

MathGov 3R / 1-2 Two-Phase Public Introduction

Integrity surface	Current requirement	Claim boundary
3R	Reality. Rights. Ripples.	A public doorway, not the complete technical specification.
1-2	Qualify through RG, RF/NCRC, TRC, and CSV; then rank selectable options through RLS.	Ranking cannot rescue prior failure.
Eligible shorthand	Eligible(a) is teaching shorthand for survival of the declared qualification sequence.	It is not a sixth gate, new authority predicate, or execution licence.

A simple new rule: bind claims to the real configuration

Do not say only, “the system is safe” or “the system can do this.” Ask:

- Which version, tools, permissions, controls, environment, and authority were tested?
- Is the action merely possible, or is it also qualified and authorized?
- Is the current state valid, is the transition admissible, and is the resulting state viable?
- What change would make the earlier approval expire?

MathGov therefore keeps alignment alive through requalification. When the material configuration or evidence changes, the affected checks run again.

AI systems: two connected governance questions

An AI system may occupy an instrumental tool-role while exhibiting substantial functional intelligence or operational agency. Tool-role tells us how a system is used; it does not establish what the system can do, whether it experiences anything, how it should be protected, or what legitimate authority it may hold.

MathGov treats AI through two connected tracks inside the existing five-level method. Capability–Consequence Governance asks what the exact configuration can do and cause, and what evidence, controls, permissions, accountability, and authority are required. Possible-Welfare Governance asks how the system should be treated if welfare-bearing states become credible or cannot safely be excluded. Capability is not sentience; possible sentience is not competence or authority; neither track creates a new gate.

This is the public doorway, not the full specification. Use it to understand the rhythm; use the Canon for conformance. Informative, non-normative public teaching companion. Canon controls if any conflict occurs.

Use the smallest profile that is safe

- **Quick / Tier 1:** a short record for low-stakes, reversible choices. A qualitative ruin-trigger screen is required; a material trigger requires escalation and governed TRC assessment.
- **Standard / Tier 2:** evidence, parameters, gates, controls, and uncertainty for material organizational or community choices.
- **Audit / Tier 3:** a full reproducible record, hashes, independent review, and replay for high-stakes or public decisions.

The controlling reading specification is

Core_15/MATHGOV_REPRODUCIBILITY_AND_USE_STANDARD_v1.7.docx. The framework stays the same at every level. Only the documentation burden changes. A user must escalate when rights, ruin, vulnerable groups, physical execution, consequential AI authority, or public validation claims become material.

The shortest version

Reality. Rights. Ripples.

1. Qualify. Reality. Rights. Ruin. Viability.

2. Rank. Ripple-score the viable options.

That is the operating rhythm.

First make sure the option belongs in the ring. Then find the one that produces the strongest conditional residual ripple among selectable options, if decisive.

The full formal cascade

RG -> RF -> TRC -> CSV -> RLS

Reality Grounding -> Rights Floor -> Tail-Risk Constraint -> Containment and Structural Viability -> RippleLogic Scoring

Teaching note: this doorway compresses the architecture into three pillars and two phases. The formal architecture has five levels: one grounding precondition, three qualification gates, and one ranking layer.

Why this matters

MathGov does not begin by scoring every imagined option. It first asks which options have enough reality, rights protection, ruin control, and viability to enter the decision field. Only then does RippleLogic score the survivors.

The MathGov answer is: ground the claim, protect the right, bound the ruin, verify the system, make exit and correction real, and rank only what survives qualification.

The three pillars

Pillar	Plain meaning	Formal location
Reality	Do we have contact with what is real, possible, grounded, testable, bounded, auditable, correctable, and viable?	RG opens reality contact. CSV tests operational reality and structural viability. PC-AEP is used when physical or causal execution claims matter.
Rights	Are protected beings, dignity, consent, thresholds, and non-collapse safeguards protected?	RF/NCRC protects rights directly. TRC protects against ruin pathways that can destroy the conditions for rights and flourishing.
Ripples	Among the options that survive qualification, which creates the best net flourishing?	RLS ranks only the survivors across affected unions and welfare dimensions.

rights and ruin hardening

Rights hardening: RF/NCRC does not discount a categorical or severe rights violation merely because it is brief or evidence confidence is low. Rights severity, categorical prohibitions, and severe-hazard exposure are checked separately from ordinary welfare timing.

