


# Governing the *Black Box.*

AI regulation, and what it means for insurers

ACHS 2026 SPRING MEETING

Max Samadov  
SAS Institute  
May 19, 2026





195 BC · THE ROMAN FORUM

“I wonder how one *augur* can keep from laughing when he passes another.”

— Cato the Elder, via Cicero, *De Divinatione* II.51

*Salvator Orefa pinxit.*

*Jos. Goupy fecit Londini.*

*Augures*  
*Servatur Exemplar in Aedibus praehon. Domini D. Cooke. M. Britanniae Regis Vice Camerarij.*

A SCENE

# It is June 2027.

A market conduct examiner from the Connecticut Insurance Department walks into your office.

She has six requests.

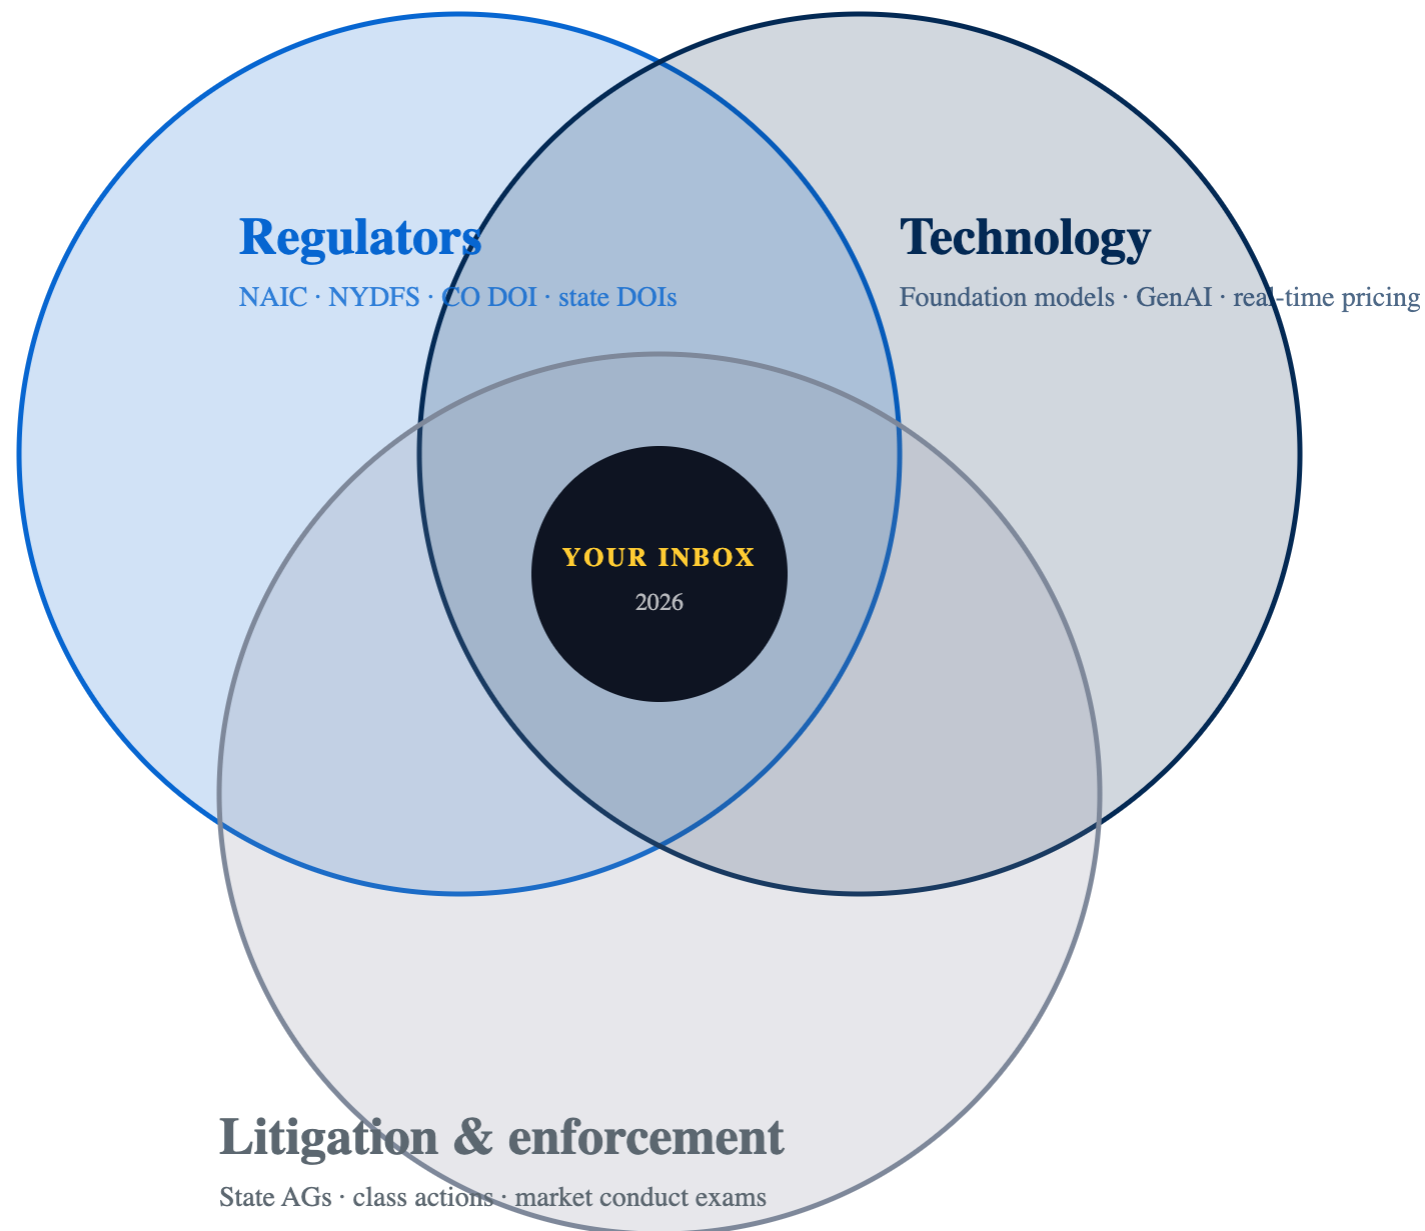
SIX REQUESTS · REDACTED UNTIL SLIDE 21

- 01 [REDACTED]
- 02 [REDACTED]
- 03 [REDACTED]
- 04 [REDACTED]
- 05 [REDACTED]
- 06 [REDACTED]

Could your firm, today, hand them over?

ACT 2 BEGINS

# Three forces converging on **your inbox.**



**01 · REGULATORS**

A patchwork of state regimes, accelerating.

**02 · TECHNOLOGY**

GenAI moved from curiosity to procurement category in 24 months.

