management process, not proof that the resulting decisions are

³⁴⁶ *Goldberg v. Kelly*, 397 U.S. 254, 262–64, 267–71 (1970); *Mathews v. Eldridge*, 424 U.S. 319, 332–35, 347–49 (1976). The constitutional inquiry concerns attributable government conduct, a protected interest, and the procedure due in context. The proposed duties against private deployers require statutory authority. See *infra* Part IV.

³⁴⁷ *Goldberg*, 397 U.S. at 267–71 (opportunity to present the recipient's case and a decision based on the evidence, without a required full opinion or formal findings); *Mathews*, 424 U.S. at 335, 348–49 (context-specific balancing and the procedures at issue). These holdings do not prescribe a universal algorithmic documentation form.

³⁴⁸ Cal. Gov't Code § 11425.50(a)–(e) (written decision, reasons, record basis, and specified findings in covered adjudications); N.Y. C.P.L.R. § 7803(3)–(4) (distinct grounds for Article 78 review). See *supra* note 330. Neither provision supplies a general private-deployer disclosure duty.

³⁴⁹ 29 C.F.R. §§ 1607.3(A), .4(A)–(E), .6(A)–(B); *Griggs v. Duke Power Co.*, 401 U.S. 424, 431–32 (1971); *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 425–27 (1975). Recordkeeping, evidence of adverse impact, validation, and ultimate liability are distinct inquiries. See *infra* Part IV.B for the separate limits governing the civil-rights authorities.

³⁵⁰ Nat'l Inst. of Standards & Tech., *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, §§ 5.1–5.2, tbls. 1–2 (2023) (including GOVERN 1.1–1.5, 2.1 and MAP 1.1, 2.2–2.3). The framework is voluntary and adaptable to context.

³⁵¹ ISO/IEC 42001:2023, *Information technology — Artificial intelligence — Management system*, cls. 5.2, 6.1.1–6.1.4 (2023). Organizational policy and risk processes are distinct from individual disclosure rights.

lawful or sound.

The existing instruments illustrate different parts of that step. Colorado’s SB 26-189 retains developer documentation and deployer obligations supporting notice, explanation, correction, and review in covered consequential decisions.³⁵² New York City requires a published bias-audit summary and specified employment notices.³⁵³ California’s risk-assessment and ADMT rules address specified processing by covered businesses, with distinct assessment and consumer-rights requirements.³⁵⁴ The NAIC Model Bulletin describes a written AI Systems Program proportionate to an insurer’s risks; its legal force depends on the adopting jurisdiction and applicable insurance law.³⁵⁵ These instruments allocate information and responsibility differently.

This Article proposes documentation of purpose, use context, inputs and outputs, validation, limitations, material changes, and responsible personnel. The strongest objection is that a deployer often did not build the system. A hospital licensing a prediction tool may have no access to its training data; an employer may receive a score without the vendor’s validation record. Appendix § 3.4(a) therefore limits upstream training-data documentation to material held or reasonably obtainable, while requiring a record of unavailable material, efforts to obtain it, and substitute assurance accepted under the implementing rules.³⁵⁶ That fallback makes the gap visible. It cannot establish the adequacy of information that remains missing.

Colorado illustrates the trade without resolving it. The 2026 successor reduced important duties in the earlier framework, including the deployer impact-assessment and risk-management requirements, but retained developer documentation and recordkeeping. It did not simply abandon the supply side of the information problem.³⁵⁷ A state choosing narrower developer duties must still ask what deployers can obtain, what contracts can secure, and when inadequate information should prevent a particular use. Those choices have operational costs as well as constitutional implications.

Seven County offers a limited analogy. Under NEPA, the Surface Transportation Board did not have to evaluate specified effects of separate projects outside its regulatory authority; the

³⁵² Colorado S.B. 26-189, §§ 6-1-1702 to -1705, -1708 to -1709, § 5 (2026). Developer information and update duties are distinct from deployer compliance records, which must be retained for at least three years. Consumer duties have coverage and exception limits and generally commence January 1, 2027. *See supra* note 88.

³⁵³ N.Y.C. Admin. Code §§ 20-870 to -874; 6 R.C.N.Y. §§ 5-300 to -304. Audit, publication, notice, and data-disclosure requirements have distinct conditions. *See infra* Part V.C.

³⁵⁴ Cal. Code Regs. tit. 11, §§ 7001(e), (ddd), 7150, 7152, 7155, 7157, 7200, 7220–7222. The January 1, 2026 effective regulations distinguish assessment obligations from ADMT consumer obligations, whose compliance begins January 1, 2027. *See supra* note 198.

³⁵⁵ Nat’l Ass’n of Ins. Comm’rs, *Model Bulletin: Use of Artificial Intelligence Systems by Insurers* §§ 3–4 (adopted Dec. 4, 2023) (risk-proportionate AIS Program, third-party systems, validation, monitoring, and examination); *see supra* note 289. The Bulletin permits appropriate alternative practices and is not itself a uniform state statute.

³⁵⁶ *Infra* Appendix § 3.4(a) (possession, custody, control, reasonable obtainability, and the record of unavailability, efforts, and substitute assurance); *id.* § 4.4(b)–(c) (useful decision explanation and specified protection of confidential details). These provisions serve different access functions.

³⁵⁷ Colorado S.B. 24-205, §§ 6-1-1702 to -1703 (2024), compared with S.B. 26-189, §§ 6-1-1702 to -1709, § 5 (2026). The successor retains developer information, change-notice, and record duties while removing the earlier general deployer assessment and risk-management architecture. *See supra* notes 9, 16, 18, 88. Reduced developer-side exposure is a design inference, not an adjudicated consequence or a finding about legislative motive.

majority applied *Public Citizen*'s account of legally relevant causation.³⁵⁸ The three Justices concurring in the judgment relied on the Board's limited statutory authority, while disagreeing with the majority's broader reasoning.³⁵⁹ Neither opinion creates a constitutional exemption for a private deployer that lacks upstream information. The useful design question is where responsibility can be made effective. A deployer controls whether to use the tool, for which decisions and populations, under which safeguards, and with what response to failure. A duty focused on those choices is substantial. Whether it is sufficient depends on the information the deployment actually requires.

2. Risk Assessment as the Operationalizing Obligation

Documentation records what the deployer knows. Risk assessment asks what that information means for people subject to the decision: what could go wrong, how serious it would be, how likely it is, and what can reduce the risk. A pre-deployment assessment cannot answer those questions permanently. Material changes and experience in use can require a new assessment, while monitoring supplies information between formal reviews.

Mathews makes the risk of erroneous deprivation relevant to the constitutional procedure due in a particular government setting; it does not require a standardized AI impact assessment in every case.³⁶⁰ Foreseeable harm is also a useful organizing question for this Article's proposed duty, without importing a general products-liability obligation onto every user of software.³⁶¹ The more concrete regulatory analogies are sector-specific. Banking model-risk guidance addresses development, use, validation, and monitoring, subject to the successor guidance's scope and exclusions.³⁶² FDA's Quality Management System Regulation incorporates a device quality-management standard and its risk-management requirements; it is not a general AI-deployer law.³⁶³

NIST's MAP and MEASURE functions help characterize risks and select evaluation methods.³⁶⁴ ISO/IEC 42001's planning provisions address organizational AI risk assessment, risk

³⁵⁸ *Seven County Infrastructure Coalition v. Eagle County*, No. 23-975, slip op. at 16–19 (U.S. May 29, 2025); *Department of Transportation v. Public Citizen*, 541 U.S. 752, 765–70 (2004). The relevant inquiry concerns federal agency authority and NEPA causation; the AI duty proposed here rests on different legal authority.

³⁵⁹ *Seven County*, No. 23-975, slip op. at 1, 7–10 (Sotomayor, J., concurring in the judgment, joined by Kagan and Jackson, JJ.). Kavanaugh wrote for five Justices; Gorsuch did not participate. The concurrence's narrower route should not be attributed to the majority or restated as a general private-actor immunity.

³⁶⁰ *Mathews*, 424 U.S. at 335, 347–49; see *supra* note 347. Error risk and the likely value of safeguards enter the constitutional balance; the particular statutory assessment proposed here is a separate design choice.

³⁶¹ Foreseeability organizes the proposed assessment: identify plausible adverse effects before deployment and reconsider them as evidence changes. This proposal does not depend on classifying every AI system as a product or treating every deployer as its manufacturer.

³⁶² Bd. of Governors of the Fed. Rsr. Sys., *Revised Guidance on Model Risk Management*, SR 26-2, at 1–3, 9–12 (Apr. 17, 2026); see *supra* note 184. The nonbinding successor is most relevant to institutions above the stated asset scale and excludes generative and agentic AI from its scope. SR 11-7 supplies the historical predecessor.

³⁶³ 21 C.F.R. §§ 820.1, .7, .10 (QMSR, effective Feb. 2, 2026). The regulation incorporates ISO 13485:2016 and specifies the covered manufacturer's quality-system obligations; design-and-development scope varies by device class. It is distinct from the particular marketing-authorization pathway.

³⁶⁴ NIST AI RMF, *supra* note 350, §§ 5.2–5.3, tbls. 2–3 (MAP context and impact characterization; MEA-

treatment, and system impact assessment. Selection of controls is part of that process; Annex A is not a command to apply every control to every deployment.³⁶⁵ The proposed statute must decide which assessment elements become enforceable and who must perform them.

The comparators differ in both timing and reach. Colorado's earlier 2024 enactment called for annual impact assessments and assessment within ninety days after an intentional and substantial modification was made available; SB 26-189 does not carry that assessment mandate forward.³⁶⁶ California requires assessments before covered processing and later review and updates under its own schedule.³⁶⁷ Virginia's vetoed HB 2094 proposed developer-supported deployer assessments, but did not become law.³⁶⁸ The NAIC Bulletin describes risk-proportionate controls and validation as supervisory expectations.³⁶⁹ EU AI Act Article 27 requires fundamental-rights assessments for specified categories of high-risk-system deployers.³⁷⁰

The proposal here uses three triggers: initial deployment, a material change, and scheduled review. It asks the deployer to identify foreseeable adverse effects, assess probability and severity, select safeguards, and explain the decision to proceed. Part V.B distinguishes the broader design choices from the precise content and intervals in the Appendix. An assessment should make a decision answerable, not turn a completed form into permission to disregard substantive law.

3. Auditability as the Condition for Backward-Looking Third-Party Review

Auditability makes the record usable by someone other than its maker. A deployer must preserve enough information to reconstruct the system's role, examine a disputed result, and test whether stated safeguards operated. This is different from a duty to commission an independent audit: a system may need to be auditable even when no periodic third-party assessment is required.

Access is the central difficulty. In *Loomis*, the Wisconsin Supreme Court permitted limited consideration of a proprietary risk score at sentencing, with restrictions and cautions. It did not establish a general right to inspect the underlying formula. Wexler's criticism exposes the

SURE methods, evaluation, and risk tracking). The framework's practices do not create a legal assessment duty without an applicable legal instrument.

³⁶⁵ ISO/IEC 42001:2023, *supra* note 351, cls. 6.1.2–6.1.4. Selected controls require justification; not every Annex A control applies to every system.

³⁶⁶ Colorado S.B. 24-205, § 6-1-1703(3)(a)–(g) (2024), particularly subsection (3)(b)'s annual and intentional-and-substantial-modification triggers. This is the earlier enacted design, rather than a claim that its original commencement text remains operative. S.B. 26-189, § 5 (2026), generally commences January 1, 2027 and omits that mandate. *See supra* note 88.

³⁶⁷ Cal. Code Regs. tit. 11, §§ 7150(a), 7152, 7155. Covered processing includes specified ADMT significant decisions; existing processing has a separate December 31, 2027 assessment deadline. *See supra* note 354.

³⁶⁸ *See supra* note 211 (Virginia H.B. 2094, 2025 session, vetoed). It is a proposed comparator, not an enacted assessment duty.

³⁶⁹ NAIC Model Bulletin, *supra* note 355, § 3 (written program, risk controls, validation and documentation proportionate to risk). Applicable law and jurisdictional adoption determine enforceability.

³⁷⁰ Regulation (EU) 2024/1689, art. 27, annex III, points 5(b)–(c), arts. 111, 113, as amended by Regulation (EU) 2026/1744. Article 27 identifies public-law bodies, private public-service providers, and specified credit and insurance deployers, with stated exceptions. Scope and application dates must be read together; it is not a universal deployer assessment rule.

cost of relying on information that the defense cannot independently test.³⁷¹ The policy question remains how to make review effective while protecting legitimate confidential material.

SOX and PCAOB provide one institutional analogy: management assessment, independent examination, defined standards, and oversight of auditors, within a financial-reporting regime that has important attestation exemptions.³⁷² Civil discovery provides another. Federal Rule 26 permits relevant, nonprivileged, proportional discovery and allows a court to protect trade secrets, including by restricting or withholding disclosure. Neither relevance alone nor the availability of a protective order guarantees access. The unpublished *Chubbs* decision illustrates the distinct barriers in criminal discovery: on that record, the defendant had not made the particularized showing necessary to obtain the proprietary code.³⁷³ Wexler argues for using protective arrangements after ordinary discovery requirements are satisfied; that is a normative position, not a uniform rule governing algorithmic evidence.³⁷⁴

Sectoral examination powers offer an administrative route, within the regulator's jurisdiction and confidentiality rules.³⁷⁵ NIST's MEASURE and MANAGE functions support evaluation, tracking, and response.³⁷⁶ Assurance about an organizational process must still be distinguished from independent testing of a particular decision system.³⁷⁷ An auditor's access to a management process does not itself give an affected person access to the evidence needed to challenge an adverse decision.

The current examples reflect those distinctions. New York City requires an independent bias audit under specified methodology, but the Comptroller identified problems in implementation and enforcement.³⁷⁸ Colorado's successor preserves compliance records without imposing the earlier assessment regime.³⁷⁹ The NAIC's 2026 evaluation-tool pilot contemplated use in

³⁷¹ *State v. Loomis*, 881 N.W.2d 749, ¶¶ 50–63, 98–105 (Wis. 2016); Rebecca Wexler, *Life, Liberty, and Trade Secrets: Intellectual Property in the Criminal Justice System*, 70 Stan. L. Rev. 1343, 1369–71 (2018). *Loomis* permitted the limited use at issue with cautions and independent sentencing reasons; Wexler criticizes the barriers to testing proprietary evidence. See *supra* notes 335, 339.

³⁷² 15 U.S.C. §§ 7211, 7214–7215, 7262; PCAOB, AS 2201, ¶¶ .01–.04, .13, app. A ¶ A5. The management assessment, covered auditor opinion, inspection and disciplinary functions are different. Section 404(b) attestation has statutory exemptions. See *infra* notes 549–550.

³⁷³ Fed. R. Civ. P. 26(b)(1)–(2), (c)(1)(G) (Dec. 1, 2025 rules); N.H. Rev. Stat. Ann. § 350-B:5 (secrecy protection in an action under that trade-secrets chapter). *People v. Superior Court (Chubbs)*, No. B258569, slip op. at 15–16, 21–24 (Cal. Ct. App. Jan. 9, 2015) (unpublished), <https://www4.courts.ca.gov/opinions/nonpub/B258569.PDF> (granting relief from disclosure on the record-specific relevance-and-necessity showing). Neither the civil rules nor that unpublished criminal ruling establishes unconditional algorithmic access.

³⁷⁴ Wexler, *supra* note 371, at 1409–12 (ordinary discovery thresholds followed by protective arrangements, with attention to restrictions that can prevent useful expert work). The argument is a proposed approach to criminal discovery, not a description of settled practice in every jurisdiction.

³⁷⁵ See *infra* Part IV.C (sectoral authority and jurisdictional limits); NAIC Model Bulletin, *supra* note 355, § 4 (regulatory requests and examination). These sources do not establish identical schedules, credentials, or access rules for insurance, banking, and hospital regulators.

³⁷⁶ NIST AI RMF, *supra* note 350, §§ 5.3–5.4, MEASURE 1.3, 2.1, 2.9, 2.13 and MANAGE 4.1–4.3. Independent assessment within the framework is not synonymous with a legally mandated external audit.

³⁷⁷ The proposed external audit and decision-specific access rights must be specified independently of organizational assurance. See *supra* note 372 and *infra* Part V.C.

³⁷⁸ N.Y.C. Admin. Code §§ 20-870 to -874; 6 R.C.N.Y. §§ 5-300 to -304; N.Y. State Comptroller, *Enforcement of Local Law 144 – Automated Employment Decision Tools*, Report 2024-N-6, at 15–19, 23, 32–35 (Dec. 2, 2025). See *infra* notes 546–548 for actual requirements, sample limits, and agency comments.

³⁷⁹ Colorado S.B. 26-189, §§ 6-1-1703, -1706, -1709 (2026); see *supra* note 352. Records and exclusive At-

market-conduct and financial examinations and reviews; pilot use is different from a generally adopted periodic audit mandate.³⁸⁰ California's assessment records and regulator-submission provisions are also distinct from its separate cybersecurity-audit requirements.³⁸¹

The proposed design therefore assigns different access routes. Regulators obtain records through defined examination or filing powers; credentialed auditors obtain material needed for the specified audit; affected persons obtain decision-specific explanations and correction rights; courts apply the governing discovery rules. Confidentiality must be addressed at each route. Appendix § 4.4 requires a useful explanation while protecting specified trade-secret details, and § 3.4(a) addresses unavailable upstream material. Neither provision makes the information problem disappear.

4. Monitoring as an Ongoing Duty

Monitoring observes performance between assessments and audits. The deployer compares actual use with the assumptions and baselines on which deployment was approved, looking for degradation, changed input populations, recurring errors, and discriminatory patterns. Its frequency should follow the consequences and pace of the decisions. An annual audit and ongoing observation serve different purposes.

Banking model-risk guidance illustrates monitoring tied to validation and changing conditions.³⁸² FDA postmarket surveillance and adverse-event reporting illustrate legally specified obligations imposed on particular device actors and products.³⁸³ Drug safety reporting supplies another bounded comparison; expedited case reporting and periodic aggregate reporting have different rules and guidance.³⁸⁴ These regimes do not impose identical duties, but each recognizes that initial review cannot settle every question arising in use.

NIST's MEASURE and MANAGE functions supply a voluntary vocabulary for tracking risk and responding to what observation reveals.³⁸⁵ Responsibility and cadence must be chosen

torney General enforcement do not constitute an independent-audit mandate or create a new private action.

³⁸⁰ NAIC, *AI Systems Evaluation Tool Pilot Project Summary*, at 1–3 (2026) (planned March–September use in examinations and reviews, feedback, and consideration of adoption). The plan does not establish completed findings or final adoption. *See supra* Part III.C.3 and *infra* Part IV.C.2 for the later supplement/exposure posture.

³⁸¹ Cal. Code Regs. tit. 11, §§ 7155, 7157 (assessment retention and submissions). On a qualifying request, the agency or Attorney General can require the assessment report within thirty days. This is distinct from a mandatory independent bias audit. *See supra* note 354.

³⁸² SR 26-2, *supra* note 362, at 9–12; Federal Reserve SR 11-7 (Apr. 4, 2011), historical model-validation discussion identifying conceptual soundness, ongoing monitoring, and outcomes analysis. Frequency and intensity follow risk and circumstances, not a uniform continuous-observation rule.

³⁸³ 21 C.F.R. §§ 822.1, .4–.7 (qualifying class II or III devices and FDA surveillance orders); *id.* § 803.10 (distinct reporting duties). Initial marketing status alone does not impose every surveillance obligation on every device actor.

³⁸⁴ 21 C.F.R. § 314.80(b)–(c), (j)–(l) (approved-application-holder safety review, reporting, and records); FDA, *E2D(R1) Postapproval Safety Data: Definitions and Standards for Management and Reporting of Individual Case Safety Reports* 1–2 (Mar. 2026). This final nonbinding guidance replaces the 2003 E2D guidance and refers periodic aggregate reporting to E2C(R2).

³⁸⁵ NIST AI RMF, *supra* note 350, MEASURE 3.1–3.3, MANAGE 3.1–3.2, 4.1–4.3 (tracking, monitoring, feedback, response, and communication practices).

for the deployment; organizational certification alone cannot settle them.³⁸⁶ The NAIC Bulletin likewise describes monitoring and validation proportionate to insurer risk.³⁸⁷ Colorado's earlier lifecycle risk-management requirement was not carried into SB 26-189.³⁸⁸ California's risk-assessment schedule is periodic and event-triggered; it does not impose the continuous deployment-baseline rule proposed here.³⁸⁹ EU Article 26 contains monitoring duties for covered high-risk-system deployers under its own conditions and application schedule.³⁹⁰

This Article proposes observation throughout the covered deployment, with records of material deviations and corrective action. The Appendix gives that proposal more specific form: protocols under § 2.4(a), at least quarterly monitoring for the highest tier under § 3.4A, and scheduled audits under § 3.5. The records should feed reassessment, incident response, and recourse. Requiring observation without specifying what follows a serious finding would leave the central duty unfinished.

5. Incident Response as Shared Regulatory Knowledge

Incident response connects a detected failure to investigation, remediation, and notification. Its distinctive contribution is to let a regulator see patterns across deployers that each deployer may see only in its own operations. It differs from the individual's right to challenge a particular decision, although the same event may trigger both.³⁹¹

Existing reporting regimes show what must be specified. California's breach-notification laws define covered information, triggering events, recipients, and exceptions.³⁹² Medical-device reporting distinguishes manufacturers, importers, and user facilities, with different events and deadlines.³⁹³ CIRCIA specifies seventy-two-hour incident and twenty-four-hour ransomware-payment reporting, but makes the statutory duties contingent on implementing final rules

³⁸⁶ A monitoring interval is a legislative choice tied to deployment risk. *See supra* Part III.D.2 and *infra* Appendix § 3.4A.

³⁸⁷ NAIC Model Bulletin, *supra* note 355, § 3, including risk-management controls, validation, and monitoring. Expectations scale to insurer and system risk; jurisdictional legal effect is separate.

³⁸⁸ Colorado S.B. 24-205, § 6-1-1703(1), (2)(a) (2024) (earlier reasonable-care and lifecycle risk-management design). S.B. 26-189 does not carry that general architecture forward. *See supra* notes 88, 366.

³⁸⁹ Cal. Code Regs. tit. 11, § 7155(a) (review at least every three years, update as necessary, material-change update as soon as feasible and no later than forty-five days). The assessment record is retained for the duration of the processing or five years after completion of the risk assessment, whichever is later; *id.* § 7155(c). These rules differ from this Article's ongoing monitoring proposal.

³⁹⁰ Regulation (EU) 2024/1689, art. 26(1)–(6), especially paragraph (5), arts. 111, 113, as amended. Monitoring follows the instructions and statutory conditions; notification, suspension, controlled logs, and financial-sector governance have distinct provisions. *See supra* note 370.

³⁹¹ Individual recourse concerns a decision and the affected person's rights. Incident reporting can reveal a pattern across decisions and deployers. The functions can overlap, and other substantive or contractual law may already require action even without a separate AI incident statute.

³⁹² Cal. Civ. Code § 1798.82(a)–(l), as amended by 2025 Cal. Stat. ch. 319, § 1 (S.B. 446), effective Jan. 1, 2026. The statute distinguishes owners or licensees from maintainers, defines covered information and acquisition, and specifies notices, timing, permitted delays, and exceptions. No national deadline range is inferred from this example.

³⁹³ 21 C.F.R. § 803.10(a)–(c). User facilities, importers, and manufacturers have different recipients, triggers, and ten-workday, thirty-calendar-day, or five-workday reporting requirements. *See supra* note 383.

taking effect.³⁹⁴ A timeline in an enacted statute should not be described as an immediately operative duty when the statute itself postpones its operation.

The earlier federal acquisition memorandum, M-24-18, directed agencies, to the greatest extent practicable, to include terms for vendor reporting of serious AI incidents or malfunctions within seventy-two hours or another severity-appropriate timely period after reasonable belief.³⁹⁵ Its replacement by M-25-22 illustrates the distinction between the continuing procurement channel and the particular terms an administration chooses to require. The successor's existence does not establish that every earlier reporting term survived.³⁹⁶ NIST's MANAGE function supports response and communication practices.³⁹⁷ The proposed notification duty needs its own trigger, recipient, and deadline; a management process cannot substitute for those legal terms.³⁹⁸

The state examples caution against assuming a complete package. Colorado's earlier law distinguished developer notice of discrimination risk from deployer notice of discrimination that had occurred; the successor removed that Attorney General self-reporting mechanism.³⁹⁹ The NAIC Bulletin supports regulatory examination and information requests, without specifying a general material-AI-failure notification deadline.⁴⁰⁰ EU Article 73 principally assigns serious-incident reporting to providers, with deployer responsibilities addressed separately in Article 26.⁴⁰¹

The proposed duty should define material failures, the point at which knowledge starts the clock, investigation requirements, recipients, and severity-based deadlines. Notification of affected persons is a further design choice, not an automatic consequence of regulator reporting. Appendix § 3.4A requires notice to the CAIO and sectoral regulator within seventy-two hours of discovery of a Reportable Incident, with remediation and documentation; it does not enact a general affected-person incident-notification rule. Materiality also matters: reporting every deviation could bury serious failures in noise and discourage candid monitoring. That is a design

³⁹⁴ 6 U.S.C. § 681b(a)(1)–(7), (b)–(c). Subsection (a)(7) ties the duties in paragraphs (1)–(4) to effective dates prescribed by final implementing rules. The seventy-two-hour clock concerns reasonable belief of a covered incident; the twenty-four-hour clock concerns ransom payment. The statutory design does not itself establish that an implementing rule is effective.

³⁹⁵ OMB Memorandum M-24-18, *Advancing the Responsible Acquisition of Artificial Intelligence in Government* § 4(d)(v), at 16–17 (Sept. 24, 2024). The provision includes the greatest-extent-practicable qualification, reasonable-belief trigger, and severity-calibrated reporting alternative; it should not be restated as an unqualified subcontract flow-down requirement.

³⁹⁶ OMB Memorandum M-25-22, *Driving Efficient Acquisition of Artificial Intelligence in Government*, overview, §§ 2, 4 (Apr. 3, 2025) (rescinding and replacing M-24-18). Continued acquisition controls do not prove retention of each prior incident term. *See infra* Part IV.D.

³⁹⁷ NIST AI RMF, *supra* note 350, MANAGE 2.3–2.4, 4.1–4.3 (response, recovery, override or deactivation, monitoring, and communication). These are voluntary practices, distinct from the proposed legal reporting duty.

³⁹⁸ A selected management practice does not by itself establish a legal notification duty. *See infra* Appendix § 3.4A for the proposed trigger and recipients.

³⁹⁹ Colorado S.B. 24-205, § 6-1-1702(5) (developer discrimination-risk notice), § 6-1-1703(7) (deployer discovery that use has caused algorithmic discrimination), (2024). S.B. 26-189 omits that Attorney General self-reporting architecture. *See supra* notes 88, 366.

⁴⁰⁰ NAIC Model Bulletin, *supra* note 355, §§ 3–4. Internal escalation, validation, regulatory requests, and examination are different from the general material-failure notification timetable proposed here.

⁴⁰¹ Regulation (EU) 2024/1689, art. 73(1)–(11), read with art. 26(5). Provider reporting includes causal-link, timing, severity, and sectoral qualifications. Article 73's location and application must not be inferred from the separate Chapter III deployer provisions. *See supra* note 370.

risk to address, not an established reason to suppress adverse findings.

6. Recourse as a Route to Contest an Adverse Decision

Recourse gives the person subject to the decision a practical way to challenge it. Notice identifies AI's role; explanation identifies relevant factors; access and correction address inaccurate inputs; human review tests the decision; and remedies address violations or harm under the applicable cause of action. A private decision becomes contestable through the rights the legislature creates. It does not thereby become administrative action.

Federal law offers distinct models. FCRA requires specified notices from users taking adverse action based on consumer reports, while consumer reporting agencies and furnishers have different dispute-related duties.⁴⁰² ECOA and Regulation B require creditors to give, or provide access to, specific reasons under the applicable notification rules.⁴⁰³ *Goldberg* and *Mathews* address government deprivations of protected interests and context-specific procedures.⁴⁰⁴ The ADA's accommodation and modification duties depend on the title, actor, disability-related need, and statutory exceptions; they illustrate individualized treatment without establishing a universal right to human re-decision.⁴⁰⁵

The CFPB's withdrawal of Circular 2022-03 did not itself repeal ECOA or Regulation B.⁴⁰⁶ The distinction supports a narrower durability claim: an agency can withdraw an interpretation without eliminating its underlying statute. The authority of a particular state enforcer still depends on applicable law and preemption; concurrent authority is not immune from legislative change.⁴⁰⁷ NIST addresses feedback, appeal, and redress practices.⁴⁰⁸ The proposed individual remedy requires its own legal terms, even when organizational governance already supports feedback and response.⁴⁰⁹

⁴⁰² 15 U.S.C. § 1681m(a) (report-user notice, reporting-agency identification, copy and dispute rights, and specified score information); *id.* §§ 1681i, 1681s-2 (different agency and furnisher functions). Regulation V § 1022.74 supplies risk-based-pricing exceptions, including an exception where an adverse-action notice was given, rather than a general adverse-action rule.

⁴⁰³ 15 U.S.C. § 1691(d)(1)–(6); 12 C.F.R. § 1002.9(a)–(b), official interpretations to paragraph (b)(2). Specific principal reasons must relate to the factors actually considered; generic internal-policy or qualifying-score statements are inadequate. The governing provisions distinguish notice, requested reasons, timing, and exceptions.

⁴⁰⁴ *Goldberg*, 397 U.S. at 262–64, 267–71; *Mathews*, 424 U.S. at 332–35, 347–49; *see supra* notes 346–347. The cases support context-specific constitutional procedures, not every element of the broader statutory recourse proposal.

⁴⁰⁵ 42 U.S.C. §§ 12112(a), (b)(5)–(7), 12182(a), (b)(2)(A)(ii)–(iii); 28 C.F.R. § 35.130(a), (b)(7). Employment accommodation, public-service modification, public-accommodation modification, and auxiliary aids have different predicates and defenses.

⁴⁰⁶ CFPB, *Withdrawn Guidance*, entry for Consumer Financial Protection Circular 2022-03 (listing May 12, 2025 withdrawal); 15 U.S.C. § 1691(d); 12 C.F.R. § 1002.9. The withdrawal concerns guidance; the statutory and regulatory requirements have their own legal basis.

⁴⁰⁷ *See supra* notes 396, 403 and *infra* Part IV. Replacement of executive guidance differs from repeal of a statute. The continuity of a regulatory channel does not establish the survival of a particular requirement or enforcer's authority against every preemption challenge.

⁴⁰⁸ NIST AI RMF, *supra* note 350, MEASURE 3.3, MANAGE 4.1, 4.3 (feedback, appeals, override, communication and response). These practices help design recourse but do not create an enforceable individual cause of action.

⁴⁰⁹ The individual's rights, duty-bearers, and remedies require statutory terms beyond an organization's governance commitments. *See infra* Part V.D and Appendix § 4.

