

Coordinated Correction Record

Package v1.1 | 10 October 2026 | DRAFT / NOT FINAL / NOT LOCKED

This successor addresses Claude's review of package v1.0. Current editions are Crosswalk v0.6.3, Pilot v0.3 and SEC-CDP-002 v0.1.2. The complete earlier package is preserved byte-for-byte. The three major findings and seven minor findings have the following dispositions.

Finding	Disposition	Location	Correction
C-1	ACCEPT WITH PRECISION	Crosswalk 14.2 and 14.4	Finite-bank primary estimator and randomization-based uncertainty; conditional bank power. Optional population-model failure does not invalidate valid bank inference.
S-1	ACCEPT RISK; QUALIFY EXPLOIT	SEC-CDP-002 R06-R08 and 4A	Withhold manifest digest; independently blind secret components; receipt leakage tests. No recovery benchmark or actual compromise claimed.
S-2	ACCEPT	SEC-CDP-002 2, 3A, R13-R15	Agent zone, per-unit access, paired-variant exclusion and plaintext asset inventory added.
Minor 1	ACCEPT	Crosswalk 14.2	Remove garbled model/endpoint sentence; missingness view follows equation.
Minor 2	ACCEPT WEIGHTING OPTION	Pilot 6 and 8	Equal weights for six lists within arm and expertise stratum for descriptive bands; empty cells make that contrast unavailable.
Minor 3	ACCEPT	Pilot 6	Instructions explicitly disclose that retracted submissions remain scored; include in ethics submission.
Minor 4	ACCEPT	Pilot 10	Actual N_R and 48*N_R denominators, including early closure and zero-N rules.
Minor 5	ACCEPT AND STRENGTHEN	Pilot 7 and scorer	Every flag requires authoritative accepted_before_closure=True; canonical deadline still enforced.
Minor 6	ACCEPT	SEC-CDP-002 R19	Human/data incident notification, session closure, debriefing and agreed payment.
Minor 7	ACCEPT	SEC-CDP-002 R15	Exact stream prompts/settings, tool access and context records; unavailable sampling seeds disclosed.

Interpretive qualifications

C-1 is a real inconsistency in the previous specification. A parametric bootstrap does not inherently require new items: a conditional version can hold items fixed. This successor explicitly selects a weighted evaluator-level fixed-bank estimator and design-consistent randomization-based inference for confirmatory work. It separates hypothetical-population modeling, keeps missingness and calibration requirements, and does not claim that a sharp-null permutation test automatically supplies an average-effect interval. The pilot's proposed full-bank participant bootstrap is exploratory and distinct; it still requires calibration.

S-1 is accepted as a potential disclosure channel. $C(48,24) = 32,247,603,683,100$ and $C(24,12) = 2,704,156$ are counts of balanced binary label patterns, not demonstrated workloads or times to recover a real manifest. Candidate checking requires sufficiently reconstructible bytes and access to the exposed digest; secret configuration or additional unknown fields can change feasibility. This package claims neither an attack run nor guaranteed compromise. Removing the digest from receipts and blinding individual secret components addresses the specified channel under stated assumptions. Public salts alone would not stop candidate guessing.

S-2 distinguishes the module from its surroundings. In the proposed exported-seed profile, signing/wrapping keys stay in the module while block secrets and labelled material enter protected software memory. The agent interface, paired-variant controls, response segregation and asset lifetimes now have explicit requirements. No HSM certificate or document hash proves those requirements are met.

The chosen Minor 2 correction standardizes descriptive band rates over six order lists within arm and expertise stratum. It can make a band contrast unavailable when cells are empty; at the proposed small pilot size this is plausible. Missing cells are reported, not silently dropped. No band confidence interval or hypothesis test is adopted. The overall fixed-bank primary summary is unchanged.

Verification and source scope

Claude reports independently reproducing the previous 133 checks, all three change replays and all 16 frozen text spans. These remain reviewer-reported results; current executed results are in 05_VERIFICATION/verification_results.json. The new scorer rejects false, absent and non-boolean event-order fields as well as late timestamps. Fixtures are software inputs, not participant observations. Receipt confidentiality, stream isolation, bank construction, order lists and statistical coverage have not been implemented or experimentally validated.

Primary sources checked this round: Stark's Statistics 240 randomization notes (M02), <https://www.stat.berkeley.edu/~stark/Teach/S240/Notes/ch3.htm> ; and Gottesman's University of Maryland cryptography Lecture 19 (R11), <https://www.cs.umd.edu/class/fall2022/cmsc456/lecture19.pdf> . They support distinctions between randomization and population inference, and between hiding and binding. Applying them to this design is our specification judgment, not their validation of it. No full re-audit of S01-S34 or R1-R10 is claimed.

Remaining gates

All eight pilot execution gates and infrastructure Gates A-F remain open. The fixed-bank pilot can develop without SEC-CDP-002, but recruitment is still not authorized. The two provenance dispositions and independent source confirmation of S01 remain pending before final promotion, along with non-AI methodological and engineering review. No Archive, Mirror Engine, Master Manifest, website, GitHub, Arweave or production system was changed. No study, model trial, power simulation, generator or consensus service was executed.

The supplied review is represented by the disposition table and qualifications here; this is an authored response, not a verbatim transcript of Claude's message. Review research circulation, recruitment, infrastructure execution and final promotion separately. No maturity score or safety certification is assigned.