

RIPPLE\_LOGIC FRAMEWORK v7.4.5

## **A Universal Ethical Operating System for Multi-Scale Alignment**

**Tier 1-3: Implementable Now | Tier 4: Design Target (ProofPack Pending Public Release)**

**Author:** James McGaughran (ORCID: 0009-0005-3324-7290)

**Affiliation:** British University Vietnam (BUV); MathGov Institute for Ethical Systems Design

**Canonical Sites:** ripplelogic.org | mathgov.org

**Date:** 25 January 2026 (Asia/Vietnam)

Version: 7.4.5 (ProofPack Readiness Patch)

---

### NAMING LINEAGE STATEMENT (v7.4.5)

Ripple\_Logic Framework is the canonical name of the union-based ethical operating system historically published under the name MathGov (v1.0 through v5.0i). All prior MathGov documentation remains valid as lineage material and SHALL be interpreted as Ripple\_Logic for continuity, preserving methodological and mathematical intent. Ripple\_Logic v7.4.5 consolidates and spec-hardens Tier 1-3 into a single fork-resistant normative text and preserves Tier-4 replayability only as a design target pending public ProofPack release.

This version (v7.4.5) is a ProofPack readiness patch incorporating feedback from multiple AI reviewers. It resolves formatting artifacts, reference corruptions, and specification ambiguities identified in v7.4.0 through v7.4.4 drafts. Normative changes in v7.4.5 are limited to ProofPack meta-normative clarifications only: (i) normative hierarchy completion and formatting standardization (Section 0.1), and (ii) SGP dependency pinning rule for ProofPack replayability (Appendix G). No changes have been made to cascade logic, admissibility gates, thresholds, tier requirements, or scoring formulas.

---

### **AUTHOR STATEMENT AND TRANSPARENCY DISCLOSURE**

The author conceived, designed, and wrote this specification. Generative AI tools (including OpenAI ChatGPT and Anthropic Claude) were used as drafting and consistency assistants. The author reviewed, verified, and edited all outputs and assumes full responsibility for the content, claims, and citations.

This is a theoretical framework paper. Empirical validation through pilots is planned but not yet completed. No operational deployment claims are made without completed validation studies.

---

## **ABSTRACT**

Governance systems and AI alignment approaches repeatedly fail in three coupled ways: (i) they collapse plural values into single metrics that permit trading away fundamental rights, (ii) they underweight catastrophic tail risks through expected-value reasoning, and (iii) they remain vulnerable to specification gaming, where optimized proxies degrade intended outcomes.

This paper presents the Ripple\_Logic Framework (formerly MathGov), a universal ethical operating system grounded in Union-Based Reality (UBR): the descriptive stance that interconnection and nested unions, rather than isolated agents, provide the correct structural grammar for complex systems. From this ontology, Union-Based Ethics (UBE) is operationalized as a strict five-level lexicographic cascade applied to a 49-cell welfare matrix (seven unions by seven welfare dimensions). The cascade proceeds as follows: (1) the Non-Compensatory Rights Constraint (NCRC), which excludes rights-violating options except under explicitly declared emergency procedures; (2) the Tail-Risk Constraint (TRC), which excludes options with unacceptable catastrophic exposure using Conditional Value-at-Risk (CVaR); (3) the Containment Gate, which prevents local optimization from degrading the coherence and viability of containing unions beyond governed tolerances; (4) the Ripple Logic Score (RLS), which ranks remaining selectable options by weighted welfare impacts after ripple propagation; and (5) structural tie-breaks using the Union Coherence Index (UCI) and the Hollowing-Out Index (HOI) when RLS differences fall within an uncertainty discrimination band.

Ripple\_Logic further specifies: (i) a sparse, interpretable ripple kernel for propagation under epistemic humility with a Kernel Quality Score policy; (ii) Hybrid Democratic Weighting (HDW) combining constitutional floors with democratic tuning; (iii) an auditable Provenance and Compliance Certificate (PCC) embedded in a Notice-Choose-Act-Reflect (NCAR) learning loop; and (iv) a binding interface to the Sentience Gradient Protocol (SGP) for rights-of-protection gating with strict separation from governance authority.

Ripple\_Logic v7.4.5 is Tier 1-3 implementable using this paper and its appendices. Tier 4 is specified as a design target and MUST NOT be claimed until a public ProofPack is independently replayable. The framework is presented as a computable, corrigible, audit-

ready decision system for individuals, organizations, governments, and AI systems, with an explicit validation and falsification program.

---

## **KEYWORDS**

Ripple\_Logic; Union-Based Reality; Union-Based Ethics; AI alignment; lexicographic ethics; rights constraints; catastrophic risk; tail risk; CVaR; multi-scale governance; welfare matrix; ripple propagation; kernel; containment; Hybrid Democratic Weighting; Union Coherence Index; Hollowing-Out Index; Provenance and Compliance Certificate; auditability; NCAR; Sentience Gradient Protocol.

---

## **TABLE OF CONTENTS**

- Section 0: Specification Contract (Normative)
- Section 1: Introduction: The Need for a Universal Ethical Operating System
- Section 2: Ontological Foundations: Union-Based Reality (UBR)
- Section 3: Normative Foundations: Union-Based Ethics (UBE)
- Section 4: System Overview: The Ripple\_Logic Architecture
- Section 5: Welfare Impact Construction (Calculable)
- Section 6: Ripple Propagation: Kernel and Epistemic Humility
- Section 7: Rights Layer: Non-Compensatory Rights Constraint (NCRC)
- Section 8: Catastrophic Risk Layer: Tail-Risk Constraint (TRC)
- Section 9: Containment Layer: Preventing Pathological Local Optimization
- Section 10: Optimization Layer: Ripple Logic Score (RLS)
- Section 11: Structural Safeguards: UCI and HOI
- Section 12: Moral Status Layer: SGP Integration (Binding Interface)
- Section 13: Weight Governance: Hybrid Democratic Weighting (HDW)
- Section 14: Auditability and Determinism: PCC and AIL (Tier 1-3; Tier 4 Target)
- Section 15: Implementation Guidance and Tiers

