AI System Evaluation Tools – Summary of Feedback Received and Regulator Responses

The following is a summary of substantive concerns raised related to the AI System Evaluation Tools and responses to each. Where possible, we have attempted to link to the groups that raised these concerns but that list may not be complete given the volume of comments received. Technical edits to the document were also made and are reflected via tracked changes.

While the drafting group made every effort to address regulatory feedback, there may be areas that will continue to be evaluated through the course of the pilot process and continuing public engagement.

Coordination within NAIC (TX, ACLI, CAI, IRI)
- Summary: Parties expressed concerns that this work isn't being coordinated properly within the NAIC including with the Market Regulation and Consumer Affairs (D) Committee and the Financial Condition (E) Committee or that it belongs under those committees. Working Group leadership will work with NAIC staff to:
    - Provide e-mail updates to relevant NAIC groups as work advances.
    - Concurrent to the pilot stage, provide a more substantive update to D and E Committee groups.
    - Concurrent to the pilot stage, request in-depth review from actuarial task forces
    - Provide verbal updates, as possible, to D and E Committee groups.
    - Note: Both the Innovation, Cybersecurity, and Technology (H) Committee and the Big Data and Artificial Intelligence (H) Working Group have charges allowing the development of guidance while encouraging work proceed in a coordinated fashion.

Confidentiality (AITC, APCIA, CAI, IRI, NALC, RAA, Consumer Reps, AHIP)
- Summary: Parties expressed concerns about information being requested and confidentiality protections. Response follows:
    - Update guidance to acknowledge that confidentiality would be the same as would be used for market or financial inquiries as this tool supports those processes, and to encourage citing exam authority/confidentiality protections.
    - The information provided will be kept confidential consistent with each state's exam authority.

Consumer Complaint Requests (AHIP, Consumer Reps)
- Summary: Parties expressed feedback on the burden or clarity of the information being requested. Response follows:
    - Regulators generally agreed to remove these columns/questions but may ask for additional input on how this information can be gathered as development of this tool continues.

Definitions (ACLI, AHIP, AITC, APCIA, MO, Monitaur, NAMIC, RAA, RRC)
- Summary: Feedback was received to clarify existing definitions and to provide additional definitions of terms used in the Tool. Response follows:
    o Multiple parties recommended to define "high-risk". Refer to "Materiality and Risk Assessment" below.
    o "Consumer Impact" – clarified to refer to a decision that is subject to regulation.
    o "Generalized Linear Models" – removed, as the intention of the scope of the Tool is to encompass all AI/ML models, as implied in the term "Artificial Intelligence".
    o "Inherent Risk" – added to refer to risk prior to mitigation controls.
    o "Internally/Externally Trained Models" – will be clarified to refer to the development of models.
    o "Residual Risk" – added to refer to risk remaining after mitigation controls.
    o "Use Case" – clarified to refer to descriptions of usage.

Intent (ACLI, AHIP, APCIA, ITC, RRC)
- Summary: Parties expressed concerns about how the tool will be used and if it would be used for data calls, and how this work would relate to financial or market exams. Response follows:
    o Instructions were updated to clarify the tool's intent expressing that this tool supports existing processes.

Materiality and Risk Assessment (CAI, PA, CT, RRC, Consumer Reps)
- Summary: Parties expressed questions on the way the tool references materiality and risk assessment and requested that guidance or a definition of "high-risk" be provided. Response follows:
    o Clarified language that the company's assessments will drive the evaluation process at this stage but that regulators will request information to help understand and evaluate the company's assessment process.
    o Clarified instructions so that based on responses to Exhibit A, it is clear that regulators can determine that no further inquiry is necessary. Will also include this guidance in pilot planning discussions.

Regulatory Burden (MO, CAI, AHIP, APCIA, IRI, ACLI, AITC)
- Summary: Parties expressed concerns about the regulatory burden this tool would create. Response follows:
    o These sorts of tools help to create consistency in questions avoiding similar concepts being scrutinized in drastically different ways.

