EIOPA Consultation: Opinion on Artificial Intelligence Governance and Risk Management

Q1 - Do you have any comments on the context and objectives of the Opinion?

Key message:

The AAE supports EIOPA's principle-based and proportionate approach. We would welcome, however, greater clarity on the boundary between AI systems and traditional actuarial models. Moreover, a broader reflection on other areas of insurance impacted by AI would help ensure the Opinion's relevance across the sector.

Detailed response:

We welcome the pragmatic and proportionate approach taken in the Opinion, which aligns well with the framework set out in the AI Act and avoids introducing unnecessary additional requirements. This provides a strong basis for the insurance sector to adopt AI responsibly while preserving flexibility across a diverse European market.

We particularly value EIOPA's commitment to a principle-based framework. Such an approach allows for appropriate tailoring across various business models and product types in the European insurance landscape, and we believe this is key to maintaining both innovation and consumer protection.

It might be helpful for EIOPA to consider expanding the scope of examples provided to better reflect the breadth of AI use across the insurance value chain. For instance, there are areas beyond life and health underwriting where AI systems may pose significant risks, such as in automated claims handling or pricing in property and casualty insurance. These can also impact fairness and customer outcomes.

Some uncertainty remains in the market regarding the boundary between AI systems and traditional statistical models, such as Generalised Linear Models (GLMs). While we recognise that traditional GLMs are generally considered outside the AI Act's definition, we believe it could be helpful for EIOPA to confirm this distinction and clarify expectations where such models are embedded in more complex or partially automated systems. This would support consistent interpretation and proportional application of governance requirements.

Looking ahead, we suggest it might be helpful for EIOPA to consider how the Opinion could remain adaptable over time. As AI systems evolve rapidly, future developments may introduce new risk dimensions or shift the context of existing applications. Ensuring that governance frameworks are sufficiently forward-compatible would help maintain their relevance and effectiveness.

Q2 - Do you have any comments on the scope of the Opinion?

Key message:
The AAE supports the scope of the Opinion, focusing on non-high-risk AI systems, and recognises the importance of aligning with the AI Act. However, it would be helpful for EIOPA to provide greater clarity on borderline or potentially high-impact use cases that may merit enhanced scrutiny. In addition, a clearer articulation of how the Opinion applies across different parts of the insurance sector, along with observations on Generative AI and third-party models, could support consistent and practical application of risk management practices.

Detailed response:
We agree with the principle-based scope of the Opinion, which appropriately focuses on AI systems not classified as high-risk under the AI Act, and allows for flexibility to reflect the specific business models and products used across the European insurance sector.

The Opinion rightly reflects the AI Act's designation of life and health underwriting as high-risk. However, further clarity would be helpful for borderline insurance use cases that may not meet the formal high-risk definition but could still have material impacts on individuals. For instance, AI used in motor or home insurance might affect access to essential services, and systems involved in non-life claims settlement may have life-altering financial consequences. Guidance on whether such cases warrant enhanced governance, or how they relate to the overall risk-based approach, would support consistent interpretation across the sector.

We suggest EIOPA could provide more detail on how the Opinion applies across various insurance product lines, recognising that the risk profile of AI applications may differ substantially between, for example, personal lines and commercial or reinsurance business. This would help undertakings assess proportionality more accurately.

Additionally, we propose that EIOPA may wish to briefly reflect on emerging issues such as Generative AI (GenAI), especially in claims fraud detection where synthetic content generation presents new challenges. Similarly, acknowledging the differing governance needs of AI systems developed in-house versus those sourced externally — including open-source solutions — may assist undertakings in tailoring controls appropriately. These considerations could be explored further in future iterations or complementary workstreams.

To improve clarity and usability, we suggest EIOPA include several explicit examples of in-scope AI systems within the introduction or annex. This could include more complex applications such as forecasting, weather or catastrophe modelling, and advanced actuarial risk assessment, particularly where deep learning or neural networks are involved. Such examples would help ground the Opinion in real-world use cases and address concerns that the scope may appear too abstract.

Q3 - Do you have any comments on the risk-based approach and proportionality section? What other measures should be considered to ensure a risk-based approach and proportionality regarding the use of AI systems?