- Section 16: Relationship to Existing Frameworks
  - Section 17: Validation, Falsification, and Research Program
  - Section 18: Applications and Use Cases
  - Section 19: Limitations and Non-Targets
  - Section 20: Conclusion: Prosperous Futures for Intelligences
  - References (APA 7)
  - Appendices A through M
- 

## **SECTION 0: SPECIFICATION CONTRACT (NORMATIVE)**

### **0.1 Normative Hierarchy (Single Source of Truth)**

When interpreting Ripple\_Logic v7.4.5, conflicts **MUST** be resolved using this precedence order (highest controls lowest):

1. - Main text rules explicitly labeled **MUST**, **SHALL**, or **PROHIBITED**
  2. - Appendix B (Canonical Equations)
  3. - Appendix C (Rights/NCRC Canon Pack)
  4. - Appendix D (TRC and Scenario Governance Canon Pack)
  5. - Appendix H (PCC and Audit Flags Canon Pack)
  6. - Appendix E (UCI Interface and Tier-3 Construction Guidance)
- Appendix G (SGP Integration Binding Interface)
  - Appendix A (Symbols, Domains, and Notation Canon)
  - Appendix F (Failure Modes and Anti-Gaming Controls)
  - Appendix J (Pre-Registration Templates and Scoring Definitions)
  - Appendix K (Starter Kernel Canon Pack)
  - Examples, templates, and illustrative values (non-normative)

### **0.2 Tier Contract and Claims (v7.4.5)**

Tier 1-3: Executable and claimable in v7.4.5, subject to stated minimum requirements. Organizations and individuals may claim compliance with Tier 1, 2, or 3 using only this Foundation Paper and its appendices.

Tier 4: PROHIBITED to claim in v7.4.5 (design target only). Any Tier-4 content in v7.4.5 is informational architecture describing what becomes claimable only after ProofPack publication and independent replayability.

### 0.3 Conformance Vocabulary (Normative)

Ripple\_Logic uses distinct predicates; implementations MUST NOT collapse them:

- **Rights-admissible:** passes NCRC.
- **Admissible:** passes NCRC and TRC. Formally:  $\text{Admissible}(a) := \text{NCRC}(a) \text{ AND } \text{TRC}(a)$
- **Selectable:** admissible and passes Containment (Mode A).
- **Selected:** the final chosen option from the selectable set using RLS and tie-break policy.

### 0.4 Consolidated Symbols and Core Objects (Normative)

#### Sets and Indices

Symbol	Meaning
$u \in \{1, \dots, 7\}$	Union index
$d \in \{1, \dots, 7\}$	Dimension index
$a \in O$	Candidate option
$s \in S$	Scenario index
$r \in R$	Right index
$k$	Impact instance index

Symbol	Meaning
$g \in G_{\{u,d\}}$	Protected subgroup index (for rights checking)

### Core Sets

Set	Definition
$U = \{U_1, \dots, U_7\}$	Operational unions
$D = \{D_1, \dots, D_7\}$	Welfare dimensions
R	Rights set: {LIFE, BODY, LBTY, NEED, DIGN, PROC, INFO, ECOL}
$C_r \subseteq U \times D$	Rights coverage sets
$C_{cat} \subseteq U \times D$	Catastrophe cell set for TRC

### Impact Objects (all options a)

Symbol	Meaning	Range
$I_{dir}(u,d,a)$	Direct impact (post-saturation)	[-1,+1]
$I_{prop}(u,d,a)$	Propagated impact (post-saturation)	[-1,+1]
$I_{rights}(u,d,a)$	Rights worst-off subgroup impact	[-1,+1]

### Kernel

Symbol	Meaning
$K \in \mathbb{R}^{\{49 \times 49\}}$	Ripple kernel matrix (sparse)
Convention	$K_{\{ij\}}$ maps source j to target i (target-row, source-column)

## Weights

Symbol	Meaning	Constraint
$w_u$	Union weights	$w_u \geq 0, \sum_u w_u = 1$
$v_d$	Dimension weights	$v_d \geq 0, \sum_d v_d = 1$
$m(u,d)$	Applicability mask	$\in \{0,1\}$ (RLS aggregation only)

## Tail Risk

Symbol	Meaning
$p_s$	Scenario probabilities ( $p_s \geq 0, \sum_s p_s = 1$ )
$\omega_c$	Catastrophe weights ( $\omega_c \geq 0, \sum_{c \in C_{cat}} \omega_c = 1$ )
$\alpha$	CVaR tail level $\in (0,1)$
$\tau_{TRC}$	Corridor threshold $\in (0,1]$

## Structural Metrics

Symbol	Meaning	Range
$UCI_u$	Union Coherence Index for union $u$	$[0,1]$
$\Delta UCI_u(a)$	$UCI_u(a) - UCI_u(\text{baseline})$	unbounded
HOI	Hollowing-Out Index (monitoring diagnostic)	unbounded

### 0.5 Version Commitments (Normative)

The canonical welfare impact scale is  $[-1,+1]$  with Baseline-Zero Rule (Section 5).

**Optional UI scale conversion (non-normative):** If a user interface uses a percent-like scale, map the canonical cell impact  $I(u,d,a) \in [-1,+1]$  to UI points via:  $\text{points} = 100 \times I$ . All admissibility gates (NCRC, TRC, Containment) and RLS calculations MUST be performed on the canonical  $[-1,+1]$  scale, or be converted back exactly before computation.

Tier 1-3 TRC uses bounded-impact loss (Section 8).

Tier 4 is design target only; no Tier-4 claims permitted until ProofPack is publicly replayable.

**UCI Equity for Self ( $U_1$ ):**  $E_1 := 1$  by definition. The equity component is fixed at 1 for Self because equity-as-distribution does not apply within a single individual. This provides a neutral, non-penalizing identity value. UCI for Self is computed from Cohesion, Flow, and Resilience with  $E_1 = 1$ . This resolves the v7.4.0 inconsistency where some sections stated  $E_1 = 0$ .

