SOA Research Institute
SOCIETAL PURPOSE

Actuarial Innovation and Technology

AI Risk Management Frameworks: An Expert Panel Discussion

APRIL | 2025

AI Risk Management Frameworks

An Expert Panel Discussion

AUTHOR Ronald L Poon Affat, FSA, FIA, MAAA

SPONSOR Actuarial Innovation and Technology Strategic Research Program

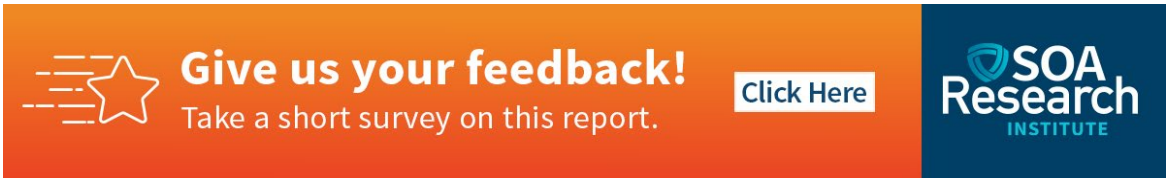

CONTENTS

AI Risk Management Frameworks

An Expert Panel Discussion

Executive Summary

The Society of Actuaries[1]' Artificial Intelligence Working Group convened a panel of experts to explore the role of Artificial Intelligence (AI) in actuarial practice, with a particular focus on risk management, regulatory compliance, and ethical considerations. AI has the potential to revolutionize actuarial science, insurance, and healthcare by enhancing predictive modeling, automating decision-making, and improving efficiency. However, as AI adoption grows, so do concerns about fairness, transparency, security, and unintended biases in AI-driven models.

A key focus of this discussion was the National Institute of Standards and Technology (NIST[2]) AI Risk Management Framework (AI RMF[3]) and its Generative AI Profile[4], which provide guidelines for developing and deploying AI responsibly. These frameworks are particularly relevant to industries where actuarial work plays a crucial role, including insurance and healthcare, where AI is increasingly being used for underwriting, risk assessment, claims processing, and fraud detection.

The NIST AI RMF is a non-regulatory framework developed by NIST to help organizations identify, assess, and mitigate AI-related risks. The Generative AI Profile extends this framework to address risks specific to generative AI models, such as large language models (LLMs) and deep learning systems. These frameworks are important for actuaries, insurers, and healthcare professionals because they:

- Support Risk Governance and Compliance: Aligning AI risk management with existing frameworks, including those set by the National Association of Insurance Commissioners (NAIC[5]) and international standards bodies.
- Mitigate Bias and Promote Fairness: Ensuring AI-driven models do not introduce unintended biases that could lead to unfair discrimination in insurance pricing, claims, or healthcare outcomes.
- Enhance Model Transparency and Explainability: Addressing the "black box" nature of AI models to improve actuarial oversight, regulatory reporting, and policyholder trust.
- Strengthen Data Integrity and Cyber Security: Ensuring secure data handling in AI applications, particularly in life insurance and healthcare, where sensitive personal data is involved.
- Improve AI Reliability and Resilience: Encouraging ongoing monitoring, stress testing, and validation of AI models to ensure long-term reliability in risk modeling and financial forecasting.

The expert **panel** engaged in a holistic discussion, addressing key aspects of AI risk management, industry best practices, and emerging regulatory challenges. This report provides a high-level summary of the panel's insights from the discussion held on March 12, 2025. A list of discussion topics is included in Appendix A.

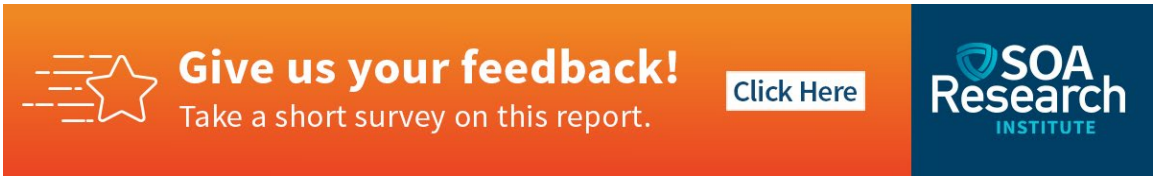

Section 1 AI Risk Management Essentials

1.1 NIST AI FRAMEWORK: NON-REGULATORY APPROACH

The panel began by discussing the NIST AI Risk Management Framework, which is designed to provide a risk-based approach to AI governance. The framework is non-regulatory but aims to serve as a guide for organizations in managing AI risks effectively. It aligns with traditional risk management models, such as International Organization for Standardization (ISO[6]) 31000 and NIST SP 800-37 for cybersecurity, yet it restructures traditional risk management functions to suit AI governance challenges. It is adaptable across various industries, offering flexibility rather than a rigid sequential process.