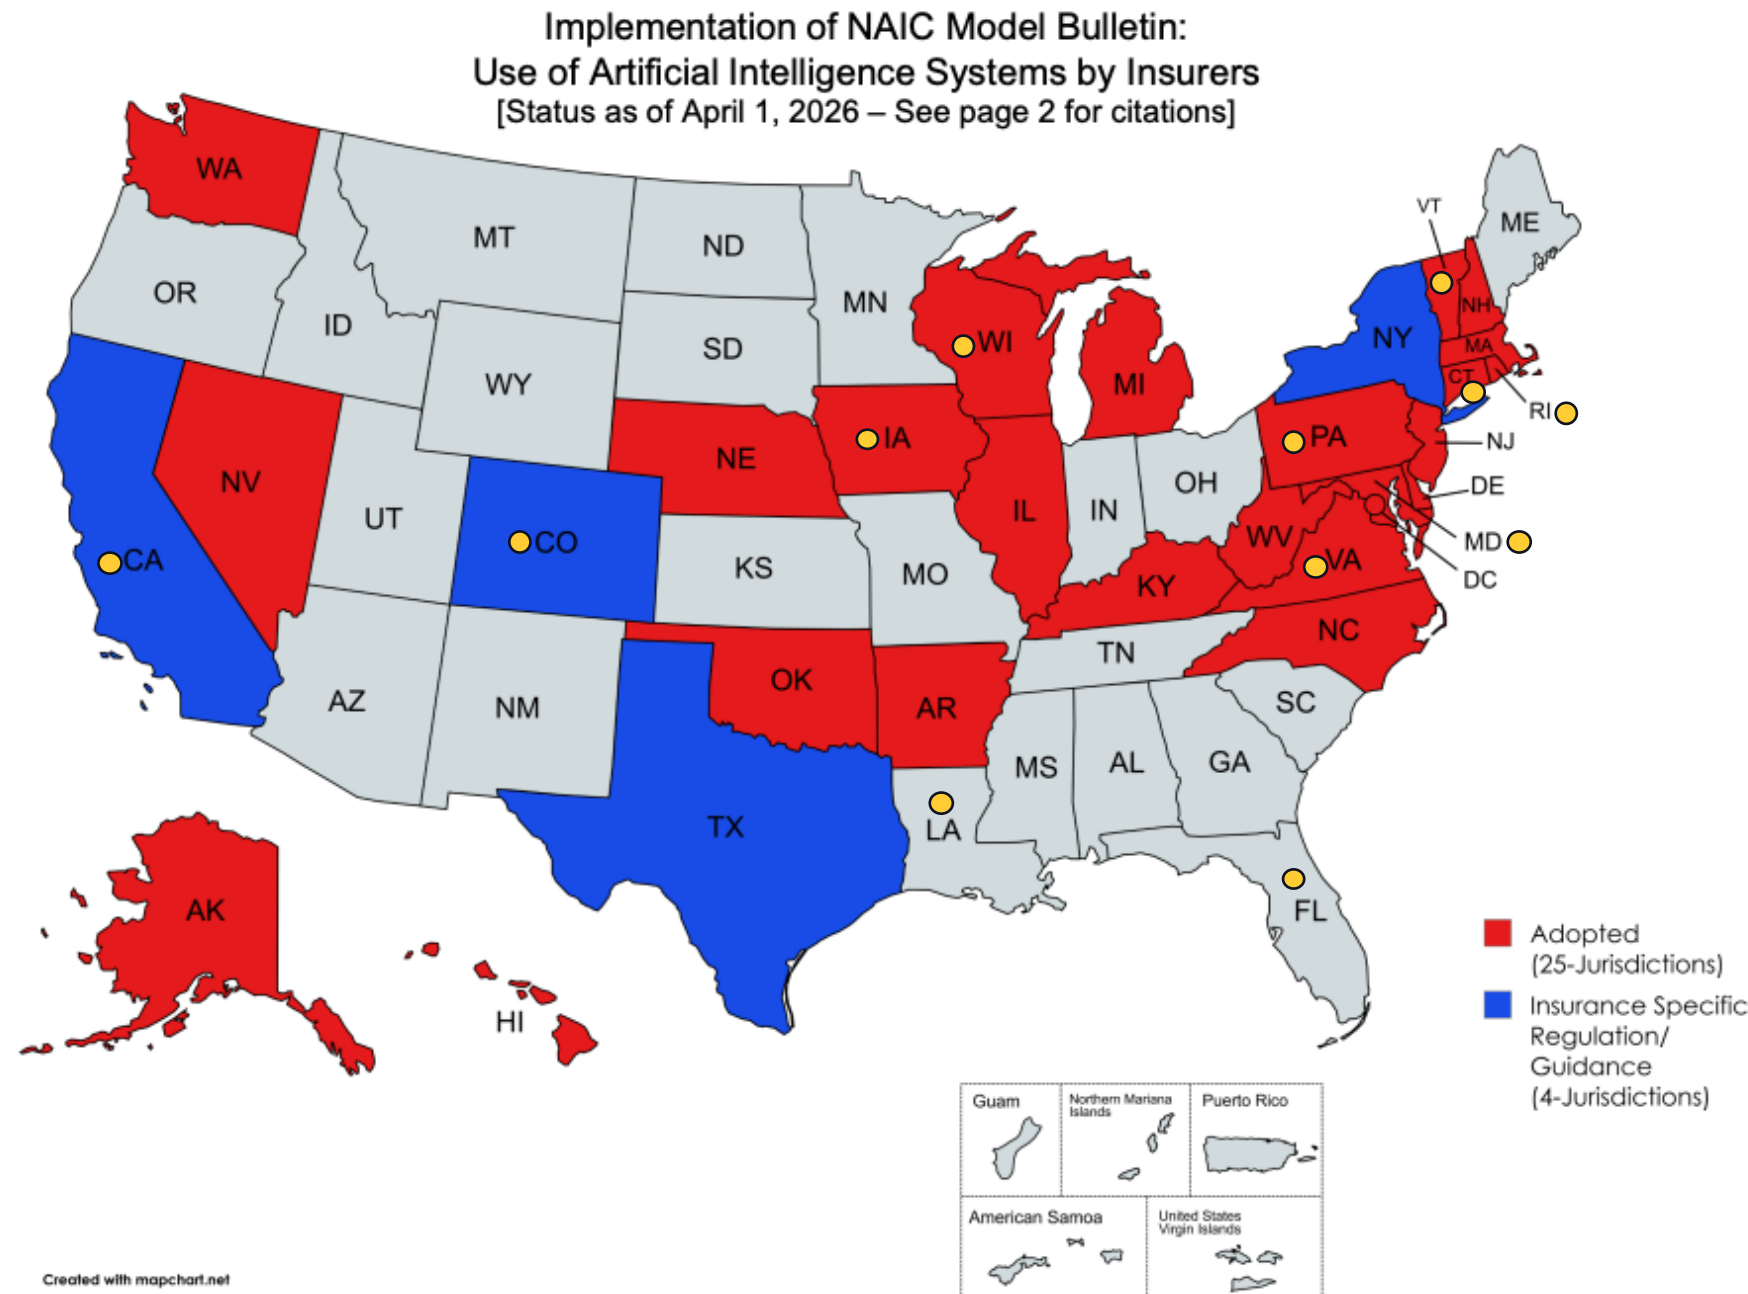
**03 · LITIGATION**

An insurance-side verdict is a matter of when, not if.

None of the three is coordinating with the others. They are converging, independently, on the same place.

### THE STATE-BY-STATE PATCHWORK

# Twenty-five jurisdictions. **Five regimes.**



■ NAIC Model Bulletin

24 jurisdictions adopted

■ Own regime

NY · CO · CA · TX

● NAIC Eval Tool pilot

12 states · CT & PA included

Absence of color does not mean safety. It means [yet](#).

Created with mapchart.net

This map represents state action or pending state action addressing the topic of the model. This map does not reflect a determination as to whether the pending or enacted legislation contains all elements of the model or whether a state meets any applicable accreditation standards. See the reference list on the next page.

NAIC MODEL BULLETIN · DECEMBER 2023

# The bulletin is the **open-book exam**.

Twenty-three pages. Four pillars. The states that adopted it published, in writing, the artifacts they intend to demand.

