

Journey to RippleLogic

v6.5 · Revision and feedback-adjudication report

1. Release decision and scope

Journey to RippleLogic: From Me to We — v6.5 is a completed, bounded revision of the supplied v6.4 reader edition. It is ready for author approval, reader distribution and onward professional copyediting. Commercial print release remains conditional on author/permissions clearance, a printer-specific full-wrap cover and a physical proof.

Prepared 28 September 2026. Cover author identity: James MathGov. The intentionally shared project name and the substantial contributions of GPT and Claude are preserved; collaboration is not represented as a sentience assessment, legal AI authorship or a transfer of publication responsibility.

What this pass did

All five supplied feedback reports were treated as hypotheses to adjudicate against the actual book and frozen technical sources, not as instructions to implement wholesale. The work revised the book, its illustrations and metadata, navigation, reflowable formats and publication notes. No technical core file was changed or reissued.

The completed book retains 13 chapters, six parts, nine appendices, five native tables and 30 embedded images. It has a 183-page 7 × 10-inch interior and a 184-page reader PDF with the unchanged front cover. The Markdown includes its media directory; the HTML embeds all images and CSS.

What the decision does not establish

This is an AI-assisted editorial and technical-alignment review. It is not an independent human referee panel, a recruited-reader study, a full empirical validation of RippleLogic, a proof of every proposition in the Canon, a copyright clearance, a physical print proof or a completed hosted publication. Checks reported here were performed for v6.5; the earlier v6.4 core-test results are not being claimed as new tests.

The 54-row Feedback_Adjudication register records 21 accepted clusters, 24 adapted clusters, seven rejected clusters and two deferred clusters. These are distinct review-claim groups, not a count of every repeated suggestion or a measure of editorial quality.

2. Source identity and corrections not adopted

The frozen pin is Release I, not Release G

The controlling baseline is RippleLogic Canon v13.0 with Sentience Gradient Protocol v8.8 in MG-RL-13.0-20260926-RELEASE-I. FINAL_FREEZE.md and VERSION_MANIFEST.json explicitly preserve unchanged earlier component masters. Both the Canon and SGP legitimately retain G-origin identities while participating in Release I. All 15 component files matched both the supplied archive bytes and the manifest hashes. The standalone Canon and SGP also matched those masters.

The source hierarchy is retained: the Canon Specification Contract governs normative precedence; the 14 Core_15 DOCX masters control their respective text, equations and tables within that hierarchy; the workbook controls only its frozen worked-run exemplar. Generated reading projections and this book are not overriding technical authorities. Frozen_Source_Identity.json records the exact hashes.

Wrong or unsupported reviewer inferences were rejected

The preparation date of 28 September 2026 was not changed to 27 September: it is the actual date of this revision. No publication event is inferred from a preparation date. A reviewer's limited upload set also does not prove the original complete v6.4 package lacked its DOCX, images, CSS, audit or worksheet.

The book is 7 × 10 inches, not the assumed 6 × 9 trim used in one legibility criticism. The comparison artwork already had four outcomes; its old image-title metadata was the defect. The pilot artwork already said Proposed. Those facts were checked on actual files rather than guessed from captions.

One small RLS score in a worked example does not show that real scores are usually small. Repeated references to a pay rise do not alone prove duplicate computation, although they can teach the rule poorly. Claims that the middle contained no personal voice overlooked the Chapter 7 confession and Chapter 8 wife scene. No new memoir was invented to solve that supposed absence.

The correct NIST AI 600-1 DOI was preserved. The dated Richardson et al. (2023) reference was not treated as a false citation. The existing truthful hardware description was not inflated merely to sound current. Later website or Canon changes do not silently override the exact meaning of this edition.

3. Technical repairs that earned their place

Rights, emergencies and definitions

Chapter 7 and Appendix I now distinguish the canonical emergency default from a universal moral ranking. Canon §7.5.1 requires the default result, plus the stated Tier 3/high-stakes alternative-order test when procedural or epistemic failure could enable the specified higher-order harms. A changed result is recorded as EMERGENCY_RIGHTS_ORDER_SENSITIVE and routed to lawful accountable authority with independent challenge and the rights vector disclosed. The text does not make the priority order arbitrarily optional.