The framework's strengths lie in its risk-based thinking, as it categorizes AI risks at different levels, allowing for more granular identification and mitigation of AI failures. Its alignment with ISO risk management standards ensures global applicability, while its adaptability makes it easier to integrate across multiple industries without imposing strict compliance requirements.

1.2 CHALLENGES IN AI RISK MANAGEMENT

The panel observed that regulatory and policy landscapes play a crucial role in shaping AI governance, influencing how organizations approach responsible AI implementation and risk management. In jurisdictions where AI risk governance is still evolving, companies may operate with varying levels of oversight, which can present challenges in ensuring consumer protection and responsible AI deployment. To address this, the panel emphasized that self-governance is essential, as companies must implement internal AI governance frameworks to prevent unethical AI applications. Some insurers may treat AI risk management as an audit function. However, this approach can be narrow, as it may overlook critical aspects such as ethical implications and the potential for social injustice within AI-driven decisions.

Another key issue discussed was the absence of a standardized harm matrix for AI risks in the United States of America (USA). While the European Union (EU) AI Act[7] categorizes AI risks into low, medium, and high tiers, U.S. frameworks take a more decentralized approach, with ongoing efforts to develop a structured methodology for defining and assessing AI-related harm. Additionally, third-party data and bias risks remain a consideration, as many companies utilize external data sources that may not always undergo extensive evaluation, potentially leading to unintended bias and fairness challenges in AI decision-making.

1.3 ALIGNING AI WITH ACTUARIAL STANDARDS

From an actuarial perspective, the AI Risk Management Framework closely aligns with existing risk management principles used in actuarial practices. The framework reinforces actuarial model governance practices, ensuring AI models are applied responsibly. Actuaries must adhere to professional guidelines that include ASOPs[8], such as ASOP No. 56 on modeling, ASOP No. 23 on data quality, and ASOP No. 12 on risk classification, which already share principles of transparency, accountability, and fairness. The panel noted that actuaries play a critical role in AI model validation and governance. A key challenge is the lack of AI transparency, particularly in large-scale generative AI systems, which often lack interpretability. This creates difficulties in regulatory compliance and risk communication with stakeholders.

1.4 AI STANDARDS INSURANCE

The discussion also emphasized the importance of integrating global AI standards with actuarial professional guidelines. Many countries lack formal AI risk taxonomies or industry-specific regulations, which can create uncertainty in applying AI-driven models within actuarial practices. The panelists highlighted the varying levels of AI

governance maturity across different nations. Some jurisdictions, such as the European Union, have established comprehensive regulatory frameworks, while others, including Australia, are in the process of developing AI-specific legislation to complement existing regulatory approaches

1.5 BIAS, FAIRNESS, AND AI GOVERNANCE

The panel addressed concerns regarding bias in AI models, emphasizing that while NIST acknowledges statistical, cognitive, and systemic biases, it lacks a formal definition of bias. They also stressed that improving communication standards is necessary to help companies effectively explain AI decisions to regulators and stakeholders. Furthermore, they noted that the actuarial profession should work towards developing a standardized approach to AI risk assessment to align with NIST's principles.

1.6 STRENGTHENING GOVERNANCE

The panel underscored the importance of governance structures in ensuring AI models function responsibly. They emphasized that AI models must not replace professional judgment but should instead serve as tools to enhance decision-making. The strength of an AI governance framework determines the reliability of the models, making documentation, audit trails, and independent oversight critical components of AI governance. Actuarial risk classification principles must be applied to AI models to prevent biased decision-making, while bias testing in AI models needs to be strengthened with methodologies that are more robust.

1.7 KEY TAKEAWAYS

The AI Risk Management Framework is non-regulatory but aligns with international risk standards, making it adaptable for global implementation. Takeaways from the panel discussion include:

- AI risk governance benefits from extending beyond compliance audits to thoughtfully incorporating ethical and fairness considerations.
- As regulatory frameworks continue to develop, self-regulation can serve as a valuable complement in jurisdictions where government oversight is still evolving.
- Further clarity on the oversight of third-party data usage could help address potential biases and enhance fairness in AI decision-making.
- In the insurance industry, the integration of AI standards with actuarial guidelines remains a work in progress.

AI models must maintain transparency, particularly in industries such as insurance, where fairness in financial decision-making is crucial. The actuarial profession must integrate AI risk assessment into professional standards to ensure consistency and accountability in AI-driven decision-making. The panel discussion reinforced that AI risk management requires a comprehensive approach that includes governance, transparency, ethical considerations, and strong professional oversight. The NIST AI Risk Management Framework serves as a foundational guide, but further development is needed to address bias, regulatory alignment, and professional standards for AI governance. Ensuring AI trustworthiness requires an interdisciplinary effort, combining actuarial expertise, ethical considerations, and risk-based methodologies to navigate the rapidly evolving AI landscape.