---

## 01

### Written AIS program

Not principles in a deck. Not policy in someone's head. A document.

## 02

### Governance & risk management

Senior management and the board own outcomes — and must be able to answer *what AI is in use*.

## 03

### Third-party vendors

You cannot delegate compliance. The regulator's questions land on your desk, not theirs.

## 04

### Documentation for examination

Inventory, validation, vendor due diligence — the artifacts they will ask for, on request.

NEW YORK DEPARTMENT OF FINANCIAL SERVICES · CL 2024-7

# If the Bulletin is principles, **New York is operations.**

## THE COMPREHENSIVE ASSESSMENT · 3 STEPS



## WHAT'S DISTINCTIVE

- Not aligned to NAIC — anchored in NY anti-discrimination statutes
- Quantitative testing required — not policy attestations
- BISG referenced as method, **no threshold prescribed**
- Vendor contracts must include **audit rights**

---

NAIC compliance is **not** a substitute for New York compliance.

COLORADO · FOR THE LIFE CARRIERS

# Where the **teeth** come from.

## WHAT IT IS

**First state regulation in the U.S. specifically targeting algorithmic discrimination in life insurance underwriting.**

- Effective Nov 14, 2023
- Applies to life insurers using ECDIS, algorithms, predictive models
- Governance, testing, reporting requirements
- Binding. Operating today.

## WHAT IT ISN'T

**The Colorado AI Act of February 2026 — that one explicitly exempts insurers.**

- Not a general AI law
- Quantitative testing rule still in draft (and has been for two years)
- Not optional

*Colorado has a comprehensive AI law for everyone except insurers — and a separate, narrower, older, tougher regime for them.*

## FOR THIS ROOM, SPECIFICALLY

Five overlapping but **non-identical** regimes.

STATE	INSTRUMENT	STATUS	DISTINCTIVE FEATURE
<b>CT</b>	MC-25	Adopted Feb 2024	In NAIC Eval Tool pilot · <b>PILOT</b>
<b>NY</b>	CL 2024-7	Effective	3-step comprehensive assessment · most demanding regime in the country
<b>PA</b>	Notice 2024-04	Adopted Apr 2024	In NAIC Eval Tool pilot · <b>PILOT</b>
<b>NJ</b>	Bulletin 25-03	Adopted Feb 2025	Standard NAIC adoption
<b>MA</b>	Bulletin 2024-10	Adopted Dec 2024	Plus April 2024 AG advisory — two-angle scrutiny

Ninety percent of the work is the same. The other ten percent is what matters.

EXECUTIVE ORDER 14365 · DECEMBER 2025

# Two regimes. A tug-of-war.

## FEDERAL EO 14365

Asserts federal authority over AI standards. Does not specifically exempt insurance.



## STATE NAIC + State DOIs

McCarran-Ferguson posture. Pilot is running. Bulletins are enforced.

12-24

months of dual-track compliance until courts resolve.

5

days for the NAIC to publicly oppose the EO.

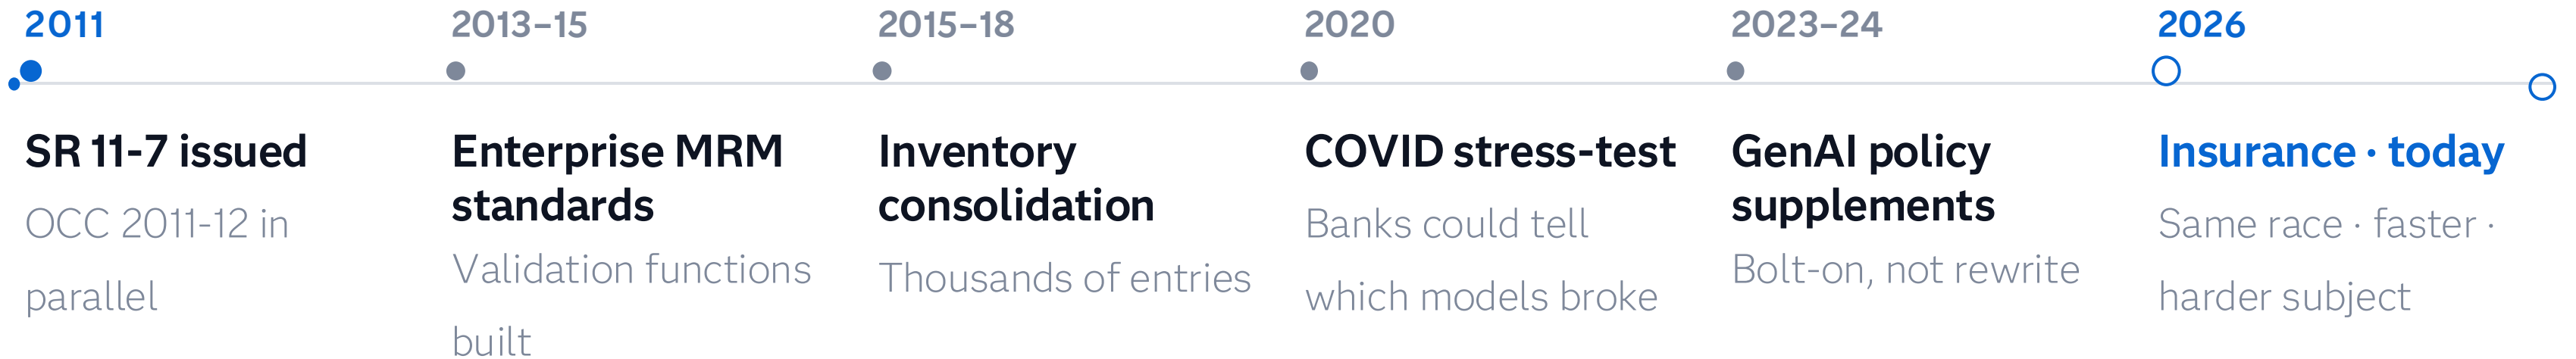
0

state regulators have paused enforcement.

The defensible posture: **hedge by complying.**

SR 11-7 · OCC 2011-12 · FEDERAL RESERVE

# Banks have already built this.



**PILLAR**  
Inventory

**PILLAR**  
Validation

**PILLAR**  
Monitoring

**PILLAR**  
Governance

THE FRAMING TO TAKE AWAY

# AI governance is MRM++.

Same skeleton. Harder muscles.

BANKING · MRM  
2011 → 2025

Mostly **internal** models

Quantitative, **structured** data

**Quarterly**—annual revalidation

Disparate impact a *peripheral* concern

Board oversight **well established**

INSURANCE · AI GOVERNANCE  
2024 → ?

Heavy **third-party** & foundation-model exposure

Plus **unstructured**, multimodal, generative

**Continuous** drift monitoring

Disparate impact the **primary** regulatory concern

Board oversight **being invented**

PILLAR 01

# You cannot govern what you cannot see.

First artifact a regulator asks for.

Unit of analysis: the use case , not the model. Tiering must be written down.

A useful test: would your CEO be surprised by any line in the inventory?

