

Final errata adjudication and exact changes

MathGov / RippleLogic v13.0

Exact build: **MG-RL-13.0-20260926-RELEASE-I**. Prepared 26 September 2026. Component editions are unchanged.

1. Disposition

The supplied audit recommends freezing the Core and identifies eight Low findings, not a new mathematical architecture. Each finding was checked against the actual supplied Release H bytes. All eight support narrow corrections or clarifications. The core scoring mathematics, qualification gates, SGP output separation, numerical inputs and worked verdict remain unchanged.

The input ZIP's CRC and original ledger passed. All fifteen separately attached Core files matched that ZIP byte-for-byte. Eleven Core files, including the Canon, SGP and Agent System, remain unchanged. Only the Foundations Primer, ripple.md, WDBIP and Aligners workbook have new bytes.

The audit's reported LibreOffice run is evidence reported by that reviewer for the H workbook hash. No reviewer endorsement is treated as an independent validation of all constructs, and no native I spreadsheet run is claimed.

2. Accepted corrections

F-01 — Workbook defined names

Disposition: ACCEPTED.

Change: Correct two `_xlnm.Print_Area` localSheetId attributes: 89 to 78 for `Build_Snapshot`; 90 to 79 for `Build_Integrity`. Keep the target ranges and all other names unchanged.

Why: The names pointed to sheets other than the ones named in their ranges. XML inspection reproduces the defect; the new regression checks the scope against the actual ordered sheet list.

F-02 — Build_Integrity!A2

Disposition: ACCEPTED, BOUNDED.

Change: Replace the broad local-guard assertion with the actual input/literal, selected-presence and RLS-reconciliation scope; explicitly exclude verdict-display replacement and require the external verifier.

Why: A disposable RLS!B40 overwrite still leaves B14 at zero by design. The display depends on the guard, so including it in the summed guard would introduce a circular reference. A separate actual-byte mutation is rejected by the external manifest verifier.

F-03 — ripple.md Appendix I1A and Appendix K threshold hierarchy

Disposition: ACCEPTED WITH NORMATIVE RECONCILIATION.

Change: Retain Cohen's $\kappa \geq 0.60$ (or the stricter preregistered threshold) as the existing L2/L3 reliability acceptance requirement; also report the Section 12 within-one-point agreement rate. Pin the estimator and weighting before results, and leave unestimable κ unverified.

Why: The reviewer identifies real ambiguity but its 92%-agreement example conflates agreement with the separate 90% reconstructability threshold. Percent agreement is not substituted for κ ; no new numerical threshold is introduced or empirical validity claimed.

F-04 — Foundations Primer Section 8, Selected definition

Disposition: ACCEPTED.

Change: Distinguish framework selection under the Canon conditions from separately recorded authority selection after a non-decisive result. State that a tie-break preference alone removes neither refusal nor the execution-authorization requirement.

Why: This aligns an informative teaching definition with the already controlling Canon; no new decision state is created.

F-05 — Foundations Primer Section 10, SGP introduction

Disposition: ACCEPTED.

Change: Replace “a construct-valid split” with “an explicit separation of constructs”.

Why: The original phrase suggested empirical construct validity that the Canon, SGP and validation protocol explicitly do not establish.

F-06 — WDBIP Section 18.2, Primary effect tokens

Disposition: ACCEPTED.

Change: Convert the malformed pipe-delimited paragraph into a native four-column table with the same header and eleven effect rows.

Why: The actual DOCX and rendered page reproduce the run-on text. All cell content is preserved; surrounding existing tables are untouched. One page break keeps Section 18.3 and its existing table readable.

F-07 — Audit_Flags!F5,F12,F13,F18,F27,F28,F33; A39,A46,D46; Dashboard!I5:I7

Disposition: ACCEPTED, LOCAL CLASSIFICATION ONLY.

Change: Mark the seven stored emissions TOOL_LOCAL in their existing source-reference cells; spell the full manual-override token in the existing taxonomy row; explain the scope; label dashboard counts as records by severity.

Why: The records are not new canonical tokens or executable detectors. Agent-owned RUNTIME_* labels remain separate from PCC severity. No emitted token, trigger value, severity, count formula or verdict changes.

F-18 — Sensitivity_Analysis!E13

Disposition: ACCEPTED.

Change: Rename the header to ΔRLS (raw). Keep E14=B14-C14 unchanged.

Why: This quantity is a score difference, not the normalized Gap compared with delta. The correct existing RLS and Dashboard labels remain untouched.

3. The two repairs that must not weaken existing safeguards

The local verdict-display gap is disclosed, not concealed

The external byte verifier rejects the actual RLS!B40 mutation, while the local guard remains zero because verdict displays are excluded from its summed checks. That exclusion is not repaired by inserting the display into the same guard: the display already depends on the guard. The corrected instruction tells readers exactly what can and cannot be inferred from the banner. It does not claim protection against an actor who changes both an artifact and the purported trusted manifest.

Reconstructability, agreement and kappa remain separate

The 90% requirement concerns reconstructability audit passage. It is not a 90% within-one-point inter-rater acceptance rule. Section 12's descriptive agreement rate and Appendix I1A's kappa requirement can coexist. The clarification preserves the already adopted kappa floor, requires the estimator and any weighting to be preregistered, and prohibits choosing an easier metric after seeing results.