Tail Emergency Mode: if every rights-safe option exceeds the ruin corridor, the default is redesign, delay, escalation, no-action review, or refusal. A temporary least-CVaR path is allowed only under a governed emergency declaration, absolute exposure cap, independent challenge, time bounds, controls, and exit criteria. It is not ordinary selection, and RLS is not used to rescue it.

Phase 1: Qualify the options

Phase 1 asks: which options are even eligible to be ranked?

Gate	Question	If the answer fails
RG - Reality Grounding	Is the option grounded in evidence, real constraints, affected unions, and a clear claim boundary? Can high-consequence claims be tested or warranted, bounded, audited, corrected, and refused?	Do not score fantasy. Narrow, investigate, mark exploratory, escalate, or refuse the stronger claim.
RF - Rights Floor	Does the option protect rights, dignity, consent, and non-negotiable safeguards?	Do not average rights away. Redesign, escalate, or refuse.
TRC - Tail-Risk Constraint	Is catastrophic, irreversible, fat-tail, or system-breaking risk bounded?	Do not gamble ruin for ordinary benefit. Bound, redesign, escalate, or refuse.
CSV - Containment and Structural Viability	Can the option actually stand, execute, be controlled, monitored, exited, corrected, and remain governable?	Do not select what cannot stand. Add controls, narrow, redesign, or refuse.

TRC is the ruin veto. CSV is the viability and control test. They are linked, but they should not be collapsed.

Reality Grounding hardening

For life-critical systems, weapons, nuclear command and control, planetary-scale cognition, critical infrastructure, medical systems, ecological interventions, or other high-consequence domains, no claim should govern action unless it can be tested or otherwise warranted, bounded to a validity domain, audited with traceable evidence, corrected through monitoring and revision, and refused when the evidence surface is insufficient.

In public language: ground the claim; protect the right; bound the ruin; verify the system; make exit and correction real; rank only what survives qualification.

Generate before ranking. Before accepting an apparent trade-off, look for no-action or continuation, delay, staged or reversible trials, redesign, constitutive mitigation, cleaner or safer alternatives, and hybrid or Pareto-improving options. MathGov does not assume every conflict can be made win-win. Where irreducible conflict remains, rights, ruin, distributional burdens, legitimacy, and residual harms remain explicit.

Phase 2: Rank the survivors

Phase 2 asks: among the options that survived qualification, which one creates the strongest conditional residual ripple among selectable options, if decisive?

RLS evaluates affected unions, welfare dimensions, evidence, uncertainty, robustness, and decision state.

Mathematical spine

Wrong method:

Decision = $\text{argmax RLS}(a)$ over all imagined options

Correct method:

$\text{Eligible}(a) = \text{RG_qualified}(a) \text{ AND } \text{RF/NCRC}(a) \text{ AND } \text{TRC_qualified}(a) \text{ AND } \text{CSV_qualified}(a)$

Here RG_qualified means RG_SUPPORTED or RG_NARROWED; TRC_qualified means TRC_PASS or a documented TRC_NOT_TRIGGERED assessment. Eligible remains teaching shorthand, not a machine-state token, sixth gate, or execution licence. CSV_qualified means CSV_PASS, CSV_PASS_WITH_CONTROLS, or a documented CSV_NOT_MATERIAL result valid for the declared claim boundary; it does not include CSV_REDESIGN_REQUIRED, CSV_ESCALATE, CSV_FAIL, or CSV_EMERGENCY_PROVISIONAL.

Define the selectable set:

$$A_sel = \{a : RG_qualified(a) \text{ AND } RF/NCRC(a) \text{ AND } TRC_qualified(a) \text{ AND } CSV_qualified(a)\}$$

Then apply the decision-state rule:

- If A_sel is empty, redesign, delay, escalate, use the governed emergency pathway where applicable, or refuse the stronger claim.
- If RLS is defined and the same leader passes the Canon's signed discrimination test against every other selectable option under every required robustness variant, the framework may return **ALLOW_FRAMEWORK_SELECTION**. A sole selectable option follows the Canon's sole-survivor rule without a fabricated pairwise test.
- If RLS is undefined, tied, close, uncertainty-overlapped, or non-decisive, return **REFUSE_DETERMINISTIC_SELECTION**. Seek more evidence, preserve multiple options, redesign, or use a separately disclosed **AuthoritySelectionRecord** within a lawful mandate.