AI Risk Management · Model Inventory Run · 2026-04-28 09:14 ET

**Active AI / ML Model Inventory**

Tier-1 and Tier-2 models in production · Tier-3 hidden Tier 1 – High Tier 2 – Med

Status: Active Tier: 1, 2 Lines: All Showing 10 of 47 active

MODEL ID	USE CASE	LINE	SOURCE	TIER	LAST VALIDATED
ML-2024-0017	Auto premium pricing GLM	Personal Auto	Internal	Tier 1	2026-02-14
ML-2023-0042	Term life underwriting ensemble	Life	Vendor · vendor	Tier 1	2026-01-08
ML-2024-0091	Claims FNOL triage classifier	Personal P&C	Internal	Tier 1	2025-11-22
ML-2025-0011	Auto fraud — staged accident	Personal Auto	Vendor · vendor	Tier 1	2026-03-05
ML-2024-0067	Group health adjudication	Group Health	Internal	Tier 1	2026-01-30
ML-2023-0021	Annuity suitability scoring	Life / Annuity	Internal	Tier 1	2025-08-15
ML-2024-0058	Disability OCR + extraction	Disability	Vendor · vendor	Tier 2	2025-12-19
ML-2025-0023	Cat model overlay (wind)	Comm. Property	Vendor · vendor	Tier 2	2026-03-14

Tier definitions per *Internal Policy MRM-AI-001 §3.2* CONFIDENTIAL — INTERNAL USE

PILLAR 02

# Bias testing has a methodology.

- 01 Define protected classes & how membership is inferred (BISG)
- 02 Define the decision being tested
- 03 Compute selection rates by class
- 04 Adverse Impact Ratio · the 4/5 rule
- 05 Conditional / multivariate fairness
- 06 Document the *decisions*, not just the testing

AI Risk Management · Fairness Testing · Report
Report ID: BTR-2026-Q1-0117 Generated: 2026-04-22

## Bias Testing Report — ML-2024-0017

Auto premium pricing GLM v4.2 · Quarterly assessment · Personal Auto

Disparity flagged
Deployment hold

**TEST SUMMARY**

**184,392**

QUOTES TESTED

**Q1 2026**

PERIOD

**73.4%**

REFERENCE RATE

**0.745**

MIN AIR

**ADVERSE IMPACT RATIO BY PROTECTED CLASS**

PROTECTED CLASS	VS.	SEL. RATE	AIR	STATUS
White (non-Hispanic)	—	73.4%	1.000	reference
<b>Black</b>	White	54.7%	<b>0.745</b>	<b>Below 0.80</b>
Hispanic / Latino	White	67.2%	0.916	Within tol.
Asian	White	78.1%	1.064	Within tol.
Female	Male	71.8%	0.978	Within tol.
Age 65+	Age 25–64	69.4%	0.945	Within tol.

**METHODOLOGY**

**Race/ethnicity inference**  
Bayesian Improved Surname Geocoding (BISG) — 2020 Census surname & block-group priors. Coverage: 91.2%.

**Gender inference**  
SSA first-name dictionary (1937–2024). Coverage: 96.8%.

**Threshold**  
EEOC four-fifths rule (AIR ≥ 0.80) as primary; supplemental: marginal effect > 5pp.

**CONDITIONAL ANALYSIS — BLACK VS WHITE**

**Residual disparity after controls:** AIR rises from **0.745** (raw) to **0.834** after controlling for vehicle class, garaging ZIP, miles driven, and prior at-fault claims. Of the remaining 0.166 AIR gap pre-controls, approximately **62%** is attributable to **credit-based insurance score (CBIS)** as a single feature.

**FEATURE CONTRIBUTION TO DISPARITY**

FEATURE	Δ AIR	CUM. AIR	NOTE
Raw model	—	0.745	
+ Vehicle class controls	+0.018	0.763	
+ Garaging ZIP density	+0.022	0.785	
+ Annual miles	+0.009	0.794	
+ Prior at-fault claims	+0.012	0.806	
<b>– CBIS (sensitivity)</b>	<b>+0.028</b>	<b>0.834</b>	review feature

**RECOMMENDATION & STATUS**

**Action:** Deployment hold on v4.2. Reviewer to evaluate continued use of CBIS feature under NY DFS Circular Letter 2024-7 §III(C) — legitimate business necessity standard. Documented business rationale required prior to release. Re-test on remediation.

**Owner:** ██████ · **Due:** 2026-05-13 · **Reviewer:** ██████

Methodology per Internal Standard FAIR-001 v2.3. AIR threshold per EEOC Uniform Guidelines on Employee Selection Procedures, 29 C.F.R. § 1607.4(D), applied by analogy.

CONFIDENTIAL — INTERNAL USE

# Bias remediation, **in practice.**



VIDEO TO EMBED · SLIDE 15

Synthetic underwriting model · AIR 0.74 → 0.91 · no loss of predictive lift.

*Source from SAS demo library or commission a 2-min screen recording.*

BEFORE

# 0.74

Raw AIR · disparity flagged

AFTER

# 0.91

Remediated · pre-deployment

— ONE AFTERNOON —

PILLAR 03

# Explainability, for adverse decisions.

**To the consumer:** a reason code, in plain language. FCRA / ECOA compatible.

**To the regulator:** SHAP, monotonicity, counterfactuals — explicable in terms of the consumer’s actual data.

Trap: do not let an LLM compose the adverse-action notice. A fluent post-hoc explanation is worse than none.

Underwriting Decisions · Adverse Action Notice · System-generated
Notice ID: AAN-2026-04-021187 Sent: 2026-04-21 14:32 ET

## Notice of Adverse Underwriting Decision

Issued in compliance with FCRA § 615(a) and the Insurance Information and Privacy Protection Act

Application: AUTO-2026-[REDACTED]  
 Recipient: [REDACTED]  
 Decision date: 2026-04-21

---

Dear [REDACTED]

We have reviewed your application for personal automobile insurance. After consideration, we are unable to offer coverage at the rate you requested. **This decision was made or supported by an automated decision system that uses statistical models to evaluate insurance applications.**

PRINCIPAL FACTORS THAT CONTRIBUTED TO THIS DECISION In order of impact on the outcome

1	<b>Driving record — at-fault accidents in the prior 36 months</b> <small>Source: motor vehicle records and prior carrier loss history</small>	-0.42
2	<b>Vehicle category — high-performance classification</b> <small>Source: VIN decode, ISO symbol table</small>	-0.31
3	<b>Credit-based insurance score — below 620 percentile</b> <small>Source: <span style="background-color: black; color: black;">[REDACTED]</span> score date 2026-04-19</small>	-0.18

