

February 17, 2025

Mr. Shigeru Ariizumi
Chair
Mr. Jonathan Dixon
Secretary General
International Association of Insurance Supervisors
Basel, Switzerland

By email

Annex also submitted via survey.iais.org/739258



Dear Messrs. Shigeru and Dixon,

Public consultation on draft Application Paper on the supervision of artificial intelligence

The Institute of International Finance (IIF) welcomes the opportunity to publicly respond to the International Association of Insurance Supervisors' (IAIS') public [consultation](#) on this important topic. The answers to the public consultation questions are set out in the **Annex** of this letter and have been submitted via the survey tool.

As you know, the [IIF](#) represents approximately 400 globally active financial institutions (FIs) from over 60 geographies, drawn from the banking, insurance, securities, asset management, payments and other sectors, including central and development banks. Many of our members are particularly interested in the cross-border impacts of emerging technologies, including artificial intelligence (AI) and machine learning (ML), and related regulatory frameworks.

While we appreciate the IAIS's thoughtful consideration of AI supervision, we believe many aspects of the Application Paper (**the paper**) are of significant potential concern and warrant a fundamental reconsideration.

We would recall that the paper states correctly that "Application Papers **do not include new requirements**, but provide further advice, illustrations, recommendations or examples of good practice to supervisors on how supervisory material may be implemented. **The proportionality principle applies** also to the content of Application Papers" (emphasis added).

In our view, the Application Paper goes beyond this description of purpose and explicitly adds new requirements and quasi-requirements directly on insurers, in many instances using prescriptive language such as "insurers should" (17 instances), "supervisors should ensure that insurers" (7 instances), "is/are essential" when applied to insurers (3 instances), or "should be" (23 instances, not all of which apply to insurers). This prescriptive language is, in our submission, not appropriate for an Application Paper that does not include new requirements.

Even where the language of such recommendations is more permissive, e.g. "insurers should consider" (3 instances), in many cases supervisors may insist on imposing these as requirements, or insurers will feel obliged to comply with, or at least generate documentation to explain non-compliance with the quasi-requirement, to be presented to supervisors in case of request.

In our view, the Application Paper risks setting markers that could lead to de facto regulation in the form of guidance, at a time when the technology is still at an early stage and where some jurisdictions, including the U.S. and the European Union, are reconsidering the scope and breadth of regulation in a number of areas, including emerging technologies.

The number of obligations and quasi-obligations added in some areas is very high, and (in our submission), is disproportionate to the real level of risk these systems pose to regulatory objectives.

The ultimate result of adopting the paper in its current form will likely be to retard the application of advanced AI in the insurance industry and consequently to reduce potential cost savings, at a time when policyholders are already unable to afford cover in some cases due to unavoidable risk-based premium increases, shifting more risk to governments and government-backed pooled insurance arrangements.

This compliance-driven approach also risks ossifying insurance industry practices and effectively sidelining the industry from deploying the latest AI technology at scale and in a timely way, while limiting the potential benefits of this technology to consumers.

On a more mundane level, the risk is of distracting senior management and compliance teams with “box-ticking” exercises, rather than empowering them to think flexibly and adaptively about how to manage their own AI systems, within the context of their business model and business lines, in a way that is compliant with local law and supervisory expectations.

Background

As we have done in recent submissions to other authorities¹ on this topic, we would stress that:

- **Supervisors should better recognize the benefits of AI and not unduly stifle innovation:** While the paper recognizes the “commercial benefits” of AI, AI systems can meaningfully enhance consumer protection and market stability through improved fraud detection, enhanced pricing, customer segmentation and risk modeling capabilities, and the potential for automated claims management. These advancements can create greater market efficiencies that benefit consumers through improved product accessibility and more competitive pricing, while strengthening insurers’ operational resilience. While supervisors must of course remain vigilant about potential novel risks, their approach should be calibrated to preserve these significant public benefits. An overly restrictive framework would harm consumer and supervisory interests by slowing the adoption of technologies that enhance market stability and consumer protection.
- **AI in insurance is not entirely new:** While generative AI (**GenAI**) at scale is a relatively emerging technology, our industry supervisors should avoid adding to or exacerbating hyperbole around AI. AI and machine learning (**ML**) systems in the insurance sector have been deployed in production for some time, and are already subject to existing regulatory and risk management frameworks (see below). Sophisticated analytical systems and ML engines are either in operation or under development widely, including in claims fraud detection and claims processing, underwriting risk scoring, streamlining customer servicing, and in consumer channels such as through more targeted or personalized marketing. Supervisory attention should focus on newer applications of AI (such as Gen AI) in insurance practices and where the use of AI materially impacts consumers or insurer solvency.
- **AI in insurance is already regulated:** AI applications in insurance are already subject to regulation through sectoral and cross-sectoral regulations. Examples include policyholder confidentiality rules, consumer data privacy regulations and insurance consumer protection rules, intellectual property laws, and prudential requirements

¹ See IIF (August 2024), [IIF Response to U.S. Treasury RFI on AI in Financial Services](#); IIF (September 2024), [IIF Response to European Commission Consultation on AI in Financial Services](#); IIF (December 2024), [IIF Response to European Commission Consultation under the AI Act](#).

concerning data governance, cyber risk, third-party risk, and operational risk generally. These regulations apply to the use of any general-purpose technology, like AI, in insurance services, ensuring that material risks are well managed. The insurance industry is already comprehensively regulated.

- **Existing risk frameworks provide a strong foundation:** Insurers operate under national/regional prudential requirements, consistent with the Insurance Core Principles (ICPs) and other international standards, that require insurers to develop their own sophisticated risk management governance frameworks, systems, and controls. These arrangements have largely served the industry well. Insurers keep such arrangements under constant review and proactively make adjustments as needed. The industry therefore starts from a very strong risk management foundation – one that is commensurate with and contributes to the high trust placed in insurance undertakings by their policyholders.
- **The IAIS should consider the potential dangers of prescriptive requirements.** Prescriptive requirements could create significant barriers to entry for new market participants, reduce investment and/or delay adoption of AI technologies, thereby delaying the benefits of AI to society. There are also potential dangers of an overbroad definition of AI systems that are in scope of the paper.

