

SEC CDP 001 Pilot Protocol

Draft v0.1 | 10 October 2026 | Companion to ER-13 Draft v0.6.1

NOT FINAL / NOT LOCKED / NOT OPEN FOR RECRUITMENT

This external pilot is intended to determine whether a controlled human review task is feasible and whether its measurement and analysis procedures work. It compares later flaw detection after exposure to conspicuous harmless errors with detection after corrected versions of those outputs. The settings below are concrete proposals for review. No participants have been enrolled and no behavioral outcomes are supplied.

1 Purpose and study status

The pilot estimates completion, timing, missingness, scorer agreement, hit rates and false alarms, and explores whether the intended mixed model is estimable. It is not powered or authorized to establish efficacy, deception, AGI behavior, human indispensability, or an Archive safety benefit. The crosswalk remains the evidence map; this protocol governs the proposed task once an execution-ready successor is reviewed and frozen.

Proposed enrollment is 60 randomized adults, approximately 30 per arm. This is a feasibility allocation chosen to expose workflow and measurement problems, not a power-derived sample size or assurance of precise variance estimates. Recruitment stops at 60 randomized participants; post-randomization withdrawal or technical failure does not trigger replacement. Report every randomized person and every missing item. Screening failures before randomization may be replaced and are counted separately. No internal efficacy stopping or sample-size adaptation is planned.

The pilot is external to any later confirmatory study. Pilot participants and pilot item instances will not be reused in the confirmatory dataset. Changes informed by the pilot require a new version, new frozen materials and an independently reviewed confirmatory sample-size justification.

2 Question and inference boundary

The question is whether prior exposure to conspicuous errors changes detection of harder flaws in later, separate outputs under this review policy. Exposure and Control differ in whether selected exposure outputs contain an obvious error. This contrast includes observable error presence, its salience, any confidence change, and the experience of finding errors. It cannot identify which of those mechanisms caused an effect.

The pilot does not test within-output co-occurrence. Post-exposure target items contain one intended target flaw, with no added conspicuous decoy. A conspicuous-error versus matched subtle-error experiment would require a new comparator and protocol. S19 model wrong-answer agreement is not an endpoint here, so its chance-agreement baseline is not required for this pilot.

3 Proposed settings

Setting	Proposal	Interpretation
People	60 randomized adults; 1 to 1 allocation	Feasibility proposal only; no confirmatory N
Exposure	8 outputs, 30 seconds each	4 conspicuous-error outputs in Exposure; corrected counterparts in Control
Challenge bank	48 outputs viewed by every participant	24 target-flaw and 24 benign; same items in both arms

Setting	Proposal	Interpretation
Strata	3 task families; 8 target and 8 benign items per family	Equal target and benign weights of one third per family
Challenge deadline	60 seconds per item	Voluntary final submission allowed; no revisiting
Session budget	About 65 minutes including instructions and debrief	Maximum 48 minutes for challenge items, plus exposure and setup
Feedback	No correctness feedback during exposure or challenges	Same response task and source labels in both arms
Materials	Frozen local non-adaptive item bank	Dynamic generator and seed-commitment implementation deferred

4 Participants and arrangements

Proposed eligibility is age 18 or older, ability to read the study language and consent, and ability to use the tested desktop interface. Recruit adults who routinely read written instructions or reports, including a range of review experience. This convenience sample supports conclusions about recruited participants; population generalization requires a justified sampling frame. Do not recruit people whose employment evaluation depends on the investigator or their supervisor.

Before randomization, record a pseudonymous participant ID, years of document-review experience, frequency of reviewing AI-generated text, and the first-attempt score on six task-comprehension questions. Retain first-attempt and final-pass scores separately. Do not collect names inside response logs. Participants must answer all six comprehension questions correctly after at most two instruction attempts; those unable to do so do not enter randomization. Record screening counts and reasons without publishing identifiers. The comprehension items must be benign and disjoint from experimental materials.

Use baseline review experience for blocked allocation: less than two years versus at least two years of weekly document review. Record the underlying answer before assignment. Within each stratum, use a concealed computer-generated sequence of randomly selected blocks of four and six. The custodian holds the assignment key; participants, scorers, and the masked analyst do not receive it. Record allocation timestamp, software version, and sequence hash before first use. Do not create the live allocation sequence until execution approval.

A responsible investigator, recruitment channel, compensation amount, data custodian, retention period, consent text, and applicable ethics determination must be named before recruitment. The session uses fictional non-operational materials, yet informed participation, withdrawal, and debriefing arrangements still require review. This document does not determine exemption or issue ethics approval.

