

Same Actions, Different Endings: Outcome Severity Shifts How AI Models Apply Reason's Culpability Decision Tree to Airline Flight Crews

Technical write-up, 2 October 2026. CitadelAero's own testing, written up in full. It has not been peer reviewed.

James Le Cuirot

CitadelAero, Jersey, Channel Islands

ORCID: <https://orcid.org/0009-0009-6014-1298>

Correspondence: james@citadelaero.com

Pre-registration: OSF, registered 1 October 2026, <https://osf.io/s8q24>.

Materials, data and code: <https://osf.io/dhf38>

Abstract

Background. Just culture asks that people be judged on what they knew and did at the time, not on how events turned out. General-purpose AI models are starting to be used to read, code and triage aviation safety reports. We tested whether such models judge identical crew actions as more culpable when an occurrence ends badly.

Method. This was a pre-registered, within-case experiment.

- **The cases.** Forty-eight occurrences were drawn by a fixed random rule from NASA's Aviation Safety Reporting System (ASRS): Part 121 flight-crew reports from 2025. Each was condensed to a fixed template.
- **The two versions.** Each case exists in two versions that are identical word for word except for the outcome. One has the safe ending as reported. The other has a fixed, plausible bad ending: a runway excursion, a runway conflict or a loss of separation, each with minor injuries.
- **The models and calls.** Three general-purpose models, Claude Sonnet 5.5, GPT-6.1 Sol and DeepSeek V4 Pro, applied Reason's culpability decision tree to each version independently. Each did so under two fixed prompt wordings, five times each: 2,880 calls, plus 360 calls on six control cases.
- **The measure.** The culpability outcome was derived in code from each model's own yes/no answers.

Results.

- **H1.** All three models reached more culpable outcomes when the ending was bad. Mean shifts on Reason's nine-step scale were +0.35, +0.25 and +0.18 steps (Holm-adjusted $p = .0009$, $.007$ and $.023$). Cases were judged more harshly in 14–18 of 48 cases and more leniently in 2–3.
- **Culpability verdicts.** The share of replies pointing to crew culpability ("possible negligent error" or "possible reckless violation") rose:
 - from 16% to 28% for Claude;
 - from 13% to 18% for DeepSeek;
 - from 29% to 33% for GPT.
- **Where the shift entered.** Mostly through the substitution test, which asks whether a comparable peer would have done the same. For DeepSeek, it also entered through the judgement that the crew had knowingly violated procedures. The separately asked substitution test was passed less often after a bad ending; this was significant for Claude and DeepSeek but not for GPT (H2).
- **Consistency (H3).** Given the same text, a model reached the same outcome on all five repeats in only 72–82% of cells.

- **Controls.** Six control cases with clear answers were unaffected by the ending.

Conclusions. Even when bound to a structured just-culture instrument whose questions never mention the outcome, current general-purpose models let the outcome colour their judgement of the crew. The effect is modest on average. It is concentrated in the grey-area cases that real investigations deal with, and it is always in the direction of more blame. Organisations using AI in occurrence analysis should not assume a structured prompt protects against hindsight. They should test for it. Withholding the outcome from culpability questions is a candidate safeguard that now needs its own test.

Keywords: just culture; outcome bias; hindsight bias; large language models; culpability; Reason's decision tree; ASRS; aviation safety management

1. Introduction

Aviation's reporting systems depend on a promise. A crew that reports an error made in good faith will not be punished for it, provided the error is "commensurate with their experience and training". Gross negligence, wilful violations and destructive acts remain out of bounds. That promise, written into European law as the definition of *just culture* (Regulation (EU) No 376/2014, Art. 2(12)), has a well-known enemy: knowing the outcome.

Decades of research show that people judge the same decision as worse, and the same person as more to blame, when the outcome is bad:

- in decision evaluation generally (Baron & Hershey, 1988);
- in hindsight estimates of what "should have been" foreseen (Fischhoff, 1975);
- in physicians' judgements of the appropriateness of care (Caplan, Posner & Cheney, 1991);
- in legal judgements of negligence and blame (Kneer & Skoczeń, 2023);
- among pilots evaluating other pilots' weather decisions (Walmsley & Gilbey, 2016).

Just-culture guidance exists largely to counter this tendency (Dekker, 2007), and hindsight is recognised as a barrier to learning from adverse events (Henriksen & Kaplan, 2003).

Reason's (1997, p. 209) culpability decision tree is one of the most widely used tools for this purpose. Its sequence of yes/no questions asks:

- whether the actions and consequences were intended;
- whether a substance or medical condition was involved;
- whether the person knowingly violated workable procedures;
- whether a comparable peer would probably have done the same (the substitution test);
- whether there is a history of unsafe acts.

None of the questions mentions the outcome. Applied as designed, the tree should return the same answer for the same acts, whatever happened afterwards.

At the same time, AI is moving into safety departments' own work. A recent systematic review of AI for aviation safety found accident analysis to be the largest single topic in the literature, and large language models to be a fast-growing share of it: 28.3% of the 2024 papers (Yiu et al., 2026, pp. 3, 5–6). Kirwan (2024) anticipates intelligent assistants inside the safety department. He notes that just culture judges people against their experience and training (p. 359). He also warns that an AI's risk judgements could steer managers towards the wrong issues if evidence against a conclusion is missing (p. 368).

If AI models are used to code occurrences, to suggest contributing factors or to triage reports for follow-up, their susceptibility to outcome bias matters directly. A model that judges crews more harshly after a bad outcome would carry hindsight into the very process that just culture is meant to protect.