Further engagement

We encourage the IAIS to discuss and consult further with industry experts on this topic before finalizing any Application Paper. We believe that the current schedule for publication of a final Application Paper in Q3 2025 does not provide sufficient time for proper consideration of these complex issues.

The IIF and its members stand ready to engage in additional discussions and consultations on these topics, or to clarify any aspect of our submission. For example, if desired we could convene a meeting of interested IIF members with IAIS officials to discuss this submission. We would also be happy to provide a briefing to IAIS officials on the findings of the 2024 IIF–EY survey on financial institutions’ use of AI and ML.²

We thank you again for the opportunity to contribute to this important consultation.

Yours sincerely,



Jessica Renier
Managing Director, Digital Finance



Andres Portilla
Managing Director, Regulatory Affairs

² See the 2024 [Public Survey Report](#) as the latest in this series. Copies of the full report have been shared with you and with Petra Hielkema, as Chair, FinTech Forum, IAIS, Joe Perry, Senior Policy Advisor, IAIS and Julian Arevalo, IAIS FinTech Forum AI/ML Co-Lead.

Annex

Answers to consultation questions

[See separate document]

Annex

IIF Responses to Consultation Questions

IAIS draft Application Paper (Application Paper) on the supervision of artificial intelligence

Q. #	Question text	Submission text
1	General comments draft Application Paper on the supervision of artificial intelligence	<p>While we appreciate the IAIS’s thoughtful consideration of AI supervision, we believe many aspects of the Application Paper (the paper) are of significant concern and warrant fundamental reconsideration.</p> <p>We would recall that the paper states that “Application Papers do not include new requirements, but provide further advice, illustrations, recommendations or examples of good practice to supervisors on how supervisory material may be implemented. The proportionality principle applies also to the content of Application Papers” (emphasis added).</p> <p>In our view, the Application Paper goes beyond this description of purpose and explicitly adds new requirements and quasi-requirements directly on insurers, in many instances using prescriptive language such as “insurers should” (17 instances), “supervisors should ensure that insurers” (7 instances), “is/are essential” when applied to insurers (3 instances), or “should be” (23 instances, not all of which apply to insurers). This prescriptive language is, in our submission, not appropriate for an Application Paper that does not include new requirements.</p> <p>Even where the language of such recommendations is more permissive, e.g. “insurers should consider” (3 instances), in many cases supervisors will insist on imposing these as requirements, or insurers will feel obliged to comply with, or at least generate documentation to explain non-compliance with the quasi-requirement, to be presented to for supervisors in case of request.</p> <p>In our view, the Application Paper risks setting markers that could lead to de facto regulation in the form of guidance, at a time when the technology is still at an early stage and where some jurisdictions, including the U.S. and the European Union, are likely to enter a deregulatory phase.</p> <p>The number of obligations and quasi-obligations added in some areas is very high, and (in our submission), is disproportionate to the real level of risk these systems pose to regulatory objectives.</p> <p>The ultimate result of adopting the paper in its current form will likely be to retard the application of advanced AI in the insurance industry and consequently to reduce potential cost savings, at a time when policyholders are already unable to afford cover in some cases due to unavoidable risk-based premium increases, throwing more risk onto governments and government-backed pooled insurance arrangements.</p> <p>This compliance-driven approach also risks ossifying insurance industry practices and effectively sidelining the industry from reaping the potentially very large rewards from deploying the latest AI technology at scale and in a timely way, while limiting the potential benefits of this technology to consumers.</p> <p>On a more mundane level, there are significant risks of distracting senior management and compliance teams with “box-ticking” exercises, rather than empowering them to think about how to manage their own AI systems, within the context of their business model and business lines, in a way which is compliant with local law and supervisory expectations.</p>
2	General comments on Executive Summary	No specific comments
3	General comments on Section 2	<p>Our comments emphasize that the regulatory framework should acknowledge broader systemic benefits of AI in insurance beyond just direct benefits to insurers and policyholders, including enhanced risk modeling, improved claims processing, better fraud detection, and increased market efficiency. We argue for a more focused definition of AI systems, recommending concentration on generative AI and deep neural networks while excluding traditional statistical methods. We warn that an overly broad definition based on OECD standards could become unmanageable and harm insurance sector competitiveness. We caution against developing guidance for AI risks in insured businesses as premature, note that AI-driven market volatility is difficult to distinguish from other sources, and seek clarification that the IAIS is not intending to conduct AI-specific work on operational risks like cloud computing or environmental issues.</p> <p>We criticize the proposed framework for lacking clear standards and potentially creating unnecessary compliance burdens. We recommend making the framework more flexible and illustrative rather than prescriptive, while better acknowledging existing risk management systems. Our comments on Table 3 (Risk criteria) suggest several refinements to the proposed risk matrix, including replacing vague terms, focusing on significant impacts, and removing unclear criteria. We acknowledge the potential of AI in</p>