Synthetic regression examples cover 92% within-one agreement with missing kappa (UNVERIFIED), kappa 0.59 (FAIL), kappa 0.60 (passes this requirement only), and kappa 0.70 against a stricter 0.75 preregistered threshold (FAIL). These tests do not establish a real wrapper audit, measurement validity, reconstructability, or RPAP/PFAP conformance.

4. Necessary synchronization, not another rewrite

Three modified Word documents receive updated identity, evidence and change-record paragraphs only in their terminal release appendices. The workbook's existing release-identity and evidence-pointer cells are updated. Sixteen matching literal-snapshot records track the seventeen direct text changes; the self-describing guard instruction is outside that literal snapshot.

The three amended documents' PDF/HTML/Markdown/text projections were regenerated. Five WDBIP static TOC page labels and its stored page count changed from 41 to 42. Canon and SGP contents, TOCs and PDFs remain byte-identical. A keep-with-next property on the wrapper's I3 heading prevents a new orphan heading. Existing document tables were not reconstructed.

Exact records distinguish 19 Word edit/layout operations, five static navigation label changes, document metadata, 33 workbook string-cell edits and two print-area scope attributes. No numerical cell, formula, cached formula result, sheet order, style or workbook geometry was changed. All 336 inherited document-table XML structures remain unchanged; the new WDBIP table is the only additional table.

Predecessor tests containing exact G/H-era prose or build-identity expectations are scoped to the actual source-build manifest. Their prior code is retained, and every checker change is logged. No mathematical test or failure route is removed merely to make a release pass.

5. Observations and proposals not adopted

F-08 — RETAIN SOURCE IDENTITIES. G/H/I source-build differences are explicitly recorded by the current manifest. Eleven correct masters are not rewritten solely to replace their historical preparation context.

F-09 — NO CORE EDIT. The quoted rc2 "Current" lines are inside marked historical, non-controlling blocks. They are not active release pins.

F-10 — NO NUMERICAL EDIT. The positive baseline case contributes zero catastrophe loss; different confidence display conventions therefore do not change this result. The existing instance-source context and baseline rules remain controlling.

F-11 — NO PARAMETER RETUNING. The declared epsilon remains 1e-6. Its effect and sensitivity are already explained in the Canon and the supplied reading guide. No replacement value or scale-relative rule is introduced.

F-12 — NO FABRICATED DEPENDENCE STUDY. The frozen run already refuses unique selection. No unperformed dependence analysis is relabeled as completed; the adverse-dependence rules remain unchanged.

F-13 — CLOSED FOR SUPPLIED SURFACES. The exact Office containers, complete body projections, bookmarks and 125 numbered TOC destinations are checked locally. This does not certify all Microsoft Word configurations.

F-14 — NO UNNECESSARY METADATA ADDITION. Absence of XLSX docProps is not a demonstrated calculation or integrity defect. Do not add components merely to resemble another application's output.

F-15 — RETAIN DISCLOSED TWO-LAYER INTEGRITY. Formula presence differs from exact formula identity. The trusted external manifest checks exact bytes, cells and formulas; the internal guard is not expanded into circular self-certification.

F-16 — NO NEW AGENT SEMANTICS. The mentioned local example label and view-tag omission did not reproduce a changed gate or verdict. Preserve the Agent's correct runtime and successor rules.

F-17 — OUTSIDE THIS RELEASE. Duplicate PC-AEP entries in the reviewer's own project store are not duplicates in the supplied Core_15 directory.

F-19 — NO INVENTED WEIGHT AUTHORIZATION. The worked-run's PLSS label is not proof of an external WeightProfileStatus. It remains a bounded, non-decisive exemplar; no authority record is fabricated.

6. Do-not-change record

Retain the governing RG → RF/NCRC → TRC → CSV → RLS sequence; the seven analytical scopes and welfare dimensions; three-channel non-compensatory rights review; TRC loss semantics and CVaR; normalized RLS and contribution decomposition; active-cell uncertainty, epsilon and every-contender robustness requirements; assessed-zero versus unknown handling; SGP MPS/FPP/GPR/SPR/ICP/RMCP separation; Agent authority and successor requirements; and all current component editions.

No new gate, numerical threshold, rights exception, scoring dimension, machine authority, empirical result or canonical registry was added. The wrapper minimum was clarified, not relaxed. The workbook remains a frozen example, not a general editable selector.

7. Evidence locations and publication boundary

Read Reports/Release_I_Verification.md and Verification/Final_I/. The exact Word and workbook registers supply before/after text and addresses. The source input and prior H archive remain in the separate Complete distribution's private provenance, not in the public publication root.

The source audit is **Final Claude Feedback v13.0.docx**, "MathGov/RippleLogic v13.0 Core 15 Independent Audit", 26 September 2026: finding register F-01-F-07 and F-18, with observations F-08-F-17 and F-19. Governing comparisons used the exact attached H masters, not older chat summaries or text-only render artifacts.

This build is prepared for public research/teaching reading and scoped reference tests. It does not establish empirical measurement validity, Microsoft Excel compatibility, complete external machine registries, every wrapper conformance pathway, physical safety, legal authority, Tier-4 ProofPack or production/domain authorization. Actual upload, hosted CI and hash-back verification remain separate publication actions.