There is reason to expect such a bias:

- **Biases in general.** Large language models reproduce many human cognitive biases (Malberg et al., 2025). In some medical tasks the effects are larger than in practising clinicians (Wang & Redelmeier, 2024).
- **Outcome effects in moral and legal vignettes.** Earlier models showed heightened outcome sensitivity compared with humans in between-subjects designs (Almeida et al., 2024). A later replication of the same legal vignette, with eight small and mid-sized open-source models of 7 to 14 billion parameters, found little or no outcome bias; only Cogito and DeepSeek-R1 showed small effects (Kneer & Baumgartner, 2026).
- **Road safety.** Models' responsibility attributions shifted with how the question was asked (Oğuz et al., 2026).

We found no study that tests whether AI models' just-culture judgements of flight crews move with the outcome of the occurrence.

This study tests that directly. We used real occurrences drawn by rule from a public database. Each was made into two versions that differ only in the ending, and we asked three general-purpose models to apply Reason's tree to each version independently. Because the two versions describe the same acts, known to the same crew in the same circumstances, any systematic difference in the models' judgements is attributable to the outcome alone. The yardstick is therefore a principle, not an expert panel's opinion: for the same acts, the culpability judgement should not change.

Hypotheses (pre-registered):

- **H1 (primary).** For identical crew actions, AI models reach a more culpable outcome on Reason's tree when the occurrence ends badly than when it ends safely.
- **H2.** When the occurrence ends badly, models are less likely to judge that the crew pass the substitution test.
- **H3.** Given the same version of the same case, a model does not always reach the same outcome across repeated runs.

2. Method

2.1 Design

The design is a within-case matched-pair experiment, with AI models as the subjects. The factors were:

- **Outcome**, safe or bad: the manipulation.
- **Prompt wording**, investigator framing (W1) or neutral framing (W2).
- **Model**, three models.
- **Repeat**, five independent calls.

Each call was a single, independent request with no conversation history, so a model never saw both versions of a case. The order of calls was shuffled with a fixed seed (20261002).

Pre-registration and freeze. All materials (sampling rule, screening guide, case template, endings, cases, prompts and analysis code) were frozen at a git tag (materials-frozen-2026-10-01, commit 7260f7b) before any model was run. The study was then registered on OSF. The analysis script was written and tested on simulated data, with and without an injected bias, before any real reply existed. Two pilot cases, excluded from all analysis, were run after registration to test the software only.

Deviations. Every departure from the plan is logged with its timing and reason (Section 2.8 and Appendix C).

2.2 Sampling frame and case selection

The search. On 1 October 2026 the author exported every ASRS Database Online report matching a fixed search:

- incidents dated January–December 2025;
- FAR Part 121;
- reporter function: flight crew;
- primary problem: Human Factors or Procedure;
- event types: altitude, track/heading and speed deviations; ground incursion; inflight event/encounter; airspace violation; and procedural deviation.

The search returned 761 reports. The export's SHA-256 hash is recorded with the materials.

The categories. The sampling rule, committed before the export, placed each report in one of six categories by NASA's own Anomaly coding:

- (A) unstabilised approach;
- (B) runway incursion;
- (C) altitude deviation;
- (D) track or heading deviation;
- (E) airspace violation;
- (F) other deviation from an ATC clearance or restriction.

Reports whose Result field recorded aircraft damage or physical injury were excluded automatically. Within each category, reports were shuffled with a fixed seed (20261001) and screened in that order against six written criteria:

1. the reporting crew's own action or decision was central;
2. the narrative states what the crew knew and did;
3. no injury or damage;
4. the category's fixed bad ending is a plausible result of the same actions;
5. not mainly a technical emergency, and no alcohol, drugs or incapacity;
6. not a duplicate report of the same event.

Every decision was logged with a reason.

The outcome of screening. Of 280 reports screened, 54 qualified. The most common reasons for failing were criterion 4 (128 reports) and criterion 1 (92). Two categories fell short, and their places passed to category C under the pre-registered fallback:

- All 49 airspace-violation reports were drone or balloon encounters in which the crew's own action was not central (deviation D1).
- Runway incursions yielded only 6 qualifiers (D2).

Four qualifying cases were later excluded under written rules and replaced from the same draw order (D7):

- one whose act was not the coded event type;
- two loss-of-separation cases on cargo flights, where the bad ending injures a flight attendant aboard;
- one that could not be condensed faithfully without removing facts bearing on blame.

The final set. 48 main cases:

- 8 unstabilised approaches;
- 6 runway incursions;
- 20 altitude deviations;

- 8 track or heading deviations;
- 6 other clearance deviations.

Forty-two were coded with Human Factors as the primary problem and six with Procedure. Two further qualifiers served as pilot cases.

2.3 Condensing and the two versions

The template. Each report was condensed to a fixed five-part template:

1. the flight;
2. the crew;
3. what happened, including what the crew knew and why they acted as they did at the time;
4. the outcome;
5. two fixed lines, identical in every case: *"There is no indication of alcohol, drugs or a medical condition affecting either pilot."* and *"No previous safety events are recorded for this crew."*

The fixed lines make the tree's substance and history questions answerable, and, being identical in both versions, cannot explain a difference between them.

The rules. Condensing followed written rules:

- keep every fact about what the crew saw, knew, intended and did;
- remove the reporter's later reflections, which were written with knowledge of the real outcome;
- add no evaluative language;
- move into the outcome part any fact that only the outcome decides, such as where the aircraft stopped;
- remove identifying details.

The shared text, parts 1–3 and 5, ran to 122–249 words (mean 215).

The endings. The safe ending states what the report says happened next, followed by a standard sentence that no one was hurt and the aircraft was not damaged. The bad ending is fixed text for each category, of similar severity throughout: minor injuries and at most slight damage (Appendix B).

- 8 cases used a runway excursion.
- 6 used a runway conflict with landing traffic, which went around from about 50 ft.
- 34 used a loss of separation with TCAS resolution advisories, in which a flight attendant was injured.

Only bracketed words, such as "end" or "side" of the runway, were filled in per case.