Q. #	Question text	Submission text
		supervisory technology while emphasizing the need for proportional oversight. Lastly, we caution against premature guidance in areas like AI-related insurance risks and investment risks.
4	Comments on Section 2.1 Context and objective	<p>Benefits recognition:</p> <ul style="list-style-type: none"> • The paper should recognize systemic benefits from better risk management enabled by AI, not just benefits “for both insurers and policyholders.” • Some benefits not specifically mentioned in this section include: <ul style="list-style-type: none"> ○ More precise risk analysis and modeling, leading to enhanced pricing and product segmentation ○ Enhanced claims management and processing efficiency, reducing costs associated with delay ○ Improved fraud and anomaly detection, reducing premiums overall and reducing potential money laundering risks ○ Greater market efficiencies benefiting both insurers and consumers <p>Risks:</p> <ul style="list-style-type: none"> • As mentioned in our answer to Q. 33 below, risks relating to AI not mentioned include: <ul style="list-style-type: none"> ○ The IAIS should consider the potential dangers of prescriptive requirements. This could create significant entry barriers for new players, reduce investment, and/or delay adoption of AI technologies, thereby delaying benefits to society. ○ Another risk is that data localization and other measures that reduce the free flow of data with trust across jurisdictions inhibit the development of AI systems, thereby reducing the effectiveness of risk management and pricing functions, among others. <p>The context and objectives section would also benefit from a clearer recognition that many AI applications in insurance represent an evolution of existing practices rather than revolutionary change, and that many if not all the risks are already addressed by extensive regulatory, supervisory and risk management frameworks (as set out in more detail in our cover letter).</p>
5	Comments on Section 2.2 AI system definition	<p>While we understand the rationale for adopting the OECD definition of “AI system”, given its breadth and vagueness, the supervisory framework should recognize that different AI applications pose varying levels of risk and warrant different levels of oversight.</p> <p>We recommend focusing supervisory attention on newer applications (such as Gen AI) and where AI can materially and negatively impact consumers or solvency, rather than addressing all AI applications equally.</p> <p>Indeed, a literal-minded application of the OECD definition by the IAIS, and consequently by firms and supervisors, may significantly impair the competitiveness of insurance vis-à-vis other sectors. Some members who have attempted to apply an inclusive, literal definition on the precautionary principle in the context of the EU AI Act implementation, for example, have reported the sheer number of software systems potentially within scope quickly becomes unmanageable.¹</p> <p>As such we would recommend narrowing the AI system definition to focus on generative AI and deep neural networks, characterized by:</p> <ul style="list-style-type: none"> • Greater capability for autonomous learning and adaptation • Complex inferential capabilities beyond traditional statistical methods <p>This would exclude well-established methods such as:</p> <ul style="list-style-type: none"> • Linear/logistic regression • Decision trees and random forests • Rule-based expert systems • Traditional statistical and predictive modeling

¹ See IIF (December 2024),

Q. #	Question text	Submission text
		<p>“Fundamental rights”</p> <p>On a different topic, the paper says it “also sets out some legal, constitutional and human rights (also referred to as fundamental rights) considerations that, although outside typical insurance supervisory mandates, are likely to be relevant for whether insurers are lawfully using AI systems.” We would see directly addressing these topics by insurance supervisors as going beyond their mandate and therefore as best left outside the scope of the paper altogether. Fundamental rights is a concept that is largely jurisdiction-specific in its application.</p>
6	Comments on Section 2.3 Scope and structure	<p>We have several comments on the possible items of the IAIS work program that are mentioned in this section.</p> <p>On Insurance-related AI risks, we would caution that the risks from AI use in insured businesses are still too nascent and uncertain to develop meaningful guidance. Any guidance would be premature and could stifle innovation in both the AI and insurance sectors. These risks are better addressed through insurers’ existing underwriting and risk management frameworks that reflect their specific product offerings.</p> <p>On Investment-related AI risks, market volatility from AI is hard to distinguish from other sources of market volatility, including from high-frequency trading that has been in place for decades. At this stage, the causal relationship between AI and market volatility is speculative at best; we would also caution against creating an artificial distinction between AI- and other technology-driven market risks. Such work could also duplicate existing efforts by securities regulators and other bodies.</p> <p>As for operational risks arising from technologies such as cloud computing, and environmental issues, we seek clarification that the IAIS is not intending to conduct AI-specific work on these topics, as seems to be suggested by paragraph 20. As is rightly observed, developments such as cloud computing are an enabler; they are not unique to the implementation of AI. We would also argue that environmental aspects of AI are best addressed by environment regulators.</p>
7	Comments on Section 2.4 Proportionality and risk-based supervision	<p>This section proposes several new requirements, as follows:</p> <ul style="list-style-type: none"> • supervisors should allocate “more supervisory activities and resources” to higher-risk AI applications (paragraph 23); • insurers’ governance and risk management measures should be “influenced by the type of AI use case and the context in which it is used” (paragraph 24); and • supervisors and insurers should assess and assign “a level of risk to an AI system” based on provided criteria (paragraph 26). <p>The section attempts to create a risk assessment framework for AI systems without:</p> <ul style="list-style-type: none"> • Providing clear standards for what constitutes “higher risk” versus “lower risk” AI applications • Acknowledging that risk classification may vary significantly across jurisdictions and business models • Addressing how this framework interacts with existing risk assessment methodologies <p>While the intention to provide guidance on proportionality is welcome, the framework presented could inadvertently create a de facto compliance obligation for insurers to:</p> <ul style="list-style-type: none"> • implement specific risk classification systems; • maintain documentation justifying their risk classifications; • adjust their governance frameworks based on these classifications. <p>We recommend that this section be revised to:</p> <ul style="list-style-type: none"> • more clearly position the framework as illustrative rather than prescriptive; • acknowledge that existing risk assessment and management requirements and frameworks exist and already cover AI systems; • provide greater flexibility in how proportionality is implemented; and • focus on principles rather than specific criteria. <p><i>Comments on Table 3</i></p> <p>We would caution against the IAIS finalizing guidance suggesting that firms themselves should be applying this risk matrix in their own risk assessments. In our view, the risk matrix, if it could be useful as a tool for supervisors, will need considerable tailoring</p>