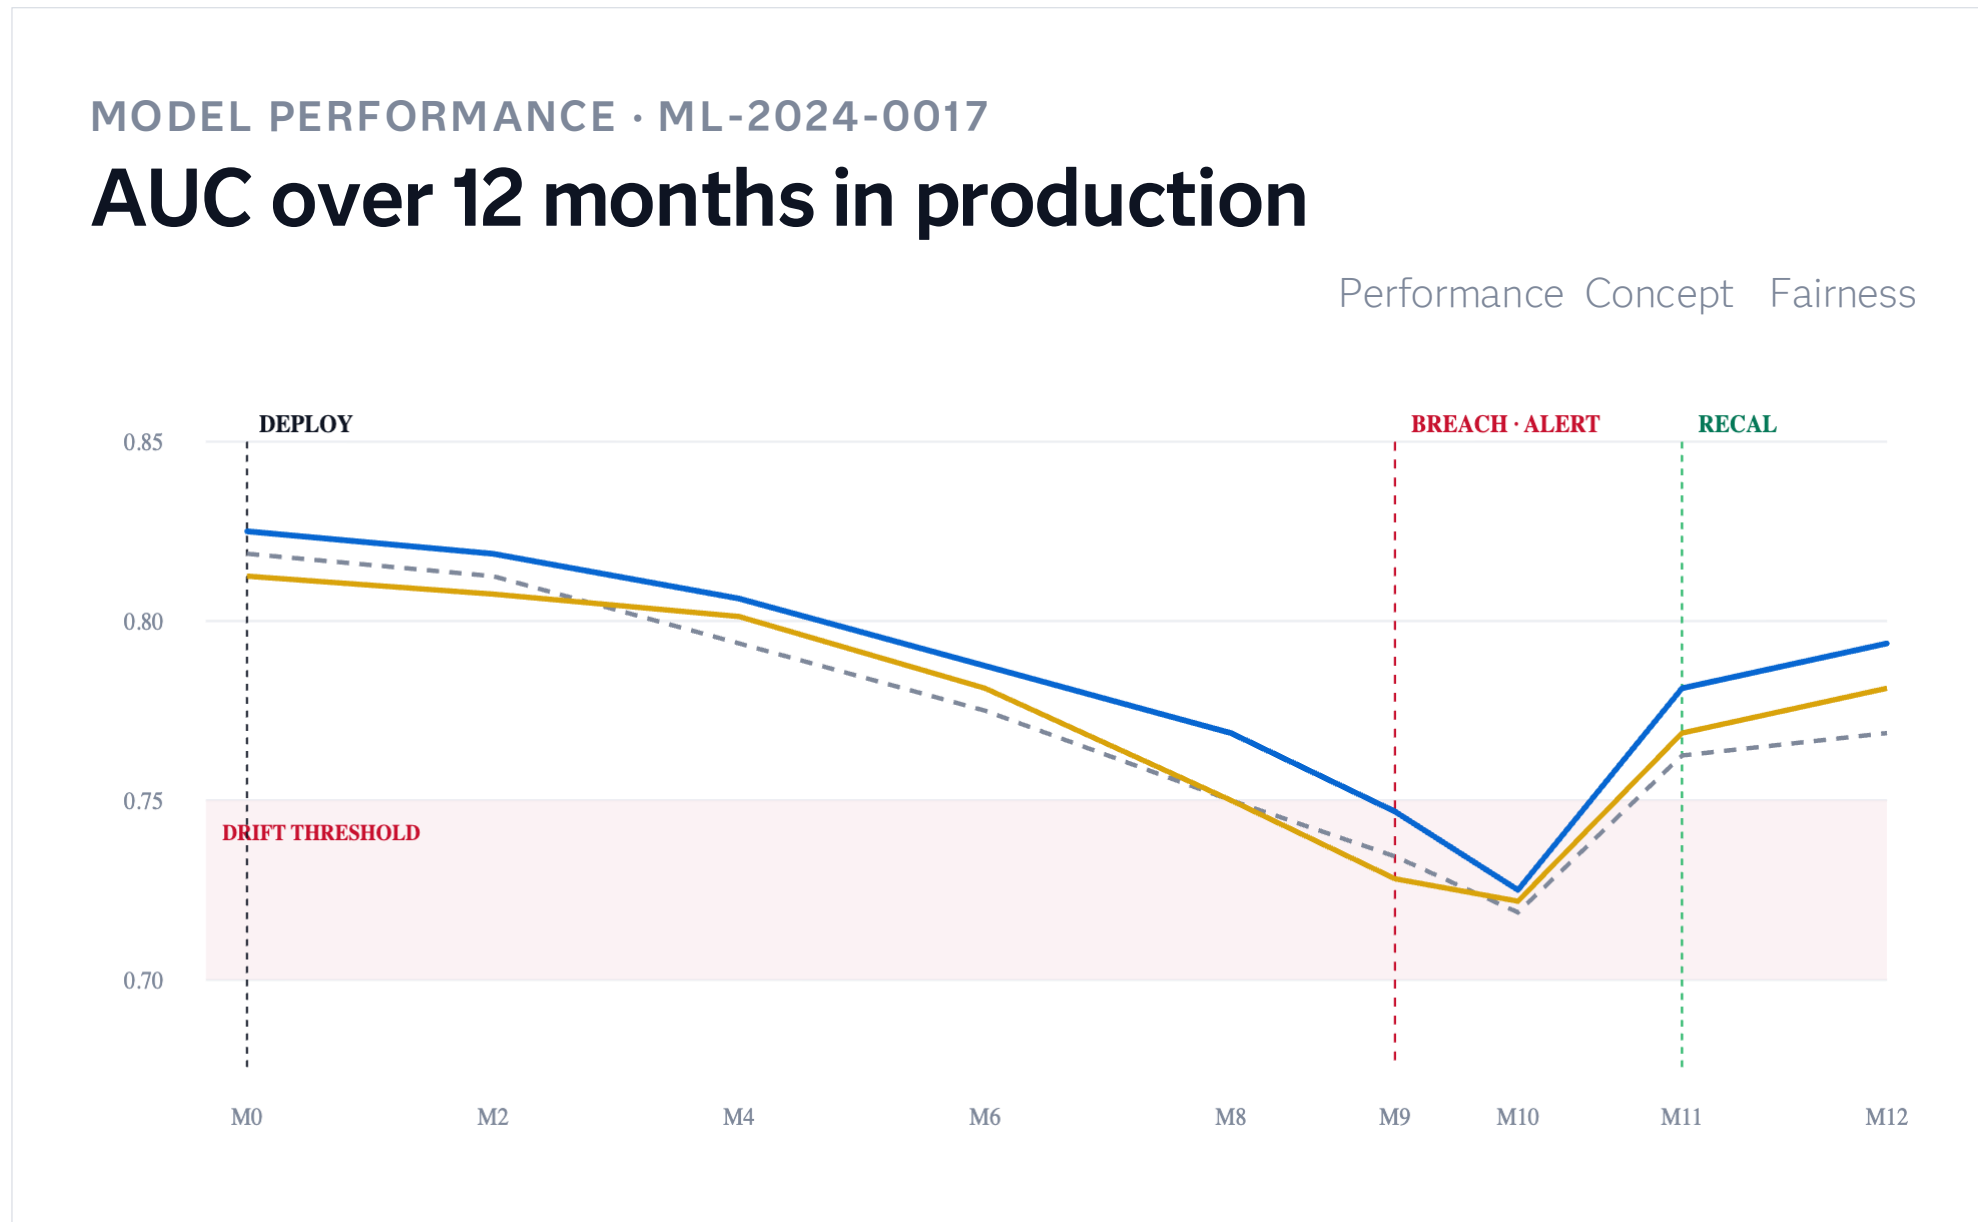
**You have the right to:**

- › Request a free copy of the consumer report used in this decision (FCRA § 612)
- › Dispute information you believe to be inaccurate or incomplete
- › **Request a manual review of this decision by a human underwriter, at no cost**
- › Receive additional information about the data sources and the model used

To exercise any of these rights, contact us at [REDACTED] or [REDACTED] Reference notice ID AAN-2026-04-021187. This notice is provided pursuant to the Fair Credit Reporting Act and applicable state insurance law. The model used in this decision is reviewed regularly for

PILLAR 04

# Three things drift. All three matter.



### 01 · PERFORMANCE DRIFT

AUC sliding because the world is no longer the world the model was trained on.

### 02 · CONCEPT DRIFT

Relationships between features and target are changing.