Colorado's successor provides notice, post-adverse-outcome information, correction, and review duties with material limits: coverage depends on the technology's role and the defined adverse outcome, correction excludes opinions, predictions, and scores, and human review is required to the extent commercially reasonable, subject to exceptions and the January 1, 2027 commencement.⁴¹⁰ New York City's employment notice and accommodation-request provisions are narrower than a general human-review right.⁴¹¹ California's ADMT access and opt-out rules have their own definitions, exceptions, and compliance dates; they should not be conflated with every right in the broader CCPA.⁴¹²

The proposal here gives an affected person an intelligible explanation, a correction process, and review by someone able to examine the record and change the result where warranted. The Appendix makes the distinction between review and re-decision express: § 4.6(a) authorizes review of a covered consequential decision; § 4.6(c) requires re-decision when review identifies error or finds the decision unsupported on the whole record. A challenge does not guarantee a different outcome. FCRA subject-matter preemption, arbitration, standing, and the remedy's legal basis remain part of the design, as Part V.D explains.

The six duties work together, but their legal footing must be assessed separately. Documentation makes conduct legible; assessment examines foreseeable risk; auditability permits scrutiny; monitoring observes use; incident response organizes action after failure; recourse lets the affected person contest a decision. Attaching them to in-state deployment and established regulatory channels can improve the fit between authority and obligation. It cannot make a duty preemption-proof or establish that it regulates only local effects. The four authorities examined in Part IV supply enactment and enforcement channels within those boundaries.

Part IV: Matching State Authorities to Duties, Institutions and Remedies

The four authorities developed here do different work. Consumer protection addresses deceptive or unfair conduct; civil rights law supplies substantive prohibitions; sectoral licensing supervises identified institutions and professionals; procurement governs what the State buys and uses. For each, the question is how an AI-related duty reaches a responsible actor, who can investigate and enforce it, and what relief is available. Existing authority, a lawful application and an effective remedy are separate steps. The portfolio is designed to make those steps explicit.

⁴¹⁰ Colorado S.B. 26-189, §§ 6-1-1701, -1704 to -1705, -1708 to -1709, § 5 (2026). The duties distinguish material influence and adverse outcome, further-information requests, inaccurate personal data, commercial reasonableness and exceptions. No nationwide ranking of recourse architectures is needed for the comparison. *See supra* note 88.

⁴¹¹ N.Y.C. Admin. Code § 20-871(b)–(c); 6 R.C.N.Y. § 5-304. Notice to covered candidates or employees, job qualifications and characteristics, and a route to request alternative selection or accommodation are specified; the law does not require an alternative selection process in every case. Data disclosure has its own conditions and exceptions.

⁴¹² Cal. Code Regs. tit. 11, §§ 7001(e), (ddd), 7200, 7220–7222. ADMT notice, opt-out and access differ from underlying CCPA correction and deletion rights. Specified human appeal is an alternative within an opt-out exception, not a universal human-review entitlement. Consumer compliance commences January 1, 2027.

A. Consumer Protection: The Representation and the Available Claim

State consumer-protection law can reach conduct involving AI without a statute named for the technology. The NCLC survey dated March 2018 found UDAP statutes in the fifty states and in D.C. Massachusetts's prohibition uses the familiar unfair-or-deceptive formulation; other states supply different texts and limits.⁴¹³ California, New York and Massachusetts each provide public enforcement and private actions, but private standing, available relief and the conduct privately actionable differ. New York's unfair and abusive limbs, for example, belong to the Attorney General; the private action specified in § 349(h) concerns deception.⁴¹⁴ A nationwide statutory presence does not establish nationwide remedial equivalence.

Federal Section 5 enforcement supplies an analogue, not a grant of state power. Nor does one provision preserving state remedies establish that the whole federal consumer-protection scheme leaves them untouched.⁴¹⁵ Operation AI Comply illustrates the distinction between conduct and technology. The FTC's September 2024 announcement collected administrative complaints concerning DoNotPay and Rytr and three federal-court business-opportunity cases. These allegations and proceedings illustrate the use of identified laws; they supply no established body of state AI doctrine.⁴¹⁶ Rytr's complaint alleged both means and instrumentalities of deception and unfairness; DoNotPay concerned representations about the performance of an AI legal service. Their theories should not be conflated.⁴¹⁷ The state examples have their own posture: Texas's assurance with Pieces Technologies resolves an enforcement matter on specified terms, while California's January 2025 advisories state the Attorney General's views about existing law. Advice, an assurance and an adjudicated holding establish different things.⁴¹⁸

The Rytr sequence also shows why an enforcement example needs its later qualification. The

⁴¹³ See *supra* note 157 (the March 2018 NCLC survey and state texts); Mass. Gen. Laws ch. 93A, § 2(a); 15 U.S.C. § 45(a)(1)-(2) (federal prohibition and enforcement coverage, including specified exclusions).

⁴¹⁴ Cal. Bus. & Prof. Code §§ 17203, 17204, 17206(a)-(b) (injunctions and restitution; private injury in fact and loss of money or property as a result; public civil penalties up to \$2,500 per violation); N.Y. Gen. Bus. Law §§ 349(b)(1), (d), (h), 350-d(a) (public unfair, deceptive and abusive authority; specified federal-compliance defense; deceptive-act private action and its damages terms; public civil penalty up to \$5,000); Mass. Gen. Laws ch. 93A, §§ 4, 9 (public enforcement and knowledge-linked civil penalties; private injury, demand and damages provisions). See Carter, *supra* note 157, at 2, 28 (dated survey, rather than current nationwide remedial parity).

⁴¹⁵ 15 U.S.C. §§ 45(a)(2), 57b(e); 12 U.S.C. § 5552 (specified state federal-law enforcement route with limits, not a general state-UDAP grant or universal saving clause). See *infra* note 428 (specific statutory limits).

⁴¹⁶ Press Release, Fed. Trade Comm'n, *FTC Announces Crackdown on Deceptive AI Claims and Schemes* (Sept. 25, 2024), <https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes> (describing DoNotPay, Rytr, Ascend Ecom, Ecommerce Empire Builders and FBA Machine). The release describes allegations and proceedings on that date, not present dispositions or state-law holdings.

⁴¹⁷ 15 U.S.C. § 45(a), (n); Order Reopening and Setting Aside Order, *In re Rytr LLC*, No. C-4806 (F.T.C. Dec. 22, 2025) (describing and rejecting the two counts as pleaded); Decision and Order § I, *In re DoNotPay, Inc.*, No. C-4812 (F.T.C. Jan. 2025) (respondent-specific substantiation terms). See *infra* notes 418-419.

⁴¹⁸ Decision and Order, *In re DoNotPay, Inc.*, No. C-4812 (F.T.C. Jan. 2025); Decision and Order, *In re Rytr LLC*, No. C-4806 (F.T.C. Dec. 16, 2024), subsequently set aside, *infra* note 419. These are consent proceedings without a merits admission. Assurance of Voluntary Compliance ¶¶ 16-19, 24, 29, *In re Pieces Technologies* (Tex. Att'y Gen. Aug. 21, 2024) (specified accuracy disclosures, no new private right and other-law qualification); Cal. Att'y Gen., *Application of Existing California Laws to Artificial Intelligence and Application of Existing California Law to Artificial Intelligence in Healthcare* (Jan. 13, 2025) (noncomprehensive advisories). These examples are not an enforcement census.

Commission entered a final order in December 2024 and set it aside in December 2025. The later order found the complaint insufficient under the asserted deception and unfairness theories; its reasoning did not repeal Section 5 or abolish those theories in every application.⁴¹⁹ State statutes continued to supply their own authority.⁴²⁰ That does not make federal posture irrelevant to them: a state may direct courts to federal interpretations or provide a federal-compliance defense. The narrower point is that a change in federal enforcement does not itself enlarge or extinguish the state's statutory grant. *AMG Capital Management* similarly limits the FTC's § 13(b) monetary-relief route, without deciding remedies under the state statutes. Relative importance is a reason to examine those routes, not a substitute for their elements.⁴²¹

FTC v. Wyndham supplies a useful legal analogy. The Third Circuit upheld the FTC's unfairness authority over the challenged data-security conduct and rejected the company's fair-notice objection.⁴²² Its account of Section 5's flexible content explains how an existing prohibition can apply to a new technological practice.⁴²³ The substantial-injury, reasonable-avoidability and countervailing-benefits criteria constrain the federal unfairness inquiry.⁴²⁴ State text must then do the work. *Wyndham* does not make California's UCL, New York's § 349 and Massachusetts's chapter 93A identical, or establish that every undesirable AI practice is unfair under each.

Process representations provide a concrete application. A company promises human review of an adverse AI-mediated decision but supplies none. Another advertises completed discrimination testing that it never conducted. In each case the representation can be compared with the service actually provided. That is a possible deception claim, subject to the governing statute's materiality, coverage, causation and enforcement requirements.⁴²⁵ Missing audit records may

⁴¹⁹ Order Reopening and Setting Aside Order, *In re Rytr LLC*, No. C-4806 (F.T.C. Dec. 22, 2025), https://www.ftc.gov/system/files/ftc_gov/pdf/Rytr-Order.pdf (2-0 vote; pleading analysis of means and instrumentalities and unfairness, together with the stated public-interest and AI-policy grounds); Decision and Order, *id.* (Dec. 16, 2024) (recording the original dissents). The set-aside concerns that complaint and order, rather than repeal of § 5 or a decision about state law.

⁴²⁰ See *supra* notes 157, 414 (sampled statutory texts). Massachusetts's express federal-interpretation provisions and New York's federal-compliance defense illustrate why independent state enactment does not make every federal development legally irrelevant.

⁴²¹ *AMG Capital Mgmt., LLC v. FTC*, 593 U.S. 67 (2021) (unanimous Court; § 13(b) does not authorize the equitable monetary relief sought); 15 U.S.C. §§ 45(l), 57b (separate remedial routes and conditions); *supra* note 414 (state remedies). A change in the relative usefulness of state routes is the Article's inference, not a measured transfer of enforcement activity.

⁴²² *FTC v. Wyndham Worldwide Corp.*, 799 F.3d 236, 247, 255, 259 (3d Cir. 2015) (Ambro, J., for a unanimous panel; interlocutory review of the denial of a motion to dismiss). The decision permits the federal unfairness case to proceed; it does not adjudicate the company's liability or a state AI-process claim.

⁴²³ *Id.* at 244 (quoting earlier decisions concerning the flexibility of the federal standard). The source is *Wyndham's* account of those authorities, rather than a separate holding that all state UDAP provisions have identical scope.

⁴²⁴ 15 U.S.C. § 45(n) (codifying the substantial-consumer-injury-not-reasonably-avoidable-and-not-outweighed-by-countervailing-benefits standard for unfair acts or practices); *Wyndham*, 799 F.3d at 255, 259 ("Wyndham is only entitled to notice of the meaning of the statute and not to the agency's interpretation of the statute. . . . In sum, we have little trouble rejecting Wyndham's fair notice claim.").

⁴²⁵ Complaint ¶¶ 13-16 and Decision and Order §§ I-IV, *In re GeoCities*, 127 F.T.C. 94 (1999) (alleged gap between privacy representations and practices; consent disposition without a merits admission). This is an analogy from federal charging practice. Whether a false representation about an AI process is actionable depends on the applicable state elements and enforcement route; the examples in the text are proposed applications, not reported holdings.

be evidence that a promised audit did not occur; they do not necessarily prove that it did not. I recommend preserving the record needed to test the representation and designing any new process duty expressly. Existing UDAP liability still requires the applicable statutory elements.

FCRA's reasonable-procedures duty supplies a limited analogy for the form of an obligation. It also shows why process and redress must be separated: the cited private-claim authorities require more than a bare procedural lapse, including the applicable inaccuracy and concrete-injury showings.⁴²⁶ The DoNotPay order likewise requires that respondent to support specified performance claims with evidence. Extending that reasoning to a false promise of assessment or review is this Article's application, to be tested under the particular state law.⁴²⁷

Preemption remains part of that test. Traditional consumer protection does not supply immunity; FCRA expressly displaces specified state requirements, and other federal statutes preserve, channel or constrain state enforcement in different ways.⁴²⁸ A later AI statute could alter the allocation again. The White House's March 2026 recommendations favor existing regulators and would preserve specified general state laws, but those are legislative recommendations. My claim about the cost of displacing established state enforcement is institutional and political. A valid preemption clause can operate without a showing about that cost. Equivalent relief through the remaining remedies is a separate question.⁴²⁹

B. Civil Rights: Substantive Protection and the Evidence to Enforce It

Civil rights law gives the process duties a substantive purpose: the record should help identify and remedy unlawful treatment of people. Title VII's deferral provision and § 708 preserve a role for state enforcement within their terms.⁴³⁰ State and local agencies, attorneys general where authorized, and private litigants act through their respective legal grants; a federal des-

⁴²⁶ Fair Credit Reporting Act, 15 U.S.C. § 1681e(b) (consumer-reporting agencies must “follow reasonable procedures to assure maximum possible accuracy” of consumer reports); *see supra* notes 402, 406 (FCRA framework). On the enforcement conditions of the private claim, *see TransUnion LLC v. Ramirez*, 594 U.S. 413 (2021) (slip op. at 19) (class members whose inaccurate internal credit files were never furnished to a third party lacked Article III standing on the reasonable-procedures claim: “The mere presence of an inaccuracy in an internal credit file, if it is not disclosed to a third party, causes no concrete harm.”); *see also Dennis v. BEH-1, LLC*, 520 F.3d 1066, 1069 (9th Cir. 2008) (amended opinion) (requiring “the prima facie showing of inaccuracy required by sections 1681e and 1681i,” and adding that the agency is not strictly liable because it “remains free to argue” at trial that it followed reasonable procedures, *id.* at 1069-70); *see also infra* note 562 (*TransUnion* and the state-court alternative).

⁴²⁷ Decision and Order § I, *In re DoNotPay, Inc.*, No. C-4812 (F.T.C. Jan. 2025) (specified evidence requirements for the respondent's claims); *supra* note 418. The order binds its respondent on its terms. It does not create a universal state-law standard for every process representation.

⁴²⁸ 15 U.S.C. § 57b(e) preserves other remedies and rights of action in its terms; that general clause must be read with any applicable specific displacement provision. *See supra* note 112 (*Morales*). Section 45f(g) displaces conflicting state online-marketplace requirements; §§ 45b(e)(4) and 45c(c)(4) channel competing actions under those federal sections rather than all independent state causes. For FCRA's specified subjects, *see* 15 U.S.C. § 1681t; *Premium Mortgage*, 583 F.3d at 106-07; *supra* notes 135-136.

⁴²⁹ White House, *National Policy Framework*, *supra* note 142, at 3-4 (existing sectoral bodies and recommended preservation of specified general state laws). The Article's argument about political and institutional cost does not establish a legal prerequisite to preemption or measure available replacement enforcement capacity.

⁴³⁰ *See supra* notes 159, 257 (Title VII concurrent-enforcement framework at 42 U.S.C. §§ 2000e to 2000e-17, with state-deferral at § 2000e-5(c) and non-displacement at § 2000e-7).

ignation does not make all state statutes or remedies equivalent.⁴³¹ The use of AI does not replace the underlying employment-law inquiry. The relevant defendant, practice, causation and justification still have to be established.⁴³²

The history of disparate-impact law explains part of this institutional resilience. *Griggs* and *Albemarle* address practices that exclude workers without the required relation to the job.⁴³³ Today § 2000e-2(k) sets out the burden framework: the claimant must demonstrate causation by the challenged practice, and the statute specifies the respondent's job-relatedness and business-necessity burden and the alternative-practice route.⁴³⁴ Congress's 1991 response to *Wards Cove* shows a legislative commitment to that framework through a substantive change in the law.⁴³⁵ Housing has a related but distinct framework under *Inclusive Communities*, including its robust-causality and valid-interest safeguards.⁴³⁶ Documentation may help test those elements. It does not establish them simply by revealing a statistical disparity.

Federal executive action has put this distinction under pressure. Executive Order 14281 directs retreat from disparate-impact enforcement and asks the Attorney General to determine whether federal authorities preempt state disparate-impact rules. The instruction to examine preemption does not itself supply the displacing law. Published agency actions include changes

⁴³¹ See, e.g., Mass. Gen. Laws ch. 151B, § 4 (specified state substantive prohibitions). FEP Agency Designation Procedures: Revising Location of FEP Agency Lists, 91 Fed. Reg. 43,539, 43,540-41 (July 16, 2026) (administrative relocation of designated and certified agency lists to the EEOC website); EEOC, *Fair Employment Practices Agencies (FEPAs) and Dual Filing*, <https://www.eeoc.gov/fair-employment-practices-agencies-fepas-and-dual-filing> (directing readers to field-office information). Designation and certification concern their own federal criteria, rather than equivalent grants of every state actor's power or remedy.

⁴³² See *supra* notes 160-162; 42 U.S.C. § 2000e-2(k). The cited doctrine and statute govern covered practices; application to a particular AI-assisted practice requires its own facts. Historical agency technical assistance is not a substitute for that showing.

⁴³³ See *supra* notes 160, 161, 162 (Part II.B.4 first-use full citations for *Griggs v. Duke Power Co.*, 401 U.S. 424 (1971); *Albemarle Paper Co. v. Moody*, 422 U.S. 405 (1975); and Civil Rights Act of 1991 codifying the *Griggs-Albemarle* burden structure at 42 U.S.C. § 2000e-2(k)(1)).

⁴³⁴ *Griggs*, 401 U.S. at 431 ("The Act proscribes not only overt discrimination but also practices that are fair in form, but discriminatory in operation. The touchstone is business necessity. If an employment practice which operates to exclude Negroes cannot be shown to be related to job performance, the practice is prohibited."); *Albemarle Paper Co.*, 422 U.S. at 426 ("[A]ny tests used must measure the person for the job and not the person in the abstract." (quoting *Griggs*, 401 U.S. at 436)). The operative standard is now statutory. Under 42 U.S.C. § 2000e-2(k)(1)(A)(i) a practice is unlawful where the complaining party demonstrates that it causes a disparate impact and the respondent "fails to demonstrate that the challenged practice is job related for the position in question and consistent with business necessity." The codification was not a simple return to pre-1989 law: § 2000e-2(k)(1)(B)(i) keeps the requirement that a complaining party identify each particular challenged practice, subject to an exception where the elements of a decisionmaking process are "not capable of separation for analysis," and § 2000e-2(k)(1)(C) fixes the alternative-employment-practice concept to "the law as it existed on June 4, 1989," the day before *Wards Cove* was decided. See *infra* note 435.

⁴³⁵ See *supra* note 162 (Civil Rights Act of 1991, Pub. L. No. 102-166, § 105(a), 105 Stat. 1071 (codified at 42 U.S.C. § 2000e-2(k)(1))). Section 2(2) of the Act named the Court's *Wards Cove* decision and stated that the decision had "weakened the scope and effectiveness of Federal civil rights protections"; § 3(2) stated the Act's purpose as codification of "the concepts of 'business necessity' and 'job related' enunciated by the Supreme Court in *Griggs v. Duke Power Co.* . . . and in the other Supreme Court decisions prior to *Wards Cove Packing Co. v. Atonio*."

⁴³⁶ *Tex. Dep't of Hous. & Cmty. Affairs v. Inclusive Cmty. Project, Inc.*, 576 U.S. 519, 539, 542-43 (2015) (recognizing FHA disparate-impact claims with causality, valid-interest and constitutional safeguards). The evidentiary value of the Article's proposed documentation is an application, not an outcome established by the case.

to particular Title VI regulations and the Energy Department's delay of an earlier direct final rule.⁴³⁷ The June 2026 OLC opinions state executive-branch positions concerning Title VII and ECOA. The EEOC then removed its affirmative-action guidelines in July. Those documents must be read with the enacted statute and the particular regulations: an OLC opinion is not a judicial holding or an amendment to § 2000e-2(k), and rescission of part 1608 does not establish removal of part 1607.⁴³⁸ The state-law provisions need their own constitutional and preemption analysis; their continued existence alone cannot settle either.

Housing and credit show why enforcement powers require separate mapping. The FHA accommodates state law within § 3615 and has federal, state and private enforcement routes.⁴³⁹ ECOA preserves qualifying state credit-discrimination protection, but a state attorney general's asserted federal enforcement route runs through the Consumer Financial Protection Act, with national-bank and notice limitations. It does not follow merely from ECOA's substantive prohibition.⁴⁴⁰ The dated fair-housing survey and HUD's substantially-equivalent-agency materials identify parts of an institutional landscape; they do not establish a current fifty-state inventory of equivalent credit or housing remedies.⁴⁴¹

For AI-assisted tenant screening or lending, the applicable prohibition and proof of the discriminatory practice remain essential.⁴⁴² Guidance withdrawn or removed from an agency's active

⁴³⁷ Exec. Order No. 14281, *Restoring Equality of Opportunity and Meritocracy*, §§ 2, 4-5, 7, 9, 90 Fed. Reg. 17,537 (Apr. 28, 2025) (signed Apr. 23). Dated regulatory examples include 90 Fed. Reg. 57,141 (DOJ, Dec. 10, 2025); 91 Fed. Reg. 20,326 (Commerce, Apr. 16, 2026); 91 Fed. Reg. 30,239 (Interior, May 22, 2026); 91 Fed. Reg. 46,733 (Education, July 24, 2026); and 91 Fed. Reg. 48,766 (Treasury, Aug. 3, 2026). The Energy notice at 90 Fed. Reg. 43,539 (Sept. 10, 2025) delays an earlier direct final rule; it does not itself establish an operative rescission.

⁴³⁸ *Constitutionality of Disparate-Impact Liability Under Title VII*, 50 Op. O.L.C. __ (June 9, 2026) (memorandum for EEOC Chair); *Whether the Equal Credit Opportunity Act Creates Disparate-Impact Liability*, 50 Op. O.L.C. __ (June 12, 2026) (memorandum for FTC General Counsel). Rescission of Guidelines on Affirmative Action Appropriate Under Title VII, 91 Fed. Reg. 40,879 (July 6, 2026) (removing 29 C.F.R. pt. 1608). Section 1607.4 remains in the eCFR text as of October 14, 2026. The OLC interpretations do not amend the statute; the rescission of part 1608 concerns a separate part.

⁴³⁹ 42 U.S.C. §§ 3613-3615 (private and Attorney General routes, their conditions and relief; preserving state rights within § 3615's terms); *see also* Mass. Gen. Laws ch. 151B, § 4 (state housing prohibitions); HUD, *Fair Housing Assistance Program*, <https://www.hud.gov/stat/fheo/assistance-program> (state/local partnership and certification). The federal Attorney General is distinct from a state attorney general; state authority requires its own law.

⁴⁴⁰ 15 U.S.C. § 1691d(f) (state-law preservation; greater protection not inconsistent), § 1691e(g)-(h) (federal referrals and United States Attorney General authority, including the reason-to-believe pattern-or-practice route); 12 U.S.C. §§ 5481(12)(D), (14), 5536(a)(1)(A), 5552(a)-(b). The state CFPA route concerns covered persons and the title's requirements; § 5552(a)(2) limits suits against national banks and federal savings associations except the specified Bureau-regulation route. Prior notice is subject to the impracticability exception and immediate postfiling notice in § 5552(b)(1)(B). These provisions do not confer identical state remedies or resolve ECOA disparate-impact interpretation. *See supra* note 438.

⁴⁴¹ Temple University Policy Surveillance Program, *State Fair Housing Protections*, <https://legacy.lawatlas.org/datasets/state-fair-housing-protections-1498143743> (updated through Aug. 1, 2019; describing the dataset's historical coverage); HUD, *Fair Housing Partners Agencies*, <https://www.hud.gov/stat/fheo/partners-agencies> (agency list examined Oct. 16, 2026); HUD, *Fair Housing Assistance Program*, *supra* note 439 (law, operation and local-resource requirements). The dataset is not a present statute census, and the agency list is not a finding of equivalent remedies for every AI claim.

⁴⁴² *See supra* notes 273-274, 403, 406 (housing and credit applications). The applications must be distinguished from the particular laws and claims actually pleaded; *see infra* note 443. Neither an algorithm's use nor a statistical disparity alone establishes every cause of action.

site cannot be presented as current policy. A proposed HUD rescission likewise has to be distinguished from a final amendment, and a supplemental notice must be read for the subjects it actually reopens. The federal RealPage amended complaint is an antitrust pleading with state claims, not a fair-housing holding about algorithmic pricing.⁴⁴³ These examples show available legal routes and changing federal positions. A particular housing claim still needs its statutory elements and supporting facts.

Connecticut v. Teal gives the record a specific job. An employer could not defeat the challenge to a discriminatory screening examination by pointing to a favorable racial balance in the final promotions.⁴⁴⁴ Title VII protects the individual excluded by the challenged practice; favorable treatment of other group members does not erase that injury.⁴⁴⁵ An AI-assisted hiring process may likewise include a screen, ranking and later selection. The record should permit the relevant practice to be identified and examined. But *Teal* does not require four-fifths testing of every component in every pipeline. The Uniform Guidelines ordinarily do not expect component evaluation when the total process has no adverse impact, subject to their exceptions, and the four-fifths measure has statistical qualifications.⁴⁴⁶ A favorable aggregate result is not a legal defense to an otherwise established violation; that is different from a universal testing command.

I recommend audit records detailed enough to examine a materially challenged selection stage and distinguish its effect from later decisions. That is a design recommendation informed by *Teal*, not a claim that the case mandates the Article's entire audit program.⁴⁴⁷ Access, data quality and the applicable causation rules will determine how much the record can establish.

Ricci v. DeStefano addresses what an employer does in response to perceived disparate impact.⁴⁴⁸ The employer's race-based rejection of completed examination results required an objective strong basis in evidence of disparate-impact liability; good-faith concern or disparity alone did not suffice. The Court also distinguished considering fair test design before administering a practice.⁴⁴⁹ An audit finding therefore needs an identified response. Class-neutral redesign, further investigation, discarding results because of race and adjusting individual out-

⁴⁴³ 24 C.F.R. § 100.500 (eCFR text as of Oct. 14, 2026); HUD's Implementation of the FHA's Disparate Impact Standard, 91 Fed. Reg. 1,475 (proposed Jan. 14, 2026); supplemental notice, 91 Fed. Reg. 51,416 (Aug. 10, 2026) (added Title VI subjects; reopening comments on those subjects). These proposals are not final repeal. Amended Complaint, *United States v. RealPage, Inc.* (M.D.N.C. Jan. 2025) (claims section: four Sherman Act and nine state claims; no FHA claim). This federal action is distinct from the later New York § 340-b litigation. The statutory and regulatory framework remains distinct from historical HUD guidance.

⁴⁴⁴ *Connecticut v. Teal*, 457 U.S. 440, 455 (1982) ("Title VII does not permit the victim of a facially discriminatory policy to be told that he has not been wronged because other persons of his or her race or sex were hired.").

⁴⁴⁵ *Id.* at 453-54 ("The principal focus of the [Title VII] statute is the protection of the individual employee, rather than the protection of the minority group as a whole."); *id.* at 455.

⁴⁴⁶ *Teal*, 457 U.S. at 450, 453 n.12, 455; 29 C.F.R. § 1607.4(C)-(D); 42 U.S.C. § 2000e-2(k)(1)(B)(i). Section 1607.4(C) ordinarily does not expect evaluation of each component where the total selection process has no adverse impact, with stated exceptions; subsection (D) qualifies the four-fifths indicator. The statute requires identification of particular practices, with its inseparability exception. None establishes universal component-by-component four-fifths compliance.

⁴⁴⁷ See *supra* notes 378-381, 387-390 (proposed auditability and monitoring); *Teal*, 457 U.S. at 453-55. Records permitting analysis of a materially challenged barrier are this Article's design recommendation. Their scope is not a universal testing mandate derived from *Teal*.

⁴⁴⁸ See *supra* note 271 (*Ricci v. DeStefano*, 557 U.S. 557, 585 (2009), articulating the strong-basis-in-evidence standard for employer voluntary compliance with disparate-impact requirements).

⁴⁴⁹ *Ricci*, 557 U.S. at 585 ("[U]nder Title VII, before an employer can engage in intentional discrimination

comes by protected class cannot be treated as one corrective act. *Ames* removes a heightened prima facie requirement for majority-group Title VII plaintiffs; it does not decide the legality of an AI audit or replace *Ricci*'s inquiry.⁴⁵⁰

The design should require documentation of the evidence and decision chain. A promise of litigation immunity for good-faith action would go further than the law supports. Part V.C proposes a state-law audit presumption. That proposal needs state-law justification and compatibility with federal law; a state cannot write a defense into Title VII or authorize conduct it forbids. Qualifying records may supply evidence relevant to the federal inquiry, which is different from satisfying that inquiry by statutory declaration.⁴⁵¹ The absence of an audit does not itself establish disparate-impact liability, and conducting one does not authorize protected-class treatment. The substantive safeguard remains the applicable law governing the practice and the response.

The federal statutes also supply different rules for coexistence. Title VII, the FHA and ECOA contain their own preservation provisions. Title II has a specified state-notice procedure, while the Civil Rights Act's § 1104 supplies an Act-wide rule against field occupation and displacement of consistent state law.⁴⁵² Title VII's own saving clause excludes state law purporting to require or permit an unlawful employment practice.⁴⁵³ These texts support concurrent state protection. Their application requires analysis of the particular federal conflict, with no com-

for the asserted purpose of avoiding or remedying an unintentional disparate impact, the employer must have a strong basis in evidence to believe it will be subject to disparate-impact liability if it fails to take the race-conscious, discriminatory action.”); *id.* at 581 (“Allowing employers to violate the disparate-treatment prohibition based on a mere good-faith fear of disparate-impact liability would encourage race-based action at the slightest hint of disparate impact.”); *id.* at 587 (the City “could be liable for disparate-impact discrimination only if the examinations were not job related and consistent with business necessity, or if there existed an equally valid, less-discriminatory alternative that served the City’s needs but that the City refused to adopt,” citing 42 U.S.C. §§ 2000e-2(k)(1)(A), (C)); *id.* at 585 (Title VII “does not prohibit an employer from considering, before administering a test or practice, how to design that test or practice in order to provide a fair opportunity for all individuals, regardless of their race”); *id.* at 579-80 (the City “made its employment decision because of race” and “rejected the test results solely because the higher scoring candidates were white”); *id.* at 584 (adopting the standard “as a matter of statutory construction to resolve any conflict between the disparate-treatment and disparate-impact provisions of Title VII”); *id.* at 587 (deciding the two points “based on the record developed by the parties through discovery”); *id.* at 585 (finding no support for the conclusion that the respondents “had an objective, strong basis in evidence to find the tests inadequate”).

⁴⁵⁰ *Ames v. Ohio Dep’t of Youth Servs.*, 605 U.S. 303 (2025), slip op. at 4-5 & n.2, 8-9 (Jackson, J., for a unanimous Court; Thomas, J., concurring, joined by Gorsuch, J.) (rejecting the Sixth Circuit’s heightened background-circumstances rule; assuming without deciding application of *McDonnell Douglas* at summary judgment). The holding concerns disparate treatment’s prima facie framework, not AI-audit legality or preemption.

