CATASTROPHISM RESEARCH ASSOCIATION

THE METHODOLOGY OF THE ASSOCIATION

Methodology Standards

VERSION 1.1 · ADOPTED 2026-07-21

Purpose: the standard every Association publication is held to. It defines how a claim is classified, how a source is graded, how a model earns its standing, and what the Association publishes when its own prior work is contradicted.

Applies to: every publication, dataset release, model, figure, forward claim, curated entry, and public statement issued under the Association's name. Start here: this is the document a reader wants if they want to know the rules.

1. Purpose and authority

The Association's standing rests on a single claim: that its research can be held to a published standard. This document is that standard. It defines how claims are classified, how sources are graded, how models are validated, how machine assistance is disclosed, how other researchers' frameworks and property are treated, how predictions are registered and scored, how contradiction is handled, and how work enters the Association's shelves. It is published openly, applied without exception, and versioned under founder change control. The Managing Partner owns the standard and its enforcement; material amendments require both founders' consent.

The methodological discipline described here is operated on the Association's internal research infrastructure, which provides claim registration, source tracking, gate-kept knowledge promotion, and calibration of the published record. The standard governs outputs; the infrastructure enforces it in practice. A companion document, the CRA Contributor Field Guide, restates this standard as working checklists and templates so that any capable contributor can apply it without further instruction; where the two differ, this document governs.

2. Epistemic classification

Every substantive claim in every Association publication carries one of three labels.

2.1 Established Finding (EF). A claim supported by direct measurement or by multiple independent, corroborating sources of Tier S1 or S2 (Section 4), with no unresolved contradicting evidence of equal grade. EF claims cite their sources at the point of use. An EF claim about a dataset must be reproducible from that dataset by a third party.

2.2 Working Hypothesis (WH). A proposed explanation or mechanism that is consistent with the evidence but not established. Every WH must be accompanied, in the same publication, by (a) the observations it explains, (b) at least one falsification condition per Section 3, and (c) its current corroboration status. A WH may not be restated as an EF in later work without the evidentiary basis being shown.

2.3 Open Question (OQ). An acknowledged unknown material to the analysis. OQ labels are obligatory, not optional humility: where the evidence does not reach, the publication says so explicitly. Suppressing an OQ that a competent reviewer would identify is a standards violation.

2.4 Labeling mechanics. Labels appear inline at first statement of each substantive claim and in a claims summary table at the end of every publication. Compound sentences mixing classifications are rewritten until each claim carries exactly one label. Background material restating the published literature may be cited without labels; anything the Association itself asserts is labeled.

2.5 Reclassification. Claims move between classes only by documented review: WH to EF requires the corroboration standard of 2.1; EF to WH or OQ requires an erratum or a revision notice per Section 12.4. Reclassifications are recorded in the claims register with date and rationale.

3. Falsification

3.1 The rule. No Working Hypothesis is published without explicit falsification conditions: observable outcomes which, if they occurred, would require the hypothesis to be revised or withdrawn.

3.2 Specification. A falsification condition names (a) the observable, (b) the instrument or dataset in which it would appear, (c) the threshold or qualitative signature, and (d) where the hypothesis is time-bound, the window within which the test applies. Conditions that no realistic observation could trigger do not satisfy this section.

3.3 Resolution. When a falsification condition triggers, the Association publishes the outcome within one publication cycle, whether or not it favors prior work. The discipline of publishing findings that contradict our prior work is the operational core of this standard, not a slogan.

4. Source discipline

4.1 Tiers. Sources are graded on intake and cited with their grade available to readers on request:

Tier	Definition	Examples
S1	Primary instrumental or archival data from an authoritative producer	IERS EOP series, IGRF, GRACE products, observatory records
S2	Peer-reviewed literature, or independently replicated analysis	Journal publications, replicated community analyses
S3	Documented secondary sources of identifiable provenance	Technical reports, curated databases, credentialed public analysis
S4	Uncorroborated material	Unattributed claims, single-witness reports, unverifiable compilations

4.2 Corroboration scaling. The evidentiary bar scales with the weight of the claim. EF status requires S1 support or at least two independent S2 sources. A WH may rest on S2/S3 material if labeled accordingly. S4 material may motivate an OQ; it may never carry an EF or serve as the sole support of a WH.

4.3 Temporal decay. Evidence ages on a stated schedule. Instrumental series are current to their latest release; a claim resting on superseded model versions or deprecated datasets loses its grade until re-based. Every publication states the data-currency date of its key series.

4.4 Independence. Sources sharing a common origin count once. Corroboration means independent derivation, not repetition.

5. Contradiction handling

5.1 Parallel presentation. Where credible evidence of comparable grade points in opposing directions, both are presented side by side with their grades. The Association does not silently reconcile, average away, or omit the weaker-fitting result.

5.2 Resolution requirements. Each documented contradiction carries an explicit statement of what evidence would resolve it, entered as an OQ in the claims register.