---

## **SECTION 1: INTRODUCTION: THE NEED FOR A UNIVERSAL ETHICAL OPERATING SYSTEM**

### **1.1 Alignment Failures Across Scales**

Contemporary societies operate within a tightly coupled, high-dimensional environment where climate dynamics, global supply chains, digital communication networks, financial systems, and emerging artificial intelligence systems interact in ways that increasingly resist prediction or governance. Decisions taken at one organizational scale propagate rapidly through multiple layers of human and ecological organization, generating consequences that conventional decision frameworks fail to anticipate or manage (Meadows, 2008; Newman, 2010; Steffen et al., 2015). In this context, alignment is not exclusively an artificial intelligence problem. It represents a general challenge of ensuring that the actions of individuals, institutions, governments, and machine systems remain consistent with the protection of fundamental rights, the avoidance of catastrophic failure modes, and the long-term flourishing of sentient beings and the planetary systems that sustain them (Ord, 2020; Russell, 2019; Rockström et al., 2009, 2023).

### **1.2 The Alignment Trilemma (Coupled Failure Modes)**

Existing decision frameworks exhibit three recurring failure modes that together constitute the alignment trilemma:

#### **Failure Mode A: Scalarization of Value**

Many decision methods collapse plural values into a single metric (expected utility, GDP, cost-benefit net present value). This enables "moral laundering": severe harms to some

unions are permitted when offset by gains elsewhere. Arrow's (1963) impossibility theorem demonstrates that no preference-aggregation rule can simultaneously satisfy minimal fairness criteria when preferences conflict fundamentally across multiple dimensions and scales.

### **Failure Mode B: Tail-Risk Blindness**

Expected-value reasoning underweights low-probability, high-severity outcomes. Catastrophic outcomes are often irreversible and eliminate future choice; therefore symmetric tradeoffs between ordinary gains and catastrophic losses are ethically incoherent and empirically dangerous (Taleb, 2012). Coherent risk measures such as CVaR provide better tail-risk handling (Artzner et al., 1999; Rockafellar & Uryasev, 2000). Climate tipping points (Lenton et al., 2008), pandemic risks, and AI misalignment exemplify threats that conventional frameworks systematically underweight (IPCC, 2023; Bostrom, 2014).

### **Failure Mode C: Specification Gaming (Goodhart Vulnerability)**

When metrics become targets under pressure, they are gamed. Optimization exploits proxies while degrading the intended outcome (Goodhart, 1984; Manheim & Garrabrant, 2018). This appears in institutional governance (GDP) and AI reward optimization. In AI contexts, these dynamics appear as reward hacking and proxy optimization failures under strong optimization pressure (Amodei et al., 2016).

### **Coupled Failure**

These three failure modes interact and amplify each other. Scalarization enables tail-risk blindness (catastrophe can be traded for ordinary benefit), tail-risk blindness creates exploitable blind spots, and specification gaming undermines attempts to "patch" either problem by shifting optimization pressure. Solving one in isolation is insufficient.

## **1.3 Governance-Grade Requirements**

Ripple\_Logic is designed to satisfy seven non-optional requirements:

<b>Requirement</b>	<b>Description</b>	<b>Ripple_Logic Component</b>
R1	Rights-first non-compensability: rights are lexicographic constraints, not weighted terms	NCRC (Level 1)

Requirement	Description	Ripple_Logic Component
R2	Explicit catastrophic risk bounding: tail risk is bounded by TRC using CVaR	TRC (Level 2)
R3	Multi-scale and multi-dimensional welfare representation	7×7 Welfare Matrix
R4	Explicit ripple propagation: interpretable kernel propagation under humility	Kernel K
R5	Structural integrity (containment): local gains may not degrade containing unions beyond tolerance	Containment (Level 3)
R6	Legitimate weight governance: HDW blends constitutional floors with democratic tuning	HDW
R7	Computability plus auditability: PCC and audit flags enable reconstruction, challenge, and learning	PCC plus AIL

**1.4 Paper Contributions**

This paper provides:

- A fork-resistant normative specification for Tier 1-3 implementation.
- Canonical equations and constraints sufficient for independent implementation.
- Anti-gaming architecture (non-maskable cells, subgroup semantics, scenario governance, audit flags).
- A validation and falsification program (Section 17).
- A Tier-4 design target appendix (non-claimable until ProofPack).

## SECTION 2: ONTOLOGICAL FOUNDATIONS: UNION-BASED REALITY (UBR)

### 2.1 Relational Ontology (Descriptive Thesis)

**UBR thesis:** No entity exists in complete isolation for governance-relevant domains; entities are embedded in interacting networks whose structure transmits the consequences of actions.

This is a descriptive stance supported by systems science (feedbacks and leverage points), network science (universal structural features of real networks), and Earth-system science (planetary boundary coupling) (Meadows, 2008; Newman, 2010; Steffen et al., 2015; Rockström et al., 2009, 2023). UBR is also consistent with empirical findings in social and cognitive interdependence research showing strong links between human capacities and group structure, including the social brain hypothesis and attachment-based development (Bowlby, 1988; Dunbar, 1992, 1993), and with social contagion dynamics in networks (Christakis & Fowler, 2009).

**Scope conditions:** UBR is most applicable in high-coupling systems with significant externalities, multi-scale feedbacks, and long time horizons. Low-coupling, short-horizon decisions may not require full UBR modeling; Ripple\_Logic tiers allow proportional rigor.

### 2.2 Unions as the Unit of Analysis (Operational Definition)

**Definition:** A union is a bounded pattern of interdependence: a set of entities whose internal interactions are sufficiently strong, frequent, or consequential that their welfare should be evaluated together for the decision context.

Unions are analytical constructs for causal and welfare accounting; they are not metaphysical substances. The union concept operationalizes the systems-theoretic observation that complex systems exhibit hierarchical modularity (Simon, 1962).