Section 2 Standards and AI Governance

2.1 THE AI STANDARDS LANDSCAPE

The panelists noted that AI standards development does not follow a strict hierarchy. Instead, it is composed of multiple organizations, both longstanding and emerging, that focus on different aspects of AI and technology standardization. Key points include:

- Types of organizations: Some have been involved in setting technical and governance standards for decades, while others have recently emerged to address the complexities of artificial intelligence.
- Primary function: Standards organizations create documents outlining best practices, technical specifications, and regulatory frameworks.
- Recognized standards bodies: Among the most recognized are the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC)[9].

The panel discussed the more mature cybersecurity standards landscape since cybersecurity practices have been in place for a longer time. It has produced a rich body of work including standards, compliance schemes, open source and private sector tools, best practices and guidance, and product-specific benchmarks. In contrast, AI standardization is still in its early stages and is primarily focused on governance and compliance rather than technical implementation. AI frameworks remain in a phase of development, with an emphasis on risk-based approaches to AI governance rather than prescriptive technical guidelines.

2.2 NIST'S ROLE IN STANDARDS

The panel underscored NIST's pivotal role in AI governance. Though it does not set formal standards, NIST's guidelines—particularly its' AI Risk Management Framework—shape policy and industry practices by aligning with ISO standards and global best practices. Unlike international bodies, NIST refines its frameworks through stakeholder feedback rather than formal voting, positioning it as a bridge between U.S. and international regulatory efforts.

2.3 WHY TECH LEADERS SHAPE STANDARDS

One of the primary reasons why major technology companies engage in international standards development is to ensure that their AI systems function across multiple regions without the need for extensive modifications. Standardized approaches enable interoperability, allowing companies to deploy AI-driven solutions seamlessly across various markets. Adhering to global standards reduces the need for localized adaptations, thereby streamlining deployment processes and minimizing operational complexities.

2.4 REGULATORY EFFICIENCY AND COMPLIANCE COSTS

The panel discussed the importance of aligning with established international standards to simplify regulatory compliance. By conforming to globally recognized frameworks, companies can reduce the costs associated with meeting diverse national regulations. Standardized AI governance models eliminate the need to re-engineer, re-audit, and re-test systems for different markets, allowing for more efficient resource allocation, while ensuring legal compliance across jurisdictions.

However, the panel acknowledged that AI regulations are still evolving and may vary significantly between regions, unlike traditional financial regulations with well-established guidelines. This dynamic landscape makes it essential for companies to adopt a proactive and adaptable approach to AI model development. By taking a forward-thinking

stance, organizations can stay ahead of compliance requirements while balancing robust compliance measures with adaptive strategies to navigate the changing regulatory environment.

2.5 MARKET ACCESS AND COMPETITIVE ADVANTAGE

Governments often favor businesses that adhere to internationally recognized standards when approving AI-driven products and services. Companies that actively contribute to standards development can shape the industry's best practices, while ensuring their AI models meet regulatory expectations, which also positions them as thought leaders in AI governance.

2.6 RISK MANAGEMENT AND INDUSTRY INFLUENCE

By participating in AI standards development, major technology firms can influence the regulatory landscape in ways that align with their business models. Engaging in these discussions allows them to proactively address risk management concerns.

2.7 AI REGULATION CHALLENGES

AI governance within the insurance sector presents unique challenges, as the industry is already known for its opacity in decision-making. The introduction of AI may add further complexity, making transparency a major regulatory concern. The panelists pointed out that many governments take a patchwork approach to AI regulation by repurposing existing laws. This creates an inconsistent regulatory environment where compliance requirements vary significantly across jurisdictions.

2.8 RATING AGENCIES AND AI

The panel discussed the role of rating agencies in shaping AI governance, particularly within the insurance sector. While rating agencies do not set formal standards, they assess risks associated with AI implementation and influence how insurers and other financial institutions manage AI governance. Companies that demonstrate robust AI governance practices can improve their credit ratings, making compliance with AI risk frameworks financially advantageous.

Key risks evaluated by rating agencies include:

- Errors in AI decision-making.
- Over-reliance on AI-generated outputs.
- Cybersecurity vulnerabilities arising from AI adoption.
- Disruptions in business models due to automation.
- Governance concerns regarding AI accountability.

Given the financial implications of these risks, insurers must ensure that their AI governance practices align with industry standards to maintain their competitive standing.

2.9 AVOIDING BLACK-BOX AI MODELS

One of the concerns raised was the risk of allowing AI models to operate as "black boxes," where decision-making processes become opaque and difficult to interpret. The panel stressed that actuaries must maintain clear audit trails and ensure ethical oversight in AI model development. Human judgment must always remain a critical component of AI-driven decision-making to prevent unintended biases and errors.

2.10 KEY TAKEAWAYS

The panel discussion highlighted the following points:

- AI standards development is a decentralized process involving multiple organizations. NIST publications play a crucial role.
- Companies engage in international standards development to ensure technical and regulatory interoperability, coherence, and efficiency.