Key message:
The AAE supports EIOPA's emphasis on a risk-based and proportionate approach. However, the Opinion could provide greater clarity on whether proportionality should relate to the nature of the business or to the potential impact of the AI use case. We also suggest acknowledging how AI may amplify certain risks and recommending additional contextual factors—such as consumer vulnerability or reputational risk—that may influence the appropriate level of governance.

Detailed response:
We welcome the Opinion's emphasis on a risk-based and proportionate approach, which is consistent with both the AI Act and existing supervisory frameworks like Solvency II. A flexible, principle-based application of proportionality is key to ensuring the framework can adapt to diverse business models and operational contexts across the insurance sector.

However, the Opinion could provide clearer guidance on what the proportionality principle should be proportionate to. There appears to be an inconsistency between different regulatory perspectives: Solvency II focuses on the nature, scale, and complexity of operations, activities and the risks inherent in the business of (re)insurance undertakings, while other frameworks—including EIOPA's earlier work on digital ethics—suggest that proportionality should relate to the risk of harm to individuals. For example, a small insurer with limited operations might deploy an AI system in claims management that could materially affect policyholders' lives. Should proportionality in this case consider the insurer's size, or the impact of the AI application?

Some of our members also believe that factors such as consumer vulnerability and financial literacy should inform proportionality assessments. These factors are relevant under the Insurance Distribution Directive and Product Oversight and Governance (POG) requirements and could be integrated here to support consistency.

We also note that AI systems may amplify traditional risks, such as those related to fairness, explainability, and data quality. As such, proportionality should take into account not only the characteristics of the AI use case, but also the extent to which these risks may be intensified.

Additional suggestions to improve the clarity and application of this section include:

- Clarifying the meaning of proportionality in Paragraph 3.1 to distinguish whether it refers to business characteristics or AI-specific risks.
- Incorporating a broader view of risk that includes not only solvency but also profitability and reputational impact.
- Considering the evolving performance of AI systems and the need for forward-looking assessments, given the speed of technological change.

- Ensuring that the discussion of image processing in Paragraph 3.6 reflects the potential for significant impacts, such as in claims handling. The current wording could be misinterpreted as suggesting such use cases are inherently low-risk.
- Providing a clearer definition of terms such as "comprehensively explain", particularly where technical methods do exist to interpret outputs from image or text-based models.

We also support the idea that insurers should regularly review and improve their AI risk management practices to avoid sudden or unexplained shifts in outcomes, thereby promoting operational stability and consumer trust.

Q4 - Do you have any comments on the risk management system section? What other measures should be considered regarding the risk management system of AI systems?

Key message:
The AAE supports EIOPA's approach to integrating AI-related risks within the existing risk management system. However, it would be helpful to provide additional clarification on terminology, accountability, and proportionality in communication and training. We suggest placing greater emphasis on embedding AI within existing governance frameworks, rather than treating it as a separate topic, and aligning references with the AI Act where appropriate.

Detailed response:
Many of our members stressed that AI risks should not be treated in isolation but rather incorporated into the broader enterprise risk management framework. This reflects the reality that AI is a tool—albeit a powerful one—and not a standalone category of risk. A more integrated approach would also support proportionality in practice.

We support the integration of AI-related risks into insurers' overall risk management systems, consistent with the principle of proportionality and alignment with Solvency II. In some jurisdictions, such as Germany, similar expectations already exist under national supervisory guidance (c.f., BaFin paper of June 2021 titled "Big Data and Artificial Intelligence: Principles for the Use of Algorithms in Decision-Making-Processes").

In that context, the Opinion could be clearer in distinguishing between when new policies or frameworks are needed versus when existing ones can be extended. For example, Paragraph 3.9 mentions "a specific AI strategy" as an illustrative policy. It may be more helpful to refer to "an AI strategy or an integrated model governance policy" to reflect the option of embedding AI governance within existing structures.

We also believe that accountability should be clearly defined and tied to the roles articulated in the AI Act. Paragraph 3.10 refers to in-house and third-party models but does not reflect the accountability distinctions between AI system providers, deployers, and other operators. Further alignment with the terminology and concepts of the AI Act would improve clarity.

On the topic of organisational responsibilities, we propose that undertakings should establish clear accountability at senior management level. One option might be to designate a Chief AI Officer (CAIO), potentially combining this role with another executive position provided the necessary expertise is present. Actuaries may be well-suited to support or contribute to this function given their skills in data analysis and risk governance.

