

ASA Evidence Standard v0.1

STATUS: FROZEN RESEARCH PROTOCOL v0.1

This standard defines what must be recorded before an ASA forecast may be presented as an experimental result. It prevents post-hoc reinterpretation and makes trajectory claims falsifiable, timestamped and auditable.

Evidence boundary

$T < T_0$: observed evidence available to ASA
 $T = T_0$: forecast generated, timestamped and sealed
 $T > T_0$: hidden continuation used only for later evaluation

Required before every run

Experiment identifier and protocol version; research question and target event; evidence source, interval and cutoff; forecast horizon; admitted and hidden evidence; baselines; evaluation measures; confirmation and invalidation conditions; known limitations and dual-use classification.

Integrity and evaluation

The frozen manifest, admitted-evidence inventory, trajectory report, alternatives, invalidators, timestamp and integrity record are preserved before outcome reveal. Completed runs retain lead time, false positives, missed events, calibration, baseline comparison, temporal error, perturbation stability and invalidator behavior. Negative and failed results remain in the ledger.

Hostile validation

ASA is tested against non-failure drift, abrupt events without useful precursors, missing or corrupted observations, matched decoys, reversals after apparent convergence and bounded input perturbations.

Role of local AI

Local AI explains the evidence packet. It does not create evidence, alter frozen analytical values or deliver the final verdict.

Public / private boundary

Public: research questions, protocol versions, synthetic experiments, aggregate results, failures, limitations and public-safe demonstrations.

Private: ASA Core source, protected algorithms, scoring logic, parameters, prompts, runtime configuration, sensitive data and security controls.

Governing principle

ASA earns credibility through increasingly difficult, timestamped and falsifiable experiments - not through increasingly confident narratives.