⁴⁵¹ *Ricci*, 557 U.S. at 581, 584-87; 42 U.S.C. §§ 2000e-2(k), 2000e-7; *infra* Part V.C (the proposed state-law presumption). Audit findings and a documented decision chain may be evidence concerning the federal standard; procedural compliance and good faith do not by themselves supply an objective strong basis. A state-law presumption cannot authorize an unlawful federal employment practice or alter the federal standard.

⁴⁵² 42 U.S.C. §§ 2000e-5(c), 2000e-7 (Title VII), 3615 (FHA), 2000a-3(c) (Title II notice procedure), 2000h-4 (Civil Rights Act § 1104, applying across titles); 15 U.S.C. § 1691d(f) (ECOA). Section 1104 disclaims general field occupation and preserves consistent state laws; absence of a Title VI-specific deferral provision does not erase that Act-wide text. Each provision retains its own conditions.

⁴⁵³ 42 U.S.C. § 2000e-5(c) (qualifying-state deferral, ordinarily sixty days unless proceedings terminate sooner, with the initial-year qualification); § 2000e-7 (preserving state liability except state law purporting to require or permit a practice unlawful under Title VII).

mon rank for the authority as a whole.⁴⁵⁴

ERISA and arbitration make the limits practical. *Pilot Life* did involve claims processing: it displaced the state-law benefits claim before it. *FMC* applies ERISA's deemer clause to a self-funded plan. *Shaw* directly addresses partial preemption of a state fair-employment law as applied to plans, preserving it only to the extent specified in the opinion's federal-law analysis.⁴⁵⁵ A proposed AI claims-handling duty needs that analysis, not an assertion that general civil rights law is always saved. The FAA can also constrain a state rule obstructing enforcement of covered arbitration agreements. *Concepcion* concerns that state-law conflict; *Epic Systems* resolves the asserted conflict with the NLRA. Neither establishes categorical abolition of substantive state protection or every public enforcement route.⁴⁵⁶

There remains a serious objection to the portfolio: good records can document harm without providing relief. A process-compliant employer may still discriminate; a person may lack an actionable substantive claim, usable evidence or access to an effective forum.⁴⁵⁷ The Article's answer is to pair process duties with the substantive law and remedy that actually cover the decision. It cannot claim that the pairing cures the gap. Causation, proof, resources and remedial limits remain; *Inclusive Communities*, *Comcast* and *Students for Fair Admissions* address particular legal constraints and supply no empirical evaluation of this approach.⁴⁵⁸ Additional substantive protections, such as specified human-decision rights or liability rules, require their own policy choice and legal analysis. The present institutional map does not adopt them by implication.⁴⁵⁹ Sectoral licensing adds powers of institutional supervision and professional discipline.

⁴⁵⁴ See *supra* Sections II.B.1-2, II.B.4; 42 U.S.C. §§ 2000e-5(c), 2000e-7. The cost of altering an established enforcement allocation is the Article's institutional argument, not an architecture-wide risk score or constitutional entrenchment of state law.

⁴⁵⁵ 29 U.S.C. §§ 1132, 1144; *Pilot Life Ins. Co. v. Dedeaux*, 481 U.S. 41, 43-44, 48, 52-57 (1987) (improper processing of a claim for benefits and ERISA's exclusive enforcement structure); *FMC Corp. v. Holliday*, 498 U.S. 52, 58-65 (1990) (insurance/deemer rules and self-funded plan); *Shaw*, 463 U.S. at 102-09 (partial preemption of the state fair-employment law, including the analysis at 105, 108). Applying these rules to a specified AI claims-handling duty remains an application, not an unexamined general-law exception.

⁴⁵⁶ *AT&T Mobility LLC v. Concepcion*, 563 U.S. 333, 339-44, 345-52 (2011) (FAA and the state class-waiver rule; § 2's generally applicable defenses remain bounded by the opinion); *Epic Sys. Corp. v. Lewis*, 584 U.S. 497, 503-19, 524-25 (2018) (enforcement of individualized arbitration agreements against the asserted NLRA conflict, rather than a separate state-law preemption holding). The cases do not decide every independent public enforcement route.

⁴⁵⁷ The objection is this Article's: production of a useful process record and availability of substantive relief are separate questions. See *supra* notes 434, 440, 446, 451 (statutory elements, authority and corrective-action limits).

⁴⁵⁸ *Inclusive Communities*, 576 U.S. at 539, 542-43 (FHA causality and safeguards); *Comcast Corp. v. Nat'l Ass'n of African American-Owned Media*, 589 U.S. 327 (2020) (§ 1981 but-for causation); *Students for Fair Admissions, Inc. v. President & Fellows of Harv. Coll.*, 600 U.S. 181, 230 (2023) (challenged Harvard and UNC admissions programs). These decisions do not establish the frequency of trade-secret objections, state enforcement capacity or the effectiveness of AI audit duties.

⁴⁵⁹ The substantive-backstops question is separate from the evidence-producing function developed here. A specified liability presumption, categorical human-decision right or proxy restriction would need its own coverage, justification and remedy. They are alternatives for further policy and legal work, not additions to the operative worked statute by implication.

C. Sectoral Licensing: Authority over the Institution and Its Use

1. Banking and Financial Services

Banking supplies a concrete model for governing an institution's use of a model it may not have built. The federal framework introduced in 2011 through SR 11-7 and the parallel OCC bulletin addressed model development, validation, use and governance inside supervised banks.⁴⁶⁰ Its treatment of purchased models is especially useful. A bank could rely on a vendor's expertise, but it remained responsible for validating its own use of the product. Proprietary code could limit access; sensitivity analysis, benchmarking and monitoring of the bank's own outcomes supplied ways to examine the product despite that limit.⁴⁶¹ The FDIC adopted the framework in 2017, with changes reflecting its supervisory population and a qualification for institutions below \$1 billion in assets.⁴⁶² This is an institutional analogue for the Article's documentation and monitoring duties. The analogy supplies no grant of state power or evidence that existing staff can implement every proposed duty.⁴⁶³

The three agencies replaced that framework with joint guidance in April 2026. The revision expressly disclaims enforceable standards and prescriptive requirements. The agencies expect it to be most relevant to banking organizations above \$30 billion in assets, while potentially relevant to smaller organizations with significant model risk.⁴⁶⁴ It also excludes generative and agentic AI from its scope, while directing organizations to use their risk management and governance practices to determine controls for tools outside the document. Traditional statistical models and non-generative, non-agentic AI remain within the stated principles.⁴⁶⁵ The exclusion establishes a boundary in this guidance. It does not establish that the supervisors cannot govern frontier AI, that other banking law stops applying, or that an institution may leave those tools uncontrolled.

The lesson I draw is about responsibility. A bank's obligation to govern a lending decision need not turn entirely on the technical family of the model used to make it. A state may likewise define an assessment around the decision, its risks and the institution's ability to examine

⁴⁶⁰ See *supra* note 165; FDIC, FIL-22-2017, *Adoption of Supervisory Guidance on Model Risk Management* (June 7, 2017); OCC Bulletin 2026-13, *Model Risk Management: Revised Guidance* (Apr. 17, 2026) (background identifying the original joint 2011 guidance and FDIC's later adoption).

⁴⁶¹ FDIC, *Supervisory Guidance on Model Risk Management* 1-2, 7-8, 13-14 (2017 version accompanying FIL-22-2017), <https://www.fdic.gov/news/news/financial/2017/fil17022a.pdf> (validation, independence, vendor products and institution-specific use); *Revised Guidance on Model Risk Management*, SR 26-2, at 5-12 (Apr. 17, 2026) (current model-use, validation, governance and third-party principles).

⁴⁶² OCC Bulletin 2026-13, *supra* note 460 (background and rescission of OCC Bulletin 2011-12); FDIC, FIL-22-2017 (adoption with technical changes and the stated qualification for institutions below \$1 billion). The 2017 adoption is historical; the agencies issued revised joint guidance in 2026.

⁴⁶³ Conference of State Bank Supervisors, *State Financial Regulation 101* (July 1, 2023), <https://www.csbs.org/state-financial-regulation-101>; CSBS, *Networked Supervision*, <https://www.csbs.org/net-worked-supervision> (describing coordination, shared platforms and examinations). These are the organization's descriptions of its arrangements, not a grant of state power or independent outcome evaluation. The proposed analogy concerns institutional responsibility.

⁴⁶⁴ SR 26-2, *supra* note 461, at 2-3; OCC Bulletin 2026-13, *supra* note 460. The guidance disclaims enforceable standards and supervisory criticism for noncompliance with the guidance; it does not prevent action for violations of law or unsafe or unsound practices. SR 26-2, at 2 n.1.

⁴⁶⁵ SR 26-2, *supra* note 461, at 3 n.3 (generative/agentic scope exclusion, residual governance and continued principles for traditional and non-generative, non-agentic models). The exclusion is not a finding of supervisory incapacity or absence of other applicable law.

them. That choice can reduce dependence on a changing technological definition; it cannot eliminate the need to update standards, obtain information from vendors or fund supervision. The 2026 revision makes the analogy more precise, without proving the superiority of a state design.⁴⁶⁶

The state's legal route depends first on the charter. State-chartered banks operate within a dual system of state and federal supervision.⁴⁶⁷ Federal Reserve membership and deposit-insurance status help determine the federal supervisor; national banks fall under a different allocation of authority, with the OCC and the National Bank Act central to the inquiry.⁴⁶⁸ State examination and enforcement powers must then be identified in the relevant banking code. California authorizes examinations of covered banks and cease-and-desist proceedings for specified violations or unsafe practices. New York provides examination powers and orders addressing unauthorized or unsafe and unsound conduct. These are powers over identified institutions and conduct, exercised under their statutory procedures.⁴⁶⁹

For AI-assisted underwriting, fraud detection or account management, the practical inquiry is whether the proposed record or assessment concerns conduct the regulator may supervise. A deceptive account-service claim may also support a consumer-protection action under Section IV.A, if that statute covers the defendant and conduct. Neither route creates a general power over the institution's software vendor, and neither guarantees a customer's access to examination records or damages.⁴⁷⁰

National-bank cases require a second distinction: the validity of the substantive state duty and the means of enforcing it. *Watters v. Wachovia Bank* held that a national bank's mortgage business, including business conducted through its operating subsidiary, was outside the state licensing and visitorial regimes challenged there.⁴⁷¹ Congress subsequently changed the subsidiary rule. Section 25b preserves the applicability of state law to subsidiaries, affiliates and agents that are not themselves nationally chartered, and states that the National Bank Act does not occupy the field. For a state consumer financial law applied to a national bank, however, its discrimination, significant-interference and other-federal-law provisions still require analysis. *Cantero v. Bank of America* rejects a categorical answer based merely on federal authorization of the banking power; it requires a practical assessment of the state law's interference, informed

⁴⁶⁶ The responsibility-centered response is this Article's inference. See *supra* notes 165, 184-185; *infra* Part V. A technology-neutral duty still needs sufficient specification, vendor information and institutional resources; the supervisory guidance does not validate the proposed state regime's outcomes.

⁴⁶⁷ 12 U.S.C. §§ 1813, 1818; CSBS, *State Financial Regulation 101*, *supra* note 463 (dual-system description). The existence of this allocation is not immunity of state-chartered bank conduct from other federal rules.

⁴⁶⁸ 12 U.S.C. §§ 1813, 1818, 25b. The specific federal supervisory allocation and state consumer financial law inquiry must be distinguished; see *infra* notes 471-473.

⁴⁶⁹ Cal. Fin. Code §§ 500, 580 (examinations and specified cease-and-desist proceedings); N.Y. Banking Law §§ 36(1)-(3), 39(1)-(2) (examinations, inquiries and orders); 205 Ill. Comp. Stat. 5/48(2)(a) (examination example). These selected powers do not establish every regulator's jurisdiction, penalties or authority over all vendors. The New York display carries alternative versions contingent on the identified buy-now-pay-later implementation provisions; the banking-organization powers cited here appear in both.

⁴⁷⁰ See *supra* Section IV.A; N.Y. Banking Law § 36(10) (confidentiality of supervisory materials, with its statutory qualifications). Existing supervisory authority, public law enforcement, disclosure to an affected customer and a private claim require separate legal bases.

⁴⁷¹ *Watters v. Wachovia Bank, N.A.*, 550 U.S. 1, 7, 20-22 (2007) (Ginsburg, J., for five Justices; Thomas, J., took no part). For the later statutory subsidiary provisions and *Cantero*, see *infra* note 472. The operating-subsidiary result is a historical holding.

by the relevant precedents. On remand in May 2026, a divided Second Circuit again held the New York interest-on-escrow law preempted. That application illustrates why rejection of a categorical test does not establish the state duty's validity.⁴⁷²

Cuomo v. Clearing House Association addresses enforcement. A state attorney general's judicial action to enforce applicable, non-preempted law is not the supervisory visitation reserved by federal law. The Court nevertheless left the injunction against the Attorney General's own-authority investigative subpoena in place. Congress incorporated the judicial-enforcement distinction in § 25b(i)(1).⁴⁷³ A state AI claim against a national bank therefore needs both a duty that remains applicable and an authorized enforcement route. The institutional portfolio is useful precisely because it makes those two questions visible.

2. Insurance

Insurance has an additional statutory protection, but the law and its operation must satisfy its predicates. The defendant's insurer status is insufficient. McCarran-Ferguson provides that a federal statute not specifically relating to insurance shall not be construed to invalidate, impair or supersede state law enacted for the purpose of regulating the business of insurance.⁴⁷⁴ *Fabe* applies that inquiry within a single state priority statute: protection of policyholders and necessary administration qualified, while preferences for other creditors did not. *Humana* allowed federal RICO remedies alongside Nevada insurance law where they did not conflict with the law, frustrate declared state policy or interfere with the administrative regime. ERISA supplies a separate limit, including the deemer clause's treatment of self-funded plans.⁴⁷⁵ The result is substantial support for state insurance regulation, with predicates that must be established for the particular AI duty and federal provision.⁴⁷⁶

The NAIC Model Bulletin illustrates how an existing institution can organize its expectations around AI. Its adoption map dated October 8, 2026 lists twenty-seven states, the District of

⁴⁷² 12 U.S.C. § 25b(b)(1)-(4), (e), (h)(2) (state consumer financial laws, subsidiary applicability and no field occupation); *Cantero v. Bank of America, N.A.*, 602 U.S. 205 (2024), slip op. at 8-14 (unanimous Court; practical nature-and-degree assessment under the *Barnett Bank* standard, vacating and remanding rather than deciding the escrow law's ultimate validity). Section 25b(b)(1) also retains the discriminatory-effect and other-federal-law routes; the significant-interference provision is not its only condition. On remand, *Cantero v. Bank of America, N.A.*, Nos. 21-400, 21-403, op. at 1-3, 27 (2d Cir. May 5, 2026), held the New York law preempted, with Pérez, J., dissenting. Opinion and dissent reprinted in the Supreme Court petition appendix, No. 25-1313, App. 1a-3a, 27a-29a, https://www.supremecourt.gov/DocketPDF/25/25-1313/409701/20260522183319993_Cantero.Appendix.pdf.

⁴⁷³ *Cuomo v. Clearing House Ass'n, L.L.C.*, 557 U.S. 519, 535-36 (2009) (Scalia, J., for five Justices); 12 U.S.C. § 25b(i)(1). The Court distinguished judicial law enforcement from supervisory visitation and left the injunction against the Attorney General's own-authority subpoena in place. The codification authorizes enforcement of applicable law and relief authorized by that law; it does not itself make a substantive duty non-preempted.

⁴⁷⁴ 15 U.S.C. § 1012(a)-(b); *see supra* notes 163, 283. Section 1012(b)'s first clause supplies the particular reverse-preemption inquiry, not an exemption for every act of an insurance company.

⁴⁷⁵ *Fabe*, 508 U.S. at 500-10 (five-Justice majority; state priority provisions separately analyzed); *Humana*, 525 U.S. at 309-14 (unanimous Court; compatible RICO remedies in the Nevada setting); 29 U.S.C. § 1144(b)(2); *FMC*, 498 U.S. at 58-65; *Pilot Life*, 481 U.S. at 48, 52-57. These cases do not establish that all federal civil rights laws are always outside reverse preemption or that all insurance AI is insulated.

⁴⁷⁶ *See supra* note 163; *Fabe*, 508 U.S. at 508-10. The law's insurance-regulating purpose and the particular federal interference matter; insurer status alone does not establish them. *See supra* Section II.B.2 (ERISA boundaries).

Columbia and Puerto Rico under the bulletin, with California, Colorado, New York and Texas separately identified as having insurance-specific regulation or guidance. That is the map's reference list, not an independent determination that every jurisdiction has adopted every model element.⁴⁷⁷ The model calls for a written AI Systems Program proportionate to the insurer's activities and risks. It addresses governance, testing, third-party systems and information regulators may request under existing law. The bulletin provides a model for state action and permits alternative means of meeting applicable requirements. The legal rule and any sanction must come from applicable state authority.⁴⁷⁸

New York's July 2024 circular letter supplies a jurisdiction-specific example. It states supervisory expectations for AI systems and external consumer data used in underwriting and pricing, connects those expectations to existing insurance law, and retains the insurer's responsibility for tools obtained from vendors. It also addresses records, consumer inquiries and disclosure of reasons for adverse action. The distinction between existing legal requirements and supervisory expectations remains relevant even within that single document.⁴⁷⁹ Colorado uses a different route: legislation authorizes rules under the insurance code. Regulation 10-1-1, as amended effective October 15, 2025, reaches the specified life, automobile and health lines when insurers use external consumer data or algorithms and models using those data. It prescribes governance, inventories, assessments and reporting, with implementation dates that differ by line.⁴⁸⁰ New York's circular and Colorado's rule should not be treated as legally interchangeable merely because both govern an insurer's use of AI.

The examination tools are developing too. NAIC's published pilot plan named twelve participating states for work from March through September 2026. Its stated purposes include understanding insurers' practices, assessing governance and improving regulatory tools. The working group's later materials identify a renamed AI Risk Evaluation Supplement and a version 5.0 exposure with comments closing September 29.⁴⁸¹ These records show work on supervisory

⁴⁷⁷ NAIC, *Implementation of the Model Bulletin on the Use of Artificial Intelligence Systems by Insurers* (map dated Oct. 8, 2026), <https://content.naic.org/sites/default/files/legal-adoption-map-ai-model-bulletin.pdf> (reference list: twenty-seven states, D.C. and Puerto Rico; separate insurance-specific regulation/guidance entries for California, Colorado, New York and Texas). The map cautions against treating it as a determination that a jurisdiction has adopted all model elements or satisfies accreditation standards. The April 1 count cited *supra* notes 289-292 is a dated earlier snapshot.

⁴⁷⁸ NAIC, *Model Bulletin on the Use of Artificial Intelligence Systems by Insurers* 1-9 (adopted Dec. 4, 2023), <https://content.naic.org/sites/default/files/cmte-h-big-data-artificial-intelligence-wg-ai-model-bulletin.pdf.pdf>. The model identifies existing-law predicates, proportionate written programs, third-party and information-request expectations, and alternative means of compliance. State adoption and authority must be established separately; the model is not an independent national duty or sanction.

⁴⁷⁹ N.Y. Dep't of Fin. Servs., Insurance Circular Letter No. 7, §§ II-IV, ¶¶ 20, 25-30, 35-45 (July 11, 2024), <https://www.dfs.ny.gov/industry-guidance/circular-letters/cl2024-07>. The letter covers AI and external consumer data in underwriting and pricing, with its coverage exclusions, and identifies the underlying statutory and regulatory provisions. Contract information and audit terms are expected where appropriate and available; the insurer retains responsibility under applicable law.

⁴⁸⁰ Colorado SB 21-169, Colo. Rev. Stat. § 10-3-1104.9; 3 Colo. Code Regs. § 702-10, Regulation 10-1-1, §§ 3, 5-9 (amended effective Oct. 15, 2025). The regulation's specified life, private-passenger-auto and health scope concerns external consumer data and models or algorithms using those data. Its governance and reporting duties, dates, confidentiality and statutory enforcement provisions do not establish a generic duty for every insurer AI system. See *supra* notes 200, 220, 287-288.

⁴⁸¹ NAIC, *Pilot Project Summary* 1-3 (2026), https://content.naic.org/sites/default/files/inline-files/Pilot%20Project%20Summary_1.pdf (twelve named states; March-September plan; understanding, governance and tool-development purposes); NAIC Big Data and Artificial Intelligence Working Group page, <https://content.naic.org/committees/h/big-data-artificial-intelligence-wg> (materials examined Oct. 16, 2026: AI Risk Evaluation Supplement version 5.0 exposure, comments ending

implementation. Completed pilot findings, an adopted examination requirement and consumer-outcome evidence would be needed to draw those further conclusions.

This example establishes several practical supports for regulation. Insurance departments have identifiable legal powers, regulated institutions and examination processes through which a state can ask about AI use. Their legal instruments differ; information may be confidential; and institutional uptake does not establish effective enforcement. The model bulletin's use of existing law also differs from a legislature's creation of a new assessment duty. The latter needs its own authority, specification and remedy.⁴⁸² For the insurance application in Section V.E.3, the proposed duties concern documentation, assessment, monitoring and recourse tied to the insurer's decision, with the federal overlap and each state's available powers left in view.⁴⁸³

3. Healthcare

Clinical AI makes the allocation of responsibility unusually visible. A hospital chooses a system and sets its deployment policy; a clinician uses its output in treating a patient; a manufacturer may market a regulated device. Those actors answer to overlapping legal regimes, but their duties are not interchangeable.⁴⁸⁴ Medical-practice law supplies grounds for professional discipline. California identifies gross negligence, repeated negligent acts and incompetence; Illinois and New York likewise specify disciplinary grounds. California's corporate-practice provisions impose a further actor boundary, with statutory qualifications. Whether a particular clinical use violates the applicable professional standard requires the facts and the jurisdiction's law.⁴⁸⁵

Hospital law addresses the institution. California's provisions assign responsibilities to the governing body and medical staff for governance, competence and quality of care. New York specifies governing-body, medical-staff and quality-assurance duties. Illinois requires medical-staff bylaws, credentialing and clinical review. These authorities identify institutions that can examine how a hospital introduces and supervises clinical tools. An AI assessment schedule or a patient action for licensing violations would require further legal authority.⁴⁸⁶

Sept. 29). A published pilot plan and an exposure are not completed findings or final adoption.

⁴⁸² See *supra* notes 477-481. Adoption, supervisory expectations, an enforceable duty under identified state law, and measured consumer effects are distinct. The pilot plan is directed at regulatory understanding and examination development, not a demonstrated consumer-outcome experiment. The proposed legislative translation of management principles in Parts III.D and V needs its own legal authority and specification.

⁴⁸³ See *infra* Part V.E.3 (insurance underwriting application). The legal map supplies conditional supports and limits; it does not pre-approve every proposed provision.

⁴⁸⁴ See *supra* Section III.C.4; *supra* notes 293-306 (clinical actors and federal overlays). The actor distinction developed here does not rank healthcare against every other domain or establish a categorical federal-state boundary.

⁴⁸⁵ Cal. Bus. & Prof. Code §§ 2234(b)-(d), 2032, 2400 (discipline, natural-person licensing, and corporate-practice limits with stated qualifications); 225 Ill. Comp. Stat. 60/22(A)(4) (gross negligence; the displayed August 2026 versions and future January 2027 version retain this ground); N.Y. Educ. Law § 6530(3)-(6) (repeated negligence or incompetence and single gross acts). These selected provisions establish professional disciplinary grounds, not the result of any particular AI-malpractice dispute or a fifty-state corporate-practice rule.

⁴⁸⁶ Cal. Health & Safety Code §§ 1250(a), 1275; Cal. Code Regs. tit. 22, §§ 70701(a), 70703 (governing body and medical staff; quarterly LII reproduction examined Oct. 16, 2026); N.Y. Pub. Health Law § 2801-a(1)-(3); N.Y. Comp. Codes R. & Regs. tit. 10, §§ 405.2, 405.4(b)-(c), 405.6; 210 Ill. Comp. Stat. 85/10.4; Ill. Admin. Code tit. 77, § 250.310(b)(1), (4)-(6) (bylaws, privileges and clinical review). New York § 405.2(i) cautions that assure/ensure language does not expand governing-body damages

A patient harmed after an AI-assisted decision might have a tort claim, a complaint to a licensing authority, or a claim concerning deceptive representations. The available defendant, elements, injury requirements, access to records and relief differ across those routes. The payer setting is different again. California's SB 1120 regulates utilization review by covered plans and insurers; the insurance provision bars a software tool from denying, delaying or modifying services on medical-necessity grounds and reserves that determination to a competent licensed professional. Its legal route is payer regulation; hospital-licensing and medical-board powers remain separate.⁴⁸⁷

Federal device law must be tested against the actual state requirement. The FDA's applicable marketing pathway and device-specific conditions matter; permission to market a device does not itself resolve whether a clinician exercised reasonable care. Nor does that separation make every duty addressed to a clinician or hospital immune from preemption. Other federal regimes, including information-blocking, payer and emergency-care requirements, can enter the same decision environment.⁴⁸⁸

Riegel establishes that state common-law duties can be requirements within § 360k(a). A state assessment, warning or use condition must therefore be examined for what it requires with respect to the device, and whether it differs from or adds to the applicable federal requirement concerning safety, effectiveness or another covered matter. Naming the hospital as obligor does not answer that question. *Buckman* adds a distinct conflict inquiry: the fraud-on-the-FDA claims there depended on federal disclosure duties and interfered with the federal enforcement scheme.⁴⁸⁹ A state cannot avoid that problem simply by calling enforcement of a federal device obligation professional oversight.

Parallel state duties provide a possible route, within their own limits. *Lohr* and *Riegel* recognize that § 360k does not itself displace damages duties that parallel federal requirements; *Riegel* did not decide the particular parallel-claim argument, which had not been preserved. *Buckman* prevents treating every FDCA violation as a privately enforceable state claim. Section 396 separately preserves practitioner authority to use legally marketed devices within a legitimate treatment relationship while retaining specified federal restrictions.⁴⁹⁰ The proposed hospital duties must be situated within those provisions, requirement by requirement. Some will concern institutional selection, training or review without adding a device requirement; others may cross the line. The hospital application in Section V.E.1 must meet those boundaries, with professional discipline, institutional supervision and patient redress treated as distinct

liability. Facility categories, exemptions, enforcement and private remedies require their own law.

⁴⁸⁷ Cal. Ins. Code § 10123.135(a), (e), (j), as amended by SB 1120, 2024 Cal. Stat. ch. 879, § 2 (effective Jan. 1, 2025); see *supra* note 297 (parallel plan setting). Subdivision (j)(2) specifies the medical-necessity restriction and competent licensed-professional determination. This is payer regulation with its own coverage and enforcement, not a new medical-board or hospital remedy.

⁴⁸⁸ See *supra* notes 301-305 (the separate device, information-blocking, payer and emergency-care overlays). The governing marketing pathway and actual federal requirement must be identified for the device and duty.

⁴⁸⁹ 21 U.S.C. § 360k(a)-(b); *Riegel*, 552 U.S. at 321-25, 327-30; *Buckman*, 531 U.S. at 347-53; 21 U.S.C. § 337(a). *Riegel* concerns applicable federal requirements and additional or different state requirements within the clause; *Buckman* concerns the fraud-on-agency claims and federal enforcement scheme. Neither is avoided by the obligor's employer label.

⁴⁹⁰ *Lohr*, 518 U.S. at 492-502; *Riegel*, 552 U.S. at 330 (parallel-duty principle; particular argument not preserved); *Buckman*, 531 U.S. at 352-53 (not every federal violation supports a state claim); 21 U.S.C. § 396 (practitioner-use preservation, retaining specified federal restrictions). Institutional duties require their own analysis and are not presumed parallel merely because NIST or ISO principles informed them.

functions.⁴⁹¹

D. Public Procurement: Duties Connected to What the State Buys

Procurement gives the State a direct way to ask for evidence before using an AI system: a specification can require the vendor to explain the product, supply records and support review of the contracted service. The legal basis needs care. *United Haulers* and *Davis* distinguish the government-favoring measures before them from discrimination favoring private local competitors.⁴⁹² *Reeves* and *Hughes* address market participation. That doctrine requires identification of the transaction and market; a contract is not enough to establish its boundary. The four-Justice downstream reasoning in *South-Central Timber* is a warning against using one transaction to control a separate market after the participation ends.⁴⁹³ These decisions do not establish a general procurement presumption or settle every residual Commerce Clause question.⁴⁹⁴

The federal policy documents have a narrower relevance. Executive Order 14365 excludes otherwise lawful state procurement and own-use laws from the legislative preemption recommendation it orders.⁴⁹⁵ The March 2026 White House framework recommends a similar preservation. This is evidence of a proposed political boundary and a useful institutional parallel. Any preservation Congress adopts will turn on the law actually enacted.⁴⁹⁶

Federal acquisition practice supplies more concrete design material. OMB's April 2025 memoranda address acquisition and agency use separately.⁴⁹⁷ M-25-22 replaces the prior acquisition memorandum and specifies controls within its coverage, including attention to vendor informa-

⁴⁹¹ See *infra* Part V.E.1 (hospital application). The proposed implementation must distinguish institutional duties, professional standards, federal device requirements and the patient's available recourse.

⁴⁹² *United Haulers*, 550 U.S. at 342-47; *Dep't of Revenue of Ky. v. Davis*, 553 U.S. 328, 341-42, 353-55 (2008); see *supra* Section II.A.5. Public-function nondiscrimination is not identical to market participation or preemption.

⁴⁹³ *Reeves, Inc. v. Stake*, 447 U.S. 429, 436-40 (1980) (state sale of cement); *Hughes v. Alexandria Scrap Corp.*, 426 U.S. 794, 808-10 & n.19 (1976) (subsidy described as purchasing in effect, with Congress's power reserved); *South-Central Timber Dev., Inc. v. Wunnicke*, 467 U.S. 82, 95-99 (1984) (Parts III-IV, plurality of White, Brennan, Blackmun and Stevens, JJ.; Powell and Burger, C.J., concurred in the judgment favoring remand without joining that reasoning). The market and transaction must be identified; contract form is not dispositive.

⁴⁹⁴ See *supra* notes 492-493, Section II.A.5. *United Haulers* Part II-D's *Pike* analysis has four votes; *Davis* Part III-B's market-participant alternative has three, and Part IV reserves whether *Pike* applies while declining the requested comparison on that record. Neither supplies a general rebuttable presumption for all procurement conditions.

⁴⁹⁵ See *supra* notes 104, 167 (Part II.A.5 and Part II.B.4 first-use full citation of Executive Order 14365 § 8(b)(iii)) (Executive Order 14365 § 8(b) directs that the federal-preemption legislative recommendation "shall not propose preempting otherwise lawful State AI laws relating to . . . (iii) State government procurement and use of AI").

⁴⁹⁶ White House, *National Policy Framework*, *supra* note 142, at 3-4. Its recommended exclusions include general laws and state procurement/use. Political preservation must be distinguished from an enacted saving clause and from the courts' separate constitutional doctrines.