The NCRC glossary now bars compensation of a supported rights-floor violation rather than claiming that harms cannot appear in an optimization input. Rights-relevant effects of surviving options remain visible. The stale Reproducibility Standard v1.4 reference was removed; the exact v1.7 pin remains in the back matter. Chapter 7 uses plain selection/pass language rather than ambiguous machine-like shorthand.

The No-Escape Question remains an empathy and attention test, not a formal NCRC predicate. Disliking the loss of a contract does not prove a rights violation. Two final readthrough repairs likewise distinguish community costs from protected-rights breaches and remove an unsupported retrospective floor-adjacent classification of the supplier memory.

Scoring, weight governance and risk

The text distinguishes informal teaching points from the exact optional displays in Canon §0.5 ($100 \times I$) and §10.1A ($100 \times RLS$). A shared numeric range does not validate a rough judgment or make an index a percentage of ethical correctness. The cell-weight explanation includes scope $\times$ dimension weighting and the qualifications concerning masks, multipliers, a common comparison basis and normalization.

Governance-lock-in risk is now named within TRC where relevant. Severe personal harms remain rights/CSV concerns even without a formal catastrophe trigger; a local-ruin extension must be declared and applied consistently rather than invented for a preferred option. Appendix D clarifies constitutional minimum weights, democratic and structural proposals, affected-party challenge and ex ante comparison declarations.

The ordered architecture, non-attenuated rights review, unknown-versus-zero distinction, effect identity, every-contender robustness, sole-survivor handling, non-decisiveness and separation of selection from execution remain intact. Initial stakeholder discovery was already present in v6.4 and was retained, not advertised as a newly discovered missing stage.

4. Narrative, teaching and the ending

Less machinery before the reader needs it

The opening interlude keeps the journey image and a plain-language compass. Other architecture diagrams now appear with the relevant method or institutional discussion, or in Appendix I. Chapter 2 no longer interrupts a colon with an audit disclaimer or answers Chapter 3's central question prematurely. Chapter 3 retains its fair treatment of intellectual predecessors while compressing a repeated closing movement.

Chapter 4 retains the rocket scene and the early horizon essay. Its full orbital worked run moves to Chapter 11, after the procedure has been taught, and the 7+2 map moves to Chapter 5. This obtains the sequencing gain without renumbering or relocating the whole chapter. Chapter 6 loses its inaccurate teaching-order note and a duplicate lengthy promotion treatment; comparison/decisiveness detail is concentrated around the Chapter 8 calculation.

Memoir preserved without invented reconciliation

The short Chapter 5 supplier account was removed because its upstream labor/environment assertions were not established by the detailed Chapter 6 account. It is replaced by a forward pointer. This does not assert that the two passages described one event or two; that remains unconfirmed. The detailed recollection and its observed/reported/inference/unknown note are textually unchanged. No medical causal link or new chronology is supplied.

Exact text comparisons confirm preservation of Chapter 1, including its exercise; the supplier recollection and evidence note; the Chapter 8 wife scene; and Appendix F. Chapter 9 also matches after its declared internal-heading and figure-number changes. Signature short lines and recurring motifs remain where they carry voice or reinforce safe use. No mechanical reduction quota was imposed.

One primary emotional return

The Epilogue now offers personal experience and hope instead of a general performance promise. Pronouns are consistent. The working-session maxim and excavation handoff move to Appendix H, leaving the lamp and next real decision as the main narrative's landing. The team contribution and communal invitation are preserved, not erased.

Appendix E is retained as a distinct five-minute teach-back rather than another full inventory. Appendix H is now Working Conversations and Shared Stewardship, with less duplicate aphorism and an explicit distinction between discovered interdependence and chosen ethical commitments. Appendix I remains a pinned interpretation aid, not a mutable online substitute.

5. Sentience, philosophy and source support

A more usable SGP chapter

Chapter 10 opens with ordinary protection questions and an animal/AI contrast, then introduces the six distinct outputs. It adds a categorical minimum-protection ladder and an explicitly hypothetical octopus-care example. Neither an animal class nor a persuasive AI conversation is assigned a numerical consciousness probability or an automatic MPS band.

