WORKSHOP COMMENTS | NOVEMBER 15, 2026 | REG-2026-00007

Making AI-assisted insurance claims open to meaningful challenge

Kerry Thornhill

Independent researcher | contestability.ai

EMAIL kerry@contestability.ai

WEB contestability.ai

Recommendation

The Department should evaluate AI safeguards partly by what an affected claimant can actually do: identify a material ground for an adverse determination, investigate its evidential and policy basis, and pursue a supported objection before someone able to address it. A stated reason, an appeal address, or a human signature does not by itself demonstrate that capability.

This response addresses workshop topics 2-5, with an evaluation proposal relevant to topic 6. It recommends a practical test for explanations and review, illustrated by a fictional property claim. It does not report an actual insurer's conduct or an empirically validated standard. [1]

Existing duties and proposed implementation

California already requires more than a bare denial label. For first-party denials, 10 CCR section 2695.7(b)(1) requires written reasons with their known factual and legal bases and an explanation of how a cited provision applies. Subsection (b)(3) requires information about Department review; subsection (b)(2) recognizes a specified fraud-investigation disclosure limitation. These provisions are reproduced in the Department's adjuster guide. [2]

The recommendations below are proposed ways to make explanations and review usable when AI contributes to a claim determination. They should be assessed alongside existing obligations, not treated as proof that those obligations are absent or optional. This response does not offer a comprehensive account of insurance law, determine coverage, or transfer California's employment-specific rules into insurance.

The Necessary Informational Conditions (NIC) framework asks what informational operations a person needs for a defined investigative task. It distinguishes those functions from any particular document or interface. Institutional duties to provide review, authority to change decisions, and implementation of remedies require separate justification and assessment. [4]

A fictional denial: correct roof age, disputed causation

Assume a policy covers storm-caused roof damage and excludes deterioration attributable solely to ordinary wear. These simplified terms are stipulated for this example; actual policy language, causation rules, and burdens would need separate analysis.

After a storm, a policyholder supplies photographs and a contractor's report identifying fresh impact damage. An automated assessment supports denial. The letter cites the wear exclusion, describes the roof as old, and offers a review route. A human adjuster signs it.

The roof's age is recorded correctly. However, the assessment used an earlier image, and the uploaded contractor's report was not incorporated into the assessment. The process also treats age as grounds for attributing this particular damage to wear. The letter does not reveal those gaps or explain the inference. These are invented facts, not findings about an insurer or product.

The claimant can already object: a storm occurred and the submitted report supports storm damage. The further investigative problem is understanding why the insurer's answer rejects that account. The claimant needs a feasible way to distinguish omitted evidence, mistaken image selection, and an unsupported causal inference. Requiring another upload without checking its actual use may reproduce the same defect.

Claimant: "You have confirmed the roof's age. Which images and report did you consider, and what supports your conclusion that this damage came solely from wear?"

Illustrative inadequate response: "The age field is accurate, and an adjuster reviewed the result."

That response addresses data accuracy but leaves the material causal objection unanswered. A substantive review would examine the correct images, the report and contrary evidence, the applicable provision, and the inference linking them to the determination. It might support reversal, a reasoned decision to maintain the denial, or further investigation. The example leaves the outcome open.

Gilinsky and Marlow's discussion of claims handling, internal records, and human input motivates this inquiry. The fictional case and proposed tests are my application, not their reported case, recommendation, or endorsement. [3]

Recommended safeguards for topics 2-5

1. Make the material grounds investigable. A claimant should be able to identify the disputed loss or amount, the provision being applied, the decisive factual or evaluative findings, and the role AI played. Explain how those findings support the determination in this claim. A general description of the model or a confidence score may leave that connection unexplained.

2. Make evidence use traceable. Provide a usable account of the material records actually considered, their relevant dates or versions, and the treatment of contrary submissions. Distinguish receipt of a document from its use in the assessment. A short evidence summary, relevant extracts, a supported conversation, or a competent intermediary may supply the function; a complete technical log need not be the default claimant-facing form. Preserve sufficient internal provenance to investigate disputes and system changes.

