AI INSIGHTS FOR ACTUARIES

SOA Antitrust Compliance Guidelines

Active participation in the Society of Actuaries is an important aspect of membership. While the positive contributions of professional societies and associations are well-recognized and encouraged, association activities are vulnerable to close antitrust scrutiny. By their very nature, associations bring together industry competitors and other market participants.
The United States antitrust laws aim to protect consumers by preserving the free economy and prohibiting anti-competitive business practices; they promote competition. There are both state and federal antitrust laws, although state antitrust laws closely follow federal law. The Sherman Act, is the primary U.S. antitrust law pertaining to association activities. The Sherman Act prohibits every contract, combination or conspiracy that places an unreasonable restraint on trade. There are, however, some activities that are illegal under all circumstances, such as price fixing, market allocation and collusive bidding.

There is no safe harbor under the antitrust law for professional association activities. Therefore, association meeting participants should refrain from discussing any activity that could potentially be construed as having an anti-competitive effect. Discussions relating to product or service pricing, market allocations, membership restrictions, product standardization or other conditions on trade could arguably be perceived as a restraint on trade and may expose the SOA and its members to antitrust enforcement procedures.

While participating in all SOA in person meetings, webinars, teleconferences or side discussions, you should avoid discussing competitively sensitive information with competitors and follow these guidelines:

- Do not discuss prices for services or products or anything else that might affect prices
- Do not discuss what you or other entities plan to do in a particular geographic or product markets or with particular customers.
- Do not speak on behalf of the SOA or any of its committees unless specifically authorized to do so.
- Do leave a meeting where any anticompetitive pricing or market allocation discussion occurs.
- Do alert SOA staff and/or legal counsel to any concerning discussions
- Do consult with legal counsel before raising any matter or making a statement that may involve competitively sensitive information.

Adherence to these guidelines involves not only avoidance of antitrust violations, but avoidance of behavior which might be so construed. These guidelines only provide an overview of prohibited activities. SOA legal counsel reviews meeting agenda and materials as deemed appropriate and any discussion that departs from the formal agenda should be scrutinized carefully. Antitrust compliance is everyone's responsibility; however, please seek legal counsel if you have any questions or concerns.

Presentation Disclaimer

Presentations are intended for educational purposes only and do not replace independent professional judgment. Statements of fact and opinions expressed are those of the participants individually and, unless expressly stated to the contrary, are not the opinion or position of the Society of Actuaries, its cosponsors or its committees. The Society of Actuaries does not endorse or approve, and assumes no responsibility for, the content, accuracy or completeness of the information presented. Attendees should note that the sessions are audio-recorded and may be published in various media, including print, audio and video formats without further notice.

AI Insights for Actuaries

Moderator:

Eric Milner, FSA
President & Actuary
Baker Benefits Administrators

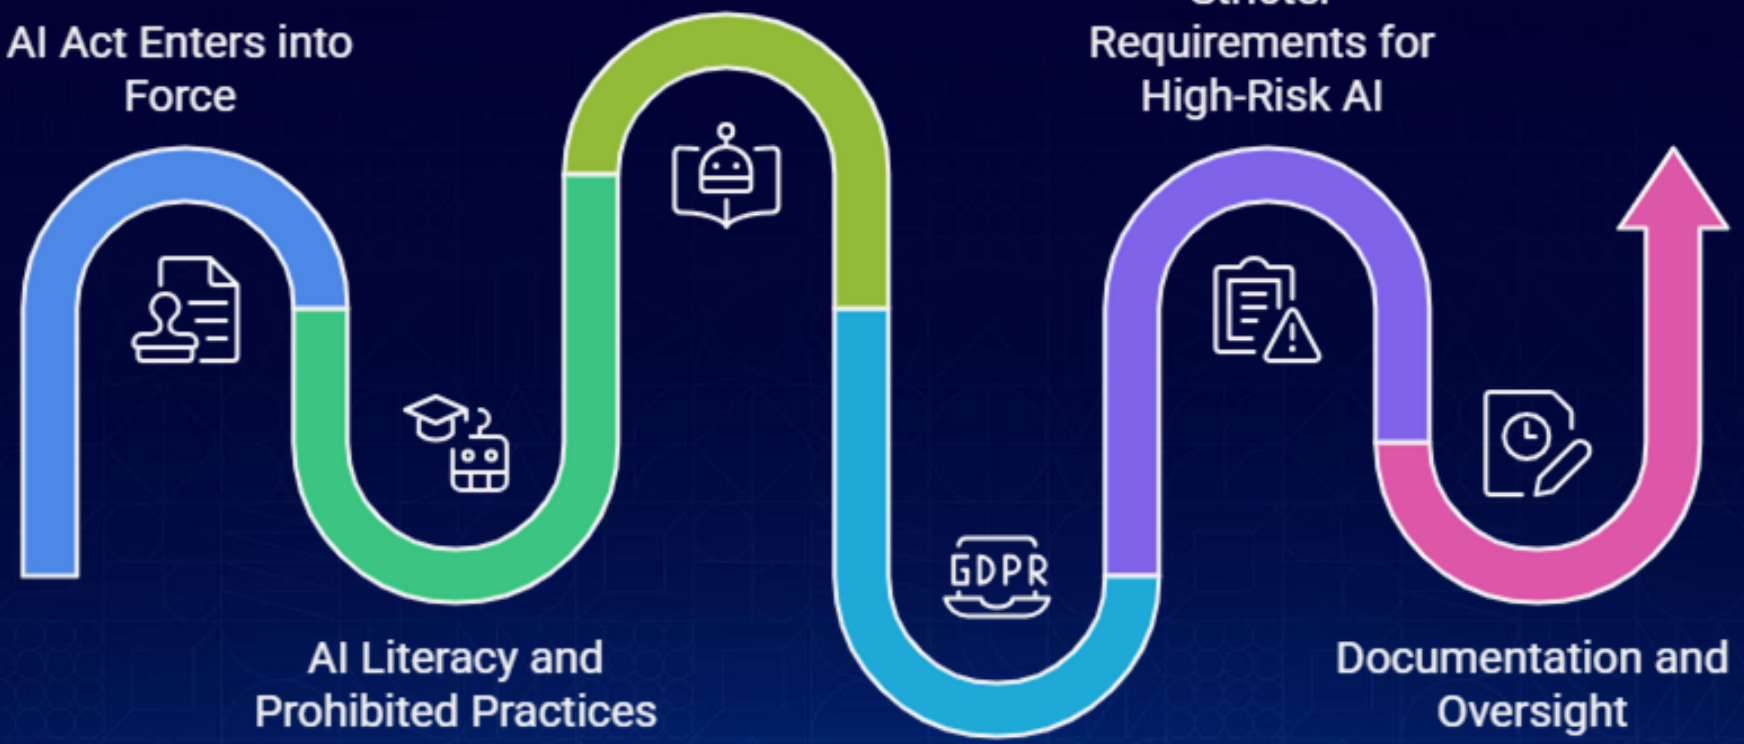
Presenters:

Carlos Arocha, FSA
Managing Partner
Arocha & Associates GmbH

Amit Jiwani
Chief Product Officer
Merit Medicine, Inc.

May 14, 2026

The EU AI Act


Aug 2024
AI Act Enters into Force
Feb 2025
AI Literacy and Prohibited Practices
Aug 2025
GPAI and Governance Rules
GDPR
General Application Date
Aug 2026
Stricter Requirements for High-Risk AI
Documentation and Oversight

Digital Operational Resilience Act (DORA)

Digital Operational Resilience Testing

Conduct regular tests to identify vulnerabilities and improve resilience.



ICT Third-Party Risk Management

Manage risks associated with third-party ICT providers to maintain resilience.

ICT-Related Incidents

Establish effective incident response plans to minimize impact.

ICT Risk Management

Implement robust strategies to mitigate ICT risks and ensure business continuity.

Information Sharing

Share information about threats and incidents to foster collective resilience.

Baseline: What was the Starting Point?

Before the **EU AI Act** and **DORA**, insurers were not operating in a regulatory vacuum. But the rules were dispersed across existing frameworks rather than centered on AI itself

- AI use in insurance was already subject to existing insurance-sector rules, including governance, risk management, conduct, and oversight expectations
- **GDPR** already mattered where automated processing could affect individuals in a material way
- ICT and cyber resilience expectations existed, but requirements were spread across multiple rules, guidelines, and supervisory practices
- Governance focused more on “models” and “systems” than on “AI” as a distinct regulatory category
- Result: firms often had partial controls, but not a unified AI accountability framework.

EU moved from indirect oversight to AI and resilience regimes

The EU has now made AI governance and digital resilience much more explicit, structured, and enforceable

- The **EU AI Act** creates a horizontal, risk-based regime for AI across sectors
- The Act entered into force on August 2024
- AI literacy obligations and prohibited AI practices started to apply on February 2025
- Rules on governance and general-purpose AI models started to apply on August 2025
- The Act becomes broadly fully applicable on August 2026, while certain high-risk AI systems embedded in regulated products have a transition until August 2027
- AI used for risk assessment and pricing for L&H insurance is treated as high-risk
- DORA has applied since January 2025 and creates a harmonized framework for ICT risk management, incident reporting, resilience testing, and third-party ICT risk oversight

This changes how actuarial AI must be governed

For actuaries, the question is no longer only whether AI performs well. It is whether the model, workflow, and operating environment are defensible