MPS, FPP, GPR, SPR, ICP and RMCP remain separate. ICP is named correctly, with informative treated as its status. RMCP-P100 is qualified as a reference plateau, not supremacy or moral worth. Human FPP-100 remains independent of performance. Removing merely from the FPP resource-conversion sentence closes a misleading reading while retaining narrow governed exceptions. Registers of capability language and detailed evidence-profile material are kept in Appendix I.

The Biosphere discussion now expressly separates ecological life support from each affected animal's welfare. Mentioning animals in an ecological chapter is not intrinsically wrong; confining their moral consideration to that row would be. The new wording avoids that implication without inventing an eighth operational scope.

The strongest objections answered, not multiplied

Appendix D now explains why the eight rights are declared protection commitments with enabling-condition and anchor-based rationales, who must be accountable for an application, and how priorities can be challenged. It admits residual aggregation rather than claiming complete separation from consequential reasoning. The We contains each Me; wider scope does not create higher human worth or unrestricted authority.

No unverified personal history of deriving the rights from the UDHR was added. No numeric Self-floor slogan is offered as proof against collectivism. Appendix F's actual failure tests, burden test and openness to simpler alternatives remain unchanged.

External checks are bounded

The retained 2023 planetary-boundaries paper remains identified as dated. A primary PIK account dated 21 September 2026 supplies a newer assessment note, with its time/evidence limitations explicit. Butlin et al. (2023) and the New York Declaration (2024) provide additional consciousness-research context, not validation of SGP bands. The source ledger separates these fresh primary-source checks from the unchanged bibliography, which was not fully reverified source by source in this pass.

Copyright/disclosure suggestions were considered without claiming legal clearance. The existing rights-qualified open-license statement is retained. The first-person collaboration disclosure credits substantive AI input and states James's manuscript/publication responsibility without claiming that every sentence underwent a formal Canon run.

6. Format, image and navigation verification

The complete current reading set

The master DOCX is genuine OOXML and was reopened successfully. It contains 2,354 body paragraphs, five native tables and 30 image assets. There are no tracked-change elements or comment parts. All images have non-placeholder descriptive metadata. Four image assets changed at the pixel level: the public architecture footer, the journey diagram's crowded labels, the zero/unknown teaching footer, and the governance matrix's clipped subtitle/U6 label.

Twenty-four numbered figures were renumbered to match reading order; all six visual interludes are now listed separately. The five native tables remain native tables and have repeat-header settings. Seventy-one contents, figure, interlude and table entries have checked target pages and internal links. The PDF bookmarks and cover page labels are separate from the printed interior numbering.

Substantive parity and completeness

DOCX visible text and the HTML generated from the distributed Markdown match exactly after NFKC/whitespace normalization and removal of the print-only navigation page numbers. The normalized comparison contains 299,315 characters in each version. The pipeline preserves the prose, headings, table text and equations; this is not a claim of byte-identical formatting.

All 2,355 non-navigation text paragraphs, including 111 native-table paragraphs, were located in the final PDF under the documented coverage normalization. Two Greek-equation paragraphs required layout-aware Poppler extraction because the first extractor reordered font runs; both were recovered and visually checked. Coverage normalization is not punctuation identity or a proof of semantic correctness.

All 30 Markdown image paths resolve, and their hashes match the DOCX assets. The HTML includes 30 base64 images whose decoded hashes also match. It contains embedded CSS and no external stylesheet dependency; every internal link target resolves. Markdown must be distributed together with its media folder. The standalone HTML can be read without that folder.

Rendered review

All 183 interior pages were rendered and visually inspected at 150 dpi. The two final text-patch pages, 64 and 70, were re-rendered and checked; the remaining pages retained their text and pagination. Programmatic checks found no off-page text/images or text-image overlap. All inspected font resources are embedded. LibreOffice produced the PDF; full-page raster review used MuPDF after the parallel Poppler raster phase stalled, and Poppler also independently recovered the equation text.

This is a screen-based layout review, not a vendor-specific print proof, accessibility certification or guarantee of identical pagination in every version of Word. The original front-cover bytes and aspect ratio are preserved; they do not constitute a paperback spine/back-cover wrap.