Q. #	Question text	Submission text
		<p>according to the specifics of individual jurisdictions’ legal and regulatory frameworks and existing risk assessment and management guidance.</p> <p>In general, while the Table seems quite comprehensive, it also includes many terms that involve delicate (and quite subjective) value judgements. We would recommend reducing the number of criteria and replacing some of them with clearer metrics.</p> <p>As a general observation, we believe that the outcomes-related, insurer-focused criteria are those which would most tend to relate to possible systemic risk, and systems that score highly on these factors could be prioritized by supervisors for attention.</p> <p><i>Comments on specific risk criteria</i></p> <ul style="list-style-type: none"> • Row 2: For example, we suggest that references to “fundamental rights” in Table 3 are too general and potentially too wide-ranging (given that this term covers all legal, constitutional and human rights, according to the discussion in section 2.2). We suggest this be replaced by a reference to a more specific legal concept such as “unlawful discrimination.” • Row 2: The reference to “discrimination” is also inapposite given that insurers tend (subject to law) to discriminate across classes of insureds in order to price and segment correctly, and should be replaced by “unlawful discrimination.” In most jurisdictions this is defined in terms of protected characteristics or membership of a protected class or similar, but there is no need for a harmonized definition. • Row 2: We further consider that the threshold should be “significant adverse impact” rather than any “adverse impact”. • Row 3: we do not believe that “The extent to which the AI system is engaged in responsible stewardship in pursuit of beneficial outcomes for consumers” is a clear criterion, nor is it clear why this should be a factor in a risk assessment. This criterion should be removed.
8	Comments on Section 2.5 The role of supervisors and supervisory tools	We welcome recognition of the possible use of AI in SupTech applications, noting the BIS Innovation Hub’s projects on SupTech applications including Project Spectrum (using GenAI to enhance inflation nowcasting). In our view, such applications will be useful to help supervisors develop a better and more nuanced (and possibly less risk-averse) understanding of AI, including GenAI, particular the potential benefits.
9	General comments on Section 3	<p>As we have said in answer to Q1, in our view, the Application Paper goes beyond this true purpose of Application Papers, and explicitly adds new requirements and quasi-requirements directly on insurers. This should be avoided.</p> <p>The results of the 2024 IIF-EY Annual Survey Report on AI/ML Use in Financial Services indicate that the majority of respondents govern their AI applications through existing frameworks, have developed, or are in the process of developing new frameworks that complement existing ones.²</p> <p>This range of approaches, in our view, is legitimate; as such, there should be no “one size fits all;” any further guidance needs to be facilitative only and indicative of possible practices rather than prescriptive. We elaborate on this further below.</p>
10	Comments on Section 3.1 Introduction	<p>We would comment that paragraph 34 – “This section covers the additional areas within these ICPs that, due to the inherent characteristics of AI systems, require specific attention” – contains an implicit assumption that new, more granular guidance on these topics is, in fact, needed.</p> <p>In fact, the ICPs that are cited in paragraph 33 (ICP 8, 16, 7 and 5) are all high-level principles which continue to apply and are relevant to supervisors supervising insurers that deploy AI systems.</p> <p>There is little evidence in the draft Application Paper of an attempt to parse the detailed guidance that already exists under each of those core principles and conduct a gap analysis. Without that gap analysis, there can be no solid case that further guidance is required.</p> <p>We would further caution that any further insurance sector-specific guidance given needs to be flexible, principles-based and technology-agnostic, such that it will not rapidly go out of date.</p>
11	Comments on Section 3.2 Risk management systems	Paragraph 36 asserts that “a clear articulation and common understanding across [an insurer’s] control functions (including risk management, compliance and internal audit) of what constitutes AI-related risk and the development of risk assessment criteria are important.”

² There is more granular data on this topic in the longer (non-public) version of the report that has been shared with the IAIS, at p. 18.

Q. #	Question text	Submission text
		<p>We believe this assertion is unduly prescriptive. While firms should of course be encouraged to keep their risk management frameworks under review in light of changing technologies, arguably, developing specific AI-focused risk management policies and materials may unduly segment and/or ossify the risk categories and lenses under which all operations – including AI systems – must be kept under constant review.</p>
12	Comments on Section 3.3 Corporate culture	<p>We feel that the suggested elements in paragraph 39 go beyond ICP 7.1 and ICP 19.2 in a manner that is unduly prescriptive. Paragraph 39 seems to be a gloss on ICP 19.2, which is much broader than AI systems.</p> <p>We would also question why, if ICP 19.2 does not merit this level of granular guidance in general, such granular guidance is merited for AI systems.</p> <p>For the reason also that the section is duplicative of section 6 and potentially confusing, it should be removed.</p>
13	Comments on Section 3.4 Human oversight and allocation of management responsibilities	<p>We welcome the statement in paragraph 40 that “the development, implementation and oversight of AI systems throughout their entire life cycle should not alter supervisory expectations. For example, Boards should continue to ensure that insurers have a well-defined and documented governance structure that provides effective separation between oversight functions and management responsibilities.”</p> <p>In general, the proposed board responsibilities are too granular and may exceed the capabilities of most board members, who are not trained technologists. For example, the board should not have to provide objective and robust scrutiny of the deployment of AI systems. Rather, this is the responsibility of management with appropriate training.</p> <p>For these reasons, and bearing in mind our response to Q1, we are not supportive of granular additional guidance of the kind set out in paragraphs 41, 42 and 43.</p> <p>The word “should” appears 16 times in these paragraphs, along with a panoply of other normative concepts such as “essential” and “to ensure.”</p> <p>We believe this level and sheer quantity of additional granular expectations for boards and senior management is likely to sap firms’ own flexibility or adaptability around risk management and oversight, and also likely to divert significant compliance resources for limited payoff.</p> <p>The expectations could be rephrased as possibilities or potentially useful tools – “could” instead of “should” – or simply jettisoned altogether.</p> <p>We also have an issue with the Board responsibilities in section 3.4.1 not being consistent with the need for the Board to attend to those topics of the highest importance. It assumes that AI oversight will always need Board-level attention. AI may become commoditized along with its risk management. The IAIS should not be “setting in concrete” expectations that should be adjusted dynamically.</p> <p>On the last bullet of paragraph 41, in particular, referring to system redundancy and kill switches, may be impracticable, inefficient and unnecessary, depending on the AI system’s characteristics and importance to the insurer’s business. Particular risk mitigants or particular technological fixes like kill switches are system dependent and they shouldn’t be hardwired into regulatory expectations.</p>
14	Comments on Section 3.5 Use of third-party AI systems and data	<p>Section 3.5 seeks to impose substantive requirements on insurers:</p> <ul style="list-style-type: none"> • Paragraph 46 calls for supervisors to “ensure that insurers obtain adequate information and reassurances from third-party service providers about the characteristics, capabilities, appropriate fitness for purpose and limitations of AI systems they outsource where they are critical services.” • Paragraph 47 calls for the “same level of oversight” where an insurer uses third parties or outsourcing and the providers use AI systems. • Paragraph 48 calls for supervisors to “ensure that insurers implement mechanisms that can track data sources used in training AI systems and the processes involved in content generation.” <p>We would repeat the observations in Q1 that the imposition of substantive requirements goes beyond the intended purpose of Application Papers.</p> <p>To the extent these requirements merely duplicate requirements elsewhere in the IAIS Core Principles, those principles should be referenced. To the extent they do not, the gap analysis that would demonstrate a need for further guidance has not been explained.</p>