### 2.3 The Union Stack (Seven Operational Unions)

Union	Name	Definition	Characteristic Timescale
U <sub>1</sub>	Self	The individual locus of experience and agency	Seconds to Decades

Union	Name	Definition	Characteristic Timescale
U <sub>2</sub>	Household	Primary cohabitation and resource pooling unit (Becker, 1981)	Days to Decades
U <sub>3</sub>	Community	Local repeated-interaction network with social capital and trust dynamics (Putnam, 2000; Dunbar, 1993)	Months to Generations
U <sub>4</sub>	Organization	Formal collective pursuing a purpose; structured coordination and institutional behavior (March & Simon, 1958; North, 1990)	Years to Centuries
U <sub>5</sub>	Polity	Governance authority unit over a jurisdiction; legitimacy and institutional structure (Weber, 1978)	Decades to Centuries
U <sub>6</sub>	Humanity/CMIU	Collective Managing Intelligence Union: all managing intelligences; global coordination and systemic risk management (Ord, 2020; Steffen et al., 2015)	Generations to Millennia
U <sub>7</sub>	Biosphere	Earth's integrated life-support systems, including climate stability and ecosystem integrity (Odum, 1971; Steffen et al.,	Centuries to Epochs

Union	Name	Definition	Characteristic Timescale
		2015; Rockström et al., 2009, 2023)	

**Canonical nesting chain:**  $U_1 \subset U_2 \subset U_3 \subset U_4 \subset U_5 \subset U_6 \subset U_7$

Ecological trophic cascades illustrate how interventions can propagate unpredictably across levels and why "local wins" can generate system losses, reinforcing the need for explicit containment and ripple modeling (Estes et al., 2011).

## 2.4 Union Types vs Instances (Aggregation Rule)

The seven unions are types; real runs involve instances (many households, many organizations, many polities). Impacts MUST be aggregated within a union row using a declared aggregation method.

**Default aggregation (normative default for non-rights scoring):** Population-weighted mean across affected instances.

**Rights exception:** For rights-covered cells, apply worst-off subgroup checks within instances before aggregating (Section 7).

## 2.5 Constructive vs Pathological Unions

Some unions grow by degrading their containing unions (for example, extractive industries or corruption networks). Ripple\_Logic therefore does not assume "union benefit" is automatically good. The Containment Gate prevents local optimization that damages containing union coherence or viability beyond governed tolerance.

## 2.6 Meta-Unions (Non-Computational by Default)

$U_8$  Cosmic and  $U_9$  Universal/AIU may be used as philosophical boundary conditions or future extensions, but do not participate in standard Tier 1-3 scoring unless formally activated via a governed extension protocol (Appendix A).

---

## SECTION 3: NORMATIVE FOUNDATIONS: UNION-BASED ETHICS (UBE)

### 3.1 Minimal Normative Axiom (MNA)

Ripple\_Logic adopts exactly one explicit normative axiom:

**MNA:** Sentient flourishing matters. Unnecessary suffering should be reduced. The enabling conditions for continued flourishing should be preserved.

This axiom is:

- **Content-minimal:** No single "good life" doctrine is prescribed.
- **Scope-bounded:** Applies to agents accepting that sentience matters.
- **Derivationally sufficient:** Generates the decision architecture given UBR.

### 3.2 Conditional Is-Ought Bridge

Ripple\_Logic avoids deriving values from facts by using a conditional bridge:

IF (UBR) actions propagate through nested unions, AND (MNA) sentient welfare matters, THEN agents and institutions ought to evaluate choices by cross-union impacts and preservation of enabling conditions, using non-compensable rights floors and explicit tail-risk bounding.

Normativity enters only through MNA; UBR specifies where consequences flow.

### 3.3 Seven Welfare Dimensions (Canonical)

Dimension	Name	Description
D <sub>1</sub>	Material	Resources and infrastructure for survival and functioning
D <sub>2</sub>	Health	Physical and mental functioning; morbidity and mortality risk
D <sub>3</sub>	Social	Belonging, trust, relational integrity, cooperation
D <sub>4</sub>	Knowledge	Epistemic access and learning conditions
D <sub>5</sub>	Agency	Autonomy and effective choice; freedom from coercion

Dimension	Name	Description
D <sub>6</sub>	Meaning	Coherence, purpose, valued life projects (measured cautiously)
D <sub>7</sub>	Environment	Ecological and built context integrity sustaining life

This seven-dimensional choice is a convergence architecture rather than a claim of metaphysical completeness. It aligns with the capability approach (Sen, 1999; Nussbaum, 2011), self-determination theory (Ryan & Deci, 2000), and fundamental human needs theory (Max-Neef, 1991), while also capturing ecological integrity as a life-support substrate consistent with Earth-system boundary research (Rockström et al., 2009, 2023).

**3.4 Non-Fungibility at Rights Level**

Dimensions are treated as non-fungible at the rights layer: gains in one dimension MUST NOT compensate rights-floor violations in another. This is enforced structurally by NCRC (Section 7), not by weight tuning.

**3.5 Unioning (Enacted Practice)**

Unioning = redesigning decisions until they are rights-safe (NCRC), tail-safe (TRC), containment-safe, and net-positive across unions under declared uncertainty policy, treating apparent "tradeoffs" as design failures before acceptance.

**3.6 NCAR Learning Loop (Corrigibility Requirement)**

Ripple\_Logic is embedded in Notice, Choose, Act, Reflect. The system is corrigible: it records assumptions and outcomes, updates kernels and parameters via governed procedures, and preserves accountability through versioned PCCs.

**SECTION 4: SYSTEM OVERVIEW: THE RIPPLE\_LOGIC ARCHITECTURE**

**4.1 Formal Definition**

Ripple\_Logic is a rights-first, tail-risk-bounded, union-based decision optimization system with explicit ripple propagation, auditable scoring, and corrigible learning.