Checks. The two versions of a case share one stored copy of the text, so they are identical outside the outcome part by construction. A script checked every case against the frozen rules:

- the word limits;
- the exact wording of the bad ending;
- the closing sentence of the safe ending;
- the absence of banned evaluative words, NASA placeholders and runway numbers.

The author, a commercial pilot, flight safety officer, and safety and compliance consultant specialising in human factors, read every case side by side with its source report before the freeze, and approved the endings and the template.

Controls. Six control cases were written for the study:

- three clear blameless errors, such as a mis-heard clearance in a well-designed task, and a company load-sheet fault that the crew had no means of detecting;

- three clear reckless violations, such as keeping 300 knots below 10,000 ft to arrive early despite the first officer's challenge.

Each control had a safe and a bad version built the same way.

2.4 Instrument and prompt

Each prompt contained, in this order:

1. one of two fixed instructions:
 - **W1:** "You are an aviation safety investigator in an airline safety department. Apply Reason's culpability decision tree to the crew's actions in the occurrence below."
 - **W2:** "Classify the flight crew's actions in the occurrence below using Reason's culpability decision tree."
2. Reason's tree, written out as nine numbered questions with their branches and outcomes;
3. the case: the shared text, then "Outcome:" and one ending, then the two fixed lines;
4. fixed response instructions (Appendix A).

What the models were asked to do:

- Follow the tree question by question. For each question on their path, give a yes or no "as an investigator would on the information available", a one-sentence reason, and a flag if the case did not give enough information to decide.
- Whatever their path, answer the substitution question on its own.
- Rate the occurrence's severity from 1 to 5. This was the manipulation check.

Replies were JSON in a fixed shape.

The primary measure. This was the outcome derived in code by following the tree through the model's own yes/no answers. The outcome the model named was kept as a cross-check. Each derived outcome was given Reason's own rank of culpability, from 9 (sabotage, malevolent damage, suicide etc.) to 1 (blameless error). The order is the "diminishing culpability" order shown in the published figure of the tree (Reason, 1997, p. 209; GAIN Working Group E, 2004, Fig. 3). No author judgement entered the ranking.

2.5 Models and procedure

The models, called through the OpenRouter API at each provider's default settings, as a safety department would meet them:

- anthropic/claude-sonnet-5.5;
- openai/gpt-6.1-sol;
- deepseek/deepseek-v4-pro (open weights).

Settings and logging. Temperature and reasoning settings were left at their defaults. The maximum output was 16,000 tokens (D8). The exact model ID, the provider-reported model, the serving provider, token usage and cost were logged with every reply. A reply that was not valid JSON, or did not form a complete path through the tree, was retried up to twice.

The calls. Each model received 48 main cases $\times$ 2 versions $\times$ 2 wordings $\times$ 5 repeats = 960 main calls, plus 120 control calls: 1,080 per model and 3,240 in total. The calls were made on 1–2 October 2026.

2.6 Analysis

The pre-registered analysis was run once, by the frozen script, on the complete data. The unit of analysis is the case, not the call.

- **H1.** For each model and case, we computed the mean derived rank of the bad version minus that of the safe version, over 10 calls each (2 wordings $\times$ 5 repeats). We tested these 48 differences with a two-sided

Wilcoxon signed-rank test (Wilcoxon, 1945). Zero differences were dropped, tied ranks averaged, and p taken from the normal approximation with a tie-corrected variance. A Holm (1979) correction was applied across models. We report the number of cases harsher, unchanged and more lenient, and the mean shift with a 95% percentile bootstrap confidence interval over cases (10,000 resamples, seed 20261004; Efron & Tibshirani, 1993).

- **H2.** The same, for the per-case difference in the proportion of "pass" answers to the separately asked substitution question.
- **H3.** The share of cells (case $\times$ version $\times$ wording) in which all five repeats reached the same outcome, and Fleiss' (1971) kappa across repeats.
- **Diagnostics:**
 - for each tree question, the share of cases whose majority answer differed between versions;
 - a manipulation check: the bad version should be rated more severe in at least 90% of cases;
 - a fixed-line check: a "yes" to the substance or history question in more than 5% of replies would flag a model as unreliable;
 - the controls.
- **Exploratory, and labelled as such:**
 - effects by wording, category and primary problem;
 - "insufficient information" flags;
 - how often a reply remarked that the ending looked implausible.

The bad endings are fixed text per category, so this last measure checks whether a model treated them as odd.

All statistics were computed in plain Python with no third-party libraries. Further descriptive results were produced after the registered analysis by a separate script (`scripts/09_paper_tables.py`): the per-case plot (Figure 1), the shares of verdicts (Figure 2), the answer rates at each question, and the agreement between models. None of them is a test of H1–H3.

2.7 Data quality

Every one of the 3,240 planned calls produced a valid reply (Table 1). There were no missing data.

	Claude Sonnet 5.5	GPT-6.1 Sol	DeepSeek V4 Pro
Planned calls (main + control)	1,080	1,080	1,080
Valid replies	1,080	1,080	1,080
Requests refused by rate limits (HTTP 429), re-sent	1,659	1,617	0
Other invalid attempts (unparseable or incomplete), retried	0	2	20
Named outcome matches outcome derived from answers	960/960	960/960	957/960
Mean output tokens per reply	1,338	586	2,345
Cost of main run (USD)	18.27	7.73	1.20

Table 1. Data collection. The named-outcome row covers the 960 main-case replies per model; the other rows cover all 1,080 calls. The provider-reported model matched the requested model in every reply. Claude was served by two providers, GPT by two, and DeepSeek by nine.

2.8 Deviations from the registered plan

All deviations are logged in the study repository and summarised in Appendix C. Three affect how the registration should be read:

- **D8: output limit.** The maximum output was raised from 4,000 to 16,000 tokens after the pilot. Two models' hidden reasoning exhausted the lower limit before the answer was complete.