⁴⁹⁷ OMB Memorandum M-25-22, *Driving Efficient Acquisition of Artificial Intelligence in Government* (Apr. 3, 2025); OMB Memorandum M-25-21, *Accelerating Federal Use of AI Through Innovation, Governance, and Public Trust* (Apr. 3, 2025).

tion and contract performance.⁴⁹⁸ M-25-21 addresses high-impact federal uses, with specified practices and qualifications such as pilots and waivers.⁴⁹⁹ The continuation of acquisition governance across the two memoranda is observable; the claim that procurement form caused political survival is not established by the documents. The statutory direction in the Advancing American AI Act also matters. State law must supply the relevant purchasing powers.⁵⁰⁰

CMMC supplies an analogy for making standards contractual and checking compliance. The operative defense-acquisition provisions use required CMMC status and affirmations, with different assessment modes, conditional statuses, exceptions and phased applicability. Applicable clauses can also require subcontract flow-downs tied to the covered information and contract terms.⁵⁰¹ The lesson for AI procurement is to specify what evidence is required, who assesses it, what must be maintained and what follows noncompliance. Whether a state chooses self-attestation, agency review or an independent assessor requires an authorized design. Assessor competence, subcontract terms and enforcement against the proper contracting party remain implementation problems in both systems.⁵⁰²

Connecticut shows an existing allocation of public responsibilities. Public Act 23-16 assigns executive-branch inventories and ongoing assessments to DAS, policies to OPM and preimplementation responsibilities to agencies. It separately assigns Judicial Department functions. Its § 4 also requires state business contracts to contain an applicable privacy-law compliance term.⁵⁰³ The later Public Act 26-15 creates developer and deployer obligations in the private setting, with staggered dates and specified enforcement. The chronology illustrates public and private governance in the same state. It does not prove that the earlier arrangement caused or enabled the later enactment, or that the two impose equivalent duties.⁵⁰⁴

⁴⁹⁸ OMB Memorandum M-25-22, *supra* note 497, §§ 2(c), 4(c)-(d) (replacement of M-24-18, covered acquisition and contract provisions). Mandatory directions, encouraged practices, implementation timing and coverage limits should be distinguished; the memorandum is not state procurement authority.

⁴⁹⁹ OMB Memorandum M-25-21, *supra* note 497, §§ 4-5 (high-impact definition, practices and specified qualifications). Coverage and the pilots/waivers matter; this is not a statement that every federal AI use receives identical controls.

⁵⁰⁰ Advancing American AI Act, Pub. L. No. 117-263, div. G, tit. LXXII, subtit. B, §§ 7222, 7224(d), 7226, 136 Stat. 3668 (2022), set out in 40 U.S.C. § 11301 note (official 2023 edition text). The cited AI-hygiene and rapid-pilot provisions carry five-year sunset terms; that does not mean every provision sunsets. *See supra* note 396 (acquisition analogy).

⁵⁰¹ *See* DFARS Acquisition Rule, 90 Fed. Reg. 43,560 (Sept. 10, 2025) (effective Nov. 10); DFARS 204.7502-7504; 252.204-7021(a), (d)-(f) (NOV 2025); 252.204-7012(b)(2), (m). Required status and affirmation, self/C3PAO/DIBCAC modes, conditional status and phasing through November 2028 prevent a universal current-certificate description. Applicable subcontract clauses depend on the contract and covered information. The December 2017 NIST SP 800-171 implementation date concerns the historical § 7012 requirement.

⁵⁰² *See* DFARS provisions cited *supra* note 501; *infra* Part V.A. State authority to select standards, credential assessors, incorporate clauses, obtain records and enforce remedies must be established under the applicable state framework. Flow-down does not by itself establish direct government-subcontractor privity.

⁵⁰³ 2023 Conn. Pub. Acts No. 23-16, §§ 1-4 (SB 1103) (DAS, OPM, agency and Judicial Department assignments; applicable privacy-law compliance term in state contracts with businesses). The AI-governance provisions and the separate privacy contract provision should be described according to their own scope. *See supra* notes 214, 238.

⁵⁰⁴ 2026 Conn. Pub. Acts No. 26-15, §§ 7-13 (approved May 27, 2026) (developer/deployer roles; sections effective Oct. 1, 2026 and principal deployment triggers Oct. 1, 2027; § 12 Attorney General enforcement and no new private action; § 13 employment-law preservation). *See infra* note 532. The enactment sequence does not establish procurement causation or equivalent public/private remedies.

NASCIO's March 2026 report recounts 2025 survey categories concerning agency policies, inventories and advisory bodies. Those figures are evidence of reported governance activity, not a verified count of enforceable procurement clauses, independent audits or successful outcomes.⁵⁰⁵ I recommend starting with the State's own purchases because a defined service gives the State a concrete opportunity to specify records, inspect performance and learn what oversight requires. Authority, staffing, access to vendor information and the consequences of breach still have to be supplied. The recommendation does not assume that an existing CIO can exercise every power the proposed State CAIO would receive.⁵⁰⁶

Preemption remains a separate limit. *Crosby* invalidated Massachusetts's purchasing restriction because it obstructed the federal Burma sanctions scheme.⁵⁰⁷ Its rejection of the spending defense, together with *Gould*'s treatment of additional procurement sanctions in the labor setting, prevents treating market participation as a general preemption exemption.⁵⁰⁸ The relevant federal law may already exist in a sector; Congress need not enact a comprehensive AI statute before a conflict can arise. Voluntary NIST or ISO alignment can inform a specification, but does not supply a federal parallel-requirements defense. The comparison must identify the purchased service, the state condition and the binding federal rule it may affect.⁵⁰⁹

The implementation proposed in Section V.A connects specifications and evidence to assessment and contract administration, distinguishing present practices from proposed authority. CMMC's annual affirmations and subcontract clauses remain analogues; they cannot be reported as controls already supplied by that Section or by every state contract.⁵¹⁰ A public purchaser may create useful information and experience. Whether those benefits reach private deployments depends on later law, contracts and institutional choices. The sequence itself guarantees no extension.

The four authorities together provide several ways to assign responsibility, but they do not guarantee substitution. A consumer-protection action, a civil rights claim, a licensing examination and a contract remedy reach different defendants, conduct and injuries. A person excluded from one route may lack another, and a single federal provision may constrain several routes at once. The Article's resilience claim is that distributing functions across identifiable legal institutions can preserve lawful and useful parts of governance as other parts change. That is an argument about design and the political cost of displacement, not a constitutional barrier to Congress or a

⁵⁰⁵ NASCIO, *Beyond Generation: The Rise of Agentic AI in State Government 2* (Mar. 2026) (recounting 2025 State CIO Survey: 88 percent responsible-use, guardrail, security or ethics measures; 84 percent inventories; 82 percent advisory committees or task forces). These are the report's survey categories, not an October 2026 procurement-clause census or measured efficacy.

⁵⁰⁶ See *infra* Part V.A (present practices and proposed procurement-first controls). The recommendation depends on lawful authority, sufficient information and resources, and appropriate administration of the contract; an existing office is not proof of those conditions.

⁵⁰⁷ *Crosby*, 530 U.S. at 373 n.7, 376-87; see *supra* notes 120, 156, 169.

⁵⁰⁸ *Crosby*, 530 U.S. at 373 n.7 (rejecting the spending-power defense); *Gould*, 475 U.S. at 287-91 (additional procurement sanctions in a field governed by federal labor law). Dormant Commerce Clause market participation and preemption remain separate inquiries.

⁵⁰⁹ An executive direction, proposed agency statement, voluntary framework and binding federal rule have different legal effects. For the identified instruments, see *supra* notes 139-146 and Section II.B.1. A relevant existing sectoral rule can constrain a purchase condition; absence of comprehensive AI legislation would not establish absence of a conflict.

⁵¹⁰ See *infra* Part V.A (acquisition specifications, evidence, assessment and proposed authority); *supra* note 501 (CMMC's own affirmation and flow-down terms). Those federal terms are analogues rather than a claim that every term appears in the proposed state implementation. Extension into private deployment needs a separately identified legal or contractual route.

prediction of equivalent redress. The implementation in Part V retains the distinction between authority already in force and powers the legislature would need to create.

Part V: Governance Primitives: A Modular Architecture for State AI Regulation

A. Procurement-First Controls as the Architectural Foundation

California already regulates its own purchases of generative AI through written procurement policy, and Appendix §§ 1.4 through 1.6 propose a particular statutory arrangement alongside it: eligibility evidence fixed by rule, a vendor assessment that state reviewers read before the contract, and a vendor warranty tied to that assessment. The California rules appear in the State Administrative Manual (SAM) chapter on generative AI procurement, which applies to state entities except as a statute or policy exempts them; in chapter 23 of the State Contracting Manual (SCM); and in the non-cloud model information-technology terms published by the Department of General Services (the model contract).⁵¹¹ These documents supply administrative policy and model contract language. Whether the model terms govern a given purchase depends on how they are incorporated and on the statement of work, which can change many of them. I treat them as existing components. The comparison locates what the Appendix adds. It does not show that the Appendix is valid or would work. The limits on the State's purchasing authority and the relevant federal overlap remain separate legal questions.⁵¹²

A hypothetical shows the existing sequence. Suppose a state benefits agency proposes to buy a tool that summarizes applicant files for caseworkers, knows that the tool includes generative AI, and plans to contract on the non-cloud model terms as published. Procurement staff must confirm with the department's chief information officer (CIO) that the purchase may proceed

⁵¹¹ Cal. Dep't of Gen. Servs., *State Administrative Manual* [hereinafter SAM] § 4986.9 (scope statement; Parts A–B) (page marked “Revised: 11/2025,” revision list entry “02/2025”), <https://www.dgs.ca.gov/Resources/SAM/TOC/4900/4986-9>; Cal. Dep't of Gen. Servs., *State Contracting Manual* [hereinafter SCM] vol. 2, ch. 23, §§ 2300–2309 (sections marked “Revised: 02/2025”; §§ 2302.1–2302.2 marked “New: 02/2025”), <https://www.dgs.ca.gov/PD/Resources/SCM/TOC/23/Chapter-View>; Cal. Dep't of Gen. Servs., Procurement Div., *Information Technology General Provisions Non-Cloud, Goods & Services*, DGS PD 403-ITGP Non-Cloud (rev. Feb. 20, 2025) [hereinafter ITGP 403], <https://www.dgs.ca.gov/-/media/Divisions/PD/Acquisitions/CMAS/2025/IT-General-Provisions-NonCloud-DGS-PD-403ITGP-Revised-02202025.pdf>. The SAM scope statement applies the chapter to “all State Entities except as specifically exempted by statute or policy”; the manual's definition of that term and any statutory exemptions are not examined here. On incorporation and variation, see SCM § 2303 (required generative AI terms “incorporated into the State's model contract terms”; amendments that include pre-2025 information-technology general provisions must incorporate §§ 22.1 and 22.2 of the “applicable and current” Cloud or Non-Cloud version), and ITGP 403 §§ 6.1, 8.1.2, 8.4.1, 8.6, 16.3 (default terms that operate unless “otherwise specified in the Statement of Work”). The Cloud general provisions are a separate document not examined here. The 2025 labels are those on the versions now served and are not an archival record of the text in force on an earlier date. SIMM 5305-F and 5310-C, SAM § 4986.2, SAM 5300 and SIMM 5300, the generative AI special provisions, statement-of-work templates, and executed contracts are separate documents not examined here; no executed contract or court enforcement is verified. In this passage, “CDT” is the California Department of Technology; “CIO” is a department's (or the state's) chief information officer; and “State CAIO” is the Appendix's defined term for the State Chief AI Officer, Appendix § 1.2(d) *infra*, which is distinct from the CIO. The model contract is cited by section for the provisions discussed.

⁵¹² See *supra* notes 93–95, 166–167, 492–496 (the DCC and preemption questions developed in Parts II and IV.D).

and must keep that confirmation in the file.⁵¹³ Before award, the agency must complete the state's generative AI risk assessment form, which classifies the risk as low, moderate, or high, and must submit it to the California Department of Technology (CDT) at every level. It must also complete a privacy assessment, which it provides to CDT on request. The SAM requires the risk assessment before award of information-technology, non-information-technology, and telecommunications procurements generally, and it exempts listed uses, such as real-property transactions and commodities with no technology or service component.⁵¹⁴ A moderate or high result requires the agency to seek CDT's consultation before award. A low result lets the purchase continue, although CDT may require consultation.

The bidder's disclosure duty applies at every risk level. The solicitation must carry a disclosure clause that the manual forbids changing. It requires a bidder to notify the state in writing if it (1) intends to provide generative AI as a deliverable, or (2) intends to use generative AI, including generative AI from third parties, to complete all or a portion of any deliverable that "materially impacts" state-system functionality, risk to the state, or contract performance. The clause takes the meaning of "materially impacts" from SAM § 4986.2.⁵¹⁵ It warns that nondisclosure may result in disqualification, and it reserves the state's right to add special provisions or to reject a bid that presents an unacceptable level of risk. If a bid discloses generative AI, procurement staff must engage the CIO before proceeding. The model contract adds a duty to furnish the information and data reasonably required for inspection and testing and to keep records of inspections and their results, available during performance and for three years after final payment. Unless the statement of work provides otherwise, the state relies on the contractor's own quality assurance system in place of state inspection and testing when it acquires goods or commercial software. The next sentence preserves inspection, testing, and acceptance of goods at their destination.⁵¹⁶

⁵¹³ SCM § 2301 (before proceeding with a generative AI purchase, releasing a solicitation that includes one, or approving a contract that includes one, procurement professionals "must confirm with their department Chief Information Officer (CIO) that the purchase may proceed"; confirmation kept in the procurement or contract file); § 2302.2 (engagement of the CIO when a bid discloses generative AI; confirmation kept in the file).

⁵¹⁴ SAM § 4986.9(A) ("[f]or all IT, non-IT, and Telecommunication procurements," before award, completion of the SIMM 5305-F generative AI risk assessment "to determine if the GenAI risk is low, moderate, or high"; submission to CDT "for all risk levels"; consultation with CDT before award for moderate or high risk; for low risk, continuation permitted, with consultation if CDT requires it; listed exemptions from the assessment, including real property transactions, commodity types with no technology or service component, and "[a]ny other GenAI use determined by CDT to be low risk"; for "all GenAI procurements," completion of the SIMM 5310-C privacy assessment, provided to CDT "for review upon request"). The forms themselves are not examined here.

⁵¹⁵ SCM § 2302 (all written solicitations must include the disclosure notification clause; "[n]o changes are permitted"; failure to report generative AI "may result in disqualification"; state's reservation of the right to incorporate special provisions or reject bids that present "an unacceptable level of risk"). See also SAM § 4986.9(A) (parallel disclosure notification language required in solicitations and contracts for information-technology and telecommunications goods and services; the SAM clause is addressed to the "Contractor" and does not contain the SCM clause's disqualification and reservation sentences). SAM § 4986.2, which both clauses use for the meaning of "materially impacts," is not examined here. See ITGP 403 § 22.1.6.

⁵¹⁶ ITGP 403 § 6.1.1 (unless the statement of work provides otherwise, for goods or commercial software "the State shall rely on Contractor's existing quality assurance system as a substitute for State inspection and testing"); § 6.1.1 (goods remain subject to inspection, test, and acceptance at destination); § 6.1.3 (contractor "shall furnish to inspectors all information and data reasonably required to perform inspection and testing," and keeps "records evidencing inspections and their result" available "during Contract performance and for 3 years after final payment") (PDF pp. 8–9).

Two further requirements are conditioned on risk. For moderate- and high-risk uses, the SAM requires the entity to evaluate vendors and suppliers “for compliance with” the cited state manual provisions, the NIST AI Risk Management Framework, and NIST 800-53, and the listed measures include ensuring that vendors provide documentation “on the development, training, and operation of their GenAI systems, including details on the origin of datasets, algorithms, and any pre-trained models used.” The SCM requires generative AI special provisions in the solicitation, award process, and contract when the assessment and CDT consultation indicate moderate or high risk, applied case by case by the CIO in consultation with CDT.⁵¹⁷ The SAM requires that documentation, but its text states no pass line for judging it, and it names the NIST framework without an edition; the documents it cites may supply more. The comparison here does not extend to the separate special provisions.

After award, the model contract turns some of these steps into obligations of the contractor, which is its term for the vendor. The contractor must notify the state in writing, immediately, when it intends to deliver generative AI or to use it to complete all or part of a deliverable that “materially impacts” state-system functionality, risk to the state, or contract performance, with the meaning given in SAM § 4986.2. At the state’s direction it must discontinue previously unreported generative AI that results in a material impact, “as determined by the State.” It must also ensure that the generative AI it supplies is equitable, non-discriminatory, and reasonably well designed to avoid harmful, offensive, dangerous, and unlawful impact. The state may, at its sole discretion, treat a failure to disclose or discontinue as a material breach when the failure results in a material impact, and its remedies include termination for default.⁵¹⁸ Contract managers must monitor, assess, and validate generative AI deliverables for equitable outcomes, inaccuracies, fabricated content, hallucinations, bias, and the need for human action in decision-making, and they must engage the CIO when previously unreported generative AI

⁵¹⁷ SAM § 4986.9(B) (for moderate and high risk GenAI use, evaluation of vendors and suppliers “for compliance with the SAM 5300, SIMM 5300 and the National Institute of Standards and Technology (NIST) Artificial Intelligence (AI) Risk Management Framework (RMF) and NIST 800-53,” including verifying adherence to state policies, privacy laws, and bias mitigation practices; mandating incident reporting, logs, and cooperation; ensuring data-protection measures; and ensuring that vendors provide documentation as quoted in the text); SCM § 2302.1 (special provisions “required to be part of a solicitation, award process, and contract when” the risk assessment and CDT consultation “indicate moderate or high risk GenAI”; “applied on a case-by-case basis as determined by the department’s CIO in consultation with CDT”). The SAM provision names the framework without stating an edition and states no pass line. SAM 5300 and SIMM 5300 are not examined here. The model contract’s security provisions are a separate subject: ITGP 403 § 21.3 requires NIST SP 800-53 controls “assuming a National Institute of Standards and Technology (NIST) MODERATE classification” unless the statement of work specifies otherwise and permits, with prior written approval of the state’s chief information security officer, a current ISO 27001 certification “in satisfaction of the foregoing requirements of this subsection” (PDF pp. 34–35). That certification is not ISO/IEC 42001, and the security classification is not the generative AI risk level set by the SAM § 4986.9(A) assessment. The inspection and audit clause, § 21.13.1, provides for testing “that controls are in place and working as intended,” with vulnerability scans as examples (PDF p. 39), and I do not treat it as a conformity check for generative AI systems generally.

⁵¹⁸ ITGP 403 § 22.2.1(a), (c) (notification; discontinuance “at the direction of the State” of “any previously unreported GenAI that results in a material impact . . . as determined by the State”), § 22.2.2 (State “at its sole discretion” may consider failure to disclose or discontinue a material breach “when such failure results in a material impact”; remedies include termination for default under § 16.3), § 22.3.2 (contractor “shall ensure that the GenAI included, or made available as part of the Deliverables is equitable, non-discriminatory, and reasonably well-designed to avoid harmful, offensive, dangerous, and unlawful impact”) (PDF pp. 42–43). The contract text refers to SAM § 4986.2 for “materially impacts.” Section 22.2.1(d) separately addresses State approval of previously undisclosed generative AI, with contract amendment that may incorporate the special provisions.

appears. The state entity separately must keep a multidisciplinary governance team to oversee changing risks over the life of the use.⁵¹⁹ The deliverables carry warranties of substantial conformity to the contract's requirements and specifications. Unless the statement of work says otherwise, the state's remedies for breach of these warranties include repair, re-performance, or replacement, or termination with a refund and the cost of cover.⁵²⁰

Those provisions carry conditions that a reviewer must keep in view. Unless the statement of work says otherwise, the contractor does not warrant that software is error-free or will run without immaterial interruption. I read that as a qualification of the software warranties, alongside the contract's express conformity warranties.⁵²¹ Default termination requires a material breach that the contractor fails to cure within the period stated in the cure notice, which may be no shorter than fifteen days unless the statement of work specifies otherwise. The contractor's liability for damages is generally limited to the purchase price, except as designated deputy directors approve, with exceptions that include liability under the statutory-compliance provision, the indemnities for intellectual-property infringement and for injury or property damage, and direct costs of a data breach attributable to the contractor. Neither party is liable to the other for consequential, incidental, indirect, special, or punitive damages, except as the statement of work or those listed exceptions provide. The generative AI section makes the contractor "liable for any Hallucination produced by the GenAI that has an adverse impact on Generated Data or a Deliverable," but the general limitation applies to liability "for any cause whatsoever" subject to its listed exceptions, and I do not read the hallucination sentence as unlimited liability.⁵²² Contract disputes run through a separate procedure: the contractor seeks a final decision from the buying department, appeals to a deputy director, and then may file a claim with the Government Claims Board. That procedure resolves contract disputes. It is not review of a standard that determines which systems may be sold to the state.⁵²³

⁵¹⁹ SCM § 2304 (contract managers "shall continuously monitor, assess, and validate GenAI contract deliverables" for the listed matters; must enforce generative AI contract terms by monitoring for previously unreported generative AI; must engage the CIO on discovery or disclosure; contracts "may include unique, contract-specific requirements"; changes in use "may require an amendment"); SAM § 4986.9(A) (State Entity "must have a multidisciplinary governance team to continuously oversee and monitor" risks "throughout the lifecycle of the GenAI use").

⁵²⁰ ITGP 403 § 8.2(a)–(d) (warranties of substantial conformity to the contract's requirements, descriptions, and specifications, and related quality warranties), § 8.6 (unless the statement of work specifies otherwise, remedy "shall include, at the State's sole discretion," re-performance, repair, or replacement, or termination with a refund and "Cost to Cover") (PDF pp. 12–13); *see also* § 8.7 (disclaimer of warranties other than those specified in the contract).

⁵²¹ ITGP 403 § 8.4.1(a) (unless the statement of work specifies otherwise, "Contractor does not warrant that any Software provided hereunder is error-free or that it will run without immaterial interruption"); § 8.4 (software tested to minimize the risk of harmful code, not a guarantee of any model-performance measure); § 8.7 (disclaimer applies "[e]xcept for the warranties specified in this Contract") (PDF pp. 12–13).

⁵²² ITGP 403 § 16.3 (default termination "only if the failure constitutes a material breach" and the contractor fails to cure within the cure-notice period, "in no event" less than 15 days "unless otherwise specified in the Statement of Work") (PDF pp. 27–28); § 16.4 (state's recovery for default "subject to Section 12"); §§ 12.1–12.2 (liability limited to the Purchase Price "for any cause whatsoever" except as approved by designated deputy directors; no liability for consequential, incidental, indirect, special, or punitive damages except as the statement of work or § 12.2 provides; exceptions for liability under §§ 13.1, 11.1, and 11.2, prevailing-party costs and fees, and direct costs of data breach attributable to the contractor) (PDF pp. 18–19); § 22.3.2 (contractor "shall be liable for any Hallucination produced by the GenAI that has an adverse impact on Generated Data or a Deliverable") (PDF p. 43). The cited text does not state how § 22.3.2 interacts with § 12 in a particular contract; that depends on the contract's terms.

⁵²³ ITGP 403 §§ 17.1–17.3 (contractor's written demand for a final decision by the buyer's department

Appendix § 1.4 arranges related functions in proposed statutory text. Under § 1.4(a), an AI system is eligible for state procurement if and only if it satisfies the criteria that the State Chief AI Officer (State CAIO) specifies by rule. The State CAIO is an officer whose office § 1.3(a) would create, and is distinct from a department's or the state's CIO.⁵²⁴ The rules must map the NIST framework's function-level outcomes (GOVERN, MAP, MEASURE, and MANAGE) and the ISO/IEC 42001 clause requirements to objectively verifiable evidence items, define the risk-profile tiers of subsection (c) by published criteria, and specify pass/fail determination, cure periods, and appeal under the state administrative procedure act. Section 1.4(b) requires the State CAIO to publish a schedule of the NIST and ISO/IEC editions adopted by rule and in force. Section 1.4(c) says the State CAIO "may" establish three tiers keyed to the system's risk profile and the deployer's decision context: vendor self-attestation for lower-risk contracts, State CAIO documentation review for moderate-risk contracts, and third-party assessor certification for the highest-risk contracts. I carry the directive in (a) and the permissive "may" in (c) as drafted and do not use either to resolve the other. This comparison does not establish that California's risk criteria and the Appendix's tier criteria are equivalent. The appeal and cure provisions belong to eligibility determinations. They are not the contract-dispute procedure and default cure described above.

The mapping asks for a choice that NIST does not make. Its GOVERN function addresses organizations that acquire AI systems as well as those that build them, but the NIST Framework is voluntary, and its Core is not a checklist: NIST says its actions "do not constitute a checklist" and that users "may apply these functions as best suits their needs."⁵²⁵ A statute that conditions eligibility on those outcomes must say what evidence counts and what result passes. The Appendix assigns that content to rulemaking, so the evidence schedule remains a policy choice that the enacting state must make. What the Appendix requires is that the editions and tier criteria be published and the evidence items fixed by rule. The point of that arrangement is that a vendor, a purchasing agency, and a reviewer of an eligibility decision can read the standard before the decision is made.

The Appendix also directs the State CAIO to address applicable sector-specific federal guidance in § 1.4(a)(3), naming the FDA Software-as-a-Medical-Device framework and SR 26-2 for covered state-government banking relationships. Those references specify the proposed procurement rule's inputs; the relevant federal overlay still has to be examined for the particular system and use.⁵²⁶ The CMMC comparison concerns a contractor-compliance and affirmation

director; appeal to the DGS Procurement Division or CDT deputy director, as applicable; claim to the Government Claims Board) (PDF pp. 29–30). The Appendix's eligibility appeal is Appendix § 1.4(a) *infra*.

⁵²⁴ Appendix §§ 1.2(d), 1.3(a), 1.4(a)–(c) *infra*. Section 1.4(a)(2) names Clauses 6, 8, and 9 of ISO/IEC 42001, the licensed text of which this discussion does not characterize. The content and current applicability of the federal guidance named in § 1.4(a)(3) are not assessed here; § 1.3(d) (political subdivisions) is not discussed.

⁵²⁵ Nat'l Inst. of Standards & Tech., *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, at 2 (Jan. 2023) ("intended to be voluntary, rights-preserving, non-sector-specific, and use-case agnostic"), at 20 (Core composed of four functions; "Actions do not constitute a checklist, nor are they necessarily an ordered set of steps"), at 21 (users "may apply these functions as best suits their needs"; GOVERN function concerns organizations "designing, developing, deploying, evaluating, or acquiring AI systems"), <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>. The Appendix defines the framework by this edition, and the State CAIO may adopt a later revision only by rule. Appendix §§ 1.2(e), 1.4(a)(3) *infra*. The comparison describes the Appendix's reference to the ISO framework without characterizing its licensed clause text.

⁵²⁶ Appendix § 1.4(a)(3) *infra*. This describes the proposed text's federal inputs; the cited federal guidance is not characterized here.

mechanism. It does not establish the content or enforceability of the Appendix's AI warranty.⁵²⁷

Section 1.5 moves the vendor's evidence into a statutory assessment. Under § 1.5(a), the vendor must deliver, before contract performance, a written assessment of the system's intended purpose and operational context; its training data, with provenance, time-period coverage, and documented gaps; its validation methodology, including testing populations and performance metrics; its known limitations and operational boundaries; its foreseeable risks of adverse outcomes for state residents in regulated domains; and the mitigation measures the vendor and the deployer will implement. Section 1.5(b) requires review by the State CIO (or the agency CIO, where applicable) and the State CAIO before the deployer enters the procurement contract, and it lets the State CAIO require revisions before execution. Section 1.5(c) requires the deployer to retain the assessment for the contract term plus five years and to make it available to the State CAIO, the State Auditor, and any sectoral regulator with authority over the deployment context. That record, its holder, and its clock differ from the model contract's, under which the contractor keeps inspection records during performance and for three years after final payment.⁵²⁸

Section 1.6 ties the contract to those documents. Every state procurement contract for an AI system must include the vendor's warranty and covenant that the system satisfies the § 1.4 standards and operates consistently with the § 1.5 assessment; that it has been tested and validated for the deployment contexts the assessment documents, using disclosed methodology and results, and performs within its documented limitations when operated under the vendor's documented conditions of use; that the vendor will promptly notify the deployer of any post-delivery discovery materially affecting the documented limitations or the assessment; and that the vendor will keep supplying the documentation the deployer needs to meet its obligations under the Act. The warranty is framed around documented limitations and conditions of use, not around error-free operation. Section 1.6(b) lists termination by the deployer or the State CAIO; a civil penalty, recoverable by the Attorney General, of up to 10 percent of the contract value or a bracketed \$500,000 that the enacting legislature completes, whichever is greater; disqualification from future state procurement for up to five years; and any remedy under the procurement code or common-law warranty doctrine. Section 1.6(c) leaves in place affected persons' rights under civil rights, consumer protection, and sectoral statutes. The solicitation clause's possible disqualification for nondisclosure is a bid-stage consequence and differs from a proposed exclusion from future procurement for up to five years. The California sections cited here do not provide for the civil penalty or for that exclusion, and I make no claim about other California authority.⁵²⁹

Return to the benefits agency. Under the California texts, the purchase would require CIO confirmation, completion and submission of the risk assessment to CDT, completion of the privacy assessment, the bidder's required disclosure, and, at moderate or high risk, the vendor evalu-

⁵²⁷ For the contractor-compliance and affirmation comparison, *see infra* note 529.

⁵²⁸ Appendix § 1.2(b) ("Deployer"), § 1.5(a)–(c) *infra*. Subsection (a) times delivery "[b]efore contract performance"; subsection (b) requires review "before the Deployer enters into the procurement contract." Both timing provisions are carried as drafted, and the discussion relies on each only as stated. *Cf.* ITGP 403 § 6.1.3 (PDF p. 9) (contractor's inspection records).

⁵²⁹ Appendix § 1.6(a)–(c) *infra*; SCM § 2302 (nondisclosure "may result in disqualification"); ITGP 403 §§ 8.6, 16.3, 22.2.2 (contract remedies). The statement about the civil penalty and the multi-year exclusion is limited to the California sections cited in this passage. For the separate contractor-affirmation comparison, *see supra* note 501; DFARS 252.204-7021(d)(3), (e)(3) (annual affirmation of continuous compliance), <https://www.acquisition.gov/dfars/252.204-7021-cybersecurity-maturity-model-certification-requirements>. This clause supplies an analogue, not the proposed AI warranty or its remedies.

ation, development and training documentation, and special provisions as the CIO and CDT apply them. Under the Appendix, once the State CAIO has issued rules, the rules would identify the evidence items for the system's risk profile, and the vendor's assessment would state its testing populations, performance metrics, and limits. Before the agency enters the contract, reviewers compare that assessment with the agency's proposed use. If the vendor validated the tool only on files unlike the agency's, or documented a limitation that the proposed use exceeds, the mismatch can appear on the face of the documents, and the warranty is then defined by the deployment contexts those documents describe. If the vendor discovers after delivery something materially affecting a documented limitation or the assessment, § 1.6(a) requires prompt notice. Each of these is a check on a paper record. None shows that the vendor can supply the underlying data, that reviewers would catch a mismatch, or that the agency would act on it. Resemblance to California's policies does not show that the Appendix is valid or effective, and the Appendix leaves appropriations and other enactment-specific machinery to the enacting state.⁵³⁰

Procurement-first is a sequencing proposal. A State can use its own purchases to examine the records, review capacity, and costs that a wider regime would require. Connecticut's 2023 act assigned inventory and ongoing-assessment functions to the Department of Administrative Services and procurement and implementation policies to the Office of Policy and Management.⁵³¹ Its 2026 act added private-sector employment-technology duties enforceable through a specified consumer-protection route and amended employment-discrimination law so that use of an automated employment-related decision technology is no defense to the specified discrimination complaint.⁵³² The chronology identifies public and private mechanisms. It does not establish that the earlier program caused the later enactment or proved its feasibility.