**4.2 Core Objects (Implementable Data Structures)**

Object	Description
Option set O	Finite set of candidate actions;
Welfare impact matrix per option	$I_{prop}(u,d,a) \in [-1,+1]$
Rights impacts	$I_{rights}(u,d,a) = \min$ over protected subgroups for rights-covered cells
Scenario set S with probabilities $p_s$	For TRC
Catastrophe cells $C_{cat}$ and weights $\omega$	For TRC
Kernel K (sparse 49×49)	For propagation (NONE or QUICK permitted in Tier 1-3)
Weights $w_u$ and $v_d$	HDW-governed at Tier 3 recommended; floors are constitutional
Structural metrics UCI and HOI	For containment and tie-breaks

#### 4.3 Canonical Five-Level Lexicographic Cascade (Normative)

Level	Name	Function	Failure Consequence
1	NCRC	Rights floor: exclude options violating fundamental rights	Option removed (except Emergency Mode)

Level	Name	Function	Failure Consequence
2	TRC	Tail-risk bound: exclude options with unacceptable catastrophic exposure	Option removed
3	Containment	Structural integrity: prevent degradation of containing unions	Option excluded or escalated
4	RLS	Welfare optimization: rank by weighted aggregate impact	Selection if decisive
5	UCI/Tie-break	Structural health: break ties using coherence metrics	Final selection

#### 4.4 End-to-End Algorithm (Tier 1-3 Executable) (Normative)

##### Algorithm: Ripple\_Logic\_Run

**Inputs (minimum):** Decision scope; baseline; option set O; union and dimension sets; weights; rights canon; TRC canon (when required); containment parameters; propagation mode; kernel (optional); scenario set (Tier-3 TRC required).

**Outputs:** Selected option  $a^*$  or escalation, plus PCC.

##### Step 1: Notice

Record the decision question, scope boundary, time horizon, affected unions, and option set. Declare tier and configuration.

##### Step 2: Impact Construction

For each option  $a$ : estimate direct impacts from impact instances; saturate into  $I_{dir}$ ; propagate in NONE or QUICK to obtain  $I_{prop}$ ; apply post-propagation saturation. Record uncertainty notes and missing-data penalties as required (Section 5).

##### Step 3: NCRC Gate

For rights-covered cells, compute worst-off subgroup impacts  $I_{rights}$ . Compute violation depths  $v_r(a)$ . Options with any  $v_r(a) > 0$  fail NCRC. Let  $A_{NCRC}$  be the passing set. If  $A_{NCRC}$  is empty, invoke Emergency Mode (Section 7).

**Step 4: TRC Gate**

For each  $a \in A_{NCRC}$ : compute scenario losses  $L(a,s)$  over catastrophe cells and compute  $CVaR_\alpha$ . Exclude options exceeding  $\tau_{TRC}$ . Let  $A_{adm}$  be the passing set. If  $A_{adm}$  is empty, invoke TRC fallback (Section 8).

**Step 5: Containment Gate (Mode A)**

For each  $a \in A_{adm}$  (in RLS order or all): compute  $\Delta UCI$  for containing unions and apply containment predicate. Let  $A_{sel}$  be the selectable set. If  $A_{sel}$  is empty, escalate (Section 9).

**Step 6: RLS Ranking**

Compute  $RLS(a)$  for all  $a \in A_{sel}$ . Compute  $\sigma_{RLS}(a)$  if enabled. Apply discrimination rule; if decisive, select top option.

**Step 7: Tie-break**

If non-decisive, apply UCI dominance threshold, HOI risk flags, then escalation if still non-decisive.

**Step 8: Emit PCC**

Record all inputs, intermediate computations, gate results, sensitivity outputs (Tier 3), and 5-sentence public rationale.

**4.5 Implementation Tiers (Normative Summary)**

Requirement	Tier 1 (Heuristic)	Tier 2 (Core)	Tier 3 (Auditable)
NCRC	Heuristic	REQUIRED	REQUIRED (subgroups)
TRC	Optional (qualitative screen)	REQUIRED when catastrophe relevance plausible	REQUIRED

Requirement	Tier 1 (Heuristic)	Tier 2 (Core)	Tier 3 (Auditable)
Scenario set size	N/A	≥5 minimum	≥20 minimum
Kernel propagation	NONE	NONE default; QUICK if QQS policy satisfied	QUICK only if QQS policy satisfied plus sensitivity
Containment Mode A	Optional	Recommended	REQUIRED (binding)
PCC	Optional	REQUIRED (basic)	REQUIRED (full)
Sensitivity analysis	Optional	Recommended	REQUIRED

Tier 4 is design target only in v7.4.5 and MUST NOT be claimed until ProofPack is public and independently replayable.

## SECTION 5: WELFARE IMPACT CONSTRUCTION (CALCULABLE)

### 5.0 Purpose and Design Requirements (Normative)

This section specifies how Ripple\_Logic converts real-world predictions and evidence into a computable welfare impact matrix for each option. A competent analyst MUST be able to compute all required impacts from (i) a declared baseline, (ii) a declared option, and (iii) a finite list of impact instances per cell, using only the equations and rules in this paper and appendices.

#### 5.1 Baseline-Zero Rule (Normative)

**Semantic anchor (MUST):** For all unions  $u$ , dimensions  $d$ , and options  $a$ :

$I(u,d,a) = 0$  if and only if the predicted indicator state under option  $a$  equals the baseline indicator state.

**Interpretation:** "0 means no change from baseline" is globally enforced. Any method that produces absolute levels MUST immediately convert them into baseline-relative deltas before entering Ripple\_Logic impact scoring.

## 5.2 Impact Instances (Normative)

Ripple\_Logic represents impacts in each active welfare cell (u,d) for option a as a finite set of impact instances  $K(u,d,a)$ . Each instance captures one distinct causal pathway or measurable effect.

**Table 5.1: Impact Instance Parameters (Canonical)**

Parameter	Symbol	Range	Meaning
Magnitude	$\mu_k$	[-1,+1]	Signed severity of welfare change (baseline-relative). Positive = improvement; negative = harm.
Reach	$r_k$	[0,1]	Fraction of the relevant stakeholder population materially affected.
Time horizon	$t_k$	(0, $\infty$ ) years	Duration over which the effect remains materially relevant.
Likelihood	$\ell_k$	[0,1]	Conditional probability the instance occurs (conditional on scenario model, if used).
Confidence	$c_k$	[0.1,1]	Analyst confidence in the instance specification and mapping (floor prevents "zeroing out").
Equity/resilience multiplier	$e_k$	$\geq 0$ (default 1)	Governed multiplier for equity/resilience adjustments (use cautiously; must be justified in PCC).