Training and communication are also key elements. We suggest that requirements to communicate AI governance policies to "all staff" (as stated in Paragraph 3.10) may be disproportionate. Instead, policies should be communicated to all relevant staff who are reasonably expected to engage with or be affected by the AI systems.

Relatedly, clarification is needed around who should receive training, and when. For example, staff implications will differ depending on whether an AI system is integrated into an existing pricing model, or whether a generative AI tool is rolled out company-wide.

We also propose expanding the list of risk areas in Paragraph 3.7. In addition to technical and financial risks, insurers should consider areas such as accountability, effective communication, and internal collaboration, all of which are essential to embedding sound AI governance.

To support risk management practice, clearer definitions of “accuracy” and “robustness” would be welcome. Moreover, the scope of material risks should extend beyond solvency to include profitability, reputational impacts, and effects on core processes such as underwriting and claims.

Finally, we agree that risk management practices should be dynamic. Paragraph 3.9 could highlight the need to review AI approaches when systems already in use begin to exhibit unforeseen behaviour, as may happen with automated learning models. This would encourage ongoing vigilance and responsiveness.

Q5 - Do you have any comments on the fairness and ethics section? What other measures should be considered to ensure a fair and ethical use of AI systems?

Key message:
The AAE supports the inclusion of fairness and ethics within the Opinion. We suggest refining the expectations around bias in data and outputs to reflect practical limitations and aligning definitions of fairness with existing regulatory and actuarial frameworks. Additionally, assigning clear organisational accountability for ethics-related risks and expanding guidance on hidden biases, redress mechanisms, and model monitoring would enhance the Opinion's practical utility.

Detailed response:
We welcome the strong focus on fairness and ethics in the Opinion. These principles are fundamental to public trust and align closely with actuarial values. We particularly endorse the need for explicit assignment of ethics risk at a senior level within the organisation. Without such clear accountability, there is a real danger that ethical issues could be overlooked despite Board-level awareness.

With regard to training data, Section 3.13 currently states that it should be "free of bias". While this is an important goal, we believe it would be more practical to require insurers to demonstrate that they have taken reasonable steps to identify and mitigate bias, including through model correction techniques. This reflects the reality that some biases may be unknown or unavoidable, particularly in third-party AI systems or those trained on large-scale internet datasets.

Hidden or indirect biases—such as those embedded in assumptions or proxies—also warrant further attention. We suggest that the Opinion explicitly call for the use of both pre- and post-modelling tests to assess fairness in outcomes, especially where AI decisions affect pricing, underwriting or claims.

The lack of access to protected characteristics (e.g. race or gender) in datasets can make bias detection more difficult. In such cases, guidance on proxy testing or indirect fairness assessments would be welcome.

On the definition of fairness itself, we note that multiple and sometimes conflicting definitions exist (e.g. equal opportunity vs. demographic parity). We recommend that EIOPA clarify which interpretation it expects undertakings to apply, to support regulatory consistency and comparability across firms.

Section 3.12 refers to a "consumer-centric" approach. Some of our members noted that this may be difficult to uphold in practice when firms rely on externally-developed AI tools, such as general-purpose chatbots. In such cases, we suggest that customer-centricity be interpreted as a responsibility to assess and manage the risks of using such tools rather than guaranteeing full control.

Additionally, it would be helpful if the Opinion referenced customer vulnerability and financial literacy explicitly, as these are relevant factors in understanding fairness and ethical risks in customer interactions.

We also encourage elaboration of Paragraph 3.15 on redress. While we support the principle of giving customers the right to information and contestation, the mechanism for redress may vary across the value chain and by the nature of the system. We suggest distinguishing between redress for customers and internal feedback

mechanisms to support model retraining and correction. These serve different purposes but both are critical to trustworthy AI.

A further enhancement would be to include a paragraph addressing model and system governance, including model drift, lifecycle testing, and performance monitoring. These are often underdeveloped in current insurance practices. Incorporating guidance on fairness metrics (e.g. F1-score, recall, precision), explainability, and robustness under a governance and testing umbrella would promote more reliable outcomes.

Lastly, it may be helpful to clarify or reconsider terms such as "meaningfully explainable", which may be open to interpretation. Definitions grounded in technical standards would support more consistent application.