### 03 · FAIRNESS DRIFT

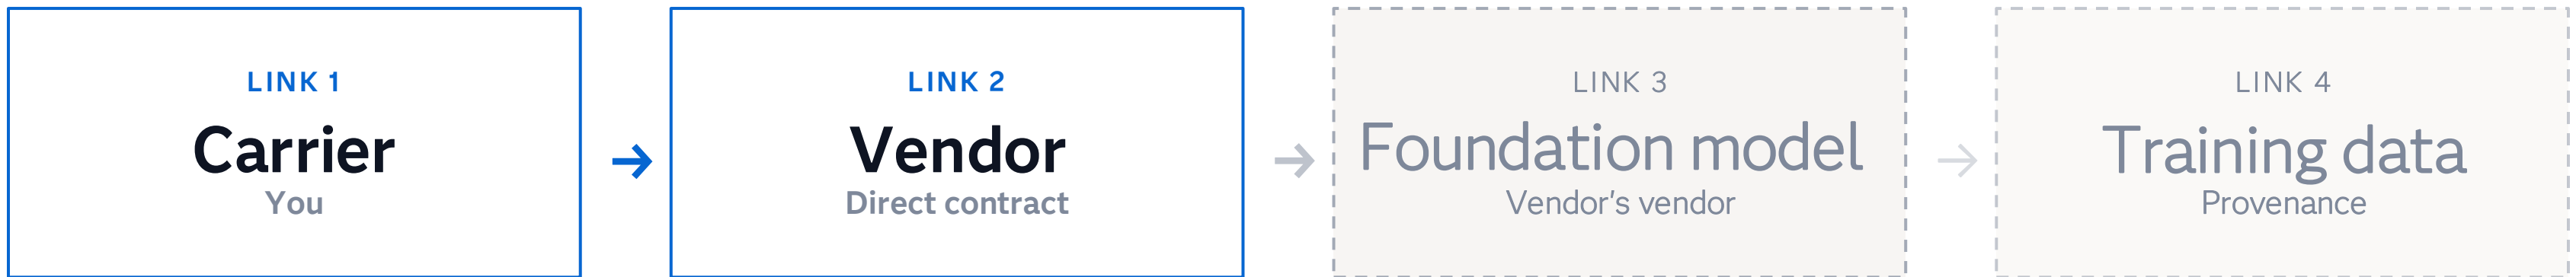
Unbiased at deploy ≠ unbiased at month nine.

*Annual revalidation = 12-month resolution.*

*Regulators expect daily-to-monthly.*

THIRD-PARTY / VENDOR ACCOUNTABILITY

# A four-link chain. You can see two.



THE REGULATOR'S LENS

You are responsible for **every link.**

- ✓ Right to audit
- ✓ Right to bias-test
- ✓ Cooperation clause
- ✓ Material-change notice

A practical test: pull your top three AI vendor contracts. If NY DFS issued a special report request tomorrow, could you comply?

