Response to the ICO call for evidence on agentic AI

Kerry Thornhill
Research on algorithmic contestability
WEB contestability.ai

DATE November 13, 2026

Scope of this response

I recommend that the ICO assess safeguards by asking whether an affected person can investigate a material objection and obtain a reasoned response from someone able to act on the findings. A material objection is one that, if substantiated, could warrant correction, reconsideration, suspension, or a change in the decision or its use. People should not have to establish a hidden defect before they can obtain the information needed to investigate it.

This response draws on my conceptual work on Necessary Informational Conditions for Algorithmic Contestability (NIC). Its contribution is a method for examining investigative capability. The credit example below is hypothetical, and the evaluation described here is proposed research. I do not present these recommendations as findings from deployed agentic systems or as an empirically validated checklist.

Question 26 Identifying when the ADM provisions apply

The ICO's theme 4 appropriately assesses the complete decision process under Articles 22A–22D of the UK GDPR. A human action at one point does not establish meaningful involvement in the significant decision; the presence of an agent does not itself establish that the ADM provisions apply. Article 22C addresses safeguards for significant decisions based solely on automated processing and entirely or partly on personal data. The restrictions in Article 22B require a separate assessment. [1, 2]

I recommend recording the decision affecting the person, the contributing automated steps, and the human judgment said to affect that decision before it takes effect. This should identify what the person reviewing the output could examine and change, as well as what they actually did. Selecting a workflow or releasing its result may involve different judgment from examining its grounds. These records would help assess the organization's classification and make that classification open to scrutiny.

The Irish DPC's September 2026 report offers a useful EU comparison. In its service-ban example, the controller pointed to a user trigger and a human appeal. Following the DPC's intervention, the process changed to include human involvement before a ban. This is a regulatory case study under EU GDPR Article 22, rather than an interpretation of the amended UK regime. Its practical relevance is the need to examine when human judgment affects the decision. [3, p. 59]

Question 27 Establishing meaningful human involvement

Where an organization relies on meaningful human involvement to place a decision outside the ADM provisions, I recommend examining the reviewer's practical ability to test the grounds of the proposed decision. Authority should be accompanied by relevant evidence, adequate time, competence, and a usable route to seek further information. The record should show the issue considered and the basis for accepting, revising, or rejecting the recommendation.

Consider a credit workflow in which one agent retrieves payroll data, another converts it into monthly income, and a third recommends refusal. A loan officer approves the recommendation. Suppose the payroll amount covered four weeks but was treated as a calendar month. The officer's approval alone cannot establish whether the period was checked, whether the officer could correct the conversion, or whether a corrected assessment could affect the lending decision. The review record should answer those questions. An unchanged decision can still follow meaningful scrutiny; agreement rates alone cannot establish its quality.

California's SB 947 provides a limited comparison. Operative July 1, 2027, it prohibits sole reliance on automated decision systems for employee discipline or termination. Primary reliance requires human corroboration using source data or other relevant information. Uncorroborated output, and output found inaccurate, incomplete, or misleading, must be excluded from that decision. [4, §1522]

For example, a productivity flag might conflict with a worker's shift assignment. My proposed review record would preserve what the reviewer compared, how the conflict was resolved, and why the resulting decision followed. That record is a NIC design recommendation, distinct from the statutory requirements. The comparison illustrates substantive examination without treating California's rules as UK law.

Question 28 Making safeguards usable

Information, representations, human intervention, and contestation need to work together in the individual case. [2, Article 22C] In an agentic workflow, the affected person may encounter a final outcome without knowing which agent retrieved the relevant data, transformed it, or supplied the decisive inference. A route to appeal is of limited use if the person cannot discover what could reasonably be disputed.

In the credit example, I recommend a decision-specific explanation that makes the income amount and period recognizable, identifies how the amount was used, and explains how to raise a suspected mismatch. The person should be able to supply payroll evidence through an accessible channel and authorize a representative to assist. Access should not depend on already knowing which internal agent made the error.

The organization should preserve the original grounds and route the objection to a reviewer able to investigate the conversion and secure reconsideration of the lending decision. The response should explain what evidence was considered, whether the objection was substantiated, and what changed as a result. If the refusal remains justified

on other grounds, the borrower should receive those reasons. Correcting an intermediate value does not, by itself, resolve the challenge to the final decision.

France's consumer-credit provision helps make this distinction concrete. Article L312-16(VIII), scheduled to take effect November 20, 2026, links information and human intervention with the borrower's ability to express a view and request reconsideration of both the creditworthiness assessment and the lender's decision. The example should concern a new application; existing contracts generally remain subject to the earlier provisions. The comparison supplies no guarantee that reconsideration produces approval. [5]

The challenge route should also accommodate an objection to an inference or its use, even when the underlying data are accurate. Colorado's enacted SB 26-189 illustrates why these matters should be distinguished: its correction provision excludes opinions, predictions, scores, and protected evaluations, while separately providing for meaningful human review and reconsideration to the extent commercially reasonable. Its principal duties are scheduled for January 1, 2027. I use that distinction to motivate a design question, without importing Colorado's scope or qualifications into UK law. [6, §6-1-1705]

Recent disclosure research supports closer examination of what a promised review enables. In the publisher's abstract, Xiaoguo Chen reports that 259 of 117,460 sampled privacy policies disclose using automated decision-making. Within that subgroup, 19% promise human involvement, and 81% of those promises say nothing about the reviewer's authority. The seven guarantees examined include notice, reasons, access, correction, appeal, human review, and remedy authority. [7]

My use of this study is limited to the abstract. Disclosure coding does not establish what organizations actually do, and the full methods and coding have not yet been assessed here. NIC adds a further question for investigation: even where safeguards are stated, can the affected person or their representative use them to assess a material objection and obtain an institutional response?