I favor beginning with purchasing because the State can define the evidence it seeks, identify the officials who must review it, and test its administrative capacity through its own contracting decisions. Vendors compete for those contracts, but competition does not establish that vendors will accept the conditions, that small suppliers can furnish the evidence, or that the State has sufficient bargaining power. Existing NIST-based controls supply components for the proposal; they do not choose the evidence threshold or show that a particular obligation is proportionate. The legislative histories discussed in Part III.B provide context for the political-economy question.⁵³³ Whether this sequence reduces opposition or improves administration is a hypothesis to test. The description alone establishes neither.

The worked Appendix extends beyond the eligibility, assessment, and warranty provisions compared here. Section 1.3 defines the proposed State CAIO's authority, and § 1.7 provides for an audit program and incident reporting. The preliminary note leaves state-specific enactment machinery, including appropriations, to the enacting legislature. Those completion choices

⁵³⁰ Appendix § 1.5(b), § 1.6(a)(3)–(4) *infra*; Appendix, Preliminary Note *infra* (appropriations and other enactment-specific machinery left to the enacting state).

⁵³¹ 2023 Conn. Pub. Acts No. 23-16, §§ 1–2 (DAS inventory and ongoing assessment; OPM policies governing procurement, implementation, and assessment), <https://www.cga.ct.gov/2023/ACT/PA/PDF/2023PA-00016-R00SB-01103-PA.PDF>. See *supra* notes 214, 238, 503.

⁵³² An Act Concerning Online Safety, 2026 Conn. Pub. Acts No. 26-15 (approved May 27, 2026). The Connecticut Unfair Trade Practices Act enforcement hook, enforced exclusively by the Attorney General with no private right of action, is § 12; the amendment to Conn. Gen. Stat. § 46a-60(b) making automated-tool use no defense to employment-discrimination liability is § 13; § 14 makes the parallel amendment to Conn. Gen. Stat. § 46a-81c. Commentary refers to the act as the Connecticut Artificial Intelligence Responsibility and Transparency (CART) Act; the act itself enacts no short title.

⁵³³ See *supra* notes 208–211 (Part III.B.1 discussion of Colorado's 2024 act and 2026 replacement, and Virginia's vetoed 2025 proposal).

matter to whether the proposed records can be obtained, reviewed, and acted on.⁵³⁴

B. Deployer-Centered Impact Assessments as the NEPA Analogue Adapted to AI

An impact assessment makes the decision to deploy answerable before adverse outcomes accumulate. It should explain the intended use, the people and interests affected, foreseeable risks, safeguards, and the reasons for proceeding. Initial assessment, material-change updates, and periodic review perform different jobs; a statute must specify each trigger.⁵³⁵ Alternatives, including a non-AI approach, can make the exercise more searching. That recommendation extends beyond some of the assessment schedules discussed here.⁵³⁶

NEPA supplies a useful analogy with defined limits. Its statutory framework continues after the rescission of CEQ's government-wide regulations, while agency procedures must be examined separately.⁵³⁷ *Robertson* distinguishes an informed decision from a mandated result: NEPA does not require an agency to select the least harmful alternative. That does not mean any completed document satisfies its procedural requirements.⁵³⁸ This Article borrows five design features for AI assessment: tiered triggers, alternatives analysis, attention to aggregate effects, public participation and disclosure, and judicial review of the process. They are proposed choices, not a description of current NEPA requirements or duties that automatically transfer to private deployers.⁵³⁹

The limitation of an assessment is that good paperwork may coexist with harmful decisions. NEPA is not categorically a one-shot exercise: *Marsh* recognizes supplementation when major federal action remains and significant new information bears on environmental effects. The proposal here goes further by coupling deployment assessments to ongoing monitoring, incident response, and substantive consumer-protection, civil-rights, and licensing law.⁵⁴⁰ That coupling gives adverse findings an enforcement route. Its effectiveness depends on the quality of the evidence and the institutions using it.

⁵³⁴ Appendix §§ 1.3, 1.7(a)–(b), Preliminary Note *infra*.

⁵³⁵ See *supra* Part III.D.2; *infra* Appendix §§ 2.3–2.6. Assessment, monitoring, and trail preservation have different functions and triggers.


⁵³⁶ Alternatives analysis is this Article's recommendation. Appendix § 2.4 incorporates the six items of § 1.5(a) and adds specified plans, analysis, and attestation; that schedule does not expressly require non-AI alternatives.

⁵³⁷ 42 U.S.C. §§ 4332, 4336–4336a; *Robertson v. Methow Valley Citizens Council*, 490 U.S. 332, 349–53 (1989); CEQ, *Removal of National Environmental Policy Act Implementing Regulations*, 90 Fed. Reg. 10,610 (Feb. 25, 2025), finalized without changes at 91 Fed. Reg. 618–19 (Jan. 8, 2026). The historical agency examples differ: Navy, 90 Fed. Reg. 29,453; Army Corps, 90 Fed. Reg. 29,461; Interior, 90 Fed. Reg. 29,498 (all July 3, 2025). They do not establish every agency's current procedures.

⁵³⁸ *Robertson*, 490 U.S. at 349–53, 350 n.13. The process/result distinction presupposes adequate identification and evaluation. The opinion's particular comment-response rule came from then-current CEQ regulations; 42 U.S.C. §§ 4332(2)(C), 4336a(c) separately address statutory comment procedures.

⁵³⁹ The five features are this Article's selection for AI legislation. Cumulative-effects analysis and an APA-style review standard are proposed design choices, not an assertion that every feature remains a general federal NEPA requirement. See *supra* note 358 for the separate limits in *Seven County* and *Public Citizen*.

⁵⁴⁰ *Marsh v. Oregon Natural Resources Council*, 490 U.S. 360, 371–76 (1989), particularly 374 (supplementation with remaining major federal action and significant new information). The proposed AI monitoring and substantive-law coupling addresses the independent design objection that an adequate assessment need not produce a less harmful result.



Sources: NIST AI100-1 §§5.1–5.4; ISO public description; Parts III.D/V; Appendix.

Figure 2. Framework inputs and proposed statutory choices. Sources: NIST AI100-1, §§ 5.1–5.4 (2023); ISO/IEC 42001:2023, official public description; Parts III.D and V and Appendix §§ 1.7(b), 3.4A(c), 4.7–4.9.

NIST’s MAP function helps identify context, system characteristics, affected populations, and impacts.⁵⁴¹ ISO/IEC 42001 supplies a related organizational planning approach.⁵⁴² Legislation must select the required content, responsible actor, access rules, and consequences. Neither a reference to a framework nor an organization’s certification establishes that a specific deployment has satisfied the statute.

Appendix § 2 gives the assessment proposal a specific form. It applies to private Covered Deployers within its domain, nexus, tier, and exemption rules. Section 2.4 incorporates the six vendor-assessment items listed in § 1.5(a), then adds risk-management, disparate-impact, recourse, and senior-official attestation requirements. Those six incorporated items concern purpose, training data, validation, limitations, risk, and mitigation; they do not expressly require

⁵⁴¹ NIST AI RMF, *supra* note 350, § 5.2, tbl. 2, MAP 1–5. These categories organize context, categorization, capabilities and expected benefits, component risks, and effects on people and society. They are not, without a specified legal adoption, a statutory checklist.

⁵⁴² ISO/IEC 42001:2023, *supra* note 351, cls. 6.1.2–6.1.4; *infra* Appendix § 2.2(e) (benchmark uses).

non-AI alternatives analysis. Section 2.5 requires assessment before deployment, an update within ninety days of a material change, and review annually for the highest tier and every three years for lower tiers. The worked text selects those intervals from a wider range of possible designs.

Section 2.6 requires filing with the CAIO and sectoral regulator and a redacted public registry summary. It reserves public comment and a response document for specified highest-tier public-benefits and criminal-justice uses. A separate provision permits review of specified final CAIO actions for reasoned consideration on the whole record. The design does not promise unrestricted public access to every assessment or a hearing on every private deployment. Section 2.7 supplies public enforcement, alongside the private remedies preserved by § 4.

Proportionality matters. Risk tiers and coverage limits structure the worked text, which uses a small-deployer exemption, with statutory exceptions; it does not enact a comprehensive size-based ladder of internal and external assessors. CMMC illustrates why assessment modes can differ with risk, but its federal cybersecurity levels do not determine appropriate AI thresholds.⁵⁴³ An exemption also carries a cost: it removes some decisions from the assessment duty. A state should evaluate administrative capacity and expected coverage as well as compliance expense before choosing the numbers.⁵⁴⁴ The bracketed defaults in the Appendix are proposals for that choice, not empirically validated cutoffs.

C. Audit Trail Mandates as the Backward-Looking Accountability Spine

An assessment explains the decision to deploy. An audit trail preserves what happened afterward: relevant inputs and outputs, the use made of them, monitoring results, and corrective action. The records must be usable by a regulator, auditor, or court under the applicable access rules.⁵⁴⁵ Preservation periods should fit the relevant examination and litigation needs. Appendix § 3.4 specifies trail content but does not itself supply a general minimum retention period for private deployments; the five-year vendor-assessment period in § 1.5(c) serves a different procurement duty.

New York City provides an important implementation record. Local Law 144 took effect on January 1, 2023, and enforcement began on July 5. It requires an independent bias audit within the preceding year, a publicly available summary, and specified employment notices. Its rules define independence and prescribe audit calculations and disclosures.⁵⁴⁶ The problem

⁵⁴³ 48 C.F.R. § 252.204-7021, definitions and paragraphs (d)–(f); *id.* §§ 204.7501–.7504. CMMC distinguishes Level 1 self-assessment, Level 2 self-assessment or C3PAO assessment, and Level 3 government DIBCAC assessment, with contract-specific and phase conditions. The Appendix’s State review and third-party tiers are its own proposal. *See supra* note 501.

⁵⁴⁴ *Infra* Appendix § 2.3(b)–(d), § 3.5(a). Proportionality trades compliance cost against coverage and assurance. The worked text does not establish an automatic exemption from external audits for every small deployer or prove that a proposed threshold removes unequal burdens.

⁵⁴⁵ *See supra* Part III.D.3; *infra* Appendix §§ 3.2(a), 3.4–3.5. The trail’s specified categories and access rights differ from an illustrative record of every operational input and output. Section 5.3’s action limitation is not a retention rule; § 1.5(c)’s contract-plus-five-year rule concerns the procurement assessment.

⁵⁴⁶ N.Y.C. Admin. Code §§ 20-870 to -874; Local Law 144 of 2021, § 2; 6 R.C.N.Y. §§ 5-300 to -304; DCWP, *Automated Employment Decision Tools: Frequently Asked Questions* (enforcement commencing July 5, 2023). The statute became effective January 1, 2023; qualifying shared or vendor-arranged audits can satisfy the applicable conditions.

documented by the Comptroller was not the absence of those legal requirements. It was the performance of the regime in practice.

The Comptroller reviewed thirty-two research-identified companies, comprising twenty-two employers and ten vendors. Its review of the twenty-two employers identified three potential independence problems, five audits with insufficiently explained tiered comparisons, and nine with insufficiently explained data; those categories may overlap. Eight posting concerns were also identified as of November 2024, subject to website and timing limitations. The report describes two demand letters to two of four employers DCWP had identified as noncompliant and records no Local Law 144 penalty. It does not establish that no penalty was imposed anywhere, and DCWP disputed parts of the report's inferences.⁵⁴⁷

The record supports closer attention to filing, inspection, and enforcement. It does not establish the causes of the observed problems, prove intentional deployer gaming, or measure whether the law improved hiring outcomes. Studies of published audits and implementers add evidence about disclosure, discretion, and practical constraints, with their own sampling limits.⁵⁴⁸ The additional mechanisms proposed here are credentialing and discipline, regulator-direct filing, false-attestation liability, and inspection of audit quality. They are responses to identifiable weaknesses, not demonstrated cures.

SOX separates several functions that an AI statute should also distinguish. Section 404 addresses management's assessment of internal control over financial reporting and, for issuers subject to subsection (b), an independent auditor's attestation. Important issuer exemptions apply. PCAOB's AS 2201 addresses an opinion on internal-control effectiveness, with reasonable assurance, which is not absolute.⁵⁴⁹ Senior-officer certifications arise separately under §§ 302 and 906. Section 906's criminal provisions address knowing or willful certification of a covered financial report that fails the statutory requirements; they are not a general false-audit offense.⁵⁵⁰ Appendix §§ 2.4(d) and 3.6 borrow the idea of responsible-officer attestation, while leaving state offense classification and mens rea to be completed at enactment.

The corrective-action problem is distinct from audit quality. *Ricci* requires a strong basis in evidence for the race-based employment action at issue there; a statistical disparity alone did not meet that standard. It did not adjudicate the constitutionality of the action or validate an AI audit safe harbor.⁵⁵¹ Appendix § 3.8 proposes a rebuttable presumption for state-law claims,

⁵⁴⁷ N.Y. State Comptroller, *Enforcement of Local Law 144 – Automated Employment Decision Tools*, Report 2024-N-6, at 15–19, 23, 32–35 (Dec. 2, 2025). Research leads were not a representative census of all users. Potential issue categories overlap; the November 2024 posting observation has use and website-timing limits. Agency-wide penalty totals do not establish a Local Law 144 penalty count.

⁵⁴⁸ Lucas Wright et al., *Null Compliance: NYC Local Law 144 and the Challenges of Algorithm Accountability*, 2024 ACM FAccT (391-employer search, with limits on noncompliance inference); Lara Groves et al., *Auditing Work: Exploring the New York City Algorithmic Bias Audit Regime*, 2024 ACM FAccT (seventeen practitioner interviews); Marissa Kumar Gerchick et al., *Auditing the Audits: Lessons for Algorithmic Accountability from Local Law 144's Bias Audits*, 2025 ACM FAccT 29 (forty-four reports containing 116 audits). These methods support bounded observations about audit practice, rather than a causal demonstration of outcomes or intent.

⁵⁴⁹ 15 U.S.C. § 7262(a)–(c); PCAOB, AS 2201, ¶¶ .01–.04, .13, app. A ¶ A5. Emerging-growth-company and nonaccelerated-filer exemptions limit auditor attestation. Management's report and the auditor's opinion on ICFR effectiveness are distinct, and reasonable assurance is not a guarantee.

⁵⁵⁰ 15 U.S.C. § 7241 (SOX § 302); 18 U.S.C. § 1350(a)–(c) (SOX § 906). The latter distinguishes knowing and willful certification of a covered noncompliant report. PCAOB disciplinary powers under 15 U.S.C. § 7215 address specified firms and associated persons, not a general senior-officer disqualification rule for AI audits.

⁵⁵¹ *Ricci v. DeStefano*, 557 U.S. 557, 584–87 (2009); see *supra* Part IV.B. Statistical disparity, potential

conditioned on methodology, independence, a qualifying finding, proportionate correction, and a documented chain. It expressly leaves federal Title VII unchanged. Satisfying a process requirement is not, by itself, a defense to federal discrimination liability. Nor does a failure to audit automatically establish disparate-impact liability.

The audit provisions combine content requirements under § 3.4, monitoring and incident response under § 3.4A, tiered audit intervals and independence rules under § 3.5, attestation under § 3.6, and direct filing and inspection under § 3.7. Its state-law presumption is a separate policy choice, with the conditions and rebuttal specified in § 3.8.⁵⁵² These mechanisms make claims of compliance more testable. Whether they produce reliable assurance depends on methods, access, auditor incentives, and regulator capacity.

The information gap remains the strongest objection. If a vendor withholds essential material, a record of unsuccessful requests may explain inadequate assurance without making the deployment safe. Procurement can help: a public buyer can require supporting material as a condition of the contract. Some evidence may be reusable for another customer, but validation and updating can remain deployment-specific. CMMC's express subcontract clauses illustrate contractual flow-down; a public contract does not automatically bind unrelated private sales.⁵⁵³

The implementing rules should make the required efforts to obtain information testable: a specific written request, escalation after refusal, and documented consideration of available alternatives are this Article's recommended minimum.⁵⁵⁴ Those details are not already enacted in § 3.4(a). The existing fallback requires recorded efforts and substitute assurance accepted under CAIO rules; it does not prove that any particular substitute is adequate. Publication of named adequacy findings would also require legal authority, confidentiality protections, and a fair process. The current redacted registry provisions should not be treated as a general public-naming power. Procurement leverage and disclosure can influence developer conduct, and their legal effect remains subject to the preemption limits developed in Part II.

ERISA supplies a particularly sharp limit. In *Gobeille*, the Court held that Vermont's reporting requirements intruded on central matters of plan administration. Parallel duties can be preempted even when they do not contradict federal requirements.⁵⁵⁵ *Rutledge* upheld a different kind of regulation, principally affecting costs. Its observation that the law reached PBMs whether or not they served ERISA plans addressed the reference-to inquiry; general applicability did not end the separate connection-with inquiry.⁵⁵⁶

disparate-impact liability, business necessity, and less discriminatory alternatives are distinct. The case decided the statutory question and did not decide constitutional sufficiency.

⁵⁵² *Infra* Appendix § 3.8. Conditional favorable treatment has different forms: DOJ Antitrust Division's Leniency Program offers conditional nonprosecution, while 17 C.F.R. § 240.10b5-1(b)–(c) supplies a specified affirmative defense. Neither analogue authorizes an AI state-law presumption or displaces federal discrimination standards.

⁵⁵³ 48 C.F.R. § 252.204-7021(f); *see supra* note 501. CMMC's contractual flow-down is tied to covered subcontract information and its exceptions. Reuse of AI evidence in unrelated private transactions depends on relevant content, currentness, and contractual access, rather than this federal clause alone.

⁵⁵⁴ The written-request, escalation, and alternatives sequence is this Article's implementing-rule recommendation. Appendix § 3.4(a) requires a record of efforts and substitute assurance under CAIO rules, without already specifying those steps or proving their sufficiency.

⁵⁵⁵ *Gobeille v. Liberty Mutual Insurance Co.*, 577 U.S. 312, 319–25 (2016); *see supra* notes 131–132. Reporting and recordkeeping can be central matters of ERISA plan administration; a different stated purpose or parallel requirement does not by itself avoid preemption.

⁵⁵⁶ *Rutledge v. Pharmaceutical Care Management Ass'n*, 592 U.S. 80, 87–92, 89 n.2 (2020); *see supra*

Calling a record an AI decision record therefore cannot separate it from plan administration when it records the adjudication of benefits. An adopting state must examine the duty's actual operation, including what records it compels and how it affects administration. The Appendix does not expressly draw a plan-administration boundary. A limit addressing that problem would reduce the information available for reviewing benefit-denial systems, and its precise terms require a substantive drafting decision. The present proposal cannot promise either complete access to those records or immunity from ERISA preemption.

D. Dispute and Remedy Pathways as the Affected Person's Recourse Architecture

The final primitive gives the affected person a route from information to challenge. Disclosure and explanation identify the system's role and the factors behind the decision. Correction addresses inaccurate inputs. Meaningful human review examines the record and can require re-decision when error or lack of support is found. Appendix § 4.6 distinguishes the right to request review from the conditional duty to re-decide; contesting an output alone does not guarantee a changed outcome.⁵⁵⁷

FCRA offers an instructive allocation of responsibilities among reporting agencies, users, and furnishers. It does not map those categories automatically onto AI developers, deployers, and vendors.⁵⁵⁸ ECOA supplies a more specific credit-decision model for reasons, under the applicable notice and request rules.⁵⁵⁹ The ADA supplies distinct accommodation and modification duties, with their statutory limits.⁵⁶⁰ The proposed general recourse right extends beyond those existing regimes. Enacting it requires legislative authority and careful coordination with the existing regimes.

FCRA coordination is particularly important. Section 1681t(b)(1) preempts specified subject matters, with different references and qualifications; subsection (B), for example, concerns the timing of action on disputed information under § 1681i. Appendix § 5.5A limits the subject matter of §§ 4.3–4.5 by reference to identified federal provisions and applies those duties to the extent not preempted. A state's additional explanation or correction duty still requires comparison with the actual federal subject matter. A savings clause does not settle that inquiry.

Enforcement presents a choice of institutions and incentives. Private actions can extend capacity beyond agency budgets, but bring litigation costs and coordination problems. Public enforce-

note 133. The general-applicability observation addresses reference to ERISA plans. The separate connection-with analysis sustained the particular cost regulation and did not decide an AI trail or the unargued overlap with *Gobeille*.

⁵⁵⁷ See *supra* Part III.D.6; *infra* Appendix § 4.2(b), § 4.6(a)–(c). Meaningful review requires the specified authority, independence, access, and written basis. Re-decision follows the finding in subsection (c), rather than every challenge.

⁵⁵⁸ 15 U.S.C. §§ 1681m(a), 1681i(a), 1681s-2(a)–(b), 1681t(a)–(b); see *supra* note 402. Reinvestigation timing, extensions, source responsibilities, and enforcement differ by provision. The proposed process action is not inferred from FCRA's existence; § 1681m(h)(8) and § 1681s-2(c)–(d) contain specified enforcement limits.

⁵⁵⁹ 15 U.S.C. § 1691(d)(1)–(6); 12 C.F.R. § 1002.9(a)–(b), official interpretations to paragraph (b)(2); see *supra* notes 403, 406. A private credit decision does not become an APA adjudication because the creditor must provide reasons.

⁵⁶⁰ 42 U.S.C. §§ 12112(b)(5), 12182(b)(2)(A)(ii)–(iii); 28 C.F.R. § 35.130(b)(7); see *supra* note 405. Disability-related accommodation or modification depends on the actor, entitlement, need, and applicable defenses.

ment can aggregate complaints and prioritize systemic conduct, but remains constrained by staffing, legal powers, and priorities.⁵⁶¹ This Article favors public enforcement for cross-cutting obligations alongside applicable private discrimination and sectoral claims. That preference does not exhaust the worked text: § 4.8(a)(1) preserves UDAP actions where applicable, and § 4.8(a)(4) and (d) create an independent action for violations of the specified process duties, with bracketed damages and timing provisions. The existence and scope of underlying causes of action remain governed by their own law.

Statutory damages do not remove federal standing requirements. *TransUnion* requires concrete harm for each class member recovering individual damages, while leaving open the timing of the standing showing at class certification. State courts apply their own justiciability rules, which require separate examination.⁵⁶² Claims against public deployers also need state-specific immunity and claims-presentation machinery. Neither an independent statutory right nor a damages schedule completes that work.

The Appendix makes the operational choices visible. Section 4.5 requires investigation within thirty days, followed by the specified correction, source notification, and re-decision using corrected data. Section 4.6 provides review within thirty days, shortened to seventy-two hours for specified urgent adverse actions, and a stay where lawful and feasible. Section 4.7 supplies administrative complaint routes. Section 4.9 preserves the FAA and channels covered private claims through enforceable arbitration arrangements, while protecting specified administrative routes and duties to the extent permitted by law. Public enforcement is valuable, but it is not an automatic equivalent to an individual's judicial or arbitral remedy. The recourse record can support a substantive claim; it cannot establish its elements or guarantee the forum in which it will be heard.

E. Three Worked Sector Applications

Applied to a hospital, an employer's hiring process, or an insurer's underwriting decision, the four primitives take different forms. Each setting identifies a concrete decision, the institution able to answer for it, the evidence review requires, and the legal limits on obtaining and using that evidence. The applications below adapt the proposed design to those differences. Existing licensing systems, employment rules, and supervisory instruments supply starting points; they do not establish that the entire design is already law.

1. Hospital AI: Clinical Responsibility and Federal Overlays

A hospital that introduces a diagnostic score makes choices the system's manufacturer cannot make for it: where the score enters the workflow, who may rely on it, when a clinician must investigate an apparent error, and when use should stop. A payer deciding whether to authorize treatment makes a different decision. Diagnostic tools, clinical decision support, coverage

⁵⁶¹ The choice between public and private enforcement is this Article's institutional recommendation. See *supra* Part III.B.3 for the distinct enforcement families and Part IV for the authority required by each channel. A family label does not establish enforcement intensity or sufficient resources in every jurisdiction.

⁵⁶² *TransUnion LLC v. Ramirez*, 594 U.S. 413 (2021) (majority slip op. at 7–15, 15 n.4, 20–27); *Spokeo, Inc. v. Robins*, 578 U.S. 330, 338–42 (2016). Each class member recovering individual damages needs Article III standing; *TransUnion* expressly reserves whether each member must demonstrate it before certification. State-court justiciability and the proposed cause of action require separate analysis.

review, and operational systems can all affect a patient, but they call for different duties and institutions.⁵⁶³

Professional regulation reaches the licensed practitioner under the applicable practice act. Physicians, nurses, physician assistants, and pharmacists do not all answer to one medical board or share one statutory scope of practice.⁵⁶⁴ Hospital licensing addresses the institution's governance, staffing, clinical privileges, and quality systems. The representative provisions developed in Part IV identify possible homes for supervision of clinical AI. Imposing the six duties requires the relevant authority and implementation.⁵⁶⁵ Payer oversight is separate again: the regulator's powers depend on the plan, insurer, activity, and applicable federal limits. A state insurance department is not a universal supervisor of Medicare Advantage plans, self-funded employer plans, or every prior-authorization contractor.⁵⁶⁶

The distinction between product and practice helps identify the issue without deciding preemption. FDA marketing authorization and professional responsibility do different work, but a rule addressed to a hospital can still impose a requirement with respect to a device. Diagnostic software and clinical decision support also require attention to the particular device category and statutory exclusions; FDA, CMS, and ONC do not supply interchangeable device rules.⁵⁶⁷ A covered payer's use of an algorithm is another regulatory setting. California's SB 1120 places limits on utilization-review software and reserves specified medical-necessity determinations to a competent licensed professional. It is not a general right to physician review of every hospital AI output.⁵⁶⁸

The details of that example matter. The California provisions govern covered plans and insurers and the specified authorization, delay, modification, or denial of services; their requirements must be read within that coverage.⁵⁶⁹ Arizona's cited medical-review law supplies a related general comparator without expressly naming AI. Maryland expressly addresses algorithmic tools in specified utilization review, while Indiana's cited provision addresses medical-necessity downcoding with its own reviewer and information requirements.⁵⁷⁰ Colorado's enacted 2026

⁵⁶³ See *supra* Part IV.C.3; *supra* notes 485–490. Device marketing, hospital participation, coverage review, and health-IT certification regulate different functions.

⁵⁶⁴ See *supra* note 485. For professional misconduct, compare N.Y. Educ. Law § 6530(3)–(6), (28), (32), with Cal. Bus. & Prof. Code § 2234(b)–(d). Those provisions concern their covered practitioners and disciplinary grounds; the scope and regulator for another profession require its own practice act.

⁵⁶⁵ See *supra* note 486; Cal. Code Regs. tit. 22, §§ 70701(a), 70703(a)–(d), (g)–(h) (governing body and organized medical staff); 42 C.F.R. § 482.12(a), (c), (e) (hospital governing body, patient care, and contracted services). These institutional duties are distinct from the Article's proposed six-duty package.

⁵⁶⁶ See *supra* Part IV.C.2; *supra* notes 474–483. Compare 42 C.F.R. § 422.402 (Medicare Advantage standards supersede state law except state licensing and solvency laws), with 29 U.S.C. § 1144(a), (b)(2) (ERISA preemption, insurance saving, and deemer provisions); *FMC Corp. v. Holliday*, 498 U.S. 52, 61–65 (1990). The relevant payer, duty, and federal provision determine the boundary.

⁵⁶⁷ 21 U.S.C. §§ 360k(a), 396; *Riegel v. Medtronic, Inc.*, 552 U.S. 312, 321–25, 327–30 (2008); see *supra* notes 489–490. Section 396 preserves specified practitioner use of a legally marketed device while retaining federal restrictions on sale, distribution, labeling, and promotion. Marketing pathway and applicable federal requirements remain material to preemption.

⁵⁶⁸ 2024 Cal. Stat. ch. 879 (SB 1120), §§ 1–2; Cal. Health & Safety Code § 1367.01(k); Cal. Ins. Code § 10123.135(j); see *supra* notes 297, 487.

⁵⁶⁹ Cal. Health & Safety Code § 1367.01(a), (e), (k)(1)–(7); Cal. Ins. Code § 10123.135(a), (e), (j)(1)–(4). The provisions distinguish individual clinical information, provider decisionmaking, inspection, and software limits. Section 1367.01(k)(7) conditions application to Medi-Cal managed care on necessary federal approval and preservation of federal financial participation.

⁵⁷⁰ 2025 Ariz. Sess. Laws ch. 165, §§ 1–3 (HB 2175) (individual medical-director review of specified

successor, generally commencing in 2027, likewise requires attention to its healthcare and device exclusions and the notices and information duties it retains, including those concerning patient financial assistance. These instruments allocate responsibility differently; they do not establish a uniform healthcare-AI regime.

For a hospital adopting a clinical tool, the proposed documentation duty should begin with the deployment: intended use, available validation, relevant patient populations, limitations, responsible clinicians, and the circumstances requiring escalation. Assessment should examine how the score changes care, including foreseeable errors and alternatives. Auditability should preserve enough evidence to examine a disputed decision; monitoring and incident response should connect what happens in use to correction, reassessment, or withdrawal. Where training or validation material remains with a supplier, the information limits developed in III.D and V.C still apply. Recording unsuccessful requests makes the gap visible without proving the system suitable for care.

Those duties need an authorized enforcement route. A professional board can investigate conduct within its disciplinary jurisdiction; a complaint does not itself establish malpractice or compel a particular result.⁵⁷¹ A hospital regulator can use the applicable inspection, deficiency, and licensing machinery, with the relevant procedural safeguards.⁵⁷² Insurance supervision, consumer protection, and civil-rights enforcement may supply additional routes, subject to their actor, coverage, remedy, and preemption limits.⁵⁷³ A patient action requires an available cause of action and its elements. Private Title VI enforcement requires the relevant funding relationship and intentional discrimination; disability claims require the applicable actor, duty, and remedial predicates. Discriminatory effects alone do not create every private damages claim. A licensing power, a funding condition, and a private damages claim should not be treated as equivalent ways to obtain relief.