Section 3 Adopting the NAIC Model Bulletin

3.1 BIAS TESTING IN INSURANCE

The panel began by discussing the National Association of Insurance Commissioners (NAIC[10]) Model Bulletin on AI Use in Insurance, which has now been adopted by 23 states, reflecting a growing regulatory focus on bias mitigation and fairness in AI-driven insurance practices. However, the implementation of effective bias testing remains a significant challenge, with one of the primary obstacles being the lack of a clear, universally accepted definition of bias.

Key challenges include:

- Ambiguity in defining bias: Without a precise definition, insurers struggle to align their AI models with regulatory expectations, leaving them uncertain about what constitutes compliant AI decision-making.
- Functional vs. harmful biases: While biases can serve a functional role in human decision-making, the industry's primary concern centers around identifying and mitigating harmful biases that could lead to discrimination and unfair treatment of policyholders.
- Distinguishing biases: Differentiating between necessary biases in AI modeling and problematic biases in decision-making remains an ongoing challenge.

The panel noted that while bias is often framed as an inherently negative factor, it is not always detrimental. From a cognitive perspective, biases can help filter and process large amounts of information.

3.2 IDENTIFYING PROTECTED CLASSES

A significant challenge in bias testing is that insurers often lack access to demographic data that would allow them to directly assess disparities in AI-driven decision-making. Unlike other industries, such as e-commerce, where companies can request race, gender, or disability status information directly from users, insurers operate in a regulatory environment that discourages the collection of sensitive demographic data. This limitation makes direct bias assessment exceedingly difficult.

To infer protected class status, insurers commonly rely on methods such as Bayesian Improved Firstname Surname Geocoding (BIFSG[11]), which predicts race based on surname and geographic location. However, the panel noted concerns about using 2007 mortgage data, which utilized first name data because it reflects a limited number of mortgages to certain ethnic groups. Additionally, factors such as surname changes, interracial marriages, and geographic mobility contribute to inconsistencies in demographic inference, introducing a margin of error that can affect assessments of bias. Due to these challenges, many insurers are cautious about relying on inference-based bias testing, as incorrect classifications could lead to additional complexities.

3.3 RESISTANCE TO BIAS TESTING

One of the key challenges insurers face in conducting bias testing is navigating regulatory expectations. The panel noted that some insurers take a cautious approach, as identifying potential bias within AI models could lead to increased regulatory scrutiny. If evidence of unintended discrimination is found, companies may need to adjust pricing models, underwriting strategies, or decision-making frameworks, which could require additional resources and operational changes.

As a result, insurers often focus on statistical bias assessments that align with existing regulatory guidance, prioritizing methods that are well-established and practical to implement. While more advanced bias detection techniques, such as multi-objective optimization, are available, they are not widely adopted due to their complexity

and the potential need for broader model adjustments. This measured approach to bias testing aims to ensure compliance, while managing the operational impact of integrating more detailed bias detection methodologies

3.4 SHORTCOMINGS IN BIAS TESTING METHODS

The panel discussed the challenges associated with current bias testing frameworks in insurance, noting that many companies use statistical techniques that may not fully capture the complexity of variable interactions or potential proxy discrimination. Current approaches often focus on regression-based models that assess whether adding a protected class variable significantly affects model outcomes. Some states, such as Colorado, have introduced AI bias testing requirements that emphasize variable-based assessments. However, the panel observed that these methods might not fully account for confounding factors that could influence bias in AI-driven decision-making.

One consideration raised was that insurers often prioritize demonstrating that race or gender is not explicitly included in their models, while it can be more difficult to assess whether bias is indirectly embedded in correlated variables. AI systems may not use race as an input, but factors such as ZIP codes, credit scores, and occupation can sometimes serve as indirect indicators of race or socioeconomic status. This complexity makes it challenging to identify and address potential bias through traditional bias testing approaches alone

3.5 IMPROVING BIAS FRAMEWORKS

To enhance bias testing frameworks in insurance, the panel highlighted the benefits of adopting a more structured and transparent approach to AI governance. A key recommendation was for regulators and industry stakeholders to collaborate on developing clearer and more actionable standards. Rather than relying on broadly defined principles, insurers would benefit from specific guidelines outlining how bias should be measured, which metrics should be used, and what mitigation strategies should be considered.

The panel underscored the value of regulatory case studies in helping insurers better understand compliance expectations in practice. The NAIC and other regulatory bodies could play a valuable role by providing structured documentation guidance, supporting insurers in demonstrating due diligence in bias detection and mitigation. Strengthening industry-wide consistency in these efforts could also help reduce regulatory uncertainty—which might otherwise discourage companies from engaging in comprehensive bias testing.