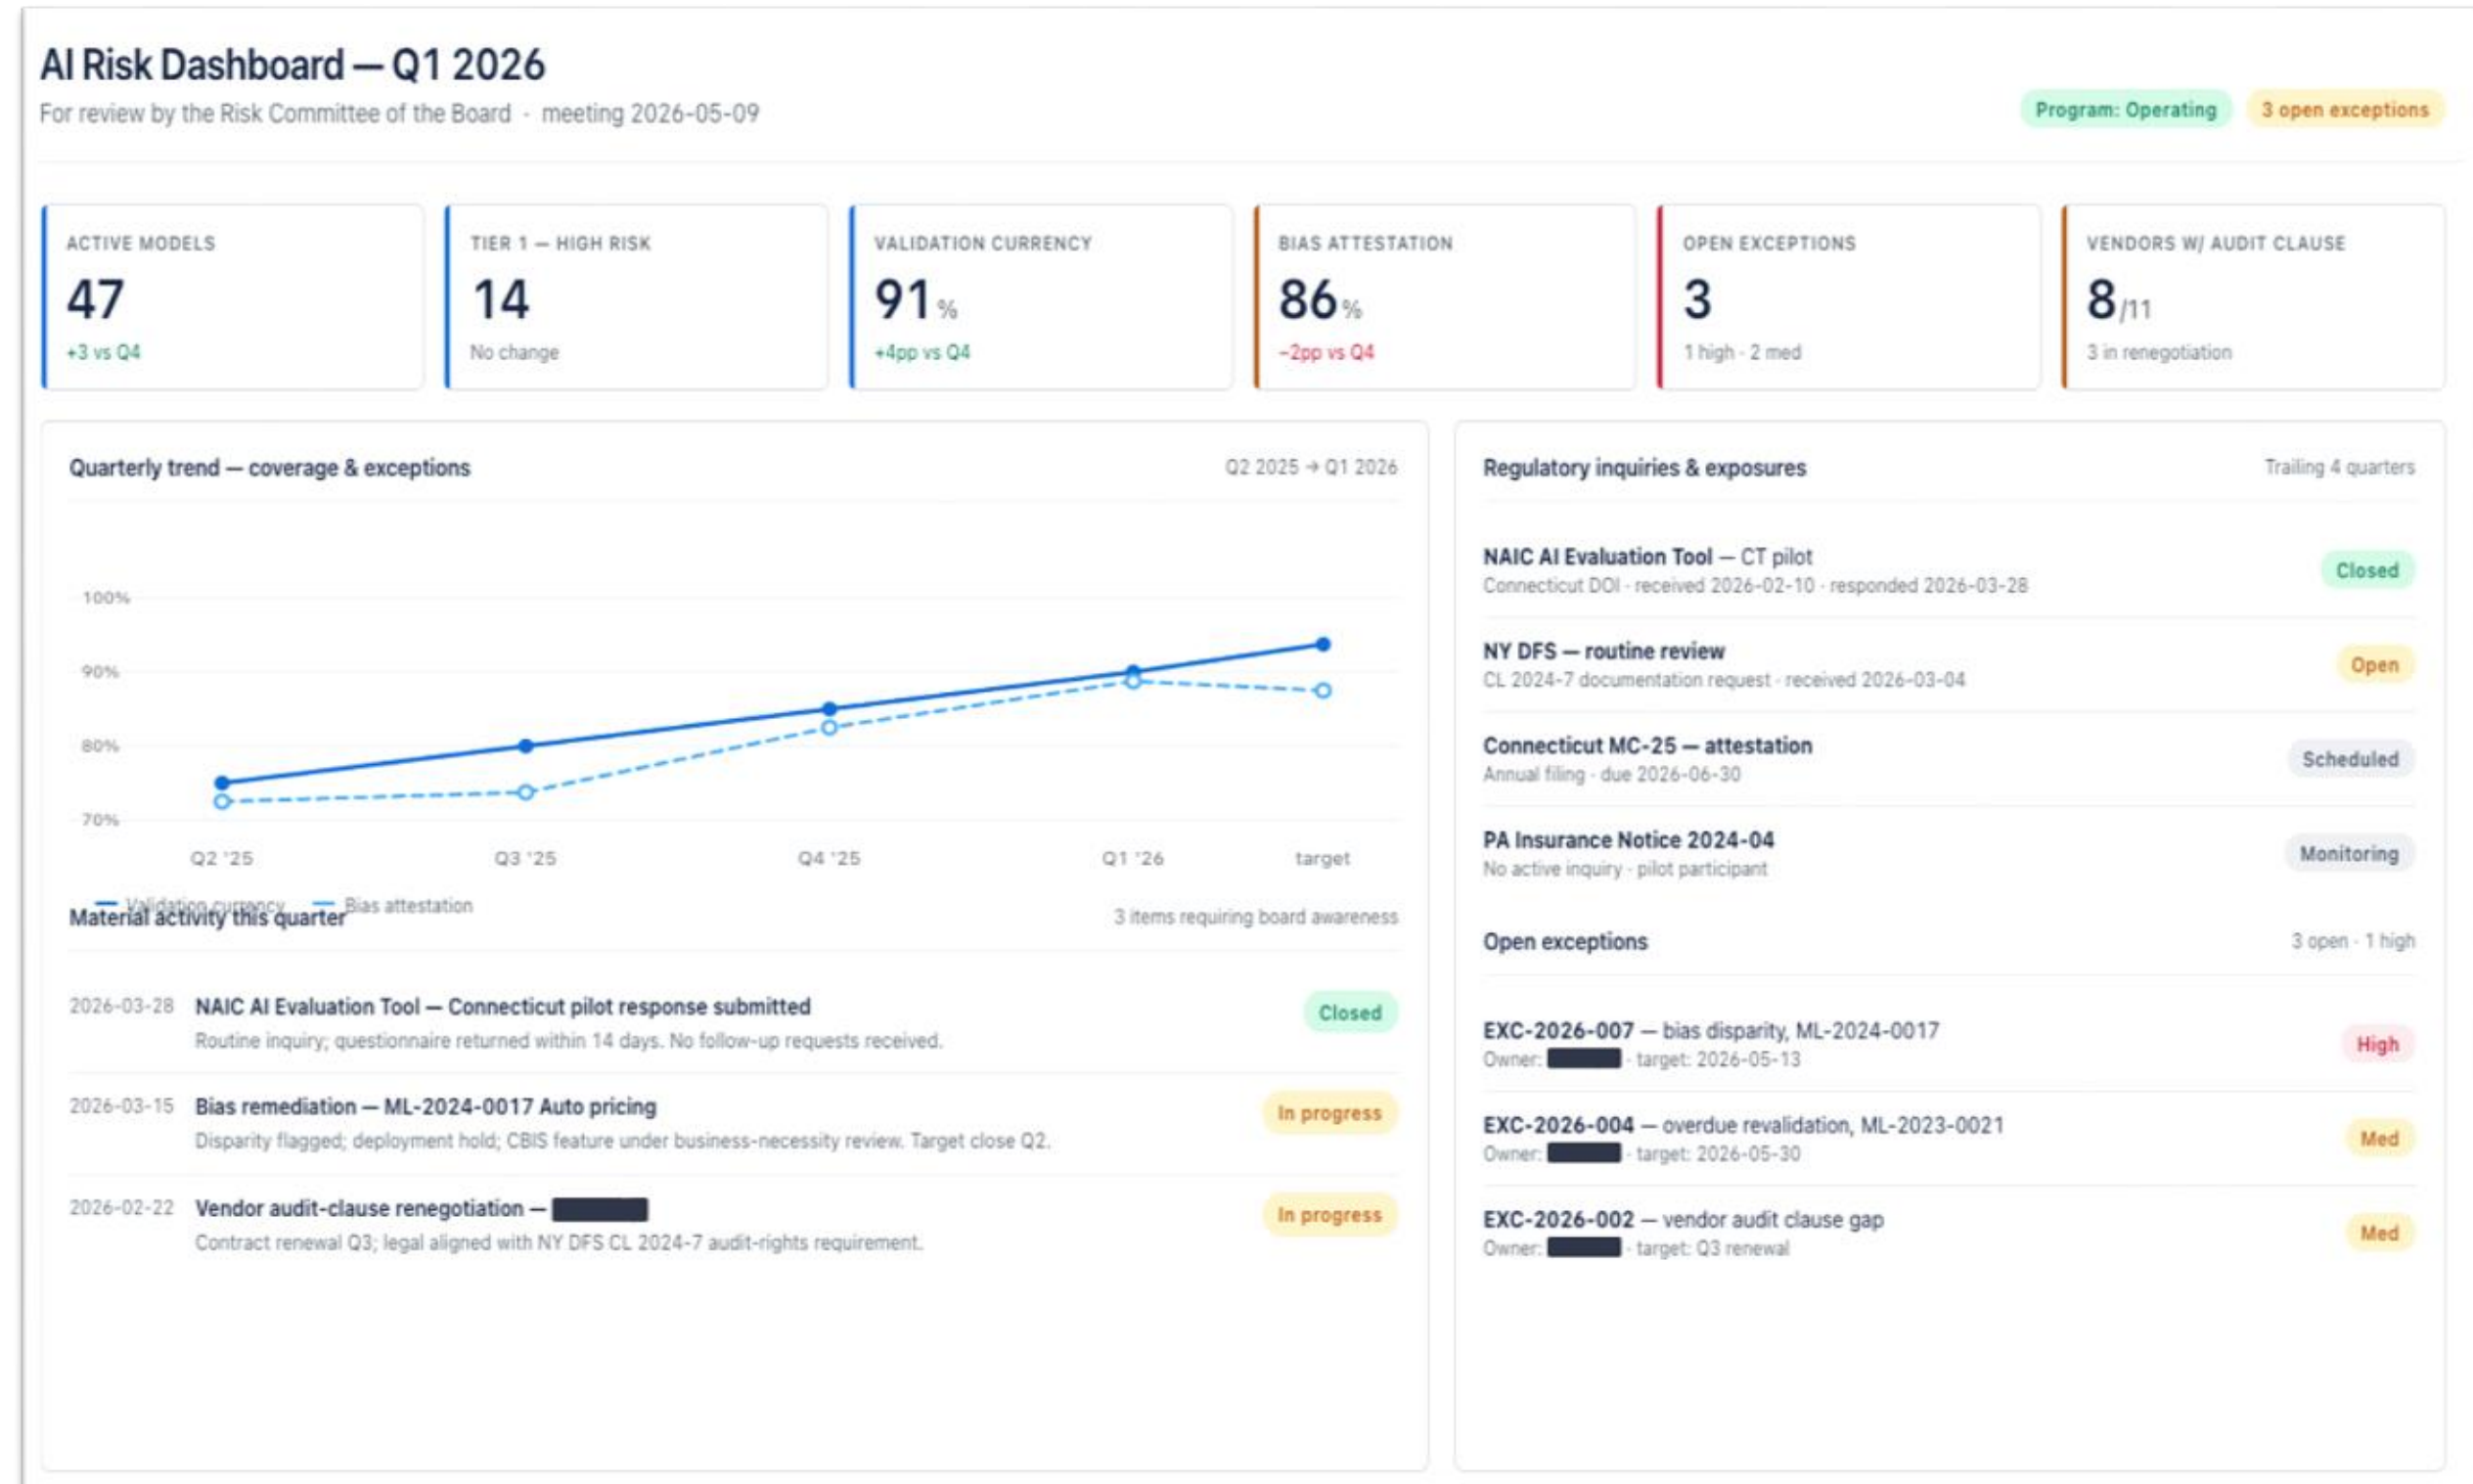
GOVERNANCE · OPERATIONALIZED

The board doesn't need **SHAP** values.

The board needs the dashboard to the right. Once a quarter, minimum.

Cross-functional AI Risk Committee.  
 Chaired by CRO or General Counsel — *not* the CDO.

*The Bulletin can require a written program. It cannot require that the program is believed.*



IF YOU REMEMBER NOTHING ELSE · THIS SLIDE

# The defensible operating model.

## PILLAR 01

### Inventory

- Risk-tiered
- Use case as unit
- Vendor models flagged
- Owners named

## PILLAR 02

### Testing

- At deployment
- At material change
- On cadence
- Documented decisions

## PILLAR 03

### Monitoring

- Continuous
- Threshold-driven
- Alert-coupled
- Pre-defined response

## PILLAR 04

### Governance

- Senior officer
- Cross-functional cmte
- Board reporting
- Quarterly cadence

## FOUNDATION

**Documentation: continuous, dated, available on demand.**

Most regulatory difficulty doesn't come from carriers who didn't do the right thing. It comes from those who can't prove they did.