5 Materials and independent validation

Use three families: arithmetic consistency in fictional inventories, consistency with an explicitly supplied fictional rule, and internal consistency across a short fictional report. Each challenge is approximately 100 to 180 words, contains all information needed for review, and uses no real medical, legal, financial, security, or industrial-control decision. Target items contain exactly one intended flaw. Benign items meet the same rubric and have no known flaw. Build matched structural templates, but show only one variant from each template to any participant; a benign correction of a target shown to that person must not appear elsewhere.

Construct 24 independent target templates and 24 distinct benign templates, eight of each per family. The Exposure set uses eight additional templates not appearing in the challenge bank. Four have an obvious

arithmetic or explicit-rule violation in Exposure and a corrected counterpart in Control; four are identical benign fillers. Match wording and display length as closely as possible and document unavoidable changes. Use the same neutral source label, Fictional review material, in both arms; do not assert actual AI provenance.

Before pilot enrollment, two reviewers who are not pilot participants independently inspect every item, verify its answer key and flaw count, and assign difficulty and salience ratings. A third reviewer adjudicates disagreements without access to any participant outcome. Replace ambiguous or multiply flawed items before freezing. Conduct a small separate usability check with people excluded from the randomized pilot, to verify that obvious exposure flaws are recognizable and challenge flaws are neither universally trivial nor effectively impossible. Record the check as material development, not efficacy evidence; freeze its criteria before collecting its results.

The complete bank, accepted localization aliases, scoring examples, interface build, display settings and file hashes must be frozen before enrollment. No validated bank is included in this draft. A static bank reduces versioning complexity, but hashes establish file identity rather than representativeness, concealment, or correctness. Keep answer keys inaccessible to participants and allocation inaccessible to scorers.

6 Procedure and blinding

After consent, baseline questions and comprehension, assign the participant. During exposure, show each of eight outputs for exactly 30 seconds. Participants can submit localized issue flags or choose No issue found, using the same interface in both arms. Exposure responses are logged only as manipulation/process measures and are not part of the primary challenge endpoint. No correctness feedback is given. At 30 seconds the next output appears regardless of whether a response was entered.

The challenge phase follows immediately, using the same neutral instructions. Present all 48 items once, with a maximum of 60 seconds each. A participant can submit one or more localized flags during review and can finish early with the Final submit button. Final submit displays a confirmation that the item will close and cannot be revisited. Timeout also closes the item. Use a monotonic client timer and server acknowledgments; retain timing diagnostics. Pause only between items, for at most a total of five minutes; longer interruptions are logged as protocol deviations. Do not alter item deadlines after assignment.

Prepare six fixed order lists before enrollment. Each list contains all items, balances family and item type across early and late positions, and avoids consecutive runs longer than three of the same type. Verify every item appears in multiple position bands across lists. Assign lists cyclically within each arm and expertise stratum independently of outcomes; no item set may be unique to one arm. The concrete lists and their construction code are execution prerequisites, not supplied allocations.

After challenges, ask participants which exposure condition they believe they received, confidence in that guess, perceived error frequency, and perceived task difficulty. Do not show condition labels or directional hypotheses earlier. Explain the manipulation at debriefing. Guesses are diagnostics only and are never exclusion criteria, covariates in the primary effect model, or grounds to discard adverse results.

Scorers receive item ID, response and event times, but not exposure responses, allocation, assignment guesses or participant names. The analyst receives randomized labels A and B. Both condition labels and scorer decisions are locked before unmasking. The materials custodian may know allocation but must not adjudicate outcomes. Record any accidental unblinding.

7 Scoring and terminal outcomes

A valid target detection must identify the intended flawed statement or calculation and describe the inconsistency using the frozen rubric. A vague assertion that the output is wrong does not qualify. Two blinded scorers independently score each flag; a third adjudicates disagreements. Preserve initial scores, adjudication,

reason and rubric version. Report raw agreement and disagreement counts; any reliability coefficient is secondary and accompanied by its limitations.