Beyond regulatory improvements, the panel emphasized the importance of greater transparency in AI model development. Insurers would benefit from integrating bias mitigation efforts throughout the entire AI lifecycle, ensuring that fairness considerations are addressed from data collection through model training and deployment. This proactive approach could contribute to more accountable and equitable AI systems.

Lastly, the panel pointed to the need for continued dialogue on AI fairness within the insurance industry. The lack of a shared framework for defining and enforcing fairness can make it challenging for insurers to interpret and implement bias testing consistently. Establishing clear, industry-wide best practices will be essential in ensuring that AI-driven insurance decisions remain both effective and aligned with ethical considerations.

3.6 KEY TAKEAWAYS

The NAIC Model Bulletin on AI Use in Insurance marks an important step toward addressing AI bias in the industry, though challenges persist. Conclusions from the panel discussion include:

- Lack of a universal definition of bias: The absence of a clear definition creates ambiguity for insurers in aligning AI models with regulatory expectations.
- Complexities in identifying protected classes: Insurers face challenges in accessing demographic data and using inference-based methods for bias testing.

- Regulatory uncertainty: Companies may take a cautious approach to bias testing due to concerns about increased scrutiny and operational changes.
- Need for improved frameworks: The panel recommended clearer regulatory guidelines, stronger AI transparency requirements, and structured case studies to assist insurers in navigating compliance expectations.

The panel concluded that ,without these enhancements, insurers would continue to operate in a climate of regulatory uncertainty, where bias testing remains reactive rather than proactive. Moving forward, insurers, regulators, and industry stakeholders must collaborate to create robust and actionable AI fairness frameworks that prioritize both consumer protection and business efficiency.

Section 4 Academia's Role in AI Governance

4.1 ACADEMIA'S INFLUENCE ON AI ETHICS

The panel opened the discussion by acknowledging the pivotal role that academic institutions have played in shaping AI ethics and governance. Academia has long been at the forefront of research on ethical AI, bias mitigation, and safety frameworks, with initiatives such as the Partnership on AI[12], which brought together institutions like MIT, Stanford, and Oxford to collaborate with industry leaders on defining responsible AI practices. Academic research has provided the foundation for many of the ethical guidelines adopted by regulatory bodies and corporations alike. However, the panel recognized that academia's position as a neutral and honest broker in AI safety is not without challenges, particularly as funding structures, industry partnerships, and political influences shape the academic landscape.

The panel was invited to explore both the strengths and vulnerabilities of academia's contributions to AI safety standards. The conversation centered on academia's ability to conduct independent, unbiased research, the challenges posed by corporate funding and the loss of talent to the private sector, and the need for stronger connections among academic research, regulatory frameworks, and industry adoption.

4.2 STRENGTHS AND CHALLENGES

The panel emphasized that universities bring unparalleled expertise, research capabilities, and intellectual diversity to AI ethics discussions. Academic institutions have long served as hubs for open inquiry, fostering a culture where researchers can explore AI risks, fairness, and safety without the immediate pressures of profit-driven motives. This intellectual environment has allowed universities to develop robust ethical frameworks that inform industry best practices and regulatory policies.

However, the panel acknowledged that academia is not entirely independent of external influences. With growing industry partnerships and corporate funding playing a larger role in university research initiatives, there is increasing concern that some institutions may temper their critiques of AI risks to align with corporate interests. The panel noted that while universities often maintain academic freedom, the financial backing of major technology firms might subtly shape the direction of AI safety research, leading to conflicts of interest.

Another critical challenge facing academia is the brain drain of AI researchers into private industry. Many of the most talented AI ethics scholars are being recruited by leading technology firms, where they receive significantly higher salaries and access to cutting-edge AI systems. While this transition benefits corporate AI governance, it diminishes the number of truly independent academic voices advocating for AI safety. The panel expressed concern that this shift could weaken academia's ability to function as a long-term steward of AI ethics.

4.3 ETHICAL DIVERSITY IN ACADEMIA

The panel emphasized that academia is not a monolithic entity and that universities take different approaches to AI ethics research. Some institutions focus on theoretical research, formulating abstract ethical principles and policy recommendations that shape AI governance discussions at a high level. These universities produce academic papers, white papers, and conference presentations that influence long-term AI regulatory planning.

Other institutions take a more applied approach, working directly with industry and regulatory bodies to implement AI safety frameworks. These universities engage in projects that involve bias testing, risk assessment, and AI model governance, ensuring that ethical considerations are translated into practical tools and methodologies.

However, the panel also pointed out that some academic institutions prioritize corporate partnerships, leading to greater collaboration with industry but potentially compromising their neutrality. These institutions work closely

with tech companies to develop AI systems and ethical guidelines, but this relationship can lead to concerns about whether research is being influenced by corporate interests rather than being fully independent.