Q6 - Do you have any comments on the data governance section? What other measures should be considered to ensure adequate data governance of AI systems?

Key message:
The AAE supports EIOPA's emphasis on high-quality data as foundational to the trustworthy use of AI systems. However, we suggest clarifying expectations around bias, accuracy, and completeness, particularly in cases involving third-party data. Greater consistency in terminology and scope across the Opinion, and consideration of the broader data lifecycle, would strengthen the practicality and coherence of this section.

Detailed response:
We welcome the focus on ensuring that data used to train and test AI systems is accurate, complete, appropriate, and relevant. However, we suggest clarifying that data should be "reasonably" complete and accurate, as full completeness may not be practical in many insurance applications, especially when unstructured or third-party data is involved.

The treatment of bias across this section would benefit from further clarification. Paragraph 3.18 refers to the assessment and correction of bias, which seems to contrast with previous suggestions that data should be entirely free of bias (e.g., Paragraph 3.13). In practice, bias is often unavoidable, especially in historical or third-party datasets. A more consistent and pragmatic expectation—focusing on reasonable mitigation efforts rather than elimination—would reflect operational realities and support feasible implementation.

Terminological consistency would also enhance the clarity of the Opinion. Paragraph 3.16 uses terms that differ from Paragraph 3.13 and from Article 82 of Solvency II. We suggest that EIOPA align the terminology used throughout the Opinion, or provide clarification on any intended distinctions. This would help users interpret requirements consistently.

In addition, we propose that EIOPA consider the principles outlined in the European Standard of Actuarial Practice 1 (ESAP1), particularly those relating to data quality and reliance on third parties (Sections 3.5 and 3.3). These principles are well-established across the actuarial profession and may offer useful reference points for insurers when assessing the adequacy and governance of data.

ESAP1: https://actuary.eu/wp-content/uploads/2019/10/ESAP1-Revised-11-10-2019-FINAL.pdf

We further recommend strengthening expectations regarding third-party data sources. Insurers may not have full transparency or control over how such data is collected or processed. Paragraph 3.20 could be expanded to address the need for consistent disclosure obligations and due diligence when engaging with external providers.

Lastly, we suggest broadening the scope of the data quality section to reflect the full data lifecycle. Areas such as data collection (e.g. via LLMs or web scraping), feature engineering, and data enrichment can significantly influence model outcomes and should be considered as part of effective AI governance. Including these dimensions would enhance the technical depth and practical applicability of this section.

Q7 - Do you have any comments on the documentation and record keeping section? What other measures should be considered to ensure adequate documentation and record keeping of AI systems?

Key message:
The AAE agrees with the importance of robust documentation and record-keeping for AI systems. However, we suggest EIOPA clarify expectations around reproducibility, the level of documentation required, and how evolving AI systems should be treated. A proportionate approach, particularly for long-term record tracking, and greater attention to post-deployment monitoring, would help ensure practicality and environmental sustainability.

Detailed response:
We support EIOPA's emphasis on appropriate documentation and record-keeping to support accountability and transparency. These are critical for effective governance and should be aligned with existing risk management practices.

A key point requiring clarification is the expectation of reproducibility and traceability. In practice, many AI systems—especially those involving automated machine learning or stochastic elements—may not allow for strict reproducibility. Some models also operate as black boxes. We encourage EIOPA to clarify whether reproducibility is an "obligation of results" (i.e., a hard requirement), or an "obligation of means" (i.e., reasonable effort to achieve it). A balance should be struck between ideal expectations and technical realities.

Also, we suggest that the level of technical detail required in documentation should be proportionate to the context and consistent with expectations for non-AI models already subject to governance under existing risk management systems. Overly detailed technical documentation may not always be practical or necessary.

The Opinion could also provide more structured guidance on post-deployment monitoring. While Annex I contains useful examples, the link between these and the monitoring obligations in the main text could be strengthened. Moreover, aspects such as model limitations, embedded assumptions, and interpretability merit more explicit treatment, ideally as a distinct documentation topic.

Regarding long-term data record tracking, we propose a proportionate approach based on the function and sensitivity of the AI system. Continuous storage of large datasets can carry a high environmental cost, and requirements should be calibrated accordingly.

We also suggest including guidance on the geographic location of data storage, given its relevance for data protection, operational resilience, and cybersecurity. Awareness of data sovereignty and regulatory obligations tied to storage location is increasingly important for insurers operating cross-border.