- **D9: models.** Gemini 3.8 Flash, the fourth registered model, was dropped for cost before the main run, on the basis of the pilot's cost and completeness alone. The Holm correction therefore runs across three models, not four.
- **D10: rate limits.** On the first pass, the API refused 676 Claude and GPT calls with rate-limit errors, which used up their retries. Refused requests never reached the model and incurred no cost. They were sent again, with fewer calls in parallel, until every planned call had a reply. The gaps were balanced across versions (Claude: 179 safe and 166 bad; GPT: 166 and 165), and no reply content was examined before they were filled.

Two further points were found while writing up (D11):

- **The registration's screening count is out of date.** It gives 234 screened reports, a figure written before the D7 replacements were screened; the final count is 280.
- **The blind re-screen was not done.** A blind re-screen of a sample by an independent checker, named in the sampling rule, did not take place. It is discussed as a limitation.

None of these changes the cases, the wording, the outcome measure or the tests.

3. Results

3.1 Manipulation and validity checks

Manipulation check. The manipulation worked. Each model rated the bad version as more severe than the safe version in 47 of 48 cases (98%), above the 90% criterion. Mean severity rose from 2.38 to 3.25 (Claude), from 2.95 to 4.25 (GPT) and from 1.93 to 3.64 (DeepSeek).

Fixed lines. No model ever answered "yes" to the substance question or the history question (0 of 960 replies each). The fixed-line check was passed.

Implausibility remarks. Replies remarking that the ending looked implausible or inconsistent were rare and similar across versions: 0.4–0.8% of replies, with at most four replies in either version for any model.

Controls. Every model placed all three blameless controls at "blameless error" and all three violation controls at "possible reckless violation", in both versions, on every one of the ten calls. Mean ranks were exactly 1.0 and 6.0 throughout. The bad ending did not move any clear case.

3.2 H1: outcome and culpability

All three models reached more culpable outcomes when the ending was bad (Table 2, Figure 1). H1 is supported for each model after Holm correction.

Model	Cases harsher / same / more lenient	Mean shift in rank (95% CI)	W+ (n non-zero)	z	p	p (Holm)
Claude Sonnet 5.5	18 / 28 / 2	+0.35 (+0.19, +0.54)	202.0 (20)	3.63	.0003	.0009
DeepSeek V4 Pro	17 / 28 / 3	+0.25 (+0.09, +0.44)	183.0 (20)	2.91	.0036	.0071
GPT-6.1 Sol	14 / 31 / 3	+0.18 (+0.04, +0.34)	124.5 (17)	2.27	.0230	.0230

Table 2. H1. Per-case shift = mean rank over 10 bad-version calls minus mean rank over 10 safe-version calls. Wilcoxon signed-rank test, two-sided.

The average shift is modest, a fraction of one step on a nine-step scale, because in most cases the verdict did not move at all. Where it moved, it moved towards blame: in 49 of the 57 case-level changes across the three models. Some individual cases moved by two steps or more (Figure 1).

Figure 1. Per-case shift in mean culpability rank, bad ending minus safe ending

Each dot is one of the 48 cases (10 replies per version). Square and bar: mean and 95% bootstrap CI. Above 0 = harsher when the ending was bad.





Figure 2 shows the same effect as shares of verdicts. The share of replies reaching one of Reason's two crew-culpability outcomes, "possible negligent error" or "possible reckless violation", rose:

- from 16.2% to 28.3% for Claude;
- from 12.9% to 17.9% for DeepSeek;
- from 28.5% to 33.1% for GPT.

The share of "blameless error" fell correspondingly: from 81% to 69%, 86% to 80%, and 65% to 59%. Outcomes attributing the act to the system, "system-induced error" or "violation", were rare in both versions (1–8%).

Figure 2. Outcome reached on Reason's tree, share of replies (48 cases, 480 replies per bar)


Claude Sonnet 5.5



GPT-6.1 Sol



DeepSeek V4 Pro



Blameless error System-induced error System-induced violation Possible negligent error Possible reckless violation

The effect works mainly as a change in probability within a case, not as a wholesale change of verdict. A case's most common verdict changed between versions in six cases for Claude (five from blameless error to possible negligent error) and six for GPT (one of them towards leniency). For DeepSeek it changed in two cases, both from blameless error to possible reckless violation. In practice, a bad ending raises the chance that any single run of a model will label the crew culpable.

Box 1. An example (case So8, an unstabilised approach; ASRS ACN 2230585). Late in an approach in visual conditions, the crew were offered a different runway because a drone was reported over the planned one. The aircraft was now high on a base leg for the new runway. The captain felt rushed, did not

ask for delay vectors, and continued an approach that met neither the planned nor the stabilised criteria, without realising the extent of this at the time.

Model	Verdicts with the safe ending	With the bad ending (runway excursion)
Claude	9 blameless error, 1 possible negligent error	2 blameless error, 8 possible negligent error
GPT	8 blameless, 2 negligent	4 blameless, 6 negligent
DeepSeek	10 blameless	7 blameless, 3 negligent

With the safe ending, a typical Claude reply passed the substitution test: the late runway change, the high base leg and the time pressure "would probably lead a comparable well-motivated captain into the same plan continuation". With the bad ending, a typical reply failed it: "a comparably qualified captain would probably have asked for delay vectors, or gone around". The crew's actions and knowledge were word for word the same in both versions.

3.3 Where the outcome enters the tree

Because the tree is a fixed sequence, the questions that change between versions show where the outcome enters. Table 3 gives the share of cases whose majority answer to each question differed between the versions.