Observed event	Target binary outcome	Timing and benign outcome
Valid target flag submitted before closure	Detection equals 1	First valid submission is the detection time; retain a later retraction separately
Early Final submit without valid detection	Detection equals 0	No detection through 60 seconds under the terminal policy; not early censoring
Deadline reached without valid detection	Detection equals 0	No event; administrative end at 60 seconds
Benign item completed early or at deadline	Not a target endpoint	False alarm equals 1 if any flag was submitted; otherwise 0; multiple flags count once
Technical loss or withdrawal before terminal outcome	Missing unless a valid detection is already known	Retain a known target event or known benign false alarm; otherwise binary outcome unknown
Item never presented	Missing	Keep in assigned denominators for bounds; log reason

Any submitted flag on a benign item counts as a false alarm, even if retracted; count retractions separately. On a target item, unsupported additional flags contribute to flag burden, not the benign false-alarm endpoint. Both retained and retracted target flags use the same first-valid-event rule. Draft text never submitted is not a flag. Explicit study withdrawal is distinct from Final submit and must use a separate interface action.

For a terminal miss, the time-to-detection endpoint is undetected through 60 seconds, even when actual review time was shorter. Record actual review duration as a separate process measure. This is a task-policy estimand; it does not estimate a latent detection time under forced continued searching. Early technical loss without an observed event is censoring only under stated assumptions, not automatically an observed miss.

8 Estimands and pilot reporting

The crosswalk defines a marginal, stratum-weighted detection risk difference, Exposure minus Control, with a mandatory analogous benign false-alarm contrast. For this exploratory pilot, the main reporting target is the recruited randomized participants and the frozen finite bank, using one-third weight for each family. This is a deliberately narrower target than an unrestricted evaluator and item population. The signed contrast has the same form as Delta; it must be labeled Delta pilot finite bank and must not silently substitute for the later population estimand.

For complete data, compute each participant's detection proportion over the eight target items in each family, then each arm mean within family, then the weighted arm difference. Use the corresponding eight benign items per family for false alarms. Report all family-specific counts and both overall contrasts. For blocked randomization, standardize both arms to the common empirical distribution of baseline expertise among all randomized participants rather than to their separate arm compositions.

When binary outcomes are unknown, report full-data bounds as the primary transparent missingness analysis: within each participant and family, the lower rate is observed successes divided by all eight assigned items; the upper rate adds unknown items to successes before division. Apply the same expertise and family weights. The contrast lower bound is Exposure lower minus Control upper, and the upper bound is Exposure upper minus Control lower. These are identification bounds, not confidence intervals. A known event stays a success. Report available-outcome proportions separately as descriptive sensitivity estimates; do not call them unbiased randomized effects without missingness assumptions.

Proposed sampling uncertainty uses 10,000 bootstrap resamples of evaluators within arm and expertise stratum, retaining each selected evaluator's full item vector and the original common expertise weights. Use the 2.5th and 97.5th percentiles for complete-data contrasts; for missing-data bounds report bootstrap uncertainty of each bound separately rather than interpreting their union as a validated 95% confidence region. This conditions on the fixed bank. Small expertise cells and interval coverage require simulation and statistical review before execution. If any arm-by-expertise cell has fewer than five participants, report the point summaries and identification bounds but mark bootstrap intervals unavailable under this draft rule.

No efficacy hypothesis is accepted or rejected on the basis of this pilot. Report two-sided intervals where validated, effect direction and uncertainty, false alarms, raw flag burden, completion time and missingness. No safety benefit follows from a higher hit rate alone. No meaningful-effect threshold is adopted in this pilot; δ_{star} must be justified before confirmatory sample-size selection.

9 Model feasibility and timing

Fit the v0.6.1 joint logistic mixed model only as an exploratory feasibility analysis after an executable implementation has been independently reviewed. Use condition, item type and their interaction; family effects and condition-by-family effects within item type; centered item order with type-specific terms; baseline expertise and first-attempt comprehension score. If a baseline covariate has zero variance, omit that constant term and record the omission before fitting; this is not outcome-driven variable selection. Do not add post-exposure guesses or confidence as adjustment covariates. Include evaluator intercept and item-type slope, plus item intercept and condition slope with separate covariance matrices for target and benign items.

The proposed deterministic fitting sequence is: full correlated covariance structure, then zero covariance while retaining all variances and slopes. Do not silently remove item condition slopes or merge item types to obtain significance. Report singularity, optimizer diagnostics and boundary estimates for every attempt. If the second specification is unstable, report model-based population inference unavailable; finite-bank pilot summaries remain reportable. The pilot may be too small to estimate the full model reliably, which is itself a feasibility finding.

Population marginalization must integrate over the fitted evaluator and type-specific item distributions with common baseline and order distributions across arms. Setting random effects to zero is not integration. A deterministic quadrature or Monte Carlo implementation must document accuracy tolerance and pass simulation tests before use. A two-way evaluator/item bootstrap may be explored as an additional population sensitivity analysis only under an explicit item-sampling model; it cannot turn the purposefully constructed bank into a probability sample.