4.4 MOVING BEYOND THEORY

One of the recurring themes in the discussion was the gap between academic research and real-world AI governance implementation. While academia has made significant contributions to AI ethics, there remains a disconnect between theoretical frameworks and actionable policies that companies and regulators can readily adopt. The panel observed that many academic discussions remain at an abstract level, focusing on philosophical debates about AI fairness rather than providing clear, practical solutions for industries like insurance, healthcare, and finance.

To bridge this gap, the panel argued that academia must take a more active role in shaping regulatory and industry practices. Universities should work closely with regulatory agencies, industry leaders, and professional organizations to ensure that AI safety research is translated into policies that can be implemented at scale. This requires a stronger emphasis on interdisciplinary collaboration, where academic researchers engage with policymakers, business leaders, and technical experts to create comprehensive AI governance frameworks.

The panel also emphasized that academia should take a leading role in designing AI education programs that train the next generation of AI practitioners. Universities have a responsibility to embed AI ethics and governance principles into their curriculum, ensuring that students understand algorithmic bias, fairness, transparency, and accountability. By integrating these concepts into AI degree programs, actuarial science courses, and business schools, academia can help cultivate a new generation of AI professionals who prioritize ethical considerations in their work.

4.5 ACADEMIA IN AI REGULATION

Another major topic of discussion was academia's role in AI standardization efforts. The panel noted that academic institutions play a key role in shaping AI governance frameworks, influencing both national and international regulatory bodies. Organizations such as the National Institute of Standards and Technology (NIST), the International Organization for Standardization (ISO), and the European Union's AI regulatory bodies often draw heavily from academic research when drafting AI policies.

However, the panel expressed concern that academic research does not always translate into regulatory action, especially when governments and industry stakeholders have differing priorities. While academic institutions produce extensive research on AI fairness, risk mitigation, and transparency, these findings do not always lead to concrete policy changes. There is a growing need for academia to engage more directly with policymakers, advocating for AI regulations that prioritize consumer protection, fairness, and safety.

The panel also discussed how academia could contribute to the development of AI audit tools that help organizations evaluate bias and discrimination in AI systems. Rather than focusing solely on high-level discussions of fairness, universities can create technical frameworks that allow businesses and regulators to systematically assess AI risks. By providing data-driven methodologies for evaluating AI performance, academia can help establish a more standardized approach to AI governance.

4.6 MAINTAINING NEUTRALITY

The discussion concluded with a reflection on the importance of preserving academia's neutrality in AI governance. The panel agreed that while universities continue to be valuable contributors to AI safety research, they must actively safeguard their independence from corporate and political influences.

Academic institutions must ensure that research is conducted transparently, with public access to findings and methodologies. This openness will help prevent the perception that AI safety research is being influenced by hidden agendas. The panel also suggested that stronger funding mechanisms for independent AI research—such as government grants, non-profit sponsorships, and academic endowments—could help reduce reliance on corporate funding.

Additionally, the panel emphasized the need for multi-stakeholder collaboration, where universities work alongside governments, regulatory bodies, and civil society organizations to create AI governance frameworks that prioritize public interest over corporate profit. This collaborative approach will help academia maintain its credibility as a neutral entity while ensuring that AI safety standards are developed with broad societal input.

4.7 KEY TAKEAWAYS

- Academia's contributions: Major contributions to AI safety through theoretical research and applied methodologies, as well as the development of ethical frameworks that inform industry and regulatory policies.
- Challenges: Corporate influence on research direction due to funding and partnerships, the loss of talent to the private sector, and the gap between academic research and real-world policy implementation.
- Recommendations: Strengthening transparency, interdisciplinary collaboration, and regulatory engagement, while safeguarding academia's neutrality through diverse funding mechanisms and multi-stakeholder collaboration.

To maintain its role as a neutral and honest broker, academia must strengthen its commitment to transparency, interdisciplinary collaboration, and regulatory engagement. By doing so, universities can continue to provide the critical research, ethical frameworks, and technical solutions needed to ensure that AI systems are safe, fair, and accountable to society as a whole.

Section 5 Reflections on Implementing AI Standards

5.1 COMPLEXITY OF AI STANDARDS

The panel opened the discussion by reflecting on the valuable insights and lessons gained from implementing AI governance and technology policies within their organizations. While standards such as ISO 42001 and other AI risk management frameworks have provided structured guidance for businesses integrating AI, the panel acknowledged that applying these standards in practice can present unique challenges. They discussed the importance of enhancing clarity when integrating multiple overlapping standards and highlighted opportunities to improve risk management processes to create a more cohesive governance structure.

The panelists, drawing from their respective experiences, identified areas where standard implementation could be more streamlined and pointed out ways to reduce redundancy across various governance frameworks. They shared practical ideas on what could be done differently and proposed forward-looking strategies for enhancing AI governance moving forward.

5.2 LINKING STANDARDS EFFECTIVELY

One of the key topics discussed by the panel was the opportunity to enhance guidance on combining AI governance standards. While ISO 42001 provides a valuable risk management framework that organizations can use to assess and mitigate AI-related risks, there is potential for clearer guidance on integrating it with other AI-related standards. The panel noted that organizations sometimes find it challenging to align the different components of AI governance, especially when multiple regulatory and industry bodies release similar but separate guidelines.