Patient recourse gives the evidence a decision-specific purpose. The proposal would explain the tool's role and principal reasons, permit challenges to inaccurate inputs and relevant external sources, and provide meaningful review by a person able to examine and change the decision. Access to a medical record does not confer control over the supplier's training corpus. Under Appendix § 4.6, review is available on request; re-decision follows identified error or a conclusion unsupported on the whole record. Urgent review, a lawful and feasible stay, correction, and available relief have separate conditions.⁵⁷⁴ The applicable claims, forum, confidentiality rules,

medical-necessity denials; effective after June 30, 2026); 2025 Md. Laws ch. 747, § 1 (HB 820), adding Md. Code Ann., Ins. § 15-10B-05.1(b)–(d) (covered carriers and contracted pharmacy benefits managers or private review agents); Ind. H.E.A. 1271, § 6 (2026), adding Ind. Code ch. 27-1-52, particularly §§ 6, 9, 11–14 (medical-necessity downcoding, reviewer and medical-record requirements, provider submissions, and AI disclosure). Indiana's § 9 reviewer need not invariably be a physician. For Colorado's successor framework, see S.B. 26-189, § 1, enacting Colo. Rev. Stat. §§ 6-1-1708(3)–(5), -1709, and § 5 (2026): the healthcare exclusions retain an employment exception, a provider-location condition, general notice, and specified financial-assistance information duties; the new framework generally applies January 1, 2027.

⁵⁷¹ See *supra* notes 485, 564. N.Y. Educ. Law § 6530(3)–(6), (28), (32) distinguishes disciplinary grounds and relevant record obligations. Investigation and discipline require the applicable jurisdiction, ground, and procedure.

⁵⁷² See *supra* notes 486, 565. The cited hospital governance and participation provisions do not themselves establish every proposed assessment interval, vendor-access right, or private remedy.

⁵⁷³ See *supra* Parts IV.A, IV.B, IV.C.2; *supra* notes 474–483. For the private Title VI limitation, see *Alexander v. Sandoval*, 532 U.S. 275, 278, 284–86 (2001). Disability duties require their distinct predicates and remedies. See 29 U.S.C. § 794(a); 42 U.S.C. §§ 12132, 12182(a)–(b), 12188. The availability of a claim against a particular provider cannot be inferred from an audit disparity alone.

⁵⁷⁴ See *supra* Part III.D.6; *supra* notes 402–412, 557–560; *infra* Appendix §§ 4.5–4.9. For medical-record

and public-entity protections still govern whether and how the patient can enforce those rights. The central federal objection must be answered at the level of the actual duty. *Riegel* requires comparison with applicable federal device requirements; a general state form or a hospital addressee is not enough. *Buckman* addresses the separate problem of using state litigation to enforce duties owed to the FDA. *Lohr* and *Riegel* leave room for genuinely parallel state duties within their holdings, while the existence of a private state claim remains a further question.⁵⁷⁵ Neither an assessment nor a trail becomes parallel to federal law simply because it concerns compliance. The relevant actor, device, marketing pathway, command, remedy, and operational effect must be compared.⁵⁷⁶

Cross-border care adds another boundary. A clinician providing telehealth must use the licensure or other lawful pathway applicable to the patient and jurisdiction; an expedited compact pathway does not itself supply a single national license.⁵⁷⁷ A hospital system operating across states must likewise identify the applicable institutional duties. Its compliance program can organize those differences, but cannot resolve a federal conflict by treating it as an administrative variation.

The hospital application therefore makes the Article's method concrete. Begin with the treatment or coverage decision, assign responsibility to the actor who can discharge it, secure evidence sufficient for the proposed review, and test the duty against the federal regime that actually governs. The opportunity for state supervision is substantial. That inquiry determines the reach of state supervision, including where a duty imposed on a deployer encounters federal product requirements.

2. Hiring AI: Testing Selection Practices and Making Review Usable

An employer using a ranking tool must decide which applicants advance, how the score affects human judgment, and how a rejected applicant can raise an error. Advertising, screening, interviews, assessments, and later promotion or dismissal decisions can present different risks. The selected state instruments discussed here show several ways to govern those uses; they do not establish a ranking of employment against every other domain.⁵⁷⁸

amendment, see 45 C.F.R. § 164.526(a)–(d) (designated record set, denial grounds, amendment and disagreement procedures). The covered entity coordinates relevant business-associate records; that responsibility does not establish control over every supplier dataset. Proposed review, correction, conditional re-decision, and relief have separate triggers.

⁵⁷⁵ *Riegel*, 552 U.S. at 321–25, 327–30; *Buckman Co. v. Plaintiffs' Legal Comm.*, 531 U.S. 341, 347–53 (2001); *Medtronic, Inc. v. Lohr*, 518 U.S. 470, 495, 500–02 (1996); see *supra* note 490. *Lohr* Part V is a majority opinion. *Riegel* declined the unpreserved parallel-claim argument in that case; its discussion does not decide the validity of a new hospital assessment requirement.

⁵⁷⁶ *Riegel*, 552 U.S. at 327–30; *Buckman*, 531 U.S. at 352–53; 21 U.S.C. § 337(a); see *supra* note 490. A duty outside § 360k(a)'s scope and a genuinely parallel duty are different arguments. *Riegel* did not accept or reject the FDA's direct-versus-incidental distinction at 330. *Buckman* rejects claims existing solely by virtue of federal disclosure requirements; chronology alone does not supply a general test for new state duties.

⁵⁷⁷ Interstate Medical Licensure Compact §§ 1, 2(d), (f), reproduced in Interstate Med. Licensure Compact Comm'n, *Commission Handbook* 29–30 (July 2025), https://imlcc.com/wp-content/uploads/2025/11/IMLCC_ComHandbook_July25.pdf. The Compact supplies a pathway to separate full member-state licenses and retains state disciplinary jurisdiction. Its patient-location rule does not decide every noncompact telehealth exception.

⁵⁷⁸ See *supra* Part III.C.1; *supra* notes 257–272. The selected instruments illustrate different employment-law mechanisms; they do not establish a comparative ranking or a complete enforcement

Employment-discrimination law supplies a substantive starting point. Title VII and state fair-employment laws reach covered actors and practices, with their own thresholds, procedures, defenses, and remedies.⁵⁷⁹ Labor agencies may have relevant recordkeeping or worker-protection powers, but the existence of a labor code does not give every agency jurisdiction over every hiring algorithm. Vendor regulation likewise depends on the vendor's role and the applicable law; selling software does not automatically make a firm a licensed employment agency.⁵⁸⁰ Consumer-protection claims require their own coverage and elements. These channels can complement one another without together supplying universal authority.

The assessment should follow the selection practice. A tool that filters candidates before an interview can exclude people whom a later, balanced hiring result never restores. *Teal* rejects a bottom-line defense to an identifiable employment barrier with the requisite discriminatory impact. It does not make each technical operation in a model a separate actionable practice.⁵⁸¹ Disability presents an additional inquiry: a video or assessment tool may screen out an applicant who needs an accommodation even where aggregate race or sex comparisons reveal no problem. The EEOC's 2022 assistance distinguished those questions; its guidance did not independently create liability.⁵⁸²

New York City's Local Law 144 offers a specific audit-and-disclosure design. Effective January 1, 2023, with enforcement beginning July 5, it conditions covered uses of an automated employment decision tool on a recent independent bias audit and a public summary, alongside specified notices. The law and rules define the tool, employment uses, independence, and audit calculations; they do not cover every use of software in every employment decision.⁵⁸³ Illinois's 2026 amendment addresses discriminatory employment uses of AI and notice under its Human Rights Act, a different form of protection.⁵⁸⁴ Maryland's 2024 automated-tool proposal supplies

history across sectors.

⁵⁷⁹ 42 U.S.C. §§ 2000e(b)–(d), 2000e-2(a)–(d), (k), 2000e-5, 2000e-7; *see supra* notes 159, 430–432. Employer, agency, and labor-organization coverage, state deferral, prohibited practices, and available procedures require their separate predicates.

⁵⁸⁰ *See supra* Part IV.B; *supra* notes 430–432. Employment-discrimination jurisdiction, labor administration, employment-agency regulation, and consumer protection are distinct channels. An assertion of authority over a particular vendor or hiring practice requires the applicable statute and actor; a citation to an entire labor code does not establish it.

⁵⁸¹ *Connecticut v. Teal*, 457 U.S. 440, 445–48, 452–56 (1982); 42 U.S.C. § 2000e-2(k)(1)(A)–(B); *see supra* notes 444–447; *infra* note 588. The statute generally requires identification and causation by a particular practice, with an exception for inseparable elements of a decisionmaking process. A technical stage is not necessarily such a practice.

⁵⁸² EEOC, *The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees*, Questions 6–12, 14–16 (May 12, 2022), https://data.aclum.org/storage/2025/01/EEOC_www_eeoc_gov_laws_guidance_americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence.pdf (reproduction). The historical, nonbinding assistance distinguishes disability screen-out and accommodation from Title VII group-impact analysis. *See* 42 U.S.C. § 12112(a), (b)(5)–(7); *supra* notes 268, 560.

⁵⁸³ N.Y.C. Local Law No. 144 (2021), §§ 1–2, codified at N.Y.C. Admin. Code §§ 20-870–20-874; 6 R.C.N.Y. §§ 5-300–5-304; N.Y.C. Dep't of Consumer & Worker Prot., *Automated Employment Decision Tools: Frequently Asked Questions* 1 (June 29, 2023) (enforcement date); *see supra* notes 546–548. The audit calculations, auditor-independence definition, historical/test-data provisions, and publication requirements already form part of the rules.

⁵⁸⁴ 2024 Ill. Laws P.A. 103-804 (HB 3773), effective Jan. 1, 2026; 775 Ill. Comp. Stat. 5/2-101, /2-102(L); *see supra* note 197. The amendment addresses the specified employment uses, prohibited effects and zip-code proxies, and notice, with implementing rulemaking. Its actor definitions and coverage differ from Title VII's.

an additional design comparator, distinct from an enacted employment duty.⁵⁸⁵ California's employment regulations and New Jersey's disparate-impact rules supply further routes through existing civil-rights law, subject to their actual scope.⁵⁸⁶

The proposed process duties would make an employer's use of a tool open to examination. Documentation should identify the decision, relevant inputs, available validation, limitations, and the role of the score in advancement. The vendor's account of its product is useful only to the extent it addresses the employer's actual use and population.⁵⁸⁷ Assessment should identify practices capable of excluding protected groups, examine the pertinent evidence, and consider lawful ways to reduce the risk. It should distinguish a technical defect, an inaccurate record, a disability accommodation, and a statistical disparity; each can call for different action.⁵⁸⁸

Auditability then preserves the evidence needed to test those conclusions. Monitoring should examine changes in use, inputs, and outcomes, with material findings feeding correction and reassessment. An independent audit supplies another form of scrutiny under specified methods and access rules. The worked statute uses risk tiers and covered decision volumes, with exemptions and exceptions; it does not place every small employer on a documentation-only regime or automatically impose annual audits on every medium or large employer. Duties and intervals must follow the worked text and implementing choices; employer-size labels alone cannot determine them.

Enforcement must connect that evidence to an institution with power to act. State civil-rights agencies and the EEOC enforce their respective laws; labor, licensing, and consumer-protection bodies need a relevant source of jurisdiction.⁵⁸⁹ The EEOC's 2022 assistance illustrates disability concerns without replacing the governing statutes, regulations, and available claims.⁵⁹⁰ A private claim likewise depends on the particular cause of action, procedural prerequisites, proof, and available relief. Maintaining an audit trail can assist that work without establishing discrimination or defeating a defense by itself.

Applicant recourse should be usable before a dispute becomes litigation. The proposal would

⁵⁸⁵ Md. H.B. 1255, 2024 Reg. Sess. (2024); Md. Gen. Assemb., *HB1255: Labor and Employment—Automated Employment Decision Tools—Prohibition*, bill history and synopsis (last updated Apr. 15, 2024), <https://mgaleg.maryland.gov/mgaweb/Legislation/Details/hb1255?ys=2024RS>. This is a dated proposal comparator. The displayed history records introduction and a hearing; it does not itself establish the proposal's final disposition.

⁵⁸⁶ Cal. Code Regs. tit. 2, §§ 11008, 11008.1, 11009(f), 11013 (2025 ADS amendments effective Oct. 1, 2025); N.J. Admin. Code §§ 13:16-1.1–1.3, -2.1–2.4, -3.1–3.3 (effective Dec. 15, 2025); *see supra* notes 264–265. California's agent and employment-agency definitions can reach persons performing the specified employer functions through an automated system. Its records provision concerns records created or received; testing evidence bears on claims and defenses without establishing an automatic defense.

⁵⁸⁷ *See supra* Part III.D; *supra* notes 365–370, 535–536; *infra* Appendix §§ 2.3, 3.3, 3.5. Deployment-specific assessment and the proposal's intervals are distinct from possession of a vendor's complete development record or from employer headcount alone.

⁵⁸⁸ *Connecticut v. Teal*, 457 U.S. 440, 445–48, 453 n.12 (1982); 42 U.S.C. § 2000e-2(k); *see supra* notes 444–447. The Article's assessment recommendation seeks evidence before harm occurs; *Teal* does not itself require an assessment of every model operation or supply the complete present statutory burden structure.

⁵⁸⁹ *See supra* notes 579, 584–586; 42 U.S.C. §§ 2000e-5, 2000e-7. Concurrent institutions retain their own coverage, charge and procedural requirements, powers, and remedies.

⁵⁹⁰ EEOC, *supra* note 582, Questions 6–12, 14–16. The assistance describes disability concerns and possible practices under the ADA. It is an explanatory resource dated 2022; the governing obligations arise under the applicable law.

provide intelligible notice and reasons, a route to correct inaccurate inputs or seek correction from an external source, and review by a person with the authority and information needed to reconsider the decision. Its proposed general rights extend beyond the particular notices required by Local Law 144 or sectoral statutes. Review and re-decision remain distinct: a challenge triggers the prescribed review; identified error or lack of support on the whole record triggers re-decision under § 4.6(c). The correction duty has its own terms and federal coordination limits.⁵⁹¹ Those rights do not guarantee access to every item in a vendor's training data or an applicant's preferred forum.

Substantive discrimination law continues to do independent work. The *Griggs* and *Albemarle* line, as codified in § 2000e-2(k), supplies burdens concerning discriminatory impact, job relatedness and business necessity, and less discriminatory alternatives.⁵⁹² *Teal* prevents aggregate success from excusing an unlawful intermediate barrier.⁵⁹³ The Uniform Guidelines' four-fifths convention is a screening aid with qualifications for statistical and practical significance and small numbers; it is not a dispositive legal threshold. A method that stops at a single ratio may therefore miss matters the governing standard requires the employer to examine.⁵⁹⁴

Corrective action raises a separate concern. *Ricci* applied the strong-basis-in-evidence standard to the race-based action before it. It neither required an AI audit safe harbor nor demonstrated that one is necessary to make an audit mandate work. Appendix § 3.8 proposes a rebuttable state-law presumption with specified conditions and expressly leaves federal Title VII standards unchanged.⁵⁹⁵ An employer cannot infer a federal defense from completion of the state process, and the Comptroller's audit does not establish that fear of disparate-treatment liability caused the problems it observed.

The hiring application gains practical value from that implementation record. Local Law 144 already requires independence and methodology; the Comptroller's selected observations concern how the regime was used and enforced. Credentialing, direct filing, inspection, and attestation are this Article's proposed responses, with their own costs and limits. The record supports testing and supervising those mechanisms. It does not prove that the proposed additions will improve hiring outcomes or that an audit can stand in for substantive employment rights.

⁵⁹¹ See *supra* Part III.D.6; *supra* notes 402–412, 557–560; *infra* Appendix §§ 4.3–4.9, 5.5A. Compare 15 U.S.C. §§ 1681b(b), 1681m(a), 1681i(a), with 42 U.S.C. § 12112(b)(5)–(7). Consumer-report procedures, disability accommodation, and the proposed general process rights have different actors and triggers; they do not supply universal control of training records or unconditional re-decision.

⁵⁹² *Griggs v. Duke Power Co.*, 401 U.S. 424, 431–33 (1971); *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 425–28 (1975); 42 U.S.C. § 2000e-2(k)(1)(A)–(C); see *supra* notes 433–435. The current statute specifies the burdens and alternative-practice rule; the historical cases should be read with that text.

⁵⁹³ *Teal*, 457 U.S. at 445–48, 452–56 & 453 n.12; see *supra* notes 444–447. A nondiscriminatory final result does not answer a challenge to the unlawful intermediate barrier. The decision does not prescribe a single significance threshold or dispense with causation and the governing burdens.

⁵⁹⁴ 29 C.F.R. § 1607.4(D); see *supra* note 269. Smaller differences may matter when statistically and practically significant or when the user's actions disproportionately discouraged applicants. Larger differences can fail to establish adverse impact with small, unreliable numbers or atypical pools. The provision permits longer-period and comparable-use evidence; four-fifths is not a dispositive liability test.

⁵⁹⁵ *Ricci v. DeStefano*, 557 U.S. 557, 582–87 (2009); see *supra* notes 448, 449, 451, 551–552; *infra* Appendix § 3.8. The proposed presumption requires its specified conditions, is rebuttable, and applies to state-law claims without altering federal Title VII standards. The Comptroller's Report 2024-N-6, *supra* notes 546–548, does not identify *Ricci* deterrence as the measured cause of its observations.

3. Insurance AI: Supervisory Coordination and Federal Limits

An insurer using a model to price a policy or decide a claim must answer for a different set of decisions from the employer or hospital. Underwriting concerns eligibility, risk classification, and price. Claims systems decide what the insurer will pay; customer-facing tools communicate with policyholders; internal models can affect reserves and solvency. The proposal should assign responsibility according to that activity and the information needed to examine it. Insurers and suppliers occupy different legal roles.⁵⁹⁶

McCarran-Ferguson supplies a distinctive statutory rule, not constitutional immunity. For a federal enactment that does not specifically relate to insurance, § 1012(b) asks whether its application would displace qualifying state law or impair that law's operation. For the regulatory route considered here, the state enactment must have that statutory purpose. The particular state law, activity, federal provision, and asserted interference matter. The fact that the defendant is an insurer does not answer those questions.⁵⁹⁷ ERISA and civil-rights coordination likewise require the actual duty and claim; a general sector label cannot settle the result.⁵⁹⁸

For underwriting, the central accountability question is whether the insurer can explain and justify the data and methods used to classify a person. External data and correlated variables can obscure an unfair-discrimination problem, and an assurance from a supplier cannot replace the insurer's own responsibilities.⁵⁹⁹ Claims review presents a related but distinct task: a proposed trail should permit examination of the basis for approval, modification, or denial, with the applicable health-coverage rules governing where relevant.⁶⁰⁰ Communications, fraud detection, and solvency models require their own risk and oversight analysis. An adopting state should choose methods for the particular function; one employment-audit metric cannot assess them all.

Colorado illustrates legislation translated into more specific supervisory duties. SB 21-169 and Regulation 10-1-1 govern covered insurers' uses of external consumer data and of models or algorithms relying on those data, with line-specific scope, exceptions, governance requirements, and dates. The rule's life, private-passenger automobile, and health-plan applications are not identical.⁶⁰¹ The NAIC Model Bulletin supplies a different instrument: a model for a risk-

⁵⁹⁶ See *supra* Part IV.C.2; *supra* notes 474–483. Underwriting, authorization, claims, communications, and internal risk management raise different questions about actors and applicable powers.

⁵⁹⁷ 15 U.S.C. § 1012(b); *Humana Inc. v. Forsyth*, 525 U.S. 299, 309–14 (1999); *U.S. Dep't of the Treasury v. Fabe*, 508 U.S. 491, 500–09 (1993); see *supra* note 283. *Fabe* distinguishes the first clause's state-law-purpose inquiry from the antitrust exemption's narrower treatment of the business of insurance. *Humana* assesses direct conflict, frustration of declared state policy, and interference with the state's administrative regime; insurer status alone establishes none of them.

⁵⁹⁸ 29 U.S.C. § 1144(a), (b)(2), (d); *FMC Corp. v. Holliday*, 498 U.S. 52, 61–65 (1990); *Shaw v. Delta Air Lines, Inc.*, 463 U.S. 85, 102–08 (1983); see *supra* note 455. Preservation of federal law does not categorically preserve every additional state civil-rights duty or remedy.

⁵⁹⁹ N.Y. Dep't of Fin. Servs., Insurance Circular Letter No. 7, §§ II.A–B, III.C–D (July 11, 2024); see *supra* notes 287–291. The circular's proxy analysis, validation, and vendor expectations concern specified uses; they do not supply a prevalence estimate or prove a particular model discriminatory.

⁶⁰⁰ See *supra* notes 568–570; Colo. Div. of Ins., Regulation 10-1-1, § 5.A.5 (provider responsibility for specified prior or concurrent health-authorizations decisions using covered external data or models). Authorization and post-loss claims handling are distinct functions.

⁶⁰¹ Colo. Rev. Stat. § 10-3-1104.9; Colo. Div. of Ins., Regulation 10-1-1, §§ 2–6, 10–11 (effective Oct. 15, 2025). Health-plan external-data coverage excludes an individual's medical records; §§ 5.C and 6 set different life and automobile/health schedules. Colorado S.B. 26-189, § 3 (2026), adds notice-and-disclosure rulemaking authority. See *supra* notes 287–288, 480.

proportionate written AI Systems Program and related supervision under applicable insurance law. Its legal force follows the adopting jurisdiction and the instrument used there.⁶⁰² New York's Circular Letter No. 7 addresses specified external data and AI uses in underwriting and pricing, with express product and data exclusions. It should not be read as covering every insurer decision or every external record.⁶⁰³

The NAIC's examination project is another step in that translation. Its 2026 pilot plan named twelve participating states for testing during March–September and contemplated later consideration of the instrument. The associated working-group materials concern development of a common evaluation supplement; plans and circulating drafts are distinct from adoption, completed examinations, or published results.⁶⁰⁴ They show an effort to organize supervision, without establishing the effectiveness of the proposed tool.

The six duties would give that supervision a deployment-specific evidence base. Documentation should identify the use, available data provenance, validation, limitations, and accountable personnel. Assessment should examine foreseeable errors and unfair discrimination in the relevant insurance decision. A usable trail, monitoring, and incident response should connect observed problems to investigation and correction. Independent assessment needs methods suited to the line and decision; an annual hiring bias audit is not an automatic template for every insurer. Nor do admitted status, surplus-lines status, and a supplier relationship constitute three size-based tiers. Coverage, exemptions, risk tiers, and audit intervals follow the worked statute.

Enforcement depends on the available institution and remedy. An insurance department can use its examination and other statutory powers within their jurisdiction. A vendor's information may be reached through the insurer's duties or applicable contractual and examination provisions, without making every technology contractor subject to the full insurer regime.⁶⁰⁵ Consumer-protection and civil-rights routes require separate coverage. An insurance score's use of credit information does not alone establish an ECOA credit transaction; Title VI requires the relevant federal-funding predicate. Dwelling insurance can also engage the Fair Housing Act, subject to its coverage and the challenged practice. The availability of a federal or state civil-rights claim cannot be inferred from the word insurance, or from a general assertion that such claims escape reverse preemption.⁶⁰⁶

⁶⁰² Nat'l Ass'n of Ins. Comm'rs, *Model Bulletin: Use of Artificial Intelligence Systems by Insurers* 1–9 (Dec. 4, 2023), particularly §§ 3–4 (written AI Systems Program, third parties, and regulatory oversight); see *supra* notes 289–290, 477–478. The model uses state-law authority placeholders and risk-proportionate expectations; an NAIC model is not itself a uniform enacted obligation.

⁶⁰³ N.Y. Dep't of Fin. Servs., Insurance Circular Letter No. 7, §§ I–IV & nn.3–4 (July 11, 2024); see *supra* notes 291, 479. The circular addresses underwriting and pricing and excludes Child Health Plus, the Essential Plan, and Medicaid managed care. Its ECDIS definition excludes specified medical-information exchange, motor-vehicle, prescription, and criminal-history data. Those exclusions and its lawful disclosure limits matter to any comparison.

⁶⁰⁴ Nat'l Ass'n of Ins. Comm'rs, *AI Evaluation Pilot Project Summary* 1–3 (2026), https://content.naic.org/sites/default/files/inline-files/Pilot%20Project%20Summary_1.pdf; see *supra* notes 292, 380, 481. The plan specifies twelve states, March–September testing, adaptation to jurisdictional needs, risk-proportionate review, and later consideration of adoption. The Big Data and Artificial Intelligence Working Group's 2026 materials identify an AI Risk Evaluation Supplement v5.0 exposure and October 8 discussion. Those materials establish planned work and draft development, without establishing completed results or uniform legal adoption.

⁶⁰⁵ Regulation 10-1-1, §§ 5.B–C, 6, 9; Circular Letter No. 7, § III.D; NAIC Model Bulletin § 4; see *supra* notes 477–482. The insurer's vendor responsibilities and the supervisor's statutory access need their actual scope; they do not automatically license every supplier as an insurer.

⁶⁰⁶ 15 U.S.C. §§ 1691(a), 1691a(b), (d)–(e) (credit, applicant, and creditor predicates); 42 U.S.C. § 2000d

Policyholder recourse should connect the challenged decision to accessible information. The proposal would supply notice and reasons, investigate identified inaccuracies in inputs or external sources, and permit meaningful review by a person with access and authority to reconsider. An insurer that does not control a record may need to seek correction from its source. Existing FCRA obligations turn on the statutory actor, information, and adverse-action predicates. The broader Appendix rights have their own federal coordination limits, and § 4.6(c) conditions re-decision on the review's findings.⁶⁰⁷ A complaint route and a private action also depend on their separate rules; neither is a general guarantee of relief.

Cross-state coordination remains a useful part of the design. The NAIC's October 8, 2026 implementation reference list lists twenty-seven states alongside Puerto Rico and the District of Columbia, with California, Colorado, New York, and Texas identified separately for other insurance-AI action. That document's categories and qualifications matter: the list is not a finding that every jurisdiction adopted every model element or made it legally effective in the same way.⁶⁰⁸ A common supervisory vocabulary can reduce repeated work and support comparison across departments. Whether the resulting practice improves decisions requires evidence beyond a listing, a model instrument, or a pilot plan.⁶⁰⁹

Compliance costs also require evidence. The Article proposes differentiated duties and the Appendix supplies tiering and exemptions with exceptions. Neither establishes the first-year or recurring cost for a particular insurer, employer, or hospital. Those estimates require a defined duty, population, method, and administrative setting; analogy to another compliance regime cannot supply the amount.⁶¹⁰

Health-benefit administration brings a particularly important federal limit. ERISA's saving and deemer provisions distinguish regulation of insurance from treating a self-funded plan as an insurer. An insurer or administrator's participation does not make every compelled record an underwriting record: a trail of benefit adjudication can concern plan administration under *Gobeille*. A civil-rights label does not itself resolve preemption or establish an available state action.⁶¹¹ Other federal statutes likewise require their own analysis, including the par-

(federal-assistance predicate); 24 C.F.R. § 100.70(d)(4) (specified discriminatory refusals of property or hazard insurance for dwellings); *Humana*, 525 U.S. at 309–14. See *supra* Parts IV.B, IV.C.2. Neither credit-based insurance scoring nor an insurer defendant alone establishes a particular civil-rights claim or its freedom from reverse preemption.

⁶⁰⁷ See *supra* notes 557–562; 15 U.S.C. §§ 1681m(a), 1681i(a), 1681t; *infra* Appendix §§ 4.5–4.9, 5.5A. An adverse-action user's notice and a consumer reporting agency's reinvestigation are distinct duties. The Appendix separately conditions re-decision and preserves federal coordination limits; access to an insurer's decision record does not imply control of every external source.

⁶⁰⁸ Nat'l Ass'n of Ins. Comm'rs, *Implementation of NAIC Model Bulletin on Use of Artificial Intelligence Systems by Insurers* 1–2 (Oct. 8, 2026), <https://content.naic.org/sites/default/files/legal-adoption-map-ai-model-bulletin.pdf>. The reference list's twenty-nine entries comprise twenty-seven states, the District of Columbia, and Puerto Rico. California, Colorado, New York, and Texas appear separately for insurance-specific regulation or guidance. The map disclaims a determination that a jurisdiction adopted every model element or met accreditation requirements; the count concerns its own listed categories.

⁶⁰⁹ See *supra* note 604. A shared instrument and planned pilot can support coordination; the cited development record does not measure examination quality, compliance, or policyholder outcomes.

⁶¹⁰ The Article supplies no per-deployer cost estimate. Its proposed duties and Appendix §§ 2.3, 3.5, 5.7 identify relevant work and allocation choices, without measuring affordability or adequacy. CMMC, NIST AI RMF, federal model-risk guidance, and the NAIC model concern different populations and functions; any empirical comparison would need to justify their relevance and define implementation costs.

⁶¹¹ See *supra* note 455; 29 U.S.C. § 1144(a), (b)(2), (d); *FMC*, 498 U.S. at 61–65; *Gobeille v. Liberty Mut.*

ticular health-coverage provision and any express or conflict-preemption rule that governs it. McCarran-Ferguson does not supply a universal exception from those inquiries.⁶¹²

The insurance application thus combines existing supervisory institutions with a proposal for more usable evidence and recourse. The distinctive opportunity is coordination through actual insurance powers and shared instruments. The limits follow the same discipline as in hospital and hiring decisions: identify the actor, command, information, enforcement route, and federal constraint before concluding that the arrangement can operate as designed.

F. Enactment and Enforcement Capacity

Each primitive creates work as well as authority. Someone must review procurement evidence, receive and examine assessments, inspect audit quality, investigate incidents, and provide recourse. A legislature should assign those functions together with the powers, staff, technical support, information systems, and funding needed to perform them. Naming an existing regulator can make use of experience and infrastructure. It also adds demands to an institution with an existing workload.⁶¹³

The New York City record shows why those demands belong in the design. In the Comptroller's 2023–2025 audit, nine of twelve test calls were not connected to DCWP. Its comparison of posted audit material identified three potential independence problems, five calculation problems, and nine data-explanation problems across the twenty-two employers drawn from a thirty-two-company review. Those categories can overlap; their sum is not seventeen independently established violations or a detection rate. The record identifies failures worth investigating without isolating their cause.⁶¹⁴

DCWP's response made resources part of the dispute. The two attorneys assigned primary responsibility for Local Law 144 had other duties, and the agency reported no additional resources dedicated to the law. It disputed aspects of the auditors' methods and conclusions and defended its allocation of effort. The Comptroller recommended improvements to intake, investigation, tracking, training, and technical assistance, including use of existing OTI resources. Neither the recommendations nor the response measures what additional funding would have achieved.⁶¹⁵ Better specification and adequate capacity should therefore be treated as complementary design requirements. This record cannot establish how much of either would have

Ins. Co., 577 U.S. 312, 320–26 (2016); *Pilot Life Ins. Co. v. Dedeaux*, 481 U.S. 41, 52–57 (1987); *Aetna Health Inc. v. Davila*, 542 U.S. 200, 208–14 (2004). Plan administration, recordkeeping, and the federal remedial scheme require separate attention. *Shaw*, 463 U.S. at 102–08, preserves Title VII's operation within its holding, without creating a general state civil-rights exception.

⁶¹² See *supra* Parts II.B, IV.B, IV.C.2–3; *supra* notes 455, 488–490, 611. Federal coverage and administration provisions have their own scope and preemption rules. The label civil rights, the use of credit-related data, or a state insurance channel does not alone establish coverage, a cause of action, or the validity of a proposed state duty.