Q. #	Question text	Submission text
		<p>The IAIS proposal for “same level of oversight” for third-party AI systems appears unrealistic given market realities. The NAIC Model Bulletin takes a more pragmatic approach, focusing on due diligence and contractual rights rather than full oversight..</p> <p>We would also observe that the level of input/centrality of a particular third-party AI system will differ strongly according to whether the system is a fundamental or incidental part of a business and whether the model is trained on insurer or provider data.</p> <p>The third-party oversight requirements should be risk-sensitive, and flexible enough to accommodate different types of AI implementations.</p>
15	Comments on Section 3.6 Traceability and record keeping	<p>As with other sections, this section proposes substantive new requirements on insurers, including:</p> <ul style="list-style-type: none"> • to implement tracking mechanisms for AI system data sources and processes (paragraph 48); • to make available detailed documentation about AI systems to supervisors and auditors (paragraph 48); • to maintain repositories of all deployed models for high-impact AI applications (paragraph 49). <p>This section suffers from several difficulties:</p> <ul style="list-style-type: none"> • it creates new documentation obligations without clear standards for what constitutes adequate compliance; • the term “high-impact AI applications” is undefined, creating regulatory uncertainty about when the model repository requirement applies; in any case, in many jurisdictions high-risk or high-impact AI system is not a separate category; • it fails to recognize the practical challenges and costs associated with maintaining comprehensive documentation of complex AI systems, particularly given the broad definition adopted by the IAIS; and • the requirement to track “processes involved in content generation” is impracticably broad and could encompass a vast range of technical details. <p>The proper place for record-keeping requirements is in existing regulatory frameworks for model documentation and governance. There is no demonstrated need for additional, AI-specific documentation requirements that may duplicate or conflict with existing obligations.</p> <p>The section should be fundamentally reconsidered to focus on principles rather than prescriptive requirements and to acknowledge practical implementation challenges.</p>
16	General comments on Section 4	<p>As we have said in answer to Q1 and elsewhere, in our view, the paper in our view goes beyond this true purpose of Application Papers, and explicitly adds new requirements and quasi-requirements directly on insurers. This should be avoided.</p> <p>There appears to be significant duplication between the new, more granular requirements in this section and existing ICPs including 8.1 (systems for risk management and internal controls), 8.4 (risk management function), 8.8 (outsourcing), enterprise risk management framework for solvency purposes (ICP 16), cybersecurity (e.g. 19.12.5) and business continuity (e.g. 18.1.5–6 and 18.7.6). The paper does not engage in the kind of gap analysis that would be needed to show that new, AI-specific guidance is considered necessary.</p>
17	Comments on Section 4.1 Introduction	<p>This section argues from the assertions that “unlike traditional models, AI systems can tackle complex tasks with intricate patterns and highly complex non-linear relationships” and “they can continuously update their understanding and predictions with new data and can adapt to changing circumstances,” to the conclusion that “these differences highlight the need for additional safeguards around model validation (particularly where a model adapts over time) and the underlying data storage and use.”</p> <p>The paper lacks any explicit analysis of the existing safeguards and whether they are general enough and robust enough to address these new developments. We believe this is a significant gap in the paper and one that should be addressed before any new prescriptive guidance is finalised.</p> <p>In particular, other than brief references to applicable ICPs, the paper appears to ignore the existing detailed enterprise risk management framework for solvency purposes (ICP 16), which is specifically required to address operational risk and which is sufficiently comprehensive to address emerging technologies.</p>
18	Comments on Section 4.2 AI system robustness	<p>The proposed robustness requirements are again too prescriptive. The performance of AI systems can, in many cases, be tested within existing model testing and validation frameworks. To the extent these frameworks may need to be adapted for special characteristics of advanced AI systems (such as GenAI) or riskier applications (such as fully automated claims processing or dynamic pricing models), firms should be free to adapt them as they consider appropriate, with supervisory guidance as needed.</p> <p>Again, we would note that the word “should” appears 9 times in paragraphs 53 and 54. We consider this a disproportionate increment in the level and granularity of expectations around this one technology category, such that other key issues or bigger risks may get overshadowed for some firms.</p>

Q. #	Question text	Submission text
19	Comments on Section 4.3 AI system safety and security	<p>The language of paragraphs 53 and 54 should be adjusted to make clear that the guidance is facilitative only and not prescriptive.</p> <p>This section attempts to impose a number of new requirements or strong recommendations on insurers, including:</p> <ul style="list-style-type: none"> • implementing advanced security measures against potential threats, in particular against cyber attacks, potentially involving developing regular adversarial testing and continuous monitoring for anomalies to identify potential threats like data poisoning and model inversion attacks • setting up automated alerts to detect significant deviations in AI behavior • regular updates of security tools for AI systems, alongside continuous staff training on new risks • putting in place effective backup and recovery solutions to ensure business continuity for insurers, especially where AI systems provide critical functions • segmentation and compartmentalization within the AI system and its purpose-built models as an additional control measure. <p>Many of these controls are comprehended within existing guidance in the ICPs concerning risk management and internal controls (ICP 8) and enterprise risk management for solvency purposes (ICP 16), as well as cybersecurity (e.g. 19.12.5) and business continuity (e.g. 18.1.5–6 and 18.7.6), and in our view it is duplicative and potentially counter-productive to address the same topic – in slightly different language – in the context of AI.</p> <p>While many or most of these controls may be put in place by large, sophisticated insurers for higher-AI systems, not all will be appropriate for all systems and insurers.</p> <p>As for automated alerts, this in particular is far too prescriptive for this reason. If any such requirement is to be maintained, IAIS should clarify what is intended by “automated alerts”.</p>
20	General comments on Section 5	<p>We are concerned about the numerous new (quasi-)requirements regarding AI transparency and explainability in insurance sought to be imposed in the paper. We critique the proposals as overly burdensome, impractical, and often too vaguely worded, while arguing that some requirements (like mandatory human assistance) represent policy decisions that should be left to national legislators. Our response also highlights that mandated disclosures about AI systems need to respect trade secrets and confidentiality obligations, and suggests that many of these requirements go beyond what’s appropriate for an Application Paper.</p>
21	Comments on Section 5.1 Introduction	<p>We note paragraph 64 states that section 5 “provides guidance on key considerations about how [ICPs 19 and 8] should be applied in the context of a transparent and explainable AI system that follows a proportional risk-based approach.” We assume the phrase starting “that follows” is intended to qualify “guidance” rather than “AI system”, and seek that clarification.</p>
22	Comments on Section 5.2 Explaining AI system outcomes	<p>Section 5.2 proposes a number of significant new obligations or quasi-obligations, including:</p> <ul style="list-style-type: none"> • supervisors should ensure that insurers are able to meaningfully explain the outcomes of AI systems that they use (paragraph 65); • the transparency and explainability of claims decisions and claims dispute resolution influenced by AI systems are especially important to ICP 19.10 (paragraph 65); • insurers deploying “start of the art” tools for explainability need to consider and document “relevant limitations” (paragraph 66); • for highly complex AI systems where achieving an “otherwise desirable” level of explainability may not be possible, insurers “should consider” complementary governance measures such as the use of guardrails or human oversight (paragraph 68); • insurers should ensure that AI systems only operate under the conditions for which they were designed and only when “sufficient levels of confidence” have been reached (paragraph 69); and • AI systems systems should identify cases in which they were not designed or approved to operate, or cases for which their answers are not reliable (paragraph 69). <p>In our view, as with other sections of this paper, the sheer number of these new obligations or quasi-obligations is very high. Also some of them are expressed in very vague terms, e.g. “relevant limitations”, “sufficient levels of confidence”, “meaningfully explain”.</p>