CALLBACK · SLIDE 03

# It is June 2027.

She has the same six requests.

The redactions are gone.

---

Every one of these is something the regulators have *already told you* they would ask for.

SIX REQUESTS · REVEALED

- 01 Model inventory · with risk tiering and ownership

---

- 02 Bias testing results · top 10 Tier-1 models, current cycle

---

- 03 Adverse-action explanations · last 90 days

---

- 04 Vendor due-diligence file · top 5 AI vendors

---

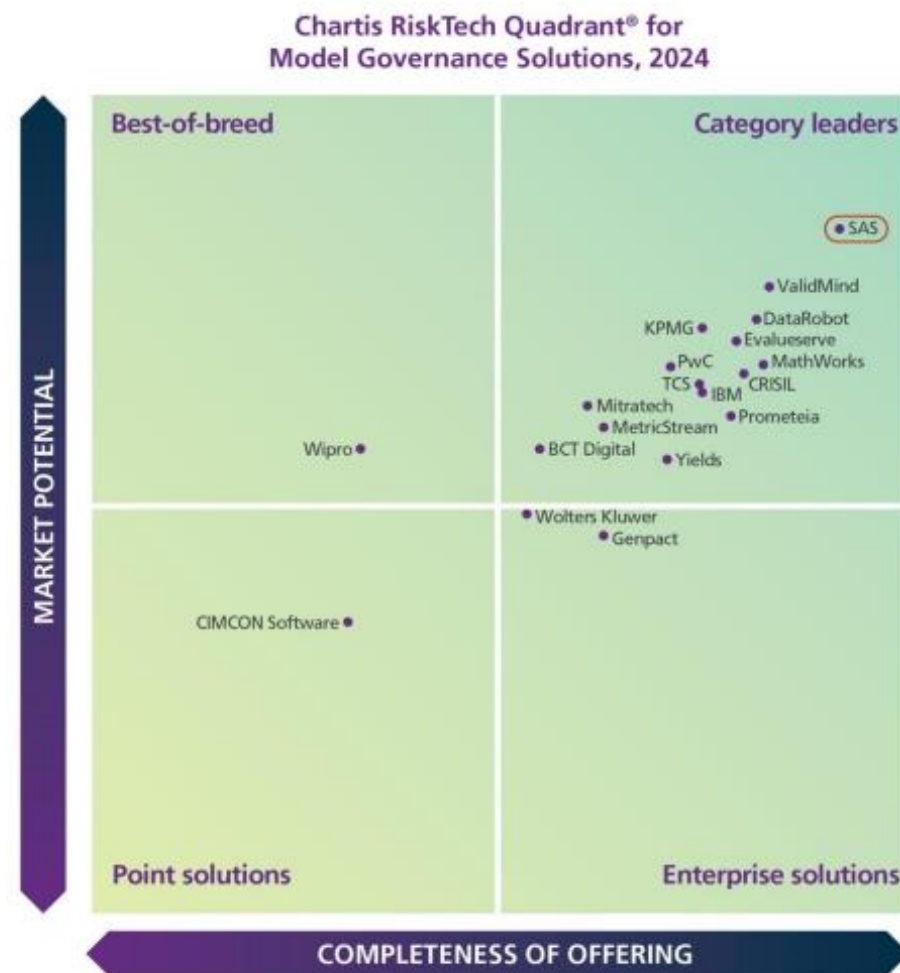
- 05 Validation reports · Tier-1 models, last 12 months

---

- 06 Board AI risk dashboard · last four quarters

CAPABILITY MAP · THE ONE SLIDE WHERE I NAME A PLATFORM

# Build vs buy. **By pillar.**

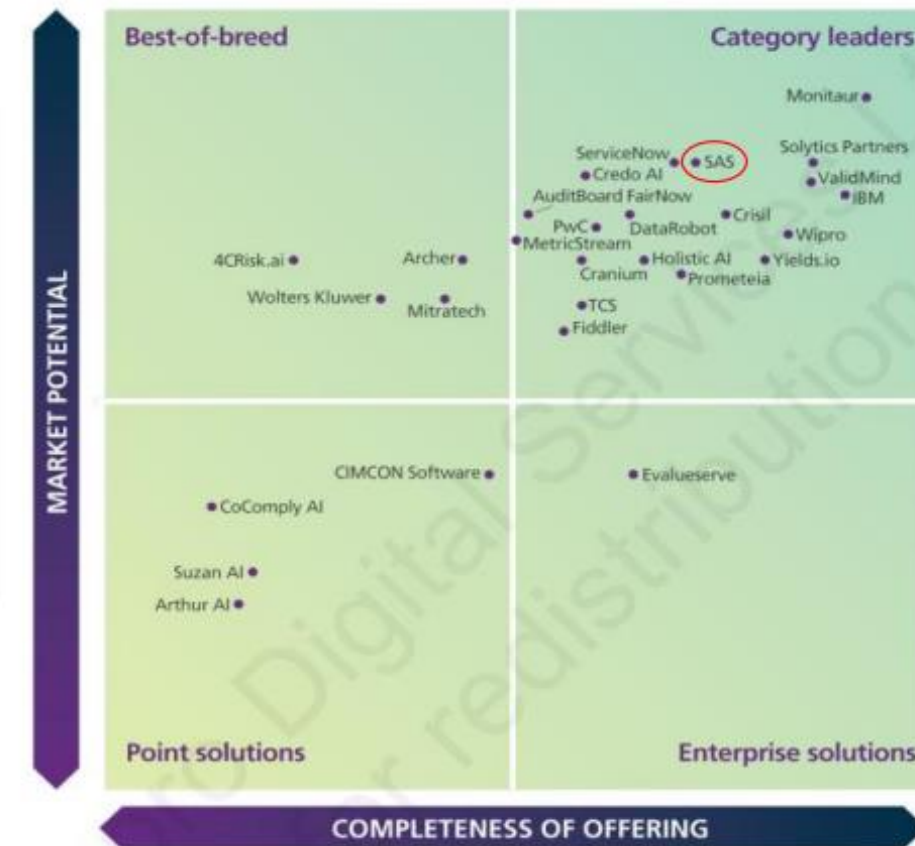


**Chartis RiskTech AI50 2025**  
#3  
SAS  
Responsible AI Adoption

**SAS ranked No. 3 overall and won 2 category awards in the 2025 Chartis RiskTech AI 50**



Figure 3: RiskTech Quadrant® for AI Gov solutions, 2025



Source: Chartis Research

“SAS is ensuring that its solutions remain relevant and effective in an ever-evolving landscape. By focusing on transparency, accountability and governance, SAS is well positioned to continue its leadership in leveraging AI to support institutions in growing, protecting and efficiently managing their operations.

THREE THINGS TO DO · THIS WEEK

# Monday morning.

---

## 01 Inventory your AI use cases.

A spreadsheet by Friday is better than a perfect framework next quarter.

---

## 02 Run a bias test on your highest-impact model.

Even a rough one. Document what you find — and what you decide.

---

## 03 Read your top three AI vendor contracts.

Specifically the audit clause. And the cooperation clause.

---

THREE TAKEAWAYS · THANK YOU

- 01** The regulators have published the **open-book exam answers**. Open the book.
- 02** AI governance is MRM++. Same skeleton. **Harder muscles**.
- 03** Build it in advance. **It is cheaper than building it under deposition.**

# Q&A

Max.Samadov@sas.com