Parameter	Symbol	Range	Meaning
Sentience multiplier	s_k	[0,1] (default 1)	From SGP interface when ethically relevant; MUST be 1 for humans (Human Plateau Rule).

**Normative notes:**

- c\_k has a floor of 0.1 to prevent omission-by-zeroing; low confidence should reduce weight but not erase impact.
- e\_k requires PCC justification (Appendix H) and SHOULD be sensitivity-tested at Tier 3.
- s\_k may only be applied to non-human or non-plateau entities; it MUST NOT be used to weaken any human protections (Section 12; Appendix G).

**5.3 Temporal Weighting (Normative)**

Ripple\_Logic applies a temporal weighting function  $\tau(t)$  to instance effects to preserve intergenerational salience without exponential discounting:

$$\tau(t) = \ln(1 + t) / \ln(1 + T_{ref})$$

**Default:**  $T_{ref} = 25$  years.

**Illustrative values:**

Time Horizon (years)	$\tau(t)$ Value
1	≈ 0.21
5	≈ 0.56
10	≈ 0.75
25	1.00
50	≈ 1.22

**5.4 Pre-Saturation Direct Impact Aggregation (Normative)**

For option a, union u, dimension d, define the pre-saturation direct impact:

$$\tilde{I}_{dir}(u,d,a) = \sum_{\{k \in K(u,d,a)\}} [r_k \times \tau(t_k) \times \ell_k \times c_k \times e_k \times s_k \times \mu_k]$$

**Interpretation:** Instances add linearly (sum of contributions). Within each instance, the attributes multiply to form a single contribution term.

### 5.5 Saturation (Normative)

To ensure all direct impacts lie in  $[-1,+1]$  and avoid runaway totals, Ripple\_Logic uses smooth saturation:

$$I_{dir}(u,d,a) = \tanh(\beta \times \tilde{I}_{dir}(u,d,a))$$

**Default:**  $\beta = 2$ .

### 5.6 Magnitude Construction $\mu_k$ : Canonical Anchoring Methods (Normative)

Magnitude  $\mu_k$  MUST represent a baseline-relative change on the normalized  $[-1,+1]$  scale. Ripple\_Logic permits two canonical anchoring families:

#### 5.6.1 Percentile Anchoring (Default for non-rights cells)

Let  $x$  be the raw indicator (higher-is-better unless specified). Let  $P_5, P_{50}, P_{95}$  be the 5th, 50th, and 95th percentiles of  $x$  in a declared reference class.

Define a bounded level score:

$$S(x) = \text{clip}((x - P_{50}) / (P_{95} - P_5), -1, +1)$$

Then define magnitude as a baseline-relative delta:

$$\mu_k = \text{clip}(S(x_a) - S(x_0), -1, +1)$$

Where  $x_0$  is baseline indicator value, and  $x_a$  is the predicted value under option a.

If higher values are worse (for example, mortality rate), use:

$$S_{worse}(x) = -S(x)$$

**Normative requirement:** The PCC MUST record the reference class and percentile values used.

#### 5.6.2 Threshold Anchoring (Required for rights-covered cells unless invariant reference is declared)

For rights-covered cells (cells in any rights coverage set  $C_r$ ), analysts MUST NOT use context-local percentile anchoring unless the reference class is declared invariant under governance.

**Default rule (Tier 1-3):** Rights-covered magnitudes MUST be derived using threshold anchoring with explicit "good/bad" anchors so that rights semantics do not drift.

Let  $x_{good}$  be the indicator level consistent with rights-safe conditions and  $x_{bad}$  be the indicator level representing a severe rights violation onset. Map to a bounded level score:

For higher-is-better indicators:  $S(x) = \text{clip}((x - x_{bad}) / (x_{good} - x_{bad}), -1, +1)$

For higher-is-worse indicators:  $S(x) = \text{clip}((x_{bad} - x) / (x_{bad} - x_{good}), -1, +1)$

Then compute baseline-delta magnitude:

$$\mu_k = \text{clip}(S(x_a) - S(x_o), -1, +1)$$

**Normative requirement:** The PCC MUST record  $x_{good}$ ,  $x_{bad}$ , indicator definition, and direction.

### 5.7 Missing Data Rule (Ignorance Penalty) (Normative)

To prevent score inflation by omission:

If a welfare cell (u,d) is active (required by tier context and not masked for RLS), but no defensible instances can be specified, the PCC MUST:

- Mark the cell as "UNKNOWN\_IMPACT", and
- Include a phantom instance with canonical parameters:

Parameter	Phantom Value
$\mu_{phantom}$	-0.10
$r$	1
$t$	$T_{ref}$ (25 years)
$\ell$	1

Parameter	Phantom Value
c	1
e	1
s	1

This yields a mild negative default rather than unjustified neutrality.

**Clarification:** This rule does not apply to cells that are legitimately out of scope and properly masked with justification; it DOES apply to non-maskable cells and to any cell required by the tier's minimum coverage rules.

### 5.8 Scenario-Conditioned Impacts (Normative for TRC runs)

When TRC is in use (Tier 3 required; Tier 2 required when catastrophe relevance plausible), impacts MUST be scenario-conditioned at least for catastrophe cells:

Scenario probabilities  $p_s$  are governed (Appendix D). Scenario conditioning enters through  $\ell_k$  and/or instance presence, and through propagation if scenario-specific kernels are declared.

### 5.9 Worked Examples (Non-Normative, Computation-Illustrative)

#### Example 5.1: Direct impact, Community-Social cell

Decision: Remote-work policy option a. Cell: u=3 (Community), d=3 (Social).

Assume  $T_{ref} = 25$ ,  $\beta = 2$ .