A point-score argmax is therefore a comparison aid inside the selectable set. It is not automatic authorization, and it does not erase uncertainty, ties, refusal, or the separation between framework selection and authority selection.

No RLS rescue

A high RippleLogic score cannot rescue failed reality grounding, a rights violation, an unacceptable ruin path, an unsupported physical-execution claim, or a structurally non-viable option.

Decision outcomes

Public labels such as SELECT are explanatory shorthand only. Canonical emitted states remain **ALLOW_FRAMEWORK_SELECTION**, **REFUSE_DETERMINISTIC_SELECTION**, **BLOCK/FAIL** states, **ESCALATE**, and separately recorded **AuthoritySelectionRecord** outcomes; no public label creates execution authority.

Outcome	Meaning
SELECT	Public teaching shorthand: the option is eligible and clearly best among survivors. It does not replace canonical machine-state or authorization tokens.
SELECT WITH CONTROLS	Public teaching shorthand: the option is eligible only with recorded safeguards, monitoring, or bounds. It is not execution approval.
NARROW	The claim, scope, domain, or time horizon must be reduced.
REDESIGN	The option fails but may be improved.
ESCALATE	Authority, evidence, expertise, or stakeholder participation is insufficient.
REFUSE	The requested claim, selection, or action is not supportable under the current evidence, authority, or gate state. Record the specific reason.
NON-DECISIVE	The evidence or scores do not justify deterministic selection.

Public script

MathGov is a decision architecture, not a claim that every option should be scored. It first qualifies options through reality, rights, ruin, and viability. Then it ranks only the surviving options by their ripple effects across affected unions. In short: qualify before you rank.

Reading the Matrix and Its Score

Union Scope tells us where and at what scale a consequence lands. Welfare Dimension tells us what kind of condition changes. One distinct effect has one primary accounting home; genuine knock-on effects are recorded separately with evidence. The dimensions are distinct lenses connected by real causal relationships, not proven independent compartments.

The seven scope rows cover Self, Family/Household, Community, Organization, Polity, Humanity / Global Coordination (including the CMIU coordination view), and Biosphere. These are overlapping analytical scopes, not a rigid hierarchy of containers or a closed list of who matters.

RLS is a weighted, normalized summary of the residual impact field after qualification. It can be decomposed by cell, scope or dimension so the overall number does not conceal its sources. Unknown is not zero; a higher score is not execution permission; a displayed index is not a percentage of moral truth or safety. Read the impact and contribution profiles, affected groups, uncertainty and dependence, and the decisive or non-decisive comparison before interpreting the scalar summary. A larger Gap is not statistical confidence or execution authority.

The 60-second checklist

1. Qualify

- Is it real?
- Are rights protected?
- Is ruin bounded?
- Can it stand, be monitored, be corrected, and be exited?

2. Rank

- Who and what is affected?
- Which option creates the strongest conditional residual ripple among selectable options, if decisive among the survivors?
- Is the result robust enough to select, or should it be non-decisive?

Claim boundary

MathGov Core v13.0 is a proposed Tier 1-3 auditable decision architecture. It is not empirical validation, legal certification, deployment certification, Tier 4 ProofPack, automated moral truth, or physical safety proof. For consequence-bearing physical systems, MathGov requires external domain evidence through PC-AEP; it does not manufacture that evidence by itself.

What This Edition Adds

v13.0 keeps the same doorway: Reality first, Rights next, then Ripples — with catastrophic risk and structural viability protecting the path before ranking.

It strengthens one simple rule: the same real harm or benefit should not become larger or smaller just because someone describes it in more boxes, more time slices, a friendlier reference class, or a weaker evidence frame.

The score is a conditional summary, not moral truth. When the result changes under a materially legitimate alternative, RippleLogic shows the disagreement and refuses to pretend there is one uniquely proven answer.

Rights, ruin and structural viability still cannot be traded away. Selection still is not authority, and authority still is not execution permission.

Final pocket form

Reality. Rights. Ripples.
Qualify before you rank.
Rank only what survives.