- The use of AI necessitates unique considerations in evaluating company operations. For instance, governance is a key concept in relation to using AI, but it's not woven into the fabric of the Market Regulation Handbook.
- Where industry had specific feedback, that was considered and regulators made several more targeted edits to cut back on the tool.
- This tool would be used on exams/analysis that would otherwise be planned – not a tool in search of an exam so the requests would come as part of a larger examination and not designed to be used for a larger information gathering exercise.
- Where companies have alternative means of providing necessary information that may be considered by regulators.
- The information requested by regulators represents the information necessary to understand a company's use of AI and therefore to assess and investigate the risk of such use.
- As part of the pilot process and continued public comment periods, regulators would ask for specific feedback on columns/rows/questions that are less relevant to the stated goals.

Regulatory Requirements (BCBSA, NALC)
- Summary: Parties expressed concern the tool had the appearance of creating requirements. Response follows:
    - The AI Systems Evaluation Tools represent the documentation regulators believe will be necessary to support their regulatory review. However, even after piloting the tool, regulators will continue to refine the tool particularly as best practices in managing AI risks evolve.
    - The Tool represents information regulators will request. Where that information is not available, regulators anticipate engaging with companies to understand the risk, the assessment of the risk, and to evaluate the mitigation of the risk. Nothing in the tool should be construed as required practice.
    - Some interested parties referenced alternative means of providing documentation including leveraging the NIST AI RMF Framework. As part of the pilot process, regulators may explore whether alternative industry standard documentation templates are available.

Risk of AI (NAMIC, AITC, Consumer Reps)
- Summary: Several parties expressed concern that the tools reflect a view that AI is riskier. Response follows:
    - Whether or not AI has more or less risk is not the focus of this tool. The tool is designed to help regulators understand a company's risk profile and how it addresses its inherent risk.

- This tool will help regulators get the information they need to understand risk.
- AI presents unique risks that require regulators ask different questions – our process already has tools and guidance to address the risks that humans can make. AI may not make the same mistakes, but it may make different mistakes or there could be process design issues in how a company implements AI. Additionally, there may be issues with the process to implement the AI which could present problems.
- The experience using this tool will inform longer-term discussions on how each process within the regulatory framework should be updated likely leading to referrals to D or E Committee groups.
- Regulators will need to understand a company's use of AI to design their exam approach. This tool provides a means of asking for information to help start that process of gaining an understanding of the company's use of AI.

Scope (AHIP, AITC, CAI, IRI, NALC, RAA)
- Summary: Parties expressed concern that this work goes beyond the Bulletin, that financial inquiries should not be in scope, that this tool might result in duplicate requests, that this tool should not scope in commercial carriers or reinsurers, or that H Committee should not be leading this work. Response follows:
  - Updated guidance to clarify that this Tool helps to identify and assess risk from the use of AI to mitigate adverse consumer outcomes and ensure the use of AI is in compliance with existing laws and regulations – which aligns with the expectations of the Bulletin.
  - This tool is not designed to decide who is subject to regulatory inquiry. This tool merely supports regulators performing an examination or other regulatory inquiry. Accordingly, scoping decisions related to who is subject to the tool's questions were not reflected in the tool. This tool will be a resource once a state has determined it is going to perform an exam or other inquiry and this tool would assist if the regulator has questions about a company's use of AI.
  - Regulators generally kept financial related inquiries included in the document.
  - This is an optional tool. E Committee and their related groups may decide not to use this tool based on how they evaluate it as fitting their solvency oversight responsibilities. However, at this stage of development, regulators wanted to explore the tool to see if it could beneficial financial regulators as well. Based on the results of the pilot, regulators may re-evaluate inclusion of financial questions in scope of the tool's contents.
  - To limit duplication of efforts in the company providing documentation, a statement referenced below was added to the guidance: "If information requested through the tool has already been provided to the department, the company's response should so state and reference when and how the information was provided."