- Higher-performing AI is not enough; firms need evidence of governance, validation, and human oversight
- For some insurance use cases, especially L&H, AI may fall into a high-risk regulatory category
- **DORA** means the control environment around actuarial systems matters more: data pipelines, cloud tools, vendors, incident response, continuity, and access controls
- GenAI raises additional issues: confidentiality, hallucinations, weak reproducibility, and vendor dependency
- Non-European firms should pay attention, because multinational insurers, global vendors, and supervisory expectations often move across borders
- Practical implication: actuarial teams need stronger model inventories, clearer intended-use statements, better documentation, richer validation, and more disciplined monitoring

How actuarial model governance must evolve



Six Pillars of Defensible AI

A Governance Framework

Pillar 1: Purpose, Scope, and Fit-for-Use

A clear, written articulation of intended use and an equally clear statement of what the model is not for is necessary to ensure model development clearly adheres to outcome expectations.

KEY PRACTICES

- Document intended use, target population, and the decisions it informs
- Explicitly name out-of-scope uses to prevent scope creep
- Confirm AI is the right tool vs. a simpler, more transparent alternative
- Verify the application population matches the training population

THE CORE QUESTION

Should this model exist, in this form, for this decision?

Pillar 2: Data Lineage, Quality, and Representativeness

You cannot defend a model whose data you cannot trace. Traceability of the data and control over the data universe used to train the model are crucial to the veracity of the output.

KEY PRACTICES

- Maintain a data dictionary mapping each feature to its source and transformations
- Test for representativeness across the dimensions that matter for your use case
- Establish a refresh cadence; ensure safe failure on stale or missing data
- Secure contractual visibility and audit rights for any third-party data

THE CORE QUESTIONS

Where did the data come from?

How was it prepared?

Who is represented and who is not?

Pillar 3: Development, Validation, and Independent Review

Defensible development requires a documented process, validation that includes adversarial testing, and review by an organizationally independent party.

KEY PRACTICES

- Follow a written model development standard
- Ensure validation is performed independently of the build team
- Test at the tails and edge cases—not just aggregate accuracy
- For generative AI, build explicit hallucination detection into validation

THE CORE QUESTION

Who is validating the model, and is it someone who didn't build it?

Pillar 4: Documentation, Explainability, and Communication

A model you cannot explain is a model you cannot defend. Explainability should be commensurate with the stakes of the decision.

KEY PRACTICES

- Produce a one-page summary readable by executives, auditors, and regulators
- Implement feature-level explainability for consequential individual decisions
- Maintain a model card that names known limitations
- Document what the model cannot do as carefully as what it can

THE CORE QUESTION

Can the right person get the right answer at the right level?

Pillar 5: Monitoring, Drift, and the Feedback Loop

Models decay. Defensible governance defines in advance what “no longer working” looks like and detects it before stakeholders do.

KEY PRACTICES

- Define triggers in advance for retraining, pausing, or retiring the model
- Assign clear ownership of monitoring and a documented escalation path
- Monitor inputs, outputs, and downstream outcomes—not just outputs
- Distinguish data drift, concept drift, and performance drift—each requires a different response

THE CORE QUESTION

How will you know when the model is no longer working?

Pillar 6: Accountability & Human-in-the-Loop

The Actuarial Code of Conduct does not have a carve-out for “the model did it.” Named human accountability is the profession’s most powerful card.

KEY PRACTICES

- Name a human accountable for each consequential model output
- Require meaningful human review at high-impact decision points
- Log overrides; treat them as a leading indicator of model issues
- Train staff on professionalism standards as they apply to AI-assisted work

THE CORE QUESTION

Whose name is on this output—and what does the Code require of them?

The Framework At A Glance

01 Purpose & Fit-for-Use

Should this model exist, in this form, for this decision?

WATCH FOR

Scope creep • wrong-population use • AI used where simpler tools fit

02 Data Lineage & Quality

Where did the data come from?

How was it prepared?

Who is represented and who is not?

WATCH FOR

Untraceable features • representativeness gaps • stale data

03 Validation & Independent Review

Who is validating the model—and is it someone who didn't build it?

WATCH FOR

Self-review • only aggregate accuracy • no hallucination checks

04 Documentation & Explainability

Can the right person get the right answer at the right level?

WATCH FOR

Black-box outputs • no per-case explanation • unstated limits

05 Monitoring & Drift

How will you know when the model is no longer working?

WATCH FOR

No retirement plan • only output monitoring • undefined drift triggers

06 Accountability & Professionalism

Whose name is on this output—and what does the Code require of them?

WATCH FOR

"The model did it" • no override logging • untrained staff

Weakness in any one pillar undermines the others. Treat governance as an operating discipline across the model lifecycle.

Questions?

Please remember to complete the evaluation:


Society of
Actuaries®