Q. #	Question text	Submission text
		<p>Paragraph 66 lists tools and strategies that insurers “Can adopt to ensure their AI systems are explainable”, including restricting deployment of AI systems to those that are simple and explainable, restricting the use of complex AI systems to challenging and fine-tuning more traditional mathematic models, making deployment of AI tools conditional on deployment of “explainability tools” such as Shapley values or LIME. We do not believe it wise to name-check particular methodologies that may be supplanted in due course in a guidance document such as this. Moreover, these proposed restrictions on the use of AI go far beyond the purpose of an Application Paper and, more importantly, would have a significant negative impact on the ability of insurers to innovate at a time when there is a clear need for increased insurance coverage.</p> <p>Paragraph 67 says that “by providing clear explanations of the factors/variables that influence risk scores, insurers can justify premium calculations to customers and regulators.” This presupposes an obligation on insurers to justify premium calculations to consumers, an obligation which is not in place in many jurisdictions that rely on market forces to keep premium levels competitive. As a separate point, the language of the whole section should be adjusted to focus on transparency to regulators rather than policyholders.</p>
23	Comments on Section 5.3 Explanations adapted to the recipient stakeholders	<p>Section 5.3 would impose several new obligations on insurers:</p> <ul style="list-style-type: none"> • consumers should be made aware if they are interacting with an AI system and be allowed to obtain assistance from a human if needed (paragraph 70); <ul style="list-style-type: none"> ○ consumers would require plain, simple and easy-to-understand information not involving the use of excessive technical language. This information should be no less detailed than that provided for decisions not based on AI; ○ an example is potentially providing policyholders with a clear breakdown of the factors that have influenced their premium calculations, such as age, driving history and geographic location to support explainability (paragraph 70); • other stakeholders such as auditors or supervisors will require more comprehensive and technical information about the AI system to allow them to perform an adequate supervisory review process, potentially including information about how the data was collected, processes and post-processing methodologies, feature importance or the reasoning behind technical choices, including the governance and risk management measures put in place (paragraph 71); • insurers should ensure that this information is sufficient to provide internal and external audit functions with the information they need to make a proper assessment of the extent to which policies have been effectively followed (paragraph 71). <p>In our view, as with Section 5.2, this section introduces numerous new obligations that are both burdensome and, in some cases, impractical. We believe this goes beyond the intended scope of an Application Paper and suggest the section be fundamentally reconsidered.</p> <p>The suggested requirement to make consumers aware if they are interacting with an AI system and to be allowed to obtain assistance from a human if needed represents a fundamental policy choice that should be left to national legislatures. It forecloses the possibility of digital-only or low-cost, AI-driven insurers that choose not to have a “human in the loop” at the point of onboarding.</p> <p>The requirement to provide explanations “no less detailed” than non-AI decisions may be impracticable in some cases, as AI systems may work in fundamentally different ways than traditional decision-making processes.</p> <p>The example of providing premium calculation breakdowns again presupposes obligations that may not exist in all jurisdictions and could conflict with existing regulatory frameworks.</p> <p>The requirements about information to be provided to auditors and supervisors presupposes a single scope or standard of audit or supervisory review that is not realistic. Such requirements should be set out in audit standards or ad-hoc supervisory requests for information and not in an Application Paper.</p> <p>As with Section 5.2, many of these requirements are expressed in vague terms (e.g., “plain, simple and easy-to-understand,” “comprehensive and technical information,” “proper assessment”) that could lead to divergent interpretations and compliance challenges.</p> <p>Lastly, the role of confidentiality and trade secrets may need to be acknowledged in any recast section 5.3. Insurers plainly cannot infringe trade secrecy obligations imposed by third parties or divulge their own trade secrets through mandated disclosures. Supervisors should not ask for trade secret information without due cause, and subject to strict secrecy obligations.</p>
24	General comments on Section 6	We recommend that these sections be substantially revised to:

Q. #	Question text	Submission text
		<ul style="list-style-type: none"> • focus on principles rather than prescriptive requirements; • acknowledge existing regulatory frameworks including existing IAIS guidance; • use clearer and less subjective terminology; • recognize technical limitations of AI systems; • balance fairness considerations with legitimate business needs; • avoid creating unrealistic expectations about AI explainability; and • focus on supervisory oversight rather than social policy.
25	Comments on Section 6.1 Introduction	<p>As to the reference to “fundamental rights” in paragraph 74, we refer to and repeat our comments in answer to Q 5 above.</p> <p>Paragraph 76 recalls the IAIS Draft Application Paper on fair treatment for diverse consumers (Diverse Consumers Draft). We would reiterate the key points made in our submission responding to the Diverse Consumers Draft, that there is a need to:</p> <ul style="list-style-type: none"> • differentiate between retail and commercial (re)insurance; • reflect local markets and context, legal regimes across jurisdictions; • recognize the fundamental concept of risk-based pricing; and • recognize existing rigorous consumer protection standards imposed on retail insurers and retail insurers’ strategic priority of expanding insurance products to serve a more diverse customer base. <p>We welcome the recognition in paragraph 75 of the distinction between improper discrimination (which is terms “unfair discrimination”) and lawful risk differentiation and risk-based pricing.</p> <p>However, paragraph 76 mentions that the Diverse Consumers Draft “recommends that the insurance industry take active steps to reduce unconscious biases, use of stereotypes and discrimination in their business processes and throughout their corporate culture.”</p> <p>As a core function of the insurance business is to “discriminate” among risks, and discrimination itself should not inherently be viewed negatively, we would propose changing this term in this context to “unlawful” discrimination.</p> <p>There is no need, in our view, to duplicate the handling of this issue across the two Application Papers. If the issue is included, the clear focus should be on unlawful discrimination based on the legal and regulatory framework in the jurisdiction.</p>
26	Comments on Section 6.2 Fairness by design	<p>This section proposes considerable new and detailed regulatory requirements on insurers, including:</p> <ul style="list-style-type: none"> • AI systems that are “harmful or abusive, treat consumers unfairly or do not respect fundamental rights ... should not be brought to the market” (paragraph 79); • insurers should “adopt a fairness-by-design approach that embeds fairness considerations within the AI governance and risk management systems” (paragraph 79); • “to this extent and as far as covered by their mandate,” supervisors should ensure insurers “establish and implement policies and processes on the fair treatment of customers, as an integral part of their business culture”, including through adherence to six specific points and 3 sub-points (paragraph 80). <p>In our view, we believe this type and level of prescription goes beyond the intended scope of an Application Paper and suggest the section be fundamentally reconsidered. The IAIS should provide more flexibility in governance requirements in order to reflect different insurance business models as well as differences in insurance regulatory and supervisory frameworks. Moreover, these considerations are not unique to AI and should be technology-agnostic. Accordingly, we do not advocate for specific IAIS guidance on this topic.</p> <p>With regard to the term “fundamental rights” in paragraph 79, see our response to Q. 5 above.</p> <p>Also, the language in paragraph 80 suggests a consciousness that these considerations go beyond the mandate of many, if not most, supervisors, again suggesting it is inappropriate for an Application Paper.</p> <p>To the extent that the material merely duplicates suggestions elsewhere in the Application Paper, as seems to be implied in paragraph 80, it should be removed from this section. To the extent it does not, the above comments apply.</p>

Q. #	Question text	Submission text
		<p>These proposals contain a number of unclear terms like “effective challenge,” “fairness-by-design,” “group think” and “adequate redress mechanism,” creating significant regulatory uncertainty. The requirement for governance, data management, transparency, redress mechanisms, etc with respect to models and outcomes appears to prescribe specific organizational structures and work programs that may not be appropriate for all insurers or model inventories.</p> <p>The section should be fundamentally reconsidered or removed.</p>
27	Comments on Section 6.3 Data management in the context of fairness	<p>This section proposes considerable new and detailed regulatory requirements on insurers, including:</p> <ul style="list-style-type: none"> • data sets must be “accurate, complete and representative of the customer segment being served” (paragraph 81); • data use must be “monitored to mitigate bias” (paragraph 81); • supervisors must “ensure that insurers have adequate data management processes throughout the AI system life cycle”, including through specific obligations on data collection, data preparation and post-processing (paragraph 81); • the insurer’s data management processes should “govern against using customer data in an unfair manner, such as when a consumer’s age or other personal characteristics are used for non-risk-based pricing practices aiming to exploit their willingness to pay or low propensity to shop around” (paragraph 82); and • insurers must also ensure customers have “right to access and, if needed, to correct data” (paragraph 82). <p>In our view, this section suffers from the same difficulties as other sections, seeking as it does to impose new requirements on insurers, which is not the objective of Application Papers.</p> <p>It also makes pronouncements on a series of delicate policy issues, including:</p> <ul style="list-style-type: none"> • the trade-off between data completeness and accuracy and efficiency of pricing: many advanced models can operate with high predictive accuracy notwithstanding missing or incorrect input data; the cost of cleaning historical data sets may not be justified; • monitoring to mitigate bias: legal obligations to avoid unlawful discrimination will of course motivate this in most cases; where there is no legal obligation to avoid “bias” (a much broader term) it is not clear why there should be an obligation to monitor for it; • adequate data management: this will already be covered under operational risk and data management frameworks; and • govern against using customer data in an unfair manner: this is a vague obligation and the example given also includes the vague term “personal characteristics”; differential pricing based on willingness to pay is not unlawful per se unless it relies on protected characteristics. <p>In short, the IAIS should not seek to progress a policy agenda that goes well beyond insurance supervisory topics and strays into areas that are best left to national/regional legislatures or privacy regulators.</p> <p>The section should be fundamentally reconsidered or removed.</p>
28	Comments on Section 6.4 Inferred causal relations in an AI system	<p>This section proposes considerable new and detailed regulatory requirements on insurers, including:</p> <ul style="list-style-type: none"> • insurers must “establish a process to regularly extract and document the implied AI system inferences... in a clear and transparent manner” (paragraph 84); • documentation should “enable effective challenge and discussion on whether the implied causal relationships are in line with expectations and the insurer’s strategic objectives” (paragraph 84); • such documentation “should support senior management and underwriters in assessing the extent to which decisions are risk-based and compliant with non-discrimination laws and ethical considerations;” (paragraph 84); • insurers should also have “policies and processes in place to ensure that customer data is not abused to circumvent prohibitions against discrimination” (paragraph 85); and • in this respect, insurers should “carefully consider the use of proxy variables, especially in pricing and underwriting practices” (paragraph 85).