SGP v8.8: protection, participation, and power are separate

When a decision may affect a welfare-bearing being, the current SGP interface does not ask one number to answer every question.

- **MPS** asks how strongly evidence supports welfare-bearing experience.
- **FPP** records full non-downgradable rights-of-protection.
- **GPR** asks how the entity can participate directly, with support, or through representation.
- **SPR** asks whether it may hold one specifically named consequential role under strict gates.
- **ICP** describes intelligence and understanding, but cannot establish sentience or authority by itself.
- RMCP profiles integrated reality-management capacity; RMCP-P100 is its open reference plateau, not a moral rank. Every human has FPP-100, while humanity collectively is the provisional current RMCP-P100 calibration anchor. Other biological, digital, hybrid, collective, or extraterrestrial intelligences may meet or exceed the same reference-plateau criteria under admissible evidence.

RMCP/P100 is not a moral-worth or authority score. The operating rule is: **unknown is not zero; capability is not welfare; protection is not authority; full protection does not require human likeness**. MPS intervals may inform residual welfare sensitivity only after rights, ruin, and viability qualification.

SGP interface clarification. MPS is not a cardinal welfare multiplier or moral-worth score. It determines protection posture and which welfare-inclusion hypotheses must be tested. Missing evidence is MPS-NE, not zero.

APPENDIX RELEASE: Identity, Source Authority and Revision Record

Framework release: MathGov/RippleLogic v13.0. Component: MathGov Public Introduction v13.0. Edition-origin preparation: 10 September 2026; the separately identified correction build is dated below. Two-part major.minor component versions are used; preserved historic identifiers are not renumbered.

Status: integrated research and teaching specification with a bounded worked-example and scoped reference implementation. Readiness is limited to the checks in Reports/Release_H_Verification.md. No empirical validation, independent human validation, full production runtime, physical-safety certification, legal authority, deployment authorization or Tier-4 ProofPack status is implied by the edition number.

Delivery identity: MathGov/RippleLogic v13.0, build MG-RL-13.0-20260926-RELEASE-H. Prepared 26 September 2026 from Release G. Only the recorded reading and release-reference corrections were made; component editions are unchanged. The manifest distinguishes corrected components from byte-identical inherited masters by their own source-build IDs and exact hashes. This is a new correction build, not a silent replacement of the earlier artifact. Live publication is not asserted.

Component/build identity. Aligners Sheet v5.9 retains its component edition but carries this correction build's identifier. Its numerical inputs, formulas and worked verdict are unchanged. Release G label, runtime-token and integrity-snapshot corrections remain preserved; this build corrects four principal string cells and their four active snapshot counterparts. See Verification/Final_H/Exact_Workbook_Changes.json. Unmodified Core masters retain their Release G source-build identity.

Verification boundary: Reports/Release_H_Verification.md and Verification/Final_H/ identify the current checks and limits. Earlier G receipts and the rejected compatibility experiment remain source-build evidence, not native-engine certification of changed H bytes. Cache-independent replay and native-application evidence are distinct. Microsoft Excel parity, empirical validation, full external-registry conformance and production authorization are not asserted.

Package source authority: Core_15 contains 14 DOCX specification masters and one XLSX Aligners frozen worked-run workbook. DOCX files control prose, equations and tables subject to Canon and SGP ownership. The workbook controls only its disclosed exemplar surfaces. Sources, Reading_HTML and Reading_PDFs are generated reading projections, not competing normative masters. The manifest and hashes identify exact bytes; supplemental code and schemas govern only their documented subset. Conflicts require recorded correction against the controlling source.

License and reuse: consult the package-level LICENSE and NOTICE and retain applicable component-specific and third-party notices. This pointer does not override a valid exception or grant rights over separately supplied private or third-party review material.

External repository publication and any absent legacy schema/validator/registry are separate authorities/evidence surfaces. The local package does not claim to update a remote repository or reproduce unavailable implementations. Current source pointers in this appendix replace prior front-matter edition pointers for this package; accurate historical source references below are retained as lineage only.

Preserved invariants: RG -> RF/NCRC -> TRC -> CSV -> RLS; rights non-compensation and unallocated rights effects; existing seven rows and seven dimensions; existing SGP MPS/FPP/GPR/SPR/ICP/RMCP separation; probability ownership, uncertainty, stability and every-contender rules; selection, authority and execution separation.