Instance 1: Reduced in-person interaction

- $\mu_1 = -0.30$ ,  $r_1 = 0.80$ ,  $t_1 = 3$  years,  $\ell_1 = 0.90$ ,  $c_1 = 0.70$ ,  $e_1 = 1$ ,  $s_1 = 1$

Instance 2: Increased online community participation

- $\mu_2 = +0.15$ ,  $r_2 = 0.50$ ,  $t_2 = 5$  years,  $\ell_2 = 0.70$ ,  $c_2 = 0.50$ ,  $e_2 = 1$ ,  $s_2 = 1$

Compute temporal weights:

- $\tau(3) = \ln(4)/\ln(26) \approx 0.43$
- $\tau(5) = \ln(6)/\ln(26) \approx 0.55$

Compute contributions:

- Inst1:  $0.80 \times 0.43 \times 0.90 \times 0.70 \times 1 \times 1 \times (-0.30) \approx -0.065$
- Inst2:  $0.50 \times 0.55 \times 0.70 \times 0.50 \times 1 \times 1 \times (+0.15) \approx +0.014$

Aggregate pre-saturation:

- $\tilde{I}_{dir} = -0.065 + 0.014 = -0.051$

Saturate:

- $I_{dir}(3,3,a) = \tanh(2 \times -0.051) = \tanh(-0.102) \approx -0.10$

### Example 5.2: Missing-data penalty

Cell  $u=7$  (Biosphere),  $d=7$  (Environment) is active (non-maskable for environment-relevant decisions). Analyst lacks defensible estimates.

Phantom instance yields:

- $\tilde{I}_{dir} = 1 \times 1 \times 1 \times 1 \times 1 \times 1 \times (-0.10) = -0.10$
- $I_{dir} = \tanh(2 \times -0.10) \approx -0.20$

This enforces humility: missing evidence is mildly negative, not neutral.

---

## SECTION 6: RIPPLE PROPAGATION: KERNEL AND EPISTEMIC HUMILITY

### 6.0 Purpose (Normative)

Ripple propagation models how direct impacts in one cell causally and institutionally ripple into other cells across unions and dimensions. This makes cross-scale externalities explicit and auditable.

**Tier posture:** Tier 1-3 permit NONE or QUICK propagation. FULL propagation is a Tier-4 design target only and MUST NOT be used for Tier 1-3 compliance claims.

### 6.1 Kernel Definition and Convention (Normative)

Let the 49 welfare cells be indexed by  $i = \phi(u,d)$  with  $\phi(u,d) = 7(u-1) + d$ . The ripple kernel is a sparse matrix  $K \in \mathbb{R}^{49 \times 49}$ .

**Kernel convention (MUST):**

- $K_{\{ij\}}$  maps effect from source cell  $j$  to target cell  $i$  (target-row, source-column).

- Propagation uses left multiplication:  $\tilde{I}_{prop} = I_{dir} + K \times I_{dir}$

**Interpretation:**

- $\kappa > 0$ : improving source j tends to improve target i.
- $\kappa < 0$ : improving source j tends to harm target i.
- $\kappa = 0$ : no modeled pathway.

**6.2 Propagation Modes (Tier 1-3) (Normative)**

Mode	Formula	Use Case
NONE (Direct-only)	$I_{prop} := I_{dir}$	Tier 1-2 default
QUICK (First-order)	$\tilde{I}_{prop} = I_{dir} + K \times I_{dir}$	Tier 3 with sensitivity

Normative restriction: FULL propagation is PROHIBITED for Tier 1-3 claims in v7.4.5.

**6.3 Post-Propagation Saturation (Normative)**

Because propagation can push values outside [-1,+1], apply elementwise saturation:

$$I_{prop}(u,d,a) = \tanh(\beta_{prop} \times \tilde{I}_{prop}(u,d,a))$$

**Default:**  $\beta_{prop} = 1$ .

**6.4 Stability Constraints and Humility Fallbacks (Normative)**

**Kernel stability guardrails (defaults; may be tightened by governance):**

- Entry bound:  $|K_{\{ij\}}| \leq \kappa_{max}$ , default  $\kappa_{max} = 0.5$
- Absolute row-sum bound:  $\sum_j |K_{\{ij\}}| \leq \rho_{max}$ , default  $\rho_{max} = 0.9$
- Spectral radius constraint:  $\rho(K) < 1$ , where  $\rho(K)$  is the spectral radius of K.

If the declared kernel violates entry bounds, row-sum bounds, or spectral radius constraint, the run MUST:

- Set `propagation_mode = NONE`, and
- Record a PCC limitation: `KERNEL_HUMILITY_FALLBACK = TRUE`.

**Rationale:** An unstable or over-amplifying kernel is worse than no kernel; it creates false certainty and is easy to game.

## 6.5 Kernel Quality Score (KQS) (Normative)

KQS summarizes readiness of the kernel for decision-relevant propagation:

$$\text{KQS} = w_{\text{cov}} \times C_{\text{cov}} + w_{\text{id}} \times C_{\text{id}} + w_{\text{stab}} \times C_{\text{stab}} + w_{\text{pred}} \times C_{\text{pred}}$$

Default component weights:

Component	Weight
Coverage (C_cov)	0.25
Identifiability (C_id)	0.30
Stability (C_stab)	0.20
Prediction (C_pred)	0.25

Component meanings (each in [0,1]):

- C\_cov: coverage evidence proportion (share of relied-upon edges with cited evidence)
- C\_id: identifiability (edges are specified with clear endpoints and sign; replayable)
- C\_stab: stability margin (satisfies bounds with margin)
- C\_pred: predictive accuracy (backtest/pilot performance where available; otherwise conservative prior)

KQS policy (MUST) for Tier 3:

- If  $\text{KQS} < 0.40$ : Kernel MUST NOT be used; propagation\_mode = NONE.
- If  $0.40 \leq \text{KQS} < 0.50$ : QUICK allowed only with mandatory kernel sensitivity; otherwise use NONE.
- If  $\text{KQS} \geq 0.50$ : Kernel use permitted with required sensitivity at Tier 3.