5.3 Internal disagreement. Disagreement between Association researchers on a published question is part of the scientific record: it is documented in the publication or its discussion note, attributed with consent, and never smoothed over for the sake of institutional voice.

6. Quantitative and reproducibility standards

6.1 Data availability. Analyses of public datasets cite the dataset version and provide the processed data underlying published figures. Member-tier releases include the cleaned datasets and analysis notebooks the publication rests on.

6.2 Uncertainty. Quantitative claims carry stated uncertainty where the underlying data supports its estimation, and figures render uncertainty envelopes rather than bare curves wherever the source provides them.

6.3 Figures. All figures conform to the CRA Design Language (40_Brand): measured data only, schematics explicitly labeled, on-figure source lines, and captions that pose open questions rather than asserting answers.

6.4 Methods sufficiency. Every quantitative publication describes its method to the standard of independent reproduction: data sources, preprocessing, algorithms, and parameter choices.

7. Forward claims and calibration

7.1 Registration. Any statement about future observable states of a monitored system is a forward claim. Forward claims are registered before publication with: statement, epistemic label (a forward claim is a WH by construction), falsification condition, resolution date or trigger, and registration date. Forward claims produced by models are additionally governed by Section 8: only models at maturity grade M2 or higher may issue them.

7.2 Resolution and scoring. At the resolution date or trigger, the claim is scored against the observable and the outcome is published. Unresolved claims may be extended only by a documented revision, never silently.

7.3 The calibration record. The Association maintains and periodically publishes its cumulative forward-claim record: claims made, resolved, and the hit rate, without selective omission. No external claim of predictive skill is made except from this record.

7.4 No unregistered prediction. Association researchers do not publish, under the Association's name, forecasts that bypass registration. Commentary on possibilities is labeled WH or OQ and framed accordingly.

8. Modeling standards

8.1 Scope. A model, for the purposes of this standard, is any quantitative construct whose outputs are presented in Association work: statistical fits, dynamical simulations, machine-learned estimators, and composite pipelines alike. Model outputs are Working Hypotheses by construction; only the measured inputs and the documented performance record of the model itself may be stated as Established Findings.

8.2 The model card. No model output appears in Association work without a published model card stating: model name and version; author(s) and their declared professional and analytical track record relevant to the modeled system; the system modeled and the question addressed; the assumptions register; input datasets with versions; parameters and the calibration method that produced them; code availability; the machine-assistance disclosure required by Section 9; the current maturity grade; and a

change log. The declared track record is subject to the misrepresentation provisions of Section 12; the gate on publication, however, is the validation record of the model, not the biography of the modeler.

8.3 Maturity grades. Every model carries exactly one grade, printed on its card and wherever its outputs appear:

Grade	Name	Requirement
M0	Conceptual	Stated mechanism or structure; no validated implementation. Outputs may be discussed only as WH or OQ and may not be plotted against data as if predictive.
M1	Verified	Implementation demonstrated correct: reproduces its stated test cases and documented sanity checks; code or sufficient specification available for review.
M2	Hindcast-tested	Documented out-of-sample performance against held-back history, with skill compared to a stated baseline or null model.
M3	Live-scored	One or more registered forward claims resolved under SOP-4, with the cumulative score published.

Forward claims may issue only from M2 or M3 models. Public statements of predictive skill may rest only on the M3 record, per Section 7.3.

8.4 Hindcast discipline. The held-out interval is declared before testing; no tuning against the holdout is permitted; skill is reported against an explicit baseline (persistence, climatology, or a stated null); and negative or null results are reported with the same prominence as positive ones. A hindcast that cannot be reproduced by a reviewer from the model card does not confer M2.

8.5 Review. Models enter Association work through SOP-5 (model intake and validation review). Grades are re-examined on every material version change and may be withdrawn on evidence of leakage, tuning on the holdout, or irreproducibility.

9. Machine-assisted work

9.1 Disclosure. Material machine assistance in producing a work (drafting, coding, analysis, literature synthesis, or figure preparation) is disclosed in the work, in a standard statement identifying what was assisted. Routine tooling (spell-checking, compilation, standard library calls) requires no disclosure.

9.2 Provenance rule. Machine output is never itself evidence. Every quantitative result must be re-derivable by a third party from the named data and the stated method, without access to the assisting system. A result that cannot be re-derived does not enter the archive, whatever produced it.

9.3 Sophistication is not evidence. Review under this standard evaluates derivations, data, and reproducibility, never vocabulary, formatting, or apparent technical fluency. Polished packaging around

an unvalidated model confers nothing; the audits of Section 12 and the ladder of Section 8 apply identically to all work.

9.4 Accountability. The human author owns every claim in their work. Fabricated citations, invented data, or unverified machine-generated results published under a contributor's name are treated as misrepresentation under Section 12, regardless of their mechanical origin.

9.5 Symmetry. The Association applies this section to itself. Its use of analytical infrastructure and machine assistance in its own operations and publications is disclosed on the methodology page of catastrophism.org and in the works concerned, under the same standard it asks of contributors.