From a usability perspective, we note that some elements in Annex I Section 1—such as "severity" and "likelihood" headings—may prompt undertakings to quantify risk factors that are better addressed qualitatively. Removing these labels or clarifying their purpose could prevent misinterpretation.

Finally, we recommend removing references to high-risk AI use cases in the documentation examples, as these fall outside the scope of the Opinion. Including such examples may cause confusion for undertakings trying to understand the applicable expectations.

Q8 - Do you have any comments on the transparency and explainability section? What other measures should be considered to ensure adequate transparency and explainability of AI systems?

Key message:
The AAE agrees that explainability is a key principle for the trustworthy use of AI. We support a proportionate and context-sensitive approach to explainability, and recommend clarifying whether this is an obligation of means or results. The Opinion should remain tool-agnostic and avoid singling out specific methods. Broader definitions and expectations for both technical and consumer-facing explainability would support consistency and practicality.

Detailed response:
We agree with EIOPA that explainability is fundamental to maintaining trust in AI systems and ensuring appropriate governance. However, further clarification is needed on whether explainability is expected as an obligation of results—meaning that AI systems must always be explainable—or an obligation of means, where insurers are expected to make reasonable efforts to achieve it. In practice, many models (especially complex or black-box systems) cannot be fully explained. Our members would favour EIOPA taking a clearer stance on this issue, potentially restricting the use of unexplainable models in high-impact contexts.

Equally important is the clarity of communication to customers. Paragraph 3.27 rightly emphasises the need for explanations in simple, accessible language. We support the suggestion that AI systems should not be deployed unless they meet this threshold, particularly for customer-facing applications. This principle reinforces consumer protection and aligns with expectations under the Insurance Distribution Directive.

The Opinion currently references specific technical tools—LIME and SHAP—as examples of explainability techniques. While illustrative, naming these methods may inadvertently give the impression that EIOPA is endorsing particular tools. We recommend revising the language to remain tool-agnostic, instead referring to "local and global model-agnostic methods," with examples given only in a footnote. This approach avoids prescribing specific methodologies and allows flexibility for firms to choose appropriate tools based on their use case.

In addition, we propose expanding the language in Paragraph 3.25 to recognise the limitations of explainability tools. These limitations should be clearly documented, including the rationale for method selection and whether any comparative assessments have been conducted. This would enhance transparency and encourage thoughtful application of multiple approaches when feasible.

A broader and harmonised definition of explainability would also be helpful. As with fairness, various definitions exist and may lead to different interpretations across jurisdictions. Regulatory convergence on what explainability means—and how it is assessed—would promote consistency across the market.

We also suggest recognising that AI explainability differs from traditional IT systems. For traditional systems, a one-to-one mapping from input to output is typical. In contrast, AI systems—especially probabilistic or learning-based models—may require statistical or stochastic explanations rather than deterministic ones. These forms of explanation should be considered appropriate, provided they support meaningful understanding and traceability of outcomes.

Lastly, we propose simplifying the language used. For instance, the phrase “meaningfully explainable” could be streamlined, as its interpretation may vary. Clarity and alignment with existing technical and regulatory frameworks would aid implementation.

Q9 - Do you have any comments on the human oversight section? What other measures should be considered to ensure adequate human oversight of AI systems?

Key message:
The AAE supports the inclusion of human oversight as a fundamental principle of AI governance. We suggest clarifying expectations on staff training, the roles of different functions in oversight, and how the "human-in-the-loop" principle should be applied in practice. Greater emphasis on Board responsibility and more precise language around function-specific roles would strengthen this section.

Detailed response:
We welcome EIOPA's emphasis on human oversight in the governance of AI systems, particularly those with customer impact. We strongly support the inclusion of a "human-in-the-loop" principle, which is essential to preserving individuals' rights to challenge or understand automated decisions.

Effective oversight depends on ensuring that Boards and senior management have the necessary skills to understand and challenge AI deployments. This may require targeted training, recruitment of personnel with relevant experience, or use of external support. We suggest that EIOPA encourage Boards to demonstrate how they are equipping themselves to meet this oversight responsibility.

Similarly, Section 3.30 could be enhanced by stating that staff training and competency requirements should be tailored based on the individual's role and the potential impact of the AI system. A one-size-fits-all approach may not be sufficient, especially where oversight responsibilities vary significantly.