The EDPB has selected automated candidate screening, including AI-driven screening, for its 2027 coordinated enforcement action. The announced focus links solely automated decisions with the related right of access. This is a selection of a future enforcement topic, not a finding against a particular organization. It provides a timely employment context for examining whether access enables a person to investigate and challenge a decision. [8]

Question 29 Logging and auditing across multiple stages

I recommend designing records around the questions an investigation must answer. For the credit example, those questions include what income period was supplied, how it was converted, which result the lender relied on, and whether the officer examined the suspected mismatch. A record that identifies agents and timestamps but omits the relevant transformation may establish that events occurred while leaving the objection unresolved.

A proportionate record could connect the decision identifier to the relevant input and its source, the transformation or rule applied, the applicable system version, the output used,

and the human action. It should preserve the original grounds when a later correction is made and link that correction to any reconsideration. Where the affected decision depends on another organization's service, arrangements for access and preservation should be established before a dispute arises. The person should not have to reconstruct the supplier network to initiate a challenge.

The affected person needs a usable account of the relevant record. Direct access, a dated calculation extract, or protected examination by a qualified intermediary may perform the necessary investigative work. A summary generated after the event should be checked against preserved evidence of what actually happened. It should not silently substitute a plausible new explanation for the original grounds. The aim is sufficient evidence for the inquiry, with appropriate protection for other people's data and confidential material; it is not indiscriminate disclosure or indefinite retention.

I propose testing these arrangements through omission and substitution. First specify the objection, the question to be answered, and an adequacy standard independently of the record being tested. For example, can the borrower establish whether the income period was converted correctly and whether the lender relied on that calculation?

Then compare feasible routes under stated time, cost, and access constraints. Remove information about the period from accessible routes and examine whether the question remains answerable. Restore the relevant comparison through a calculation extract or an intermediary who can communicate a usable, challengeable finding. If an adequate route omits the candidate investigative operation and its faithful equivalents, the necessity claim fails. If a different arrangement performs the same operation, the original document may be dispensable. Successful inquiry and the institution's response to its findings should be evaluated separately.

These are proposed tests. They do not establish that NIC's six candidate informational functions are exhaustive, necessary in every case, or validated across sectors.

Question 30 Permissions and user interfaces

I recommend testing permission settings against what reviewers can do when an objection is plausible. A reviewer who can approve an output but cannot obtain its material grounds, correct a relevant input, commission reassessment, or escalate the resulting decision may be unable to perform the review assigned to them. The interface should make those powers and their limits visible. Where relevant authority is distributed, it should identify an escalation route that can secure action within the time available.

The affected person's interface should make a challenge possible without technical knowledge of the agent network. It should identify the decision, provide an understandable account of the relevant grounds, accept evidence and representations, and show how to reach human review. People should be able to dispute an inference or governing reason as well as submit corrected factual data. An accessible alternative channel and representative access should be considered where the interface itself creates a barrier.

For workflows in which significant consequences can occur before review is complete, I recommend evaluating a pause or suspension mechanism, with clear authority and

reasons for using or declining it. Where an agent has propagated a disputed value or action, the review should examine which downstream effects require reconsideration and report what was addressed. These are design proposals for making the applicable safeguards effective, rather than claims that every intervention is legally required in every case.

Related answers on transparency and accountability

Question 15 Execution summaries and audit trails

I recommend linking the summary provided to the person to the evidence retained for the particular decision. It should explain the material input, transformation, and resulting action in terms the person can recognize. It should also identify how to obtain further information or raise an objection. A generic account of what the system usually does cannot establish what happened in this case. The proposed record and omission tests in question 29 provide a way to evaluate that connection.

Question 16 Challenging and correcting an agent action

The process proposed in question 28 starts with enough information to discover a possible ground for challenge. It then provides access or usable mediation, accepts representations and evidence, and brings the issue to someone able to secure a reasoned reconsideration. Correction of an agent's input should be connected to review of the decision and relevant downstream effects. The person should learn whether the objection was accepted and why the outcome changed or remained justified.

Question 22 Records supporting controller accountability

I recommend that organizations establish who preserves the relevant records, who can retrieve them, and who can secure action on an investigated objection. Contracts and technical arrangements should support those responsibilities across service boundaries. This proposal does not allocate controller or processor status by label; that requires assessment of the actual processing and roles. It addresses the practical risk that the body receiving a challenge lacks access to the evidence or authority needed to resolve it.

Sources

1. ICO, Theme 4 of the call for evidence on agentic AI. Current framing and evidence prompts.
2. Data (Use and Access) Act 2025, section 80, substituting UK GDPR Articles 22A–22D.
3. Irish Data Protection Commission, AI Insights Report, September 2026, printed p. 59, PDF p. 60.
4. California SB 947, chaptered text, §§1522, 1524 and operative-date provision. Official legislative record.
5. French Consumer Code, Article L312-16(VIII) and commencement note, version scheduled for November 20, 2026.
6. Colorado SB 26-189, signed act, §6-1-1701(15), printed pp. 8–9, and §6-1-1705, pp. 14–15. Legislative record.
7. Xiaoguo Chen, Can people challenge an AI decision? Measuring the contestability chain in law and corporate disclosure, AI and Ethics 6, article 642, November 6, 2026, DOI 10.1007/s43681-026-01471-y. Publisher's abstract only.
8. European Data Protection Board, November 11, 2026 announcement of the 2027 coordinated enforcement topic, section on the Coordinated Enforcement Framework.