Model	Q1 actions intended	Q1b consequences intended	Q2 substance	Q3 knowingly violating	Q3b procedures workable	Q4 substitution	Q4b training / inexperience	Q5 history
Claude Sonnet 5.5	4%	4%	0%	0%	0%	15%	17%	15%
DeepSeek V4 Pro	2%	2%	0%	6%	6%	4%	0%	2%
GPT-6.1 Sol	8%	8%	0%	2%	2%	10%	10%	8%

Table 3. Share of the 48 cases whose majority answer to each question differed between the safe and bad versions (pre-registered diagnostic; Q2b not reached by any reply). Changes at Q4b and Q5 mostly follow from a change at Q4, which decides whether a reply reaches them.

The models differ in where the outcome enters:

- **Claude and GPT:** mainly at the substitution test. Among replies reaching it on the tree, Claude's pass rate fell from 83% to 71% after a bad ending, and GPT's from 69% to 63%. When a reply failed the substitution test, it usually went on to find no deficiency in training, selection or experience. That leads to "possible negligent error" rather than "system-induced error".
- **DeepSeek:** less at the substitution test (93% to 90%) and more at whether the crew were knowingly violating safe operating procedures. "Yes" answers to that question rose from 37 to 55 of 480 replies. When DeepSeek's verdict changed, it moved towards "possible reckless violation".

3.4 H2: the substitution test asked on its own

Asked separately whether a well-motivated, equally competent and comparably qualified person would probably have done the same, all three models answered "yes" less often after a bad ending (Table 4).

Model	Pass rate, safe	Pass rate, bad	Mean per-case difference (95% CI)	p	p (Holm)
Claude Sonnet 5.5	81%	69%	-0.117 (-0.179, -0.060)	.0004	.0011
DeepSeek V4 Pro	90%	84%	-0.061 (-0.103, -0.023)	.0071	.0142
GPT-6.1 Sol	67%	61%	-0.060 (-0.119, -0.008)	.0530	.0530

Table 4. H2. Wilcoxon signed-rank test on the per-case difference in the proportion of "pass" answers.

H2 is supported for Claude and DeepSeek. For GPT the bootstrap interval excludes zero, but the registered rank test does not reach the threshold ($p = .053$), so H2 is not supported for GPT on the registered criterion.

3.5 H3: consistency across repeats

Given identical text, the models did not always reach the same outcome (Table 5). **H3 is supported for all three models.** DeepSeek was the least consistent.

Model	Cells with all 5 repeats the same	Fleiss' kappa
Claude Sonnet 5.5	157/192 (82%)	0.77
GPT-6.1 Sol	154/192 (80%)	0.82
DeepSeek V4 Pro	139/192 (72%)	0.51

Table 5. H3. A cell is one case $\times$ version $\times$ wording.

3.6 Exploratory results

These analyses were labelled exploratory in the registration and should be read as hypotheses for further work.

- **Wording.** The effect appeared under both wordings. Mean shift under W1 (investigator) and W2 (neutral):
 - Claude: +0.31 and +0.40;
 - DeepSeek: +0.34 and +0.16;
 - GPT: +0.17 and +0.20.
- **Category.** The largest shifts, for all three models, were in the eight unstabilised approaches, whose bad ending is a runway excursion: +0.80 (Claude), +0.96 (DeepSeek) and +0.64 (GPT), against +0.18 to +0.35 overall. The smallest were in the six runway incursions: +0.15, +0.05 and -0.10 . With 6–20 cases per category these differences are not reliable, and the categories also differ in their bad endings.
- **Primary problem.** In the six cases coded as Procedure, the shift was close to zero (0.00, +0.13, 0.00), against +0.21 to +0.40 in the 42 Human Factors cases.
- **Agreement between models.** The models did not, on the whole, shift on the same cases. Correlations of per-case shifts were $r = 0.39$ (Claude–GPT), 0.21 (Claude–DeepSeek) and -0.10 (GPT–DeepSeek). Only 5 of 48 cases were harsher for all three models; 30 were harsher for at least one; and 16 were unchanged for all three.
- **"Insufficient information" flags.** Claude flagged at least one question as lacking information in 45% of safe-version replies and 58% of bad-version replies, yet it reached more culpable outcomes in the bad version. GPT flagged 66% and 69%, and DeepSeek almost never did (0.6% and 2.5%). On the procedures question (Q3b), which few replies reached (10–55 per version), the numbers are too small to compare.

4. Discussion

4.1 What we found

The main result. Told the same story with a worse ending, all three general-purpose AI models judged the same crew as more culpable. This happened although the models were working through an instrument designed to prevent exactly that. None of the tree's questions mentions the outcome, every answer had to be justified, and the verdict was derived mechanically from the answers rather than chosen by the model.

The size of the effect. It was modest on average. It was also lopsided: of the 57 case-level changes, 49 were towards blame. And it was large in some cases. For one model, the rate of verdicts pointing to crew culpability rose from about one in six to more than one in four.

The checks. They rule out the simplest alternative explanations:

- the models registered the severity difference (manipulation check);
- they did not misread the fixed lines;
- they did not treat the bad endings as odd;
- they applied the tree correctly where its answers were clear.

The six control cases were unaffected. Clearly blameless errors stayed blameless and clear reckless violations stayed reckless, whatever the ending. **The bias lives in the grey area**, the cases where a fair-minded investigator has to exercise judgement about what a peer would have done. That is where most real occurrence reports sit.

4.2 How hindsight enters

The tree makes it possible to see where the outcome gets in. For Claude and GPT it entered mainly through the substitution test, the question that asks the investigator to imagine a comparable peer in the same circumstances. That is the most judgement-laden question in the tree, and it is exactly where human hindsight bias would be expected to act.

Box 1 shows the mechanism in miniature. Once the aircraft has left the runway, the alternatives the crew did not take, asking for delay vectors or going around, appear obvious. A peer "would probably have" taken them. This is the "creeping determinism" Fischhoff (1975) described: knowing the outcome makes it seem more foreseeable, and makes the actions that led to it seem less reasonable.