⁶¹³ *Infra* Appendix §§ 1.3, 2.6–2.7, 3.7–3.9, 4.7, 5.2, 5.7. The worked allocation of authority and costs leaves state-specific appropriation, claims, offense, and implementation machinery to be completed.

⁶¹⁴ N.Y. State Comptroller, *Enforcement of Local Law 144 – Automated Employment Decision Tools*, Report 2024-N-6, at 9–10, 16–18, 23 (Dec. 2, 2025); see *supra* notes 546–548. The comparison concerns potential issues in selected published material, with overlapping categories and timing limits, not a representative count of violations or an experiment on enforcement resources.

⁶¹⁵ *Id.* at 21–22, 26–28, 34–37 (recommendations, resource explanation, methodological objections, and responses). The agency and Comptroller disagreed about the adequacy of existing efforts and proposed additional work. The report does not estimate the effect of an added appropriation.

changed the result.

New York's assessment mechanism supplies a concrete funding analogy. Financial Services Law § 206 allocates specified departmental expenses among regulated insurance, banking, and virtual-currency businesses, with rules for attributable and shared costs, allocation bases, exceptions, and estimated payments reconciled to actual expenses. That is a legislated financing structure for a defined supervisor and regulated populations. It does not show that licensing agencies generally fund themselves, that an AI assessment should use the same formula, or that assessment funding makes authority legally durable.⁶¹⁶

The worked Appendix already makes a narrower financing choice. Section 5.7(a) places an annual, tier-apportioned assessment on Covered Deployers for the CAIO's costs of administering §§ 2 and 3. Subsection (b) leaves the deposit of credential fees, if charged, to the enacting state's treasury and appropriations law; subsection (c) separates the assessment from enforcement decisions. Subsection (d) leaves the required balance for the enacting state to supply.⁶¹⁷ The account itself and the appropriation and collection machinery still need completion. The stated cost allocation also does not fund every procurement, sectoral, court, or individual-recourse function elsewhere in the proposal. It should not be described as a complete budget for the stack.

A state adapting the text should estimate workload before fixing the assessment schedule: expected filings, review intensity, examinations, complaints, urgent decisions, and the technical expertise needed for each. Tier can help allocate costs, but its relationship to supervisory work needs examination. A credential fee can help pay for an assessor program without supplying the resources to supervise every deployment. Training, access to evidence, independence, coordination, and fair procedures remain necessary even where money is available. Assessment formulas and public reporting should also make the funding relationship open to scrutiny; separating a payment from an enforcement decision addresses one incentive without removing every risk of dependence on the regulated population.

Private enforcement can complement that capacity. The worked text preserves applicable substantive actions and proposes an independent process-right action, subject to the limits developed in Section D. Those routes depend on eligible claimants, usable information, available remedies, counsel, and the governing forum. They do not supply routine supervision of the entire covered population. Enactment should therefore connect public oversight and individual recourse without assuming that either will perform the other's work.

⁶¹⁶ N.Y. Fin. Serv. Law § 206(a)–(f), particularly subsection (a)'s distinct assessed populations, allocations, and exclusions and subsection (b)'s estimated payments and reconciliation. The comparison concerns that statutory mechanism, not all state licensing budgets.

⁶¹⁷ N.Y. Fin. Serv. Law § 206(a), (d-1) (virtual-currency operating and examination expenses); *infra* Appendix § 5.7(a)–(d). The Appendix limits the annual assessment to CAIO administration of §§ 2–3 and leaves the credential-fee deposit destination to state-law completion. The credential-fee deposit rule supplies no amount or complete fee schedule; the balance remains bracketed.

Part VI: Conclusion: State Governance as Necessary, and the Path Forward

A. States Should Build Duties They Can Administer

States choosing to govern AI should begin with the consequential decision and the organization responsible for using the system. A deployer duty can identify the evidence needed to explain a result, preserve records for examination and give an affected person a route to review. Its value depends on the institution able to administer it and the relief available when it is breached. A familiar authority or a completed process cannot supply those conditions by itself.

Constitutional constraints make the design inquiry more exact. Territorial connections, discriminatory treatment, interstate burdens and federal requirements must be evaluated for the duty that would actually operate. Process and deployment can be useful starting points, while remaining subject to the same inquiry. The recommendation is a design judgment supported by particular comparisons. Each proposed duty still requires its own constitutional analysis.

Regulatory resilience requires the same discipline. Consumer protection, civil rights, licensing and procurement offer different routes. A legislature should test whether a second route can reach the same conduct and provide the needed obligation or remedy before treating it as a substitute. Coalition cost describes a possible consequence of building across several institutional arrangements. It does not establish a constitutional barrier to broad federal preemption or a measured political cost. The portfolio is useful where its channels can do concrete legal work.

B. Accountability Depends on Access to Evidence and Power to Act

The throughline returns to the person subject to the decision. The organization using AI may control the records, the supplier relationship and the means of review. The affected person may control none of them. Disclosure matters when it identifies a decision that can be questioned; a record matters when someone can obtain and examine it; review matters when the reviewer can change the result. These connections determine whether process obligations redistribute practical power or merely produce an orderly account of what the organization has already done.

The deployer is an important place to attach that responsibility, without being the only possible regulated actor. The EU AI Act's deployer obligations sit alongside provider and other obligations, and its fundamental-rights assessment has a defined scope.⁶¹⁸ Selected state laws also distribute responsibilities across users, suppliers and public institutions. Their existence supplies antecedents and alternatives. It does not establish that every system has converged on a single design or that the allocation has improved outcomes.

The contribution is the connection among constitutional analysis, an institutional portfolio and worked statutory duties. The Appendix makes that connection available for examination: who is covered, what evidence is required, what happens when supplier information is unavailable, how review proceeds and which institution can impose a consequence. Its explanations also identify where operative choices and state-specific machinery remain unfinished. That visibility is part of the proposal's value. It permits a drafter to test a clause against an actual state, decision and federal overlay before adopting it.

⁶¹⁸ Regulation (EU) 2024/1689, arts. 26–27. The deployer duties and fundamental-rights assessment have distinct coverage and exclusions, alongside provider and other obligations. *See supra* note 187.

The method may be useful beyond AI wherever organizations use a technical system to make consequential decisions. That extension requires examination; the analysis establishes no universal historical sequence for general-purpose technologies. The constitutional and institutional questions must be answered again for the activity and duty at issue.

C. Enactment Requires Specific Choices About Authority, Evidence and Relief

The structural-federalism analysis leaves further constitutional work for particular laws. First Amendment, due-process and equal-protection questions can remain even where a duty has a defensible Commerce Clause or preemption position. State constitutions add their own requirements. The Colorado litigation illustrates why procedural posture and the grounds actually alleged must be kept distinct from a merits conclusion. Later federal legislation and implementation likewise require dated attention to the authority and duties then in question.⁶¹⁹

Frontier capabilities, compute and critical infrastructure also require their own institutional analysis. Federal energy, cybersecurity, standards and assurance programs address parts of that field.⁶²⁰ Their presence does not by itself establish field preemption, exclusive federal competence or the absence of state authority. A model evaluation and a hospital's use of that model pose different questions, while particular duties may interact. New York's RAISE framework illustrates an expressly developer-facing state choice whose coverage and dates must be assessed on its own terms.⁶²¹ The proposed design concentrates on consequential uses; it does not settle the validity or desirability of every other form of AI governance.

State legislators can put the method to work through a concrete sequence. Identify the decision and actor, the state's authority and territorial connection, and the harm the proposed duty would address. Compare the duty with actual federal requirements and account for interstate burdens. Specify the records the organization controls, what it must seek from suppliers, lawful withholding and the consequence of inadequate evidence. Then identify the reviewer, enforcement institution and available relief, with resources and transition rules adequate to the work.

Procurement can begin that process where the state controls a purchase or use. It can support contractual evidence, supplier accountability and experience with oversight. Extending duties to private organizations requires its own coverage, authority and implementation decisions;

⁶¹⁹ See *supra* note 6 (the dated *Weiser* order and federal pleading). First Amendment, substantive rights, state constitutional requirements and later federal legislation require their own analysis; their reservation defines the scope of this Article.

⁶²⁰ U.S. Dep't of Energy, *AI for Energy: Opportunities for a Modern Grid and Clean Energy Economy* (Apr. 2024); 6 U.S.C. § 681b (CIRCIA reporting and implementing-rule framework); NSA, CISA et al., *Principles for the Secure Integration of Artificial Intelligence in Operational Technology* (Dec. 3, 2025); NIST, *Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations*, NIST AI 100-2 E2023 (2024); MITRE, *Adversarial Threat Landscape for Artificial-Intelligence Systems (ATLAS)*; GAO, *Artificial Intelligence: Generative AI Training, Development, and Deployment Considerations*, GAO-25-107651 (Oct. 22, 2024). Statute, guidance, technical framework, knowledge base and descriptive research have different legal force. Their existence alone does not establish exclusive federal authority.

⁶²¹ 2025 N.Y. Laws ch. 699, as amended by 2026 N.Y. Laws ch. 96, including amended N.Y. Gen. Bus. Law § 1421(1)(a). The standards clause requires a description of incorporation of national standards, international standards and industry consensus practices; it does not establish that any particular federal standard has been incorporated as law. See *supra* note 12.

procurement does not confer a general power to regulate every supplier's activity. Similarly, the existing assessment provision in Appendix § 5.7 covers a defined part of CAIO administration. The account, collection, appropriation and work outside that scope still require completion. An enforcement plan must account for them before the statute is treated as ready for enactment.

Attorneys general and sectoral regulators should ask which existing powers can address a particular deceptive claim, discriminatory decision or licensed activity, and which actor and relief those powers reach. The NAIC materials and New York City's implementation record supply useful questions about institutional design, evidence and capacity. They do not provide an effectiveness certificate for the proposed statute. Public oversight and private relief must be connected through the actual rules governing access, standing, timing, arbitration, immunity and remedies.

The recommendation remains affirmative: states should place accountable organizations behind consequential uses of AI and give regulators and affected persons the means to test their decisions. The worked design supplies a way to specify that responsibility under constitutional constraint. Its promise is realized only when the legal duty, available evidence and power to act meet in the same decision environment.

Appendix A: Worked Statutory Text for the Four Governance Primitives

Preliminary Note

This Appendix supplies worked text for four governance primitives: procurement controls (§ 1), deployer impact assessment (§ 2), an audit trail (§ 3), and recourse (§ 4), followed by general provisions (§ 5). It is a template for adaptation to a state's existing law. Adoption requires more than filling in a title: definitions, implementing authority, appropriations, claims procedures and any sovereign-immunity waiver, offense classification and mens rea, and rules-first effective dates must fit the enacting state.⁶²²

The allocation of coverage matters. Section 1 reaches State agencies, authorities, and political subdivisions through procurement; its vendor-assessment and standards provisions are not the same as direct application of §§ 2 and 3 to government. Sections 2 and 3 reach private Covered Deployers under their definitions. Section 4 reaches covered public and private deployments. State Nexus attaches to the consequential decision and does not itself establish jurisdiction. Framework Standards serve the particular functions assigned to them: procurement eligibility under § 1.4 is different from the benchmark uses under § 2.2(e).

The sections support one another, but legal severability and functional completeness are different questions. Section 5.1 expresses independent-enactment intent and preserves applications that can operate after invalidation or preemption. Cross-references, access to evidence, staffing, and remedies still need examination if a legislature adopts only part of the stack.⁶²³ Bracketed amounts, periods, and state-law references are choices to complete, not established optimal defaults. Section 5.7 allocates specified administrative costs through assessments and fees, while

⁶²² See *supra* Part V.A–D. The worked text supplies a common design and state-specific completion points; it is not a claim that all necessary enactment machinery is already present.

⁶²³ See *supra* Part V.A–D; *infra* § 5.1. Independent-enactment intent does not resolve every cross-reference, resource, or operational dependency in a partial adoption.

leaving a balance to be supplied; it does not establish that the administering institutions will have adequate resources.

§ 1. Procurement-First Controls (Anchoring Statute)

§ 1.1 Findings and Purpose

The Legislature finds that: (a) artificial intelligence systems acquired and deployed by State agencies make decisions affecting State residents in regulated domains the State has traditionally governed under its police powers, including consumer protection, civil rights enforcement, sectoral licensing, and public benefits administration; (b) the market-participant doctrine of *Reeves, Inc. v. Stake*, 447 U.S. 429, 436-37 (1980), separates a State that takes part in a market from a State that regulates one and rests on the absence of any constitutional plan to limit a State's own freedom to operate in the market, so the State acts as a participant when it fixes the terms on which it buys AI systems for its own use; *United Haulers Ass'n v. Oneida-Herkimer Solid Waste Mgmt. Auth.*, 550 U.S. 330 (2007), and *Dep't of Revenue of Ky. v. Davis*, 553 U.S. 328 (2008), reach less far, holding that a law favoring the government itself while treating every private business alike is not discriminatory for Commerce Clause purposes and that courts should be slow to interfere with a traditional government function, and neither removes the State's regulation of private conduct from review; (c) Executive Order 14365 § 8(b)(iii) (Dec. 16, 2025) expressly excludes State government procurement and use of AI from the order's preemption recommendations; and (d) procurement-anchored AI governance can provide a foundation for developing deployer-centered governance through conditions connected to the State's own acquisition and use of AI; the preemption of a particular procurement condition depends on its operation and the applicable federal law, and use of spending power does not itself reduce the potential for conflict, as *Crosby v. National Foreign Trade Council*, 530 U.S. 363 (2000), explains. The purpose of this Act is to establish the procurement-anchored foundation, including the authority of the State Chief AI Officer, the substantive framework standards governing AI procurement eligibility, the pre-deployment impact assessment obligation, the vendor warranty terms, and the oversight architecture.⁶²⁴

§ 1.2 Definitions

- (a) "AI system" means a machine-based system that, for explicit or implicit objectives, infers from inputs it receives how to generate outputs that can influence physical or virtual environments, consistent with the National Institute of Standards and Technology AI Risk Management Framework (NIST AI RMF 1.0) and ISO/IEC 22989:2022.
- (b) "Deployer" means a State agency, State authority, or political subdivision of the State that deploys an AI system in furtherance of agency operations.
- (c) "Vendor" means a person or entity that supplies an AI system to a Deployer under a State procurement contract.
- (d) "State Chief AI Officer" or "State CAIO" means the officer appointed under § 1.3(a); "Office of the State Chief AI Officer" means the office established under § 1.3(a).

⁶²⁴ See *supra* Parts II.A.5, II.B.4 and IV.D for the distinct market-participant, public-function and preemption analyses. Executive Order 14365 § 8(b)(iii) concerns a recommended federal framework, not a self-executing exemption.

- (e) “NIST AI RMF” means the National Institute of Standards and Technology, *AI Risk Management Framework* (NIST AI 100-1, Jan. 26, 2023). The State CAIO may adopt a later NIST revision only by rule under the State Administrative Procedure Act, with public notice and comment and a transition period of not less than twelve months.
- (f) “ISO/IEC 42001” means *ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system*. The State CAIO may adopt a later edition only by rule under the State Administrative Procedure Act, with public notice and comment and a transition period of not less than twelve months, and shall maintain a lawful means of reference access to the incorporated standard for regulated parties [read-only access, per state convention and the publisher’s license].
- (g) “Pre-deployment impact assessment” means the document described in § 1.5 delivered by a Vendor to a Deployer before contract performance.⁶²⁵

§ 1.3 Authority of State Chief AI Officer

- (a) There is established within [the Office of the Governor or the Department of Administrative Services, per state convention] the Office of the State Chief AI Officer. The State CAIO shall be appointed by the Governor and shall serve at the pleasure of the Governor. The State CAIO shall report to [the Governor or the head of the Department of Administrative Services, per state convention].
- (b) The State CAIO shall have authority to: (1) issue binding procurement-eligibility standards for AI systems acquired by State agencies, consistent with § 1.4; (2) review pre-deployment impact assessments delivered under § 1.5; (3) exercise audit authority over Deployer use of AI systems acquired under State procurement contracts; (4) receive incident reports concerning AI systems acquired under State procurement contracts; (5) promulgate implementing regulations under the State Administrative Procedure Act; (6) coordinate with the State Chief Information Officer and with sectoral regulators concerning AI systems used by Deployers; and (7) receive and investigate administrative complaints filed under § 4.7(4) concerning compliance by Covered Deployers subject to §§ 2 and 3 of this Appendix, with authority to compel production of the filed impact assessment and Audit Trail, to issue written findings, and to refer violations to the Attorney General, the State Civil Rights Commission, or the sectoral regulator for enforcement.
- (c) The State CAIO shall not have authority to displace or supplant any sectoral regulator’s authority over AI systems used in the sector the regulator supervises.⁶²⁶
- (d) For political subdivisions, the standards issued under § 1.4 apply to procurements funded in whole or in part by State funds or made through State contract vehicles. For other procurements, a political subdivision shall either adopt the State CAIO’s standards or adopt standards no less protective. Nothing in this Act limits the authority of a home-rule subdivision beyond what [the state constitution] permits.

⁶²⁵ NIST AI RMF, *supra* note 350; ISO/IEC 42001:2023, *supra* note 351. A named framework edition is different from a state’s selected eligibility criteria and evidence requirements.

⁶²⁶ See *supra* Part IV.C; *infra* §§ 1.3(c), 5.4. Coordinating authority does not displace the sectoral jurisdiction preserved by the worked text.

§ 1.4 Substantive Framework Standards

- (a) An AI system is eligible for State procurement if and only if it satisfies the eligibility criteria the State CAIO specifies by rule. The rules shall map NIST AI RMF function-level outcomes and ISO/IEC 42001 clause requirements to objectively verifiable evidence items; define the risk-profile tiers of subsection (c) by published criteria; and specify pass/fail determination, cure periods, and appeal under the State Administrative Procedure Act. At a minimum the criteria shall address: (1) the NIST AI RMF GOVERN, MAP, MEASURE, and MANAGE function-level requirements applicable to the AI system's risk profile; (2) the ISO/IEC 42001 AI Management System requirements, including Clauses 6 (Planning), 8 (Operation), and 9 (Performance evaluation); and (3) any sector-specific federal guidance applicable to the AI deployment, including the FDA Software-as-a-Medical-Device framework for clinical AI and Federal Reserve / FDIC / OCC SR 26-2 for AI in covered banking operations within State-government banking relationships.
- (b) The State CAIO shall publish a schedule identifying the NIST AI RMF and ISO/IEC 42001 editions adopted by rule and in force for purposes of subsection (a).
- (c) The State CAIO may establish tiered procurement-eligibility standards calibrated to the AI system's risk profile and the deployer's decision context. The tiers are: Level 1 (Vendor self-attestation, lower-risk contracts); Level 2 (State CAIO documentation review, moderate-risk contracts); and Level 3 (third-party assessor certification, highest-risk contracts). The tiering architecture parallels the Cybersecurity Maturity Model Certification (CMMC) Level 1 through Level 3 architecture under 48 C.F.R. § 252.204-7021.⁶²⁷

§ 1.5 Pre-Deployment Impact Assessment

- (a) Before contract performance, the Vendor shall deliver to the Deployer a written pre-deployment impact assessment documenting: (1) the AI system's intended purpose and operational context; (2) the training data, including provenance, time-period coverage, and any documented gaps; (3) the validation methodology, including testing populations and performance metrics; (4) the AI system's known limitations and operational boundaries; (5) the AI system's foreseeable risks of adverse outcomes affecting State residents in regulated domains; and (6) the mitigation measures the Vendor and the Deployer will implement to address the foreseeable risks.
- (b) The pre-deployment impact assessment shall be reviewed by the State Chief Information Officer (or the agency CIO, where applicable) and by the State CAIO before the Deployer enters into the procurement contract. The State CAIO may require revisions to the pre-deployment impact assessment before contract execution.
- (c) The pre-deployment impact assessment shall be retained by the Deployer for the duration of the procurement contract plus five years, and shall be available for inspection by the State CAIO, the State Auditor, and any sectoral regulator with authority over the AI system's deployment context.⁶²⁸

⁶²⁷ 48 C.F.R. § 252.204-7021, definitions and paragraphs (d)–(f); *see supra* note 543. Federal Level 2 can use self or C3PAO assessment; Level 3 uses government DIBCAC assessment. The State's three proposed assurance modes are analogous tiering choices, not identical CMMC modes.

⁶²⁸ *See supra* Part V.B; *supra* notes 350–351, 365. Section 1.5(a)'s six elements are the worked procurement assessment schedule; subsection (c) supplies its contract-plus-five-year retention rule. No complete correspondence with every NIST outcome or ISO control is claimed.

§ 1.6 Vendor Warranty Terms and Remedies

- (a) Every State procurement contract for an AI system shall include a Vendor warranty and covenant that: (1) the AI system satisfies the substantive framework standards § 1.4 specifies; (2) the AI system operates consistent with the pre-deployment impact assessment § 1.5 documents; (3) the AI system has been tested and validated for the deployment contexts the pre-deployment impact assessment documents, using disclosed methodology and disclosed results, and performs within its documented limitations when operated in accordance with the Vendor's documented conditions of use; (4) the Vendor will notify the Deployer promptly of any post-delivery discovery materially affecting the documented limitations or the § 1.5 assessment; and (5) the Vendor will supply to the Deployer, on an ongoing basis, the documentation the Deployer requires to satisfy its obligations under this Act, including the § 1.5 impact-assessment and the § 1.7 audit and incident-reporting obligations.
- (b) The remedies for Vendor warranty breach include: (1) contract termination by the Deployer or the State CAIO; (2) a civil penalty of up to ten percent of the contract value or [\$500,000, per state convention], whichever is greater, recoverable by the Attorney General; (3) disqualification from future State procurement for a period of up to five years; and (4) any other remedy available under the State procurement code or under common-law warranty doctrine.
- (c) The Vendor warranty remedies are not exclusive and do not displace any affected person's right to recover under State or federal civil rights statutes, consumer protection statutes, or sectoral statutes for harms the AI system caused.⁶²⁹

§ 1.7 Oversight Architecture

- (a) The State CAIO shall maintain an audit program covering Deployer use of AI systems acquired under State procurement contracts. The audit program shall include risk-based audit selection, audit methodology consistent with auditing standards promulgated by the State CAIO, and audit-finding remediation tracking.
- (b) Deployers shall report to the State CAIO any incident involving an AI system acquired under a State procurement contract that: (1) materially affects the AI system's compliance with § 1.4 standards; (2) produces an outcome that may constitute a violation of State or federal civil rights, consumer protection, or sectoral statutes; or (3) involves a security breach affecting the AI system's training data, model weights, or output integrity. Incident reports shall be filed within seventy-two hours of the Deployer's discovery of the incident.
- (c) The State CAIO shall publish an annual report documenting: the number of pre-deployment impact assessments reviewed; the number of procurement contracts executed under § 1.6; the number of incidents reported under subsection (b); and the audit findings and remediation tracking under subsection (a).⁶³⁰

⁶²⁹ See *supra* Part IV.D; *supra* note 543 for CMMC status and affirmation. The Vendor warranty, state contract remedies, and ongoing documentation duties in § 1.6 are this Appendix's design, rather than remedies automatically supplied by CMMC.

⁶³⁰ 6 U.S.C. § 681b(a)(1), (7). The federal seventy-two-hour reporting design is an analogy; the federal statute ties operation to implementing final-rule dates. Section 1.7 states its own contractual recipients and trigger.

Commentary on § 1

Procurement gives a State a direct relationship with its supplier and a means to specify evidence, warranties, and remedies. The market-participant and public-function cases discussed in Parts II and IV provide distinct, bounded Commerce Clause arguments; neither establishes general immunity from preemption. Executive Order 14365’s recommended procurement carve-out is a direction about a proposed federal framework, not an enacted exemption. CMMC illustrates translation through defined contract requirements, assessment modes, and affirmations, without requiring third-party certification for every procurement.⁶³¹

Section 1.4’s risk tiers are this Article’s proposed allocation of assurance burdens. Lower-tier attestation can reduce compliance expense; higher-tier assessment can provide more scrutiny. Neither establishes the quality of a particular system. The vendor assessment under § 1.5 has six specified elements and is retained for the contract term plus five years. Evidence produced for one acquisition may help another, but private customers receive no automatic entitlement to it merely because a public buyer required it.⁶³²

§ 2. Deployer-Centered Impact Assessment

§ 2.1 Findings and Purpose

The Legislature finds that: (a) private-sector AI systems deployed in regulated domains (employment, housing, insurance, healthcare, education, public benefits, criminal justice) make decisions affecting State residents whose recourse depends on the AI system’s documentation, validation, and operational transparency; (b) the National Environmental Policy Act’s environmental impact statement architecture, established by 42 U.S.C. § 4332 and judicially anchored by *Robertson v. Methow Valley Citizens Council*, 490 U.S. 332 (1989), supplies a doctrinal analogue for process-based impact assessment without dictating substantive outcome; and (c) State authority over private-sector AI impact assessment operates within the State’s traditional police-power authority over consumer protection, civil rights enforcement, and sectoral licensing. The purpose of this Act is to require Deployers of AI systems in the regulated domains to conduct and document pre-deployment impact assessments and to make the assessments available to State enforcement authorities and to affected persons.⁶³³

§ 2.2 Definitions

The definitions of § 1.2 apply, except that “Deployer” for purposes of this Section, and “Covered Deployer” as that term is used elsewhere in this Act, mean any person or entity, other than a State agency, State authority, or political subdivision of the State, that deploys an AI system in a Covered Domain to make or materially influence a Consequential Decision. Additional definitions:

⁶³¹ See *supra* Part IV.D; *supra* notes 501, 543. Procurement eligibility, assessment, affirmation, and contractual flow-down depend on the specified instrument and scope; none establishes general preemption immunity.

⁶³² See *supra* Part V.A–B. Tiering can reduce some compliance burdens while demanding more assurance for high-consequence uses. Costs, coverage, and actual reliability remain separate questions.

⁶³³ 42 U.S.C. § 4332; *Robertson*, 490 U.S. at 349–53; see *supra* Part V.B. The federal agency process is a design analogy, not the source of a private deployer’s duty under this Appendix.

- (a) “Covered Domain” means employment, housing, insurance, healthcare, education, public benefits, criminal justice, and any other domain the State CAIO designates by rule under § 2.3(b), limited to domains in which Consequential Decisions materially affect access to housing-, livelihood-, health-, education-, or liberty-related services or opportunities.
- (b) “Consequential Decision” means a decision affecting access to, terms of, conditions of, or termination from any service, benefit, or opportunity within a Covered Domain.
- (c) “Affected Person” means a natural person whose Consequential Decision was made or materially influenced by an AI system subject to this Section or, for purposes of § 4, by a Deployer as defined in § 4.2.⁶³⁴
- (d) “Materially influence” means to supply an assessment, score, classification, recommendation, or other output that functions as a substantial factor in the making of a Consequential Decision. An AI system’s output is presumed to materially influence a Consequential Decision where no review with the attributes of Meaningful Human Review (§ 4.2(b)) intervenes between the output and the decision, or where the Deployer’s decision-makers adopt the output in more than [80, per state convention] percent of affected decisions; the presumption is rebuttable.
- (e) “Framework Standards,” as that term is used in this Section and in § 3, means the requirements identified in § 1.4(a)(1) and (a)(2), in the editions the schedule under § 1.4(b) identifies as in force, referenced as benchmarks for tiering, incident-trigger, and audit-finding purposes and not as independent conformity obligations of a Covered Deployer.

§ 2.3 Coverage and Tiering

- (a) This Section applies to any Deployer that uses an AI system to make or materially influence a Consequential Decision in a Covered Domain, except as exempted by subsection (c), and applies only where the Consequential Decision has a State Nexus. A Consequential Decision has a “State Nexus” if it concerns a natural person who resides in the State. The State Nexus limits which Consequential Decisions this Section reaches; it is not a test of the Deployer’s presence in the State, which the enacting State’s own jurisdictional law governs.
- (b) The State CAIO shall establish tiered impact-assessment requirements by rule, and may by rule designate additional Covered Domains meeting the criteria of § 2.2(a), with tiering calibrated to: (1) the Deployer’s size, measured by number of Consequential Decisions made annually within the State; (2) the Covered Domain’s substantive risk profile; and (3) the AI system’s documented risk profile under the Framework Standards.
- (c) The State CAIO may exempt by rule: (1) small Deployers, defined as Deployers making fewer than [200, per state convention] Consequential Decisions in the State annually; (2) AI systems used solely for back-office operations that do not materially influence Consequential Decisions; and (3) other categories the State CAIO determines do not warrant impact-assessment coverage consistent with the purposes of this Section.⁶³⁵

⁶³⁴ See *supra* Part III.A.3. The operative object is a consequential AI-mediated decision in the defined context; the definitions and exclusions here govern coverage.

⁶³⁵ The 200-Decision exemption is an author-proposed placeholder. See *supra* Part V.B and § 2.3(b)–(d) for tiering, exceptions, and calibration. It is not derived from an identified CCPA or Colorado numerical threshold.

- (d) No exemption under subsection (c) applies to Consequential Decisions concerning health-care treatment or coverage eligibility, pretrial release or sentencing, termination of public benefits, or eviction from housing. For those decisions the State CAIO may prescribe simplified compliance procedures for small Deployers but may not exempt the duty.

§ 2.4 Impact Assessment Content

The pre-deployment impact assessment under this Section shall document the six elements required for State-procurement impact assessments under § 1.5(a), and additionally:

- (a) The Deployer’s risk-management plan for the AI system, including the governance structures, monitoring protocols, and incident-response procedures the Deployer will operate;
- (b) The Deployer’s assessment of the AI system’s foreseeable disparate impact across protected classifications under State and federal civil rights statutes, including the methodology used and the assessment’s findings;
- (c) The Deployer’s plan for the Affected Person recourse architecture § 4 of this Appendix establishes; and
- (d) The Deployer’s attestation, signed by a senior official with operational authority over the AI system’s deployment, that the impact assessment is accurate to the best of the senior official’s knowledge.⁶³⁶

§ 2.5 Trigger Events

The impact assessment under this Section shall be: (a) conducted before the Deployer’s first deployment of the AI system in a Consequential Decision context; (b) updated within ninety days after any material modification to the AI system’s training data, model architecture, or operational deployment context; and (c) reviewed not less than annually for AI systems in the highest tier of risk classification under § 2.3(b), and not less than every three years for AI systems in lower tiers.⁶³⁷

§ 2.6 Public Filing and Trade-Secret Protection

- (a) The Deployer shall file the impact assessment with the State CAIO and with any sectoral regulator with authority over the Covered Domain in which the AI system operates.
- (b) The Deployer may designate specific portions of the impact assessment as containing trade secrets or commercial information that the Deployer reasonably believes would, if publicly disclosed, cause competitive injury. The State CAIO and any receiving sectoral regulator shall treat the designated portions as confidential under State public-records-exemption authority, subject to the agency’s authority to redact-and-publish the non-trade-secret portions, and subject to discovery in any civil enforcement action under § 2.7 or in any

⁶³⁶ See *supra* Part V.B–C; 15 U.S.C. § 7241; 18 U.S.C. § 1350. Section 2.4 incorporates § 1.5(a)’s schedule and adds its own requirements. The officer-certification analogy is distinct from § 404’s management and auditor ICFR provisions.