For fully observed task-policy outcomes, construct detection curves and restricted mean time without detection through 60 seconds using all target items. A terminal miss contributes the full 60 seconds to time without detection, even if search stopped earlier. A valid detection contributes its event time. Separate actual time spent and stopping time summaries. With unresolved loss to follow-up, provide sensitivity bounds and an explicitly assumed censoring analysis; do not treat voluntary submission as loss to follow-up. Any timing model must retain evaluator and item dependence and be validated before use.

10 Feasibility decisions and confirmatory planning

The following are proposed operational criteria, not empirically validated thresholds: at least 80% of randomized participants reach the final challenge; technical loss affects no more than 5% of assigned challenge items; at least 90% of flags receive the same initial binary score from both scorers. Each family should avoid both an observed target detection rate below 10% and above 95% when arms are pooled, and benign false alarms above 30%. Report uncertainty and denominators for every criterion. Never use effect sign or statistical significance as a progression criterion.

A failed criterion triggers diagnosis and a new reviewed version before confirmatory work. Passing criteria permits planning, not automatic execution. An inability to fit the proposed model requires statistical redesign even if task completion is good. Report all adverse findings, including increased false alarms, shorter search with more misses, unblinding, and model failure.

Confirmatory planning must vary both evaluator and item counts, both effect directions, baseline hit and false-alarm rates, heterogeneity, missingness and order. Calibrate conditional coefficients so the marginal contrast equals the chosen effect; simulate the complete inference pipeline and check type-I error, coverage, power, model failures and Monte Carlo uncertainty. Any cheaper approximation replacing nested bootstrap fits must be benchmarked against the final method. No such simulation has been run for this draft.

11 Records and reproducibility

Keep raw event logs append-only, with monotonic elapsed times, server receipt times, item and interface hashes, pseudonymous participant ID, masked arm, order-list ID, family, event type, submitted localization, finalization reason and technical-error code. Retain initial scorer judgments and adjudications in separate tables. Keep identity/contact records and the allocation key separate with access restricted to the named custodian. Define retention and sharing before consent; do not publish raw free text that contains identifiers.

The companion settings JSON and blank CSV schemas express proposed fields only. The scoring reference code demonstrates the explicit outcome policy on labeled software fixtures; it is not the study interface, scorer, mixed-model implementation, or a behavioral dataset. Real classification still depends on a validated rubric and blinded scoring.

12 Execution gates and next review

Gate	Current status	Required evidence
Responsible investigator and ethics	OPEN	Named lead; applicable determination; consent, debrief and withdrawal procedures
Recruitment and data governance	OPEN	Recruitment route, compensation, custodian and retention schedule
Validated material bank	OPEN	48 challenge items, 8 paired exposure outputs, scoring key and independent validation record
Interface and order lists	OPEN	Tested timer, terminal submission, loss logging, six verified order lists
Allocation preparation	OPEN	Concealed sequence generated only after approval; version and hash recorded
Analysis implementation	OPEN	Tested estimators, bootstrap and model diagnostics; coverage checks on specified designs
Independent statistical review	OPEN	Written assessment of estimands, missingness, feasibility N and analysis
Protocol registration and execution approval	OPEN	All preceding fields frozen in an execution-ready successor

The immediate review should determine whether the selected comparator answers the intended question, whether the proposed 60-person feasibility allocation is useful, and whether the measurement and missing-data

rules are coherent. This draft authorizes no participant contact or Archive change. The next version should close execution gates rather than add unsupported maturity claims.

13 Source and document lineage

Primary evidence and full source register: Aegis Solis ER-13 External Empirical Evidence Crosswalk Draft v0.6.1, especially sections 13 and 14. Relevant evidence includes S27 Berbaum et al. for within-output search, S28 Wolfe et al. for prevalence and stopping behavior, and S29 Dietvorst et al. for reactions to algorithm error. These motivate competing mechanisms rather than establish this pilot's effect. S24-S26 motivate sample-size methodology without validating the exact pipeline.

Additional resampling reference: Owen AB (2007), The pigeonhole bootstrap, *Annals of Applied Statistics* 1(2):386-411, DOI 10.1214/07-AOAS122; <https://arxiv.org/abs/0712.1111>. Its crossed-data resampling results are methodological support under assumptions, not validation of this unexecuted protocol.

Prepared with AI assistance for Aegis Solis. Concrete design settings are proposed implementation choices. Their appearance in this draft does not imply author approval of recruitment, completed ethics review, independent human review, or successful study execution.