Component	Current edition
RippleLogic Canon	v13.0

Component	Current edition
Sentence Gradient Protocol	v8.8
ripple.md Standard	v5.8
RippleLogic Agent System	v13.0
CSV Gate Standard	v2.7
RippleLogic Cascade Standard	v2.9
MathGov Reproducibility and Use Standard	v1.7
Welfare Dimension Boundary and Interaction Protocol	v1.9
RLS Validation Protocol	v2.9
RippleLogic Foundations Primer	v4.7
MathGov Public Introduction	v13.0
Physical/Causal Admissibility Evidence Profile	v2.6
Methodological Falsifiability and Dependency Integrity Standard	v2.6
Source-Coupling Integrity Standard	v2.6
RippleLogic Aligners Sheet	v5.9

Change locations and rationales: Reports/Release_H_Adjudication.md; exact edits: Verification/Final_H/Exact_Document_Changes.json and Verification/Final_H/Exact_Workbook_Changes.json. Previous change records remain source-build history. Prior releases and feedback are preserved in the private provenance archive.

Current proposals are not universal truths. Taxonomy maximality, continuous-time propagation, cross-substrate cardinal welfare, generic susceptibility/shield formulas, autonomous recovery, vendor infrastructure, zero-knowledge circuits and hardware meshes remain unvalidated unless independently demonstrated under a scoped implementation profile.

Preserved Baseline Release Material (Historical; Non-Controlling)

The following blocks are relocated intact from the recovered baseline. Their edition numbers, release-readiness wording and external-source limitations describe that historical candidate, not current component identity or newly executed verification. Governing current metadata is the matrix above.

The historical, unverified implementation reference is docs/reproducibility/MATHGOV_REPRODUCIBILITY_AND_USE_STANDARD_v1.5.md; its candidate source parity remains to be established. The controlling reading specification is MATHGOV_REPRODUCIBILITY_AND_USE_STANDARD_v1.7.docx in Core_15; the old path is not a claim of current schema or implementation compatibility.

v12.8-rc2 Historical Release Integration

HISTORICAL (NON-CONTROLLING): Release: MathGov Core Release 2026.09 — RippleLogic Canon v12.8-rc2 / SGP v8.7-rc2 — Rights Non-Dilution, Representation Integrity, and Robust Selection Release

Historical companion pins: Canon v12.8-rc2; SGP v8.7-rc2; ripple.md v5.7-rc2; Agent System v12.7-rc2; CSV v2.6-rc2; Cascade v2.8-rc2; Reproducibility v1.6-rc2; WDBIP v1.8-rc2; RLS Validation v2.8-rc2; Primer v4.6-rc2; Public Introduction v12.8-rc2; PC-AEP/MFDI/Source-Coupling v2.5-rc2; Aligners Sheet v5.8-rc1.

Release control	Historical value
Component	MathGov 3R / 1–2 Public Introduction v12.8-rc2
Release	HISTORICAL (NON-CONTROLLING): MathGov Core Release 2026.09 — RippleLogic Canon v12.8-rc2 / SGP v8.7-rc2
Architecture	RG → RF/NCRC → TRC → CSV → RLS
Role	Shortest informative public doorway to the current RippleLogic Core.
Claim boundary	Release candidate for Tier 1–3 research; source/mirror parity and implementation conformance remain separately reviewable; not empirical validation, legal authority, physical-safety certification, deployment authorization, Tier 4, or moral truth.
Source/render parity	Versioned semantic source and DOCX/PDF mirrors must agree. Filename, internal version, active pins, manifest, and hashes must agree; mismatch is release-integrity failure.

Historical Release Integration (v12.8-rc2)

HISTORICAL (NON-CONTROLLING): This candidate advances MathGov 3R / 1–2 Public Introduction from the audited v12.7 baseline to v12.8-rc2 and binds it to MathGov Core Release 2026.09 — RippleLogic Canon v12.8-rc2 / SGP v8.7-rc2. Historical references remain lineage only; unmarked current-release references use the current component matrix in Appendix RELEASE.

v12.6 Release Note (Lineage)

v12.6 adds configuration-bound assurance and qualification-continuity teaching language while preserving the 3R and MathGov 1-2 structure. The 3R and MathGov 1-2 teaching structure is unchanged.