We agree with the need for sufficient staff training before deployment of AI systems. This should be a precondition for use, particularly in high-impact contexts. Elaborating this point more explicitly would reinforce its importance.

Some of our members have questioned whether the establishment of a distinct "AI function" is necessary. We suggest clarifying that existing key functions—such as the risk management, actuarial, or compliance functions—may be equipped to address AI-related issues where appropriate. AI should be seen as a methodological enhancement, not necessarily a structural shift in governance. Nonetheless, where AI has cross-cutting impact, coordination across functions remains essential.

In that context, it might be helpful to acknowledge that actuaries may be well-suited to support or contribute to AI-related oversight roles, given their skills in data analysis, risk management and regulatory compliance. This could include supporting the responsibilities of a designated AI Officer or contributing to cross-functional AI committees, where established. Their involvement may help ensure that actuarial and technical considerations are appropriately reflected in the governance of AI systems.

We also suggest clarifying the potential overlap between roles. For example, can the responsibilities of an AI function be held by the same individual as the Data Protection Officer? Where this is permitted, expectations around role separation or safeguards should be stated. More generally, we propose further emphasis on the accountability of all business functions and the Board for assessing the impact of AI within their respective

domains. This holistic perspective reinforces the idea that AI governance is not confined to technical teams but is a firm-wide responsibility.

To support role clarity, we propose expanding the list of responsibilities noted under the actuarial function (in Paragraph 3.29) to include the function's involvement in model governance policies. This would better reflect the broader oversight role actuaries play beyond pricing and reserving.

Finally, in Paragraph 3.31, the wording around bias mitigation could be improved. We recommend simplifying "contribute to the removal of possible biases in line with the policy of the undertaking" to: "contribute to the removal of biases, consistent with the undertaking's governance framework." This makes the guidance clearer and more action-oriented.

Q10 - Do you have any comments on the accuracy, robustness and cybersecurity section? What other measures should be considered to ensure adequate accuracy, robustness and cybersecurity of AI systems?

We support the inclusion of accuracy, robustness, and cybersecurity as key principles for the governance of AI systems. However, the use of fairness metrics should be approached with caution. These tools may not always reflect real-world fairness and, if applied uncritically, can lead to misleading or counterproductive outcomes. We suggest that the Opinion encourage the use of fairness metrics only where appropriate, and as part of a broader, context-sensitive assessment supported by professional judgement. Further emphasis on robustness over time, including monitoring for model drift and data degradation, and clearer expectations around cybersecurity practices, would also strengthen this section.

Q11 - Do you have any comments on the possible risks identified for customers and undertakings?

We strongly support the application of the fairness principle in managing AI risks. Unfair treatment can arise in areas beyond high-risk use cases, such as in automated claims handling and non-life pricing, including natural catastrophe cover. These applications may carry significant consumer and reputational risk. The AAE has previously highlighted the importance of fairness in its 2024 discussion paper on Social Sustainability in Insurance. In addition to consumer harm, unfair AI outcomes can pose financial risks to insurers via litigation and reputational damage, reinforcing the prudential relevance of this principle.

Q12 - Do you have any comments on the analysis of costs and benefits?

The AAE has no specific comments on the expected costs and benefits at this time. However, we support the principle of proportionality and encourage EIOPA to monitor the practical impact of implementation, particularly for smaller undertakings and in relation to resource-intensive requirements such as documentation and data governance.

Q13 - Do you have any comments on the policy option chosen?

The AAE does not have additional comments on the proposed policy option. We support EIOPA's principle-based and proportionate approach and agree that providing supervisory convergence through non-binding guidance is an appropriate next step.

Q14 - Do you have any comments on the proposed approach?

We welcome the annexes as a useful practical supplement to the Opinion. In particular, we support the inclusion of examples to guide proportional implementation. To further enhance their usefulness, we suggest incorporating references to established model validation techniques, including quantitative testing methods, robust documentation practices, and indicators of system performance, including, where appropriate, commonly used indicators of model performance (e.g. Precision, Recall, or F1-score), recognising that these should be selected based on the use case and validation objectives.

We note that some national approaches, such as the principles outlined by BaFin in Germany, offer similar guidance in relation to data governance and algorithmic decision-making.