7. Measured changes and regression boundaries

A restrained reduction, not an imposed percentage

The full book text, including native tables but excluding navigation entries and image metadata, changes from 44,848 to 43,978 Unicode word tokens: 870 fewer, or 1.94%. The main narrative paragraph text changes from 34,211 to 32,490 tokens, approximately 5.03% shorter. That latter comparison includes the effect of moving the former excavation coda from the main narrative to Appendix H; it is not all deletion.

Page count falls from the v6.4 interior's 187 to 183. The reader PDF adds the preserved front cover for 184 pages. Page-count reduction is not itself proof of better writing; the rationale for each material relocation or cut is recorded in the adjudication register.

Arithmetic and semantic regression

The Chapter 8 synthetic comparison retains $A = 0.075$, $B = 0.240$ and $C = 0.245$; $C - B = 0.005$. The cell-weight illustration gives $0.40 \times 0.30 = 0.12$ under its stated conditions. These are teaching arithmetic, not empirical welfare measurements, calibrated uncertainty or a warranted unique framework selection.

The emergency default remains LIFE, BODY, ECOL, LBTY, NEED, DIGN, PROC, INFO, with the required sensitivity condition now present. Rights-token non-dilution, three independent rights-violation channels, TRC/CSV separation, qualified-survivor ranking, active-mass normalization, incomplete-robustness refusal, every-rival comparison, sole-survivor handling and separate authority/execution remain the governing interpretive checks.

The two provisional weight profiles in Chapter 8 were not declared an arithmetic error: different clearly illustrative contexts may use different provisional priorities. The real repair was to explain the weight construction and prohibit outcome-driven changes, not to pretend every example must have identical preferences.

Limits of the technical verdict

The review found and repaired the identified book-level conflicts and ambiguities against the supplied sources. It does not establish complete external-schema conformance, native Excel recalculation, independently replayable ProofPack status, real-world safety or comparative decision effectiveness. No core test suite was rerun in this book-only pass. Core identity checks are intentionally distinguished from those stronger claims.

The edition should not undergo another broad rewrite simply to accumulate more changes. New corrections should address a specific remaining error, evidence gap or demonstrable reader problem. The protected passages, exact source pins and identified semantic boundaries are the regression baseline for any subsequent copyedit.

8. Handoff and remaining publication decisions

What is complete now

The full revised manuscript, reader PDF, interior PDF, Markdown with assets and self-contained HTML are delivered. The report, 54-item adjudication register, exact edit records, source-verification ledger, source identities, pagination map, validation results, changelog, author checklist, file manifest and checksums accompany them. The archive contains current book deliverables and their audit records, not obsolete drafts or a modified technical core.

Author confirmation, not another model rewrite

James should approve the autobiographical facts, quotations and privacy treatment that remain: supplier chronology and later reports; the partnership and rocket memories; the colleague's personal circumstances; family acknowledgments; and any identifiable third-party implications. The conflicting Chapter 5 précis has already been removed, so publication need not invent an answer about whether two supplier events were the same.

Confirm the intended copyright/licensing and publisher-specific AI disclosure on the actual publication channel. The team name James MathGov is preserved. No ISBN, registration, publisher acceptance, permission or hosted release event has been invented. The date in the book records preparation, not proof of commercial issuance.

Print and ebook production

The 7 × 10-inch interior is a reading/production proof. A paperback upload still needs the chosen printer's trim, bleed, gutter, paper, spine-width and color requirements, a matching full-wrap cover and a physical proof. The supplied 1024 × 1536 front image is not that wrap. Dense illustrated pages should be checked on the actual paper and print process.

No EPUB is included and no EPUB validation is claimed. The self-contained HTML provides a complete reflowable reading alternative, while a later EPUB conversion must preserve navigation, tables, diagrams, alt text and source boundaries. A new format conversion must not silently drop book material.

Recommended stopping point

Use the reader PDF for final author review and the DOCX for a bounded copyedit. A fresh human reader can test the lightened opening and the decision examples; their observations should be treated as new evidence, not guaranteed by this audit. Absent a concrete error, the next work is approval and publication production rather than another expansion of the book.

The result remains a human journey into a usable ethical method: the lamp is preserved, the formal machinery arrives later, and the book's claims stay within the evidence and authority they actually have.