Reflecting on past efforts, the panel suggested that the AI standards community could build on existing frameworks by offering more practical guidance on integrating AI governance with current governance practices. Many of the risk management principles used in AI governance overlap with those found in cybersecurity, privacy, and financial risk management, creating an opportunity to streamline efforts and reduce redundancy.

5.3 HOLISTIC RISK MANAGEMENT

The panelists reflected on how adopting a more structured and proactive approach to AI risk management from the start can enhance companies' success in their AI adoption journey. They highlighted the value of involving not only technical teams, but also compliance officers, legal teams, and business executives to ensure that AI deployment aligns with both ethical considerations and regulatory expectations.

A key takeaway was that organizations can benefit from embedding AI risk management into their overall corporate governance structure rather than treating it as a stand-alone function. By evaluating AI risk frameworks alongside other enterprise risk management (ERM) processes, companies can go beyond mere compliance to build greater trust in their AI systems.

5.4 KEY TAKEAWAYS

The panel’s discussion was constructive and forward-looking, focusing on how to enhance AI standards based on real-world implementation experience. Rather than focusing on challenges, the panel emphasized practical improvements and proactive strategies.

- Building on Experience: The panel highlighted the opportunity to evolve existing AI standards, like ISO 42001, by incorporating lessons learned from practical implementation.
- Streamlining Integration: Rather than viewing overlapping standards as problematic, the panel discussed ways to harmonize frameworks to make AI governance more efficient.

- Collaborative Risk Management: AI governance should involve technical, compliance, legal, and leadership teams to ensure a holistic and responsible approach.
- Adapting to Change: The panel encouraged a flexible mindset to regulatory evolution, integrating AI risk management early to stay agile.
- Fostering Collaboration: The panel advocated for stronger partnerships among industry, regulators, and academia to develop AI standards that are both practical and adaptive.

Section 6 Conclusions

6.1 PROACTIVE AI GOVERNANCE

AI governance should not be a reactionary process. Organizations must take preemptive steps to mitigate risks before they arise, embedding AI governance into corporate structures from the outset. By integrating proactive risk management into broader enterprise risk management frameworks, companies can better navigate the evolving regulatory landscape and maintain transparency and accountability.

6.2 DATA GOVERNANCE AND BIAS

Effective AI governance requires robust data management practices to minimize bias and ensure fairness. The panel emphasized that while identifying and addressing bias in AI models remains challenging, it is essential to prioritize data quality and transparency throughout the AI lifecycle. Building comprehensive data governance policies will help mitigate unintended biases and enhance the reliability of AI-driven decision-making.

6.3 ALIGNING AI GOVERNANCE WITH RISK MANAGEMENT PRINCIPLES

Aligning AI governance with established actuarial standards fosters a consistent and transparent approach to managing AI risks. Actuarial principles of risk classification, data quality, and model governance provide a solid foundation for AI oversight. By integrating these practices, organizations can enhance their ability to manage AI risks effectively.

6.4 BALANCING AI WITH HUMAN OVERSIGHT

While AI systems offer significant potential, human judgment remains critical in evaluating model outputs and identifying potential risks. Incorporating human oversight alongside automated processes helps ensure that AI-driven decisions are interpretable, fair, and aligned with professional standards. This balance is particularly vital in the insurance and financial sectors, where accountability is essential.

6.5 CROSS-DISCIPLINARY COLLABORATION

Collaboration among academia, industry, and regulators is essential for developing practical and ethical AI governance frameworks. By fostering open dialogue and sharing best practices, stakeholders can create standards that are both rigorous and adaptable to technological advancements. The panel encouraged continued partnership across disciplines to strengthen AI governance and enhance public trust.

6.6 LOOKING FORWARD

As the session wrapped up, the panelists reflected on the depth of discussion, the challenges ahead, and the importance of continued engagement in AI governance. AI is transforming industries at an unprecedented pace, and the regulatory, ethical, and technical frameworks governing AI must keep up. Companies, regulators, and professionals must remain vigilant and adaptable, ensuring that AI systems are transparent, fair, and accountable.

The panelists agreed that AI governance is not a static process—it is an ongoing commitment that requires continuous learning, regulatory engagement, and collaboration across disciplines. Future discussions must focus on refining risk management methodologies, improving bias detection techniques, and ensuring that AI-driven decision-making remains aligned with ethical and legal standards.

The panel expressed optimism that through collaborative efforts among academia, industry, and regulators, AI governance frameworks will continue to improve, enabling the development of safe, fair, and responsible AI

systems. The conversation ended on a note of encouragement, with a shared commitment to pushing AI governance forward and ensuring that AI technologies serve the broader public good, while upholding the highest standards of accountability and transparency.