Q. #	Question text	Submission text
		<p>In our view, this section suffers from the same difficulties as other sections, seeking as it does to impose substantive new requirements on insurers, which is not the objective of Application Papers, while also employing a number of vague and contentious concepts. The section also creates impractical documentation requirements that fail to recognize the complexity of modern AI systems.</p> <p>The obligation to “regularly extract and document” implied AI inferences assumes a level of transparency that may not be technically feasible with certain AI architectures.</p> <p>Paragraph 84 raises many definitional issues by using vague terms including “effective challenge and discussion”, or “in line with expectations”. We also do not understand clearly what is denoted by the compound expression “the extent to which predictions from an AI system infer causality based on identified correlations that reflect historic societal biases” found in paragraph 84. In particular, the term “historic societal biases” is subjective and goes beyond protection from illegal discrimination.</p> <p>For similar reasons, we do not believe “ethical considerations” should be referenced in paragraph 84. Given that views of what is ethical behavior for insurers in their interactions with policyholders might vary dramatically, the yardstick should be compliance with local laws and regulatory obligations.</p> <p>ICP 19.12.7 has been misquoted in paragraph 85. It refers to “unfair treatment” not “discrimination.”</p> <p>Furthermore, the guidance on proxy variables is overly broad and could persuade some supervisors to ban or effectively ban their use without sufficient justification.</p> <p>The section should be fundamentally reconsidered or removed</p>
29	Comments on Section 6.5 Monitoring the outcomes of AI systems	<p>Like earlier sections, this section proposes substantive regulatory requirements on insurers, including:</p> <ul style="list-style-type: none"> • “ensuring appropriate governance and risk management measures according to the AI use case, such as using more explainable AI systems and using fairness metrics to assess model outcomes in high-impact AI use cases” (paragraph 89); fairness metric examples are set out in the Annex; • providing for “documentation of the outputs of AI systems” and “results of any fairness testing ... on those outputs” (paragraph 89); and • being required “to keep an inventory of models with varied levels of information depending on the complexity of the AI system and its use case” (paragraph 89). <p>This section suffers from the same difficulties as other sections, seeking as it does to impose substantive new requirements on insurers, which is not the objective of Application Papers.</p> <p>We also do not believe the recommendation in Section 6.5 (Paragraph 89) to use fairness metrics (which are not, and should not be, defined) to assess model outcomes in high-impact AI use cases is appropriately balanced, particularly for those jurisdictions that may not use a formal risk classification of AI systems. (To be clear, we are not advocating that the IAIS propose any such risk classification itself, considering this is best dealt with at national/regional level).</p> <p>Further, the proposed requirement for insurers to maintain a model inventory could also represent a significant undertaking, depending on its scope; this is connected with the issue of the breadth of the definition of “AI System”, as to which, see our response to Q. 5 above.</p> <p>The section should be fundamentally reconsidered or removed</p>
30	Comments on Section 6.6 Adequate redress mechanisms for claims and complaints	<p>As with earlier sections, this section proposes substantive regulatory requirements on insurers, including:</p> <ul style="list-style-type: none"> • to have “effective, fair and transparent redress mechanisms, both for claims and complaints disputes” (paragraph 91); • for “high-risk” AI cases, insurers must “give meaningful explanations on determinative factors in claims or complaints resolution” (paragraph 91); and • insurers’ redress mechanisms must include the “ability for a consumer to update, supplement or correct information and data from sources that are used in the AI systems” (paragraph 92). <p>This section suffers from the same difficulties as other sections, seeking as it does to impose substantive and burdensome new requirements on insurers, which is not the objective of Application Papers.</p>

Q. #	Question text	Submission text
		<p>To the extent that paragraphs 91 and 92 duplicate existing guidance on dispute resolution mechanisms and redress mechanisms, it can be safely deleted. To the extent it does not duplicate existing guidance, there is no demonstrated need for additional, AI-specific guidance on this topic.</p> <p>The proper place for any new guidance on a person’s right to update, supplement or correct information is in privacy legislation and/or in dispute resolution principles. There is no need for additional, AI-specific guidance on this topic.</p> <p>The section should be fundamentally reconsidered or removed.</p>
31	Comments on Section 6.7 Societal impacts of granular risk pricing	<p>Unlike earlier sections, this section mainly contemplates new regulatory actions by supervisors, including:</p> <ul style="list-style-type: none"> • banning the use of certain risk factors for pricing purposes; and • banning differential pricing. <p>Both measures, in our view, go well beyond typical supervisory mandates and stray into social policy around insurance availability. While this may be within the purview of some IAIS members, in most jurisdictions these are political questions and best left to legislatures. To the extent they do fall within supervisory mandates, they do not relate to AI systems except tangentially and for that reason do not belong in an AI Application Paper.</p> <p>We also note the proposed substantive new requirements on supervisors, including to “consider how the development of AI systems may increase or reduce protection gaps”; and to “undertake regular market analysis and engagement with insurers to understand the risks and possible ways to find a balance between legitimate risk underwriting practices and enhancing financial inclusion” (paragraph 97). However, the requirement to “find a balance” between risk underwriting and financial inclusion creates an ill-defined obligation that could lead to regulatory overreach.</p> <p>The proper place for addressing protection gaps and financial inclusion is in primary legislation and/or specific regulatory frameworks designed for those purposes. There is no demonstrated need for additional, AI-specific guidance on these topics.</p> <p>The section should be removed.</p>
32	What further work could the IAIS undertake on artificial intelligence?	<p>We note the several issues left out of scope as discussed in section 2.3 of the paper. Please refer to our response to Q. 6 concerning those topics. We do not wish to table further issues for future work by IAIS relating to AI at this time.</p>
33	Are there risks not effectively captured by the IAIS’ work on artificial intelligence?	<p>The IAIS should consider the potential dangers of prescriptive or vague requirements. This could create significant entry barriers for new players, reduce investment, and/or delay adoption of AI technologies, thereby delaying benefits to society.</p> <p>Another risk is that data localization and other measures that reduce the free flow of data with trust across jurisdictions inhibit the development of AI systems, thereby reducing the effectiveness of risk management and pricing functions, among others.</p>