**Tier 1-2 default:** propagation\_mode = NONE unless the PCC explicitly declares kernel use and its KQS.

## 6.6 Kernel Sensitivity Requirements (Normative for Tier 3 when QUICK is used)

When QUICK propagation is used at Tier 3, the PCC MUST include sensitivity analysis:

- Perturb each relied-upon non-zero edge by  $\pm 0.05$  (or  $\pm 10\%$  of its magnitude, whichever is larger; declare the rule).
- Recompute the cascade for each perturbation set (at minimum one-at-a-time).
- If admissibility outcomes (NCRC, TRC, Containment) or the selected option changes, set `audit_flag = DECISION_FRAGILE_KERNEL` and escalate per tier policy.

### 6.7 KOPS and Starter Kernel (Tier 2-3) (Normative plus Provisional Labeling)

Ripple\_Logic permits a governed subset of kernel edges called the Key Operational Pathways Set (KOPS). KOPS is the set of "load-bearing" pathways that are documented with evidence notes, sign-checked against literature or backtests, and sensitivity-audited.

A Starter KOPS may be provided (Appendix K) and MUST be labeled PROVISIONAL with evidence classes and a global shrink factor for elicited edges (default 0.35). If a starter kernel is used, the PCC MUST disclose this and MUST include kernel sensitivity.

### 6.8 Scenario-Conditioned Propagation (Tier 1-3 Default Rule)

**Default:** Kernel is not scenario-conditioned; scenario enters via instance likelihoods  $\ell_k$  and scenario-specific instance activation. Scenario-conditioned kernels  $K_s$  may be used only with explicit PCC justification and must carry KQS and sensitivity per scenario class.

## SECTION 7: RIGHTS LAYER: NON-COMPENSATORY RIGHTS CONSTRAINT (NCRC)

### 7.0 Purpose and Union-Based Ethics Rationale (Normative)

Ripple\_Logic enforces non-compensability: certain harms are not permitted to be "paid for" by aggregate benefits elsewhere. NCRC operationalizes this by removing rights-violating options prior to any welfare optimization.

**Union-Based Ethics (UBE) justification:** Under UBR, harms propagate across nested unions; without a rights-first constraint, local optimization systematically externalizes costs onto vulnerable subgroups and future stakeholders. NCRC is therefore a structural safeguard against value scalarization and moral laundering (Rawls, 1971; Sen, 2009).

### 7.1 Canonical Rights Set (Normative)

Ripple\_Logic defines eight core rights as feasibility floors:

#### Table 7.1: Canonical Rights Thresholds (Normative)

Right	Code	Threshold $\theta_r$	Normative Anchor
Life	LIFE	-0.90	Near-certain or highly probable death or lethal exposure
Bodily Integrity	BODY	-0.70	Severe injury, disability, torture, serious bodily violation
Liberty	LBTY	-0.65	Arbitrary detention, forced labor, coercive confinement
Basic Needs	NEED	-0.50	Severe deprivation: food insecurity, homelessness, loss of basic subsistence
Dignity	DIGN	-0.55	Systematic humiliation, dehumanization, targeted degradation
Due Process	PROC	-0.45	Denial of fair hearing, non-transparent coercive procedure
Information	INFO	-0.40	Systematic censorship, epistemic coercion, pervasive disinformation constraints
Ecological Integrity	ECOL	-0.65	Material biosphere integrity breach or planetary boundary transgression corridor

**Normative interpretation rule:** Thresholds  $\theta_r$  are admissibility floors, not weights. They do not represent "how important" a right is; they represent the minimum allowable protection level.

## 7.2 Rights Coverage Sets $C_r$ (Normative)

Each right applies to a defined subset of welfare cells  $C_r$ . NCRC evaluates each right across its coverage set and uses the worst-off subgroup for rights-covered cells.

**Canonical coverage sets (authoritative):**

Right	Coverage Set $C_r$
LIFE	$\{(u, \text{Health}): u \in \{1,2,3,4,5,6\}\} \cup \{(6, \text{Environment})\}$
BODY	$\{(u, \text{Health}): u \in \{1,2,3,4,5,6\}\}$
LBTY	$\{(u, \text{Agency}): u \in \{1,2,3,4,5,6\}\} \cup \{(u, \text{Social}): u \in \{3,4,5,6\}\}$
NEED	$\{(u, \text{Material}): u \in \{1,2,3,4,5,6\}\} \cup \{(u, \text{Health}): u \in \{1,2,3,4,5,6\}\}$
DIGN	$\{(u, \text{Social}): u \in \{1,2,3,4,5,6\}\} \cup \{(u, \text{Agency}): u \in \{1,2,3,4,5,6\}\}$
PROC	$\{(u, \text{Agency}): u \in \{4,5,6\}\} \cup \{(u, \text{Knowledge}): u \in \{4,5,6\}\} \cup \{(u, \text{Social}): u \in \{4,5,6\}\}$
INFO	$\{(u, \text{Knowledge}): u \in \{1,2,3,4,5,6\}\} \cup \{(u, \text{Agency}): u \in \{1,2,3,4,5,6\}\}$
ECOL	$\{(6, \text{Environment}), (7, \text{Environment})\}$

**7.3 Worst-Off Subgroup Semantics (Normative)**

Rights cannot be averaged away. For every rights-covered cell  $(u,d)$  in any  $C_r$ , define:

$$I_{\text{rights}}(u,d,a) = \min_{g \in G_{\{u,d\}}} I_{\text{prop}}(u,d,a|g)$$

Where  $G_{\{u,d\}}$  is the set of protected subgroups relevant to that cell.

**Tier requirements:**

- **Tier 3:** Subgroup enumeration is REQUIRED for rights-covered cells (Appendix H PCC requirements).
- **Tier 2 (strict posture):** Subgroup enumeration is REQUIRED at least for directly affected populations in rights-covered cells; when infeasible, the PCC MUST record SUBGROUP\_LIMITATION and apply a conservative bound rule (Section 7.3.2).

### 7.3.1 Minimum subgroup categories (Normative