⁶³⁷ See *supra* Part III.D.2; § 2.5. Pre-deployment assessment, a ninety-day material-change update, and annual or three-year review are separate triggers in the worked text.

Affected Person action under § 4 of this Appendix. The State CAIO may reject a trade-secret designation that is not reasonably grounded, subject to review under the State Administrative Procedure Act.

- (c) The State CAIO shall maintain a publicly accessible registry identifying Deployers that have filed impact assessments under this Section, the Covered Domain and AI system identifier for each filed assessment, and the redacted-and-published non-trade-secret summary of each assessment.⁶³⁸
- (d) For AI systems in the highest tier under § 2.3(b) deployed in public benefits or criminal justice, the State CAIO shall provide a thirty-day public-comment period on the published summary before first deployment and shall publish responses to material comments.
- (e) Final actions of the State CAIO under this Section are subject to judicial review under the State Administrative Procedure Act; review of an impact assessment's adequacy is for reasoned consideration on the whole record.

§ 2.7 Enforcement

- (a) A Deployer's failure to conduct, document, file, or update an impact assessment this Section requires is an unlawful practice under [State UDAP statute], enforceable by the Attorney General without proof of individual consumer reliance, with the remedies the UDAP statute provides, including civil penalties, restitution, and injunctive relief.
- (b) Where deployment of an AI system produces a disparate impact cognizable under [State civil rights act], the State Civil Rights Commission (or the State human rights agency, per state convention) may proceed under that act; the filed impact assessment is admissible in the proceeding, and failure to conduct a required assessment is relevant to willfulness.
- (c) The sectoral regulator with authority over the Covered Domain in which the AI system operates may bring an action to enforce this Section under the sectoral regulator's existing enforcement authority.
- (d) The enforcement authorities of subsections (a) through (c) are concurrent and do not displace each other. Affected Person rights under § 4 of this Appendix are not displaced by any enforcement action under this Section.⁶³⁹

Commentary on § 2

The NEPA analogy supplies a reason to require an informed, documented decision without prescribing a single substantive outcome. It does not make every completed assessment adequate or every resulting decision lawful. Section 2.4 adds safeguards and responsible-officer attestation to the six incorporated vendor-assessment elements. The officer-certification analogy

⁶³⁸ 42 U.S.C. §§ 4332(2)(C), 4336a(c); *Robertson*, 490 U.S. at 350 n.13; *see supra* notes 537–539. Section 2.6's specified public comment and response are this Appendix's selected procedure, not a general transfer of former CEQ regulations to private deployments.

⁶³⁹ *See supra* Parts IV.A–C; §§ 2.7, 4.7–4.8. Concurrent agency routes retain their subject-matter predicates, and the worked private action has its own terms.

belongs to SOX §§ 302 and 906, while § 404 addresses internal-control assessment and, for covered issuers, auditor attestation. Section 2.6 distinguishes confidential filing, a redacted public summary, a limited public-comment process, and review of specified final CAIO actions.⁶⁴⁰

Risk tiers and the small-deployer exemption allocate the burden and reach of the assessment duty. The 200-Decision figure is an author-proposed placeholder for calibration, not a threshold derived from the CCPA or Colorado’s successor law. The State Nexus rule applies decision by decision and carries into §§ 3 and 4 by their terms; it is not a substitute for jurisdiction or a guarantee against a Commerce Clause challenge. Expanding it to nonresident decisions would be a substantive change requiring its own analysis.⁶⁴¹

§ 3. Audit Trail Mandate

§ 3.1 Findings and Purpose

The Legislature finds that: (a) the New York State Comptroller’s December 2025 audit, *Enforcement of Local Law 144 – Automated Employment Decision Tools*, Report 2024-N-6, identifies implementation and enforcement concerns under a law and rules that already require independent bias audits and specify audit methodology; the report supports examining additional oversight mechanisms but does not establish that their absence caused the observed concerns or that the mechanisms proposed here will improve outcomes; (b) the Sarbanes-Oxley Act’s § 404, codified at 15 U.S.C. § 7262, supplies an institutional analogue for management assessment of internal control over financial reporting and, for covered issuers, independent-auditor attestation under PCAOB standards; officer certifications arise separately under § 302, 15 U.S.C. § 7241, and § 906, 18 U.S.C. § 1350, with their own report, actor, and mental-state requirements; and (c) the Public Company Accounting Oversight Board’s inspection authority over independent auditors under 15 U.S.C. § 7214 supplies the second-tier review architecture required to maintain audit-quality integrity over time. The purpose of this Section is to establish a deployer audit trail mandate combining specified records, independent review, officer attestation, regulator-direct submission, and inspection of audit quality.⁶⁴²

§ 3.2 Definitions

The definitions of §§ 1.2 and 2.2 apply, and “Deployer” has the meaning § 2.2 gives it for purposes of this Section. Additional definitions:

- (a) “Audit Trail” means the documented record of an AI system’s training data, validation outputs, deployment decisions, monitoring outputs, and incident records, maintained by the Deployer in accordance with this Section.

⁶⁴⁰ *Robertson*, 490 U.S. at 349–53; *Marsh*, 490 U.S. at 371–76; 15 U.S.C. §§ 7241, 7262; 18 U.S.C. § 1350. Process adequacy, supplementation, management assessment, auditor opinion, and officer certification are distinct functions. *See supra* Part V.B–C.

⁶⁴¹ *See supra* Part V.B and § 2.3. Size calibration changes coverage and cost; a State Nexus limit does not itself establish personal jurisdiction or determine every Commerce Clause application.

⁶⁴² N.Y. State Comptroller, *Enforcement of Local Law 144 – Automated Employment Decision Tools*, Report 2024-N-6, at 15–19, 23, 32–35 (Dec. 2, 2025); 6 R.C.N.Y. §§ 5-300 to -302; 15 U.S.C. §§ 7214, 7262. New York’s law and rules already specify independence and methodology. The report supplies bounded implementation observations, not proof that their statutory absence caused deficient outcomes. *See supra* notes 546–549.

- (b) “Independent Auditor” means a third-party assessor credentialed by the State CAIO under § 3.5(b) and free of the prohibited relationships § 3.5(c) specifies.
- (c) “Senior Official Attestation” means the attestation required under § 3.6(a), signed by an officer with operational authority over the AI system’s deployment.⁶⁴³

§ 3.3 Coverage and Tiering

The coverage and tiering provisions of § 2.3 apply to this Section, with the State CAIO’s tiering rules specifying audit-trail content requirements, audit frequency, and audit-methodology rigor calibrated to the Deployer’s tier.

§ 3.4 Audit Trail Content

The Audit Trail under this Section shall include, at a minimum:

- (a) The AI system’s training data, including provenance, time-period coverage, and any documented gaps, sufficient to permit independent reconstruction of the AI system’s training-data lineage, to the extent the documentation is within the Deployer’s possession, custody, or control or is reasonably obtainable from the Vendor or developer; where lineage documentation is unavailable, the Audit Trail shall document the unavailability, the efforts made to obtain it, and the substitute assurance accepted under the State CAIO’s rules, which may include Vendor attestations and published model documentation;
- (b) The AI system’s validation outputs, including testing populations, performance metrics, and any documented disparate-impact patterns across protected classifications;
- (c) The Deployer’s deployment decisions, including the AI system’s operational context, the population it was deployed against, and any operational restrictions the Deployer imposed;
- (d) The AI system’s monitoring outputs, including performance drift detection, disparate-impact pattern monitoring, and any operational anomalies; and
- (e) The AI system’s incident records, including any Reportable Incident under § 3.4A(b) and any incident reportable under sectoral incident-reporting requirements applicable to the AI system’s deployment context.⁶⁴⁴

§ 3.4A Monitoring and Incident Response

- (a) A Deployer shall operate the monitoring protocols its impact assessment documents under § 2.4(a), including performance-drift detection and disparate-impact-pattern monitoring, on the schedule the State CAIO specifies by rule and not less than quarterly for AI systems in the highest tier under § 3.3.

⁶⁴³ See *supra* Part V.C; *supra* notes 546–548 for observed audit and enforcement concerns, sample limits, and the distinction between a proposed mechanism and demonstrated effectiveness.

⁶⁴⁴ See *supra* Part III.D.3–5; NIST AI RMF, *supra* note 350, §§ 5.3–5.4. Section 3.4 specifies its own five trail categories, and § 3.4A adds monitoring and incidents. A framework comparison does not itself prescribe retention or every operational record.

- (b) “Reportable Incident” means an incident that: (1) materially degrades the AI system’s conformity with the Framework Standards or with the impact assessment § 2 of this Appendix requires; (2) produces an outcome that may constitute a violation of State or federal civil rights, consumer protection, or sectoral law; or (3) involves a security breach affecting the AI system’s training data, model weights, or output integrity.
- (c) A Deployer shall report a Reportable Incident to the State CAIO and to the sectoral regulator with authority over the Covered Domain within seventy-two hours of discovery, shall remediate, and shall document the incident and remediation in the Audit Trail.

§ 3.5 Independent Audit Requirement

- (a) The Deployer shall obtain an independent audit of the Audit Trail not less than annually for AI systems in the highest tier of risk classification under § 3.3, and not less than every three years for AI systems in lower tiers.
- (b) The State CAIO shall credential Independent Auditors. Credentialing shall require: (1) demonstrated technical competence in AI-system auditing; (2) demonstrated independence from the Deployers the Independent Auditor audits; and (3) commitment to the State CAIO’s audit methodology. Credentialing, denial, suspension, and revocation shall proceed under the State Administrative Procedure Act, on written findings, with appeal as the APA provides. The State CAIO shall promulgate rules governing competence examination or equivalent demonstration; continuing education; rotation of the lead auditor for a Deployer not less often than every [five, per state convention] audit cycles; limits on the share of an Independent Auditor’s annual revenue attributable to a single Deployer [fifteen percent, per state convention]; prohibition of non-audit services to audited Deployers; and reciprocity for credentials issued by other states under substantially equivalent standards.
- (c) An Independent Auditor shall not audit a Deployer with which the Independent Auditor has: (1) a current or recent (within five years) employment relationship; (2) a financial relationship (other than the audit fee for the current audit engagement); or (3) any other relationship the State CAIO determines to compromise independence.
- (d) The Independent Auditor’s audit findings shall include: (1) a finding on the Audit Trail’s compliance with § 3.4 content requirements; (2) a finding on the AI system’s compliance with the Framework Standards; and (3) any disparate-impact patterns the audit identified, with the methodology used and the patterns’ statistical significance.⁶⁴⁵

§ 3.6 Senior Official Attestation

- (a) For each audit required under § 3.5(a), the Deployer shall obtain Senior Official Attestation that: (1) the Audit Trail is accurate and complete to the senior official’s best knowledge; (2) the AI system operates consistent with the impact assessment § 2 of this Appendix requires; and (3) any deficiencies identified in the Independent Auditor’s findings have been documented in a remediation plan.

⁶⁴⁵ 15 U.S.C. §§ 7214–7215 (PCAOB inspections and specified disciplinary authority). State credentialing and discipline under § 3.5 require the state procedures and capacity the provision contemplates.

- (b) False attestation under subsection (a) constitutes: (1) a violation of the State UDAP statute, with civil penalties available under the State UDAP statute; (2) a civil violation of this Section, with a civil penalty of up to [\$1,000,000 per false attestation, per state convention], recoverable by the Attorney General; and (3) for AI systems deployed in healthcare, insurance, employment, or public benefits, a criminal violation under [State false-statement statute or willful-violation criminal-penalty provision].⁶⁴⁶

§ 3.7 Regulator-Direct Submission and Second-Tier Review

- (a) The Deployer shall submit the Audit Trail, the Independent Auditor’s audit findings, and the Senior Official Attestation directly to the State CAIO and to any sectoral regulator with authority over the AI system’s deployment context. The Deployer shall not satisfy this submission requirement by publication on the Deployer’s website alone.
- (b) The State CAIO shall maintain a second-tier inspection program over Independent Auditors, modeled on the Public Company Accounting Oversight Board’s inspection architecture under 15 U.S.C. § 7214. Second-tier inspection shall include: (1) review of a sample of the Independent Auditor’s completed audits each inspection cycle; (2) inspection findings published consistent with PCAOB inspection-publication architecture; and (3) credential-revocation authority for Independent Auditors whose audit quality the second-tier inspection identifies as deficient.⁶⁴⁷

§ 3.8 Litigation Safe Harbor

- (a) A Deployer that conducts an Independent Audit under this Section, identifies a disparate-impact pattern in the audit’s findings, and takes corrective action in good-faith reliance on the audit’s findings obtains a statutory presumption that the corrective action satisfies the strong-basis-in-evidence standard of *Ricci v. DeStefano*, 557 U.S. 557, 585 (2009), in any disparate-treatment claim arising under State law. Nothing in this Section alters the standards, presumptions, or burdens of proof applicable to any claim arising under federal law, including Title VII of the Civil Rights Act of 1964.
- (b) The presumption under subsection (a) is conditional on the Deployer’s satisfaction of four conditions: (1) the audit was conducted under statutorily specified methodology, with credentialed-auditor independence verified under § 3.5; (2) the audit’s finding identified a disparate-impact pattern with statistical significance under the methodology’s threshold; (3) the Deployer’s corrective action was proportionate to the identified disparate-impact pattern and limited to the AI system components causing the disparate impact; and (4) the Deployer documented the audit-finding-and-corrective-action chain in the Audit Trail.
- (c) The presumption under subsection (a) applies in any action arising under State law and is rebuttable by clear-and-convincing evidence that the audit finding or corrective action was pretextual or discriminatory in effect on the non-affected class.⁶⁴⁸

⁶⁴⁶ 15 U.S.C. § 7241; 18 U.S.C. § 1350(a)–(c). The federal certification provisions have defined reports, actors, and mental states. Section 3.6’s bracketed state offense and its elements remain to be completed; these federal provisions do not supply them automatically.

⁶⁴⁷ PCAOB, AS 2201, ¶¶ .01–.04; 15 U.S.C. § 7214. Inspection of audit quality is an institutional analogy, with confidentiality, reporting, and remediation qualifications under the federal regime. *See supra* Part V.C.

⁶⁴⁸ *Ricci*, 557 U.S. at 584–87; *see supra* Part IV.B and Part V.C. Section 3.8’s presumption is a proposed

§ 3.9 Enforcement

The enforcement provisions of § 2.7 apply to this Section; for purposes of § 2.7(a), a Deployer's failure to perform an obligation this Section imposes is an unlawful practice under [State UDAP statute]. Additionally, the Attorney General may bring an action for false-attestation civil penalties under § 3.6(b), and the [State criminal prosecutor, per state convention] may bring a criminal action for false attestation in the sectors § 3.6(b)(3) identifies.

Commentary on § 3

Section 3 combines a preserved trail, monitoring and incident response, credentialed audits, officer attestation, direct filing, and inspection of audit quality. The Comptroller's New York City findings support attention to implementation and enforcement, but do not show that Local Law 144 lacked independence or methodology rules, establish intentional gaming, or prove that these proposed additions will improve outcomes. SOX and PCAOB supply institutional analogies with different scope, standards, and exemptions. Section 3.4A's seventy-two-hour notification runs to the CAIO and sectoral regulator, not generally to every affected person.⁶⁴⁹

Section 3.8 proposes a state-law-only presumption for qualifying audit findings and proportionate corrective action. Its four conditions, rebuttal, and federal-law reservation must be read together. Ricci does not require this presumption or make procedural compliance sufficient under federal Title VII. The trail's other limits also matter: § 3.4(a)'s unavailable-information record is not proof of adequate assurance; § 3.4 states no general private-deployment retention interval; and the text does not expressly separate its reporting duties from ERISA plan administration.⁶⁵⁰

§ 4. Dispute and Remedy Architecture

§ 4.1 Findings and Purpose

The Legislature finds that: (a) Affected Persons subject to Consequential Decisions made or materially influenced by AI systems require notice, explanation, correction, human review, and remedy rights to make the upstream process duties of §§ 2 and 3 operationally meaningful; (b) the Fair Credit Reporting Act's adverse-action notice architecture, 15 U.S.C. § 1681m, and the Equal Credit Opportunity Act's adverse-action notice and statement-of-reasons architecture, 15 U.S.C. § 1691(d), supply the federal analogue from which the State's dispute and remedy architecture draws; and (c) the State's traditional civil rights enforcement, consumer protection enforcement, and sectoral-licensing enforcement architectures supply the remedy channels

state-law rule, not a federal holding; subsection (a) preserves federal standards, presumptions, and burdens.

⁶⁴⁹ See *supra* Part V.C; *supra* notes 547–550; 15 U.S.C. §§ 7214–7215. The proposed oversight additions respond to implementation concerns without establishing their necessity or sufficiency as an empirical matter.

⁶⁵⁰ *Ricci*, 557 U.S. at 584–87; *Gobeille*, 577 U.S. at 319–25; *Rutledge*, 592 U.S. at 87–92. See *supra* Part V.C for the separate state-law presumption and ERISA reporting boundary. Neither audit compliance nor a general-applicability label settles federal liability or preemption.

through which Affected Person recourse operates. The purpose of this Section is to establish the four-channel recourse architecture for Affected Persons.⁶⁵¹

§ 4.2 Definitions

The definitions of §§ 1.2, 2.2, and 3.2 apply, and “Deployer” for purposes of this Section means any person or entity, including a State agency, State authority, or political subdivision, that deploys an AI system in a Covered Domain to make or materially influence a Consequential Decision having a State Nexus as § 2.3(a) defines that term. The duties and rights of §§ 4.3 through 4.9 apply as to each Consequential Decision that has a State Nexus, and do not apply to a Consequential Decision that has none. Additional definitions:

- (a) “Adverse Action” means a Consequential Decision that denies, limits, or terminates an Affected Person’s access to, terms of, conditions of, or benefit from the service, benefit, or opportunity at issue.
- (b) “Meaningful Human Review” means review of an AI-mediated Consequential Decision by a natural person who: (1) has authority to change the decision; (2) is not evaluated on throughput or on rate of agreement with AI outputs, and is organizationally independent of the AI system’s procurement and operation; (3) has access to the decision inputs, the explanation § 4.4 requires, and any materials the Affected Person submits; and (4) documents the basis of the review in writing.⁶⁵²

§ 4.3 Notice Right

- (a) A Deployer shall provide an Affected Person with notice that an AI system was used to make or materially influence the Consequential Decision affecting the Affected Person. The notice shall be provided: (1) at the application or pre-decision stage, as general notice that AI may be used; (2) at the Adverse Action stage, as specific notice that AI was used in the particular Consequential Decision; and (3) on Affected Person request, with explanation under § 4.4.
- (b) The notice required under subsection (a)(2) shall include: (1) identification of the AI system used; (2) the Covered Domain and decision context; (3) the principal factors driving the Adverse Action; and (4) instructions for exercising the rights under §§ 4.4 through 4.6.
- (c) The notice required under this Section shall be provided in plain language, in the Affected Person’s preferred language to the extent the Deployer has reasonable access to the preferred-language information, and in the medium through which the Deployer communicated the Adverse Action.⁶⁵³

⁶⁵¹ 15 U.S.C. §§ 1681m(a), 1681i, 1681s-2; *id.* § 1691(d). FCRA user notice, CRA reinvestigation, furnish duties, and ECOA reasons are distinct analogies; *see supra* Part III.D.6 and Part V.D.

⁶⁵² Regulation (EU) 2016/679, art. 22(1)–(4). Paragraph (3)’s minimum intervention, expression, and contest safeguards apply to the specified paragraph (2)(a) and (c) exceptions; the statutory exception in (2)(b) has its own safeguards requirement. This Appendix’s materially-influenced trigger and reviewer safeguards are proposed extensions, not identical GDPR coverage.

⁶⁵³ 15 U.S.C. § 1681m(a). The federal report-user notice is an analogy; § 4.3 adds its own AI notice content and timing, subject to § 5.5A.

§ 4.4 Explanation Right

- (a) On Affected Person request following Adverse Action, the Deployer shall provide an explanation of: (1) the principal factors driving the AI system's output that resulted in the Adverse Action; (2) the principal data inputs the AI system used; and (3) the AI system's documented limitations relevant to the Adverse Action.
- (b) The explanation under subsection (a) shall be specific enough to permit the Affected Person to identify potentially erroneous data inputs, contest the AI system's output, or seek Meaningful Human Review under § 4.6.
- (c) The Deployer's compliance with subsection (a) does not require disclosure of trade-secret algorithmic details. The Deployer shall, however, disclose the factors and data inputs in subsections (a)(1) and (a)(2) at a level of specificity sufficient to satisfy the purpose of subsection (b).⁶⁵⁴

§ 4.5 Correction Right

- (a) An Affected Person may request correction of: (1) data inputs the AI system used that the Affected Person identifies as inaccurate; and (2) external-consumer-data sources the AI system used that the Affected Person identifies as inaccurate.
- (b) On receipt of a correction request, the Deployer shall: (1) investigate the correction request within thirty days; (2) correct any data input the Deployer determines to be inaccurate; (3) notify any external-consumer-data source the Deployer relies on of the correction request and follow the external-consumer-data source's correction procedures; and (4) re-decide the affected Consequential Decision using corrected data inputs.
- (c) The correction right under this Section is in addition to, and does not displace, any correction rights the Affected Person has under State consumer-data statutes, the federal Fair Credit Reporting Act, or sectoral statutes applicable to the data source.⁶⁵⁵

§ 4.6 Human Review Right

- (a) An Affected Person whose Consequential Decision was made or materially influenced by an AI system may request Meaningful Human Review of the Consequential Decision.
- (b) The Deployer shall provide Meaningful Human Review within thirty days of the Affected Person's request, or within seventy-two hours where the Adverse Action concerns emergency healthcare, eviction from housing, or termination of public benefits. Where a stay is lawful and feasible, the Adverse Action is stayed pending review. The review shall be conducted by a natural person with the attributes § 4.2(b) requires.

⁶⁵⁴ 15 U.S.C. § 1691(d); 12 C.F.R. § 1002.9(b)(2). The Appendix's explanation content extends beyond the creditor's specific-reasons duty. Sections 4.4(b)–(c) must be read together: functional specificity and confidentiality are both part of the proposed rule.

⁶⁵⁵ 15 U.S.C. § 1681i(a) (reasonable reinvestigation, qualified timing, and correction or deletion); *id.* § 1681s-2 (separate furnisher responsibilities). Section 4.5 supplies its own deployer investigation and correction process, subject to § 5.5A; it does not guarantee control of external source records.

- (c) If the Meaningful Human Review identifies error in the AI system’s output or determines that the Consequential Decision is not supported on the whole record, the Deployer shall re-decide the Consequential Decision and notify the Affected Person of the re-decision and any consequent action.⁶⁵⁶

§ 4.7 Administrative Recourse

An Affected Person may file an administrative complaint with: (1) the State Attorney General, where the AI-mediated Adverse Action implicates State UDAP-statute violations; (2) the State Civil Rights Commission, where the AI-mediated Adverse Action reveals disparate-impact or disparate-treatment patterns across protected classifications; (3) the sectoral regulator with authority over the Covered Domain in which the AI system operates; and (4) the State CAIO, where the AI-mediated Adverse Action implicates impact-assessment or audit-trail violations under §§ 2 or 3 of this Appendix.

§ 4.8 Private Right of Action

- (a) An Affected Person who has been harmed by an AI-mediated Adverse Action may bring a civil action against the Deployer for: (1) violation of the State UDAP statute, where applicable; (2) violation of State or federal civil rights statutes, where applicable; (3) violation of sectoral statutes applicable to the Covered Domain; and (4) violation of §§ 2, 3, or 4 of this Appendix.
- (b) The private right of action under subsection (a) is not displaced by any administrative complaint or enforcement action under § 4.7 or § 2.7 of this Appendix.
- (c) Statutory damages, attorney’s fees, and injunctive relief are available consistent with the underlying statutory cause of action.⁶⁵⁷
- (d) In an action under subsection (a)(4), the court may award actual damages; statutory damages of [\$1,000, or \$5,000 for intentional or reckless violations, per state convention] per violation where actual damages are difficult to quantify; reasonable attorney’s fees and costs to a prevailing Affected Person; and injunctive relief. An action under subsection (a)(4) shall be commenced within [two years, per state convention] of the date the Affected Person discovered or reasonably should have discovered the violation.

§ 4.9 Class Action and Arbitration Considerations

- (a) Nothing in this Section abrogates the Federal Arbitration Act, 9 U.S.C. §§ 1-16, or the Supreme Court’s interpretation of the FAA in *AT&T Mobility LLC v. Concepcion*, 563 U.S. 333 (2011), and *Epic Systems Corp. v. Lewis*, 584 U.S. 497 (2018).

⁶⁵⁶ See *supra* Part III.D.6 and Part V.D; § 4.6(a)–(c). The review entitlement, thirty-day or specified urgent seventy-two-hour interval, lawful-and-feasible stay, and conditional re-decision are distinct requirements.

⁶⁵⁷ See *supra* Parts IV.A–C and Part V.D. Section 4.8 distinguishes underlying applicable causes from the independent process-duty action and its remedies. Federal and state standing, the cause’s own harm requirement, immunity, arbitration, and merits remain separate.

- (b) Where an Affected Person has agreed to an arbitration agreement covering the underlying transaction, the Affected Person's rights under §§ 4.3 through 4.8 of this Appendix operate within the arbitration framework, subject to the State court's authority to police arbitration-agreement enforceability under State contract-formation doctrine consistent with the FAA. Nothing in an arbitration agreement limits, delays, or channels the Affected Person's right to file administrative complaints under § 4.7, any agency's enforcement authority under § 2.7, § 3.9, or § 4.7, or the notice, explanation, and correction duties of §§ 4.3 through 4.5; those rights and duties are not waivable by agreement to the extent permitted by law.
- (c) For Adverse Actions in Covered Domains where the underlying transaction is not subject to a binding arbitration agreement, the Affected Person's State-court remedies operate without the constraints subsections (a) and (b) reference.⁶⁵⁸

Commentary on § 4

Section 4 joins notice, explanation, correction, and review to administrative and private enforcement. Section 4.8 preserves applicable underlying causes of action and separately authorizes the process-duty action in subsection (a)(4), with the remedies in subsection (d). It therefore extends beyond an AG-only model for cross-cutting duties. Under § 4.6, review is available on request for a covered consequential decision, while re-decision follows the specified finding of error or lack of whole-record support. Public claims still require the state-specific immunity and claims machinery identified in the Preliminary Note.⁶⁵⁹

The federal comparisons do not erase federal limits. FCRA allocates different duties to reporting agencies, users, and furnishers; § 5.5A limits the subject matter reached by §§ 4.3–4.5. Section 4.9 preserves the FAA and enforceable arbitration arrangements, while protecting specified administrative routes and information duties to the extent permitted by law. Administrative enforcement can address conduct that private proceedings may not reach, but cannot guarantee an equivalent individual remedy. Statutory damages likewise remain subject to the forum's standing requirements.⁶⁶⁰

§ 5. General Provisions

§ 5.1 Severability

- (a) If any provision of this Act, or its application to any person or circumstance, is held invalid or preempted, the invalidity or preemption does not affect other provisions or applications of the Act that can be given effect without the invalid or preempted provision or application.

⁶⁵⁸ *AT&T Mobility LLC v. Concepcion*, 563 U.S. 333, 339–44, 351–52 (2011); *Epic Systems Corp. v. Lewis*, 584 U.S. 497, 503–08, 524–25 (2018); 9 U.S.C. § 2. Coverage, lawful contract defenses, and statutory exceptions matter; the cases do not make every purported arbitration agreement enforceable.

⁶⁵⁹ See *supra* Part V.D; §§ 4.7–4.8, 5.5A. The worked text preserves applicable underlying claims and adds the process-duty action; its federal subject-matter limitation is distinct from its allocation of actors or forums.

⁶⁶⁰ See *supra* Part V.D; *Concepcion*, 563 U.S. at 339–44, 351–52; *Epic Systems*, 584 U.S. at 503–08, 524–25; *TransUnion*, 594 U.S. 413 (majority slip op. at 15 & n.4). Public complaint routes remain subject to actual agency authority and the worked text's legal qualification; they do not establish equivalent relief or eliminate standing limits.

- (b) The legislature declares that it would have enacted each section of this Act independently of every other, and specifically that it would have enacted §§ 2, 3, and 4 had § 1 been held invalid or preempted, and would have enacted each of §§ 1 through 4 had any other been held invalid or preempted as applied to any class of Deployer, Vendor, or AI system.
- (c) A duty under this Act that references documentation originating outside the Deployer's possession, custody, or control remains operative in the form § 3.4(a) prescribes for documentation that cannot be obtained.

§ 5.2 Effective Dates

Section 1 takes effect [twelve months, per state convention] after enactment. Sections 2 and 3 take effect [twelve months, per state convention] after the State CAIO's implementing rules become final, and not later than [thirty-six months, per state convention] after enactment. Section 4 takes effect concurrently with § 2.

§ 5.3 Limitations

Except as otherwise provided in this Act, an action under this Act shall be commenced within [two years, per state convention] of the date the violation was discovered or reasonably should have been discovered.

§ 5.4 Lead Agency

The State CAIO is the lead coordinating agency for the implementation of this Act. Nothing in this section modifies § 1.3(c).

§ 5.5 Savings

This Act does not displace or diminish any right or remedy available under other State or federal law. This section preserves rights and remedies and imposes no requirement or prohibition of its own; it does not qualify § 5.5A.

§ 5.5A Fair Credit Reporting Act Preemption

- (a) Sections 4.3 through 4.5 impose no requirement or prohibition with respect to subject matter preempted by 15 U.S.C. § 1681t(b)(1)(B), (C), (D), (F), or (I), within the scope and subject to the qualifications of each cited subparagraph. The reference to § 1681i in subparagraph (B) concerns the time by which a consumer reporting agency must take action in a procedure related to the disputed accuracy of information in a consumer's file; it does not exclude all subject matter regulated by § 1681i. The duties of §§ 4.3 through 4.5 apply to the extent they are not preempted.
- (b) Subsection (a) limits the subject matter §§ 4.3 through 4.5 reach and does not limit the persons, the AI systems, or the Covered Domains those Sections reach. A duty under §§ 4.3 through 4.5 remains operative to the extent subsection (a) does not exclude the subject matter it regulates, and nothing in this section withdraws a right an Affected Person holds under federal law.

§ 5.6 Security of Filings

The State CAIO shall protect filed impact assessments and Audit Trails with administrative, technical, and physical safeguards no less protective than those [the state data-security statute] requires.

§ 5.7 Appropriation and Assessment

- (a) The State CAIO's costs of administering §§ 2 and 3 shall be defrayed by an annual assessment upon each Covered Deployer, apportioned according to the tier assigned under § 2.3 and limited to expenses solely attributable to administering those Sections.
- (b) Fees the State CAIO charges for credentialing Independent Auditors under § 3.5(b) shall be deposited [as directed by the enacting State's treasury and appropriations law].
- (c) An assessment under subsection (a) shall not vary with, and shall not be contingent upon, any enforcement action the State CAIO takes or declines to take against the assessed Deployer.
- (d) [Any balance required, per state convention.]