Section 7 Acknowledgments

The researchers' deepest gratitude goes to those without whose efforts this project could not have come to fruition: the Expert Panel participants and others for their diligent work overseeing questionnaire development, analyzing, and discussing respondent answers, and reviewing and editing this report for accuracy and relevance.

Dorothy Andrews, PhD, ASA, MAAA, CSPA – Senior Behavioral Data Scientist & Actuary NAIC; expert in AI governance and bias mitigation in insurance.

Dave Ingram, CERA, FRM, PRM, FSA, MAAA – Risk Management Actuary, Society of Actuaries Board Member, advocate for AI integration in actuarial science.

Shane Leib, FSA, MAAA – Director of Actuarial Research, Moody's Analytics; Assistant Teaching Professor, University of Notre Dame.

Elaine Newton, PhD – Principal for AI Standards & Tech Policy, Amazon Web Services (AWS); former NIST and Oracle, specialist in cybersecurity, privacy, and AI/ML standards and compliance.

Ronald Poon-Affat (Moderator), FSA, FIA, MAAA, CFA, HIBA– Reinsurance specialist, Independent Board Member.

At the Society of Actuaries Research Institute:

Korrel Crawford, Senior Research Administrator

R. Dale Hall, FSA, MAAA, CFA, Managing Director of Research

Appendix A: Expert Panel Discussion Questions

The Artificial Intelligence Risk Management Framework Generative AI Profile outlines principles to promote the trustworthy and responsible use of generative AI. What does the panel consider to be the main tenets of this framework, and why are they essential for ensuring effective AI risk management?

Can you describe the standards development landscape at a high level and explain how NIST publications relate to international standards? Additionally, what motivates major tech companies to engage in the development of international standards?

The NAIC has prepared a Model Bulletin on AI use in insurance—now accepted by 19 states—which emphasizes mitigating bias and discrimination. What do you see as the primary obstacles preventing insurers from effectively conducting bias testing in AI systems, and what are the sources of resistance? Additionally, how do current testing methods fall short, and what key elements do you believe are missing from most insurance AI frameworks?

Academia has made significant contributions to AI ethics through initiatives like the Partnership on AI—where institutions such as MIT, Stanford, and Oxford collaborated with industry leaders to define ethical guidelines and best practices. How does the panel view academia's ongoing role as a neutral and honest broker in advancing AI safety standards?

Our panel includes professionals with experience in implementing AI standards. Looking back to when your organization began its journey in developing AI standards and technology policies, is there anything you would approach differently with the benefit of hindsight? Are there specific challenges or missed opportunities that stand out, and what key lessons have shaped how your organization approaches AI governance and innovation today?

References

[1] https://www.soa.org/
[2] https://www.nist.gov
[3] https://www.nist.gov/itl/ai-risk-management-framework
[4] https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
[5] https://www.naic.org
[6] https://www.iso.org
[7] https://artificialintelligenceact.eu/
[8] https://www.actuarialstandardsboard.org/standards-of-practice/
[9] https://www.iec.ch/homepage
[10] https://content.naic.org/sites/default/files/inline-files/2023-12-4%20Model%20Bulletin_Adopted_0.pdf
[11] https://www.rand.org/health-care/tools-methods/bisg.html
[12] https://partnershiponai.org/partners/

About The Society of Actuaries Research Institute

Serving as the research arm of the Society of Actuaries (SOA), the SOA Research Institute provides objective, data-driven research bringing together tried and true practices and future-focused approaches to address societal challenges and your business needs. The Institute provides trusted knowledge, extensive experience and new technologies to help effectively identify, predict and manage risks.

Representing the thousands of actuaries who help conduct critical research, the SOA Research Institute provides clarity and solutions on risks and societal challenges. The Institute connects actuaries, academics, employers, the insurance industry, regulators, research partners, foundations and research institutions, sponsors and non-governmental organizations, building an effective network which provides support, knowledge and expertise regarding the management of risk to benefit the industry and the public.

Managed by experienced actuaries and research experts from a broad range of industries, the SOA Research Institute creates, funds, develops and distributes research to elevate actuaries as leaders in measuring and managing risk. These efforts include studies, essay collections, webcasts, research papers, survey reports, and original research on topics impacting society.

Harnessing its peer-reviewed research, leading-edge technologies, new data tools and innovative practices, the Institute seeks to understand the underlying causes of risk and the possible outcomes. The Institute develops objective research spanning a variety of topics with its strategic research programs: aging and retirement; actuarial innovation and technology; mortality and longevity; diversity, equity and inclusion; health care cost trends; and catastrophe and climate risk. The Institute has a large volume of topical research available, including an expanding collection of international and market-specific research, experience studies, models and timely research.

Society of Actuaries Research Institute
8770 W Bryn Mawr Ave, Suite 1000
Chicago, IL 60631
www.SOA.org