3. Test the reviewer's response to the objection. Assign a person with relevant competence, time, access, and authority to investigate and change the determination within the applicable process. Review should engage the claimant's specific evidence and disputed inference, including an objection to a criterion's relevance or application. A second execution of the same model or a check of input accuracy alone should not be counted as resolution. Where an issue exceeds the reviewer's authority, identify who can address it and preserve the unresolved issue during escalation.

4. Preserve insurer accountability across vendors. Procurement and operational arrangements should give the insurer sufficient access to relevant inputs, outputs, versions, evidence handling, and known limitations to investigate the individual determination. Establish who resolves a vendor-related defect and how affected claims are identified. A vendor's general performance assurance does not answer whether a particular claimant's report was considered. Claimants should have a clear institutional contact rather than having to locate the responsible technology supplier themselves.

5. Use protected access that still permits challenge. Security, privacy, and legitimate confidentiality concerns can justify tailored access arrangements. Consider redaction, targeted explanations, or qualified intermediaries with appropriate access. Evaluate what the claimant can investigate through the substitute and how errors in its account can be challenged. A conclusory assertion that the system was checked is not evidence that the substitute works. These proposals do not displace lawful restrictions or require universal disclosure of source code, model weights, or other people's data.

6. Record disposition and follow-through separately. For each material objection, record the evidence examined, reasoning, decision owner, result, and any outstanding inquiry. If a correction is warranted, check that it reaches the claim determination and relevant downstream records. Individual correction, general system remediation, and the claimant's agreement are distinct outcomes. A changed decision does not by itself establish that the person had a usable investigative opportunity.

Evaluation before prescribing a universal form

The Department could invite insurers, policyholder representatives, and researchers to test these functions with controlled cases before adopting a universal disclosure template. Begin with tasks defined independently of the documents supplied: can a claimant or advocate identify the material defect, investigate the response, and formulate a supported objection within realistic limits of time, cost, expertise, language, and assistance?

Compare a conventional reason-and-review letter with a usable evidence-and-reason account and with a supported-review alternative. Include omitted records, incorrect versions, accurate data used to support a disputed inference, and a defensible determination that survives challenge. Examine accessibility and performance across participants rather than assuming an expert advocate's success establishes ordinary claimant usability. Use synthetic or appropriately protected records.

The strongest challenge to a proposed requirement is a successful substitute. If submitting the report to a competent reviewer reliably enables the same investigation without an evidence inventory, the inventory itself has not been shown necessary. Conversely, possession of an inventory is insufficient if the review process will not address the causal objection. NIC's conditional necessity claims remain open to such tests. [4]

Measure the person's investigative capability, the institution's response, and implemented consequences separately. Neither reversal rates nor the presence of a human reviewer alone establishes meaningful contestability. Document limitations and failed routes as well as successful cases.

I recommend that future guidance ask insurers to demonstrate at least one feasible route from a material objection to an evidence-responsive reconsideration, with defined responsibility for unresolved issues and any warranted correction. This provides a concrete way to evaluate transparency, human involvement, and vendor oversight while allowing different implementations.

Sources

[1] California Department of Insurance, Invitation to Discussion Workshop: Exploring the Responsible Use of Artificial Intelligence in Insurance, REG-2026-00007, November 13, 2026, pp. 3-6. These are workshop comments; a later formal rulemaking may require resubmission.

[2] California Department of Insurance, Guide for Adjusting Property Claims in California After a Major Disaster, January 9, 2025, pp. 11-12, reproducing 10 CCR section 2695.7(b)(1)-(3).

[3] Marshall Gilinsky and Madison Marlow, Protecting Policyholders as AI Is Developed for Insurance Claims Handling: Ensuring 'Decency and Humanity' in the Digital Age, Journal of Emerging Issues in Litigation 4(2), 2024, pp. 111-118, especially pp. 112-113 and 115-116.

[4] Kerry Thornhill, From Notice to Contestable Grounds: Necessary Informational Conditions for Algorithmic Contestability, research framework. The insurance example and implementation proposals are offered for scrutiny; no empirical validation is claimed.