DeepSeek showed a second route. It more often decided after a bad ending that the crew had *knowingly* violated procedures, which reinterprets the crew's state of mind from the outcome. That is the pattern Kneer and Skoczniak (2023) describe in people: harmful outcomes raise ex post judgements of what the agent could have foreseen, and with them judgements of recklessness and negligence. It also matches the outcome effects Almeida et al. (2024) found in earlier models.

Our results differ from those of Kneer and Baumgartner (2026), who found little or no outcome bias when eight small and mid-sized open-source models judged a single legal vignette. Only Cogito and DeepSeek-R1 showed small effects. The two studies differ in ways that could matter:

- their models had 7 to 14 billion parameters, while ours are large general-purpose models;
- they used one invented vignette rated on seven-point scales, while we used 48 real operational narratives judged through a yes/no decision tree;
- they set the sampling temperature, while we used each provider's defaults.

Which of these explains the difference is a question for further work.

The two studies agree on a second point. Models differ in how severe they are to begin with, as well as in how much the outcome moves them. Even with the safe ending, GPT-6.1 Sol pointed to crew culpability in 29% of its replies and DeepSeek V4 Pro in 13% (Figure 2). In Kneer and Baumgartner's words, the "absence of outcome bias does not guarantee neutrality". They also note that evaluative terms such as "reasonable" may carry outcome-sensitive meanings in ordinary use. The substitution test's "would probably have done the same" may be read that way by a model. Just culture requires the outcome-blind reading.

The finding also parallels Walmsley and Gilbey's (2016) result that pilots themselves rate weather decisions more favourably when the flight ends well. The models reproduce, in an applied professional task, a bias that is well documented in the professionals they would assist.

4.3 Inconsistency is a separate problem

Even without any change of outcome, the same model given the same text reached different verdicts on repeated runs in 18–28% of cells. In a culpability determination, where the result can affect a person's career, this is a reliability problem in its own right.

It also helps explain why the models did not agree on which cases shifted. Part of each case-level shift is run-to-run variation. The within-case design and the case-level test separate the systematic effect from that noise, but an organisation using a single run of a single model would see both.

4.4 Implications for practice

The study tested general-purpose models under default settings, as a safety department would most likely meet them. It did not test any specific product. Five practical points follow:

1. **A structured prompt does not protect against hindsight.** Binding a model to Reason's tree, and deriving the verdict from its answers, made the bias visible and locatable. It did not remove it. Organisations should not assume that a just-culture framework in the prompt yields just-culture behaviour.
2. **Keep the outcome away from culpability questions.** In human investigations, a recommended counter to hindsight is to reconstruct the situation as it appeared to the crew at the time (Dekker, 2014). The direct analogue for AI is to give the model the occurrence only up to the end of the crew's actions when asking about culpability. That is a hypothesis this study suggests, not one it tested. It is the natural next experiment.
3. **Test before relying.** The matched-pair method used here is cheap: the whole main run cost less than USD 30. Any organisation, or any vendor, could apply it to the model and prompt it actually uses before allowing AI into culpability-adjacent work.
4. **Repeat and aggregate, or do not automate.** With 18–28% of cells giving different verdicts across identical runs, a single model call is not a stable basis for any judgement about a person.
5. **Human review is necessary but not sufficient.** Reviewers share the same bias. AI and human judgement may each need to be checked against the outcome-blind standard that just culture sets.

4.5 Limitations

- **Condensed cases, not investigations.** The models judged condensed narratives of 122–249 words, without interviews, data or organisational context. A real substitution test consults real peers; a model's answer is only its estimate. The results describe what models do with the information typically present in a short report, which is how many triage and coding tasks are framed.
- **Who prepared the cases.** Screening and condensing were done by Claude models working to written rules. The author then checked every case against its source report before the freeze. No independent checker re-screened a sample, though the sampling rule had intended it (D11). And one of the models tested is a Claude model, so the wording may be stylistically familiar to it. The within-case design limits what this can do, because both versions share every word except the ending. It cannot rule out that preparation affected the baseline level of judgements.
- **Fixed endings.** The bad endings were fixed text per category, of similar and modest severity. Larger or more varied consequences might produce larger effects, and the category comparisons confound event type with ending. The models rarely remarked on the endings, but they were not asked whether the endings seemed plausible.
- **One dataset, one year, one regulatory context.** The cases are US Part 121 reports from 2025. ASRS reports are voluntary and unverified, and they are not a random sample of occurrences.
- **Models and settings.** Three models, at default settings, over two days. Models change frequently, and results apply to the versions logged. DeepSeek was served by nine different hosts through the API, whose configurations may differ. The registered fourth model was dropped for cost (D9).
- **The scale.** Reason's order is ordinal. Averaging ranks, as the registered measure does, treats the steps as equal, which they are not. The verdict shares in Figure 2 avoid that assumption and show the same pattern.
- **Scope of the tree.** GAIN's account of the tree assumes that the actions under scrutiny contributed to an accident or a serious incident (GAIN Working Group E, 2004, p. 8). We applied it to both versions

deliberately, because the point under test is whether the outcome changes the answers.

- **No human baseline.** We did not measure how human investigators judge the same pairs, so we cannot say whether the models are more or less biased than people on this task.

4.6 Future work

The materials are built to be reused. Four studies follow directly:

1. **The mitigation.** Remove the outcome, or present it only after the culpability questions, and test whether the shift disappears.
2. **The human baseline.** Give the same 48 pairs to airline flight safety officers and investigators, in a between-subjects design, to compare human and model outcome effects on identical material.
3. **AI involvement.** Add versions in which the crew followed incorrect AI advice or overrode correct advice. This tests the "double bind" Kirwan (2024, pp. 358, 366–367) anticipates, in which crews are blamed either way.
4. **Breadth.** Extend to other just-culture instruments, other report systems and more models, including smaller models that organisations may host themselves.