10. Framework representation and third-party property

10.1 Topic fields. Association copy, figures, and metadata describe research by topic field. A researcher's named theory, framework brand, or property name appears in Association work only with the written consent of its originator, per the Research Partnership Standards Section 5.

10.2 Fidelity requirement. Work that claims to model, test, extend, or refute a named third-party framework must demonstrate fidelity to the originator's published specification of that framework. Before the name is used, the originator is offered a fidelity review with a stated response window. Consent to the name and confirmation of fidelity are recorded with the work.

10.3 The rename rule. If consent is withheld or the originator does not respond, the work may still publish on its scientific merits: it is renamed to describe, in neutral topic-field terms, what it actually models, and it makes no claim to represent the named framework. Consent governs the use of the name, never the publication of the analysis.

10.4 Straw-man prohibition. Attaching a framework's name to work that materially departs from that framework's published specification is a standards violation, handled as misrepresentation under Section 12. This applies equally to sympathetic and critical treatments.

10.5 Derived and joint work. Extensions of a partner's framework, and work built on partner-contributed material, follow the intellectual property terms of the Research Partnership Standards: the originator retains their property, attribution is explicit, and joint work is owned as agreed in writing by its authors.

11. The Archive and the Collection

11.1 Two shelves. The Association maintains two clearly separated shelves. The

Archive holds work reviewed under this standard (SOP-1 and SOP-2): Association publications, member submissions accepted after review, datasets prepared under the repository standard, and model cards. The **Collection** holds curated external material: outside analyses, threads, posts, talks, and documents that inform the field. Collection entries are held by link, metadata, and attribution; they are graded at intake under Section 4; and they are labeled so that no reader can mistake them for Association claims.

11.2 No absorption. The Collection references; it does not reproduce. Entries link outward to the owner's own property, carry the owner's name, and are removed at the owner's request under SOP-6. Curation implies neither endorsement nor review; it asserts only that the material is relevant and its provenance is stated.

11.3 No laundering. Collection material is cited in Association work under its source grade, exactly as any external source. Presence in the Collection confers no elevation. The path from Collection to Archive is commissioning or submission under SOP-2, with full review.

11.4 Compensation. Attribution and linking are free and carry no payment in either direction. Work commissioned from a Collection contributor into the Archive is compensated under the Research Partnership Standards honoraria schedule. A finder's honorarium of $100 to $250, at the Managing Partner's discretion, is paid when a Collection entry leads directly to a commissioned Archive work, to the entry's original author or to the member who surfaced it.

12. Review, publication, and correction

12.1 Internal review. Every Association publication passes a standards review before release: label audit (Section 2), falsification audit (Section 3), source audit (Section 4), figure audit (Section 6), and, where models are involved, model audit (Section 8) and framework-name check (Section 10). The Managing Partner or a designated reviewer signs the review; the sign-off is recorded.

12.2 Member submissions. Member work enters the Archive through the review process in SOP-2 (intake, screen, peer review, resolution), applying this standard in full. Acceptance is a statement that the work meets the standard, not that the Association endorses every conclusion; attribution remains with the author.

12.3 Conflicts. Reviewers with a material conflict on a submission recuse. Sponsored research discloses the sponsor in the publication itself.

12.4 Errata and retraction. Errors are corrected by dated errata appended to the work; claims that fail are reclassified visibly. Retraction removes a work from the active archive but preserves a citation stub and the reason. The archive is versioned; history is never silently rewritten.

13. Change control

This standard is versioned (v1.0, v1.1, v2.0). Amendments require founder consensus and are logged with date and rationale. Publications state the standard version they were reviewed under. The operating procedures implementing this standard are maintained separately as the CRA Operating Procedures (SOP-1 through SOP-6) under the same change control, and the CRA Contributor Field Guide restates it for working use.

Document record

Field	Entry
Document code	CRA-STD-METH-1.1
Version	1.1
Adopted	2026-07-21, under Managing Partner sign-off
Ratification	Co-founder ratification recorded per Section 13
Impression	Second impression, 2026-08-03. Front matter re-set for readers; no normative text altered.
Companion documents	Operating Procedures SOP-1 to SOP-6; Contributor Field Guide
Published	Openly at catastrophism.org/methodology
Change control	Section 13. Amendments issue as successive versions with a logged rationale; superseded versions are retained.

Amendment log.

Version	Date	Change
1.0	2026-07-20	Initial draft: Sections 1 to 7 core, review, change control.
1.1	2026-07-21	Added modeling standards (8), machine-assisted work (9), framework representation (10), the Archive and the Collection (11); renumbered review to 12 and change control to 13. Adopted under Managing Partner sign-off; co-founder ratification to be recorded on receipt of written approval.

◇

Version 1.1 · Catastrophism Research Association. Published openly as the Association's methodology standard and cited from the methodology page of catastrophism.org.