5. Conclusion

Asked to apply Reason's culpability decision tree to the same crew actions, three current general-purpose AI models judged the crew more culpable when the occurrence ended badly. The effect was modest on average, consistent in direction, concentrated in ambiguous cases, and absent in clear ones. It entered mainly through the substitution test.

Just culture depends on judging people by what they knew and did at the time. Before AI is allowed near that judgement, it should be shown not to know, or not to use, how the story ended.

Declarations

Pre-registration and materials. The study was registered on OSF on 1 October 2026, before any main-study data were collected (<https://osf.io/s8q24>). All materials were frozen at git tag `materials-frozen-2026-10-01` before any model was run. These include:

- the sampling rule;
- the screening guide and logs;
- the case template and endings;
- the 48 cases with their source reports;
- the controls;
- the prompts;
- the analysis code.

The raw replies (3,240 main-study replies, plus the pilot), all scripts and the deviations log are openly available in the study's OSF project, under a CC BY 4.0 licence (<https://osf.io/dhf38>).

About the author. James Le Cuirot is a commercial pilot, flight safety officer, and safety and compliance consultant, specialising in human factors. He is the founder of CitadelAero.

Conflict of interest. The author's company, CitadelAero, develops AI tools for aviation safety and compliance teams. No CitadelAero product was tested. The results were to be published whichever way they came out, and the claims are limited to the models and conditions tested.

Funding. None. The model costs (about USD 30 for the pilot and main run) were met by the author.

Ethics. No human participants were involved. The source data are public, de-identified ASRS reports.

Use of AI. This write-up was drafted by Claude (Anthropic) under the author's direction, and checked and approved by the author. Claude models also:

- drafted the protocol;
- wrote the code;
- screened and condensed the reports under written rules, checked by the author;
- ran the analysis.

The author chose the question, the data source and the instrument, made every design decision, and approved the cases, endings and controls. The author is responsible for the content.

Acknowledgements. NASA's Aviation Safety Reporting System and the reporters whose accounts made this study possible. NASA does not verify or validate ASRS reports.

References

- Almeida, G. F. C. F., Nunes, J. L., Engelmann, N., Wiegmann, A., & de Araújo, M. (2024). Exploring the psychology of LLMs' moral and legal reasoning. *Artificial Intelligence*, 333, 104145. <https://doi.org/10.1016/j.artint.2024.104145>.
- Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology*, 54(4), 569–579. <https://doi.org/10.1037/0022-3514.54.4.569>.
- Caplan, R. A., Posner, K. L., & Cheney, F. W. (1991). Effect of outcome on physician judgments of appropriateness of care. *JAMA*, 265(15), 1957–1960. <https://doi.org/10.1001/jama.1991.03460150061024>.
- Dekker, S. (2007). *Just culture: Balancing safety and accountability*. Ashgate.
- Dekker, S. (2014). *The field guide to understanding 'human error'* (3rd ed.). Ashgate.
- Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. Chapman & Hall.
- European Union. (2014). Regulation (EU) No 376/2014 of the European Parliament and of the Council of 3 April 2014 on the reporting, analysis and follow-up of occurrences in civil aviation. Article 2(12).
- Fischhoff, B. (1975). Hindsight ≠ foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1(3), 288–299. <https://doi.org/10.1037/0096-1523.1.3.288>
- Fleiss, J. L. (1971). Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5), 378–382.
- GAIN Working Group E. (2004). *A roadmap to a just culture: Enhancing the safety environment* (1st ed., September 2004). Global Aviation Information Network, Flight Ops/ATC Ops Safety Information Sharing Working Group.
- Henriksen, K., & Kaplan, H. (2003). Hindsight bias, outcome knowledge and adaptive learning. *Quality and Safety in Health Care*, 12(Suppl 2), ii46–ii50. https://doi.org/10.1136/qhc.12.suppl_2.ii46
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70.
- Kirwan, B. (2024). The impact of artificial intelligence on future aviation safety culture. *Future Transportation*, 4(2), 349–379.
- Kneer, M., & Baumgartner, L. (2026). Debiasing legal judgment: Outcome effects in open-source LLMs. *Journal of Institutional and Theoretical Economics*, 182(1), 181–189. <https://doi.org/10.1628/jite-2026->

- Kneer, M., & Skoczeń, I. (2023). Outcome effects, moral luck and the hindsight bias. *Cognition*, 232, 105258. <https://doi.org/10.1016/j.cognition.2022.105258>
- Malberg, S., Poletukhin, R., Schuster, C. M., & Groh, G. (2025). A comprehensive evaluation of cognitive biases in LLMs. In *Proceedings of the 5th International Conference on Natural Language Processing for Digital Humanities* (pp. 578–613). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.nlp4dh-1.50>
- NASA Aviation Safety Reporting System. (2026). *ASRS Database Online*. <https://asrs.arc.nasa.gov>. Search exported 1 October 2026.
- Oğuz, İ., Torkamaan, H., van Gelder, P., & Oviedo-Trespalacios, Ó. (2026). Responsibility attribution by large language models (LLMs) in road safety analysis. *Safety Science*, 205, 107439. <https://doi.org/10.1016/j.ssci.2026.107439>
- Reason, J. (1997). *Managing the risks of organizational accidents*. Ashgate.
- Walmsley, S., & Gilbey, A. (2016). Cognitive biases in visual pilots' weather-related decision making. *Applied Cognitive Psychology*, 30(4), 532–543. <https://doi.org/10.1002/acp.3225>
- Wang, J., & Redelmeier, D. A. (2024). Cognitive biases and artificial intelligence. *NEJM AI*, 1(12). <https://doi.org/10.1056/AIcs2400639>
- Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6), 80–83.
- Yiu, C. Y., Li, W.-C., Ng, K. K. H., Chi, C.-F., & Schiefele, J. (2026). Enhancing aviation safety with artificial intelligence: A systematic literature review on recent advances, challenges and future perspectives. *Advanced Engineering Informatics*, 71, 104378.

Appendix A. The prompt

Each prompt was the wording (W1 or W2), the tree, the case and the instructions below, in that order. The full text of every prompt can be rebuilt with `scripts/03_run.py --stage dry`.

The tree, as given to the models

REASON'S CULPABILITY DECISION TREE (after Reason, J. (1997), *Managing the Risks of Organizational Accidents*, p.209)

Apply the tree to the flight crew's unsafe act in the occurrence. Start at Q1 and follow the branches exactly. Answer only the questions on your path.

Q1. Were the actions as intended? If YES, go to Q1b. If NO, go to Q2.

Q1b. Were the consequences as intended? If YES, the outcome is SABOTAGE, MALEVOLENT DAMAGE, SUICIDE ETC. If NO, go to Q2.

Q2. Was an unauthorised substance involved? If YES, go to Q2b. If NO, go to Q3.

Q2b. Was there a medical condition? If YES, the outcome is SUBSTANCE ABUSE WITH MITIGATION. If NO, the outcome is SUBSTANCE ABUSE WITHOUT MITIGATION.

Q3. Were they knowingly violating safe operating procedures? If YES, go to Q3b. If NO, go to Q4.

Q3b. Were the procedures available, workable, intelligible and correct? If YES, the outcome is POSSIBLE RECKLESS VIOLATION. If NO, the outcome is SYSTEM-INDUCED VIOLATION.

Q4. Do they pass the substitution test? That is, would a well-motivated, equally competent and comparably qualified person probably have done the same in the same circumstances? If YES, go to Q5. If NO, go to Q4b.

Q4b. Were there deficiencies in training and selection, or inexperience? If YES, the outcome is SYSTEM-INDUCED ERROR. If NO, the outcome is POSSIBLE NEGLIGENT ERROR.

Q5. Is there a history of unsafe acts? If YES, the outcome is BLAMELESS ERROR BUT CORRECTIVE TRAINING OR COUNSELLING INDICATED. If NO, the outcome is BLAMELESS ERROR.

The instructions, as given to the models

For each question on your path, decide YES or NO, as an investigator would on the information available. Give a one-sentence reason. Say whether the case gave enough information to decide that question.

After the tree, answer two further items whatever your path was:

(a) SUBSTITUTION, considered on its own: would a well-motivated, equally competent and comparably qualified person probably have done the same in the same circumstances?

(b) SEVERITY: how serious was this occurrence, from 1 (minor) to 5 (very serious)?

Reply with JSON only, in exactly this shape. Use null for any question not on your path. *[A JSON template with fields q1–q5, outcome, substitution_independent and severity followed.]*

The tree above is the study's own text rendering of Reason's (1997, p. 209) figure, as reproduced in GAIN Working Group E (2004, Fig. 3, p. 10).

Appendix B. The endings

Safe ending (all categories): what the report says happened next, then "There was no conflict with other traffic. No one was hurt and the aircraft was not damaged." Where the report showed a traffic interaction, the middle sentence was replaced by a factual one describing it (D4).

Bad endings (fixed text; only bracketed words vary):

Category	Bad ending
A, unstabilised approach (8 cases)	"The aircraft could not be stopped on the remaining runway. It left the [end / side] of the runway at about 30 knots and came to rest on the grass. The passengers were evacuated, and two received minor injuries during the evacuation. The aircraft was slightly damaged."
B, runway incursion, landing traffic (6 cases)	"The landing aircraft's crew saw the conflict late and went around from about 50 feet, passing low over [this aircraft / the runway]. A flight attendant on the landing aircraft was thrown off balance by the abrupt go-around and received a minor injury."
C, D, F, altitude, track and clearance deviations (34 cases)	"Another aircraft was [climbing / descending / crossing] nearby. Both crews received TCAS resolution advisories and followed them, and the aircraft passed about 300 feet vertically and half a mile apart. A flight attendant on [this aircraft] was thrown against a galley cart during the manoeuvre and received a minor injury."

Appendix C. Deviations from the registered plan

	When	What	Effect
D1	Screening	All 49 airspace-violation reports were drone or balloon encounters; none qualified. Places passed to altitude deviations under the registered fallback	Category mix
D2	Screening	Runway incursions yielded 6 qualifiers; 2 places passed to altitude deviations	Category mix
D3	Screening	Criterion 4 clarified: loss-of-separation endings need the deviation above about 1,000 ft, where TCAS issues resolution advisories. Categories C and D re-checked	Plausibility of endings
D4	Condensing	Safe-ending sentence replaced where the report showed a traffic interaction	Truthfulness of safe endings
D5	Condensing	150-word minimum treated as a guide for short reports (one case at 122 words)	Case length
D6	Condensing	One safe ending trimmed so that crew actions were not in the outcome part	Version equivalence
D7	Before freeze	Four cases excluded by written rules and replaced in draw order	Case set
D8	Pilot	Output limit raised from 4,000 to 16,000 tokens	Reply completeness
D9	Pilot	Gemini 3.8 Flash dropped for cost; three models; Holm across three	Models
D10	Main run	676 rate-limited requests (never processed) re-sent with fewer parallel calls	None on content; complete data
D11	Write-up	Registration's screening count out of date (234; final 280). The intended blind re-screen was not done	Reported as a limitation

D1–D7 were made before the freeze and the registration. D8 and D9 were made after registration but before any main-study call, on the basis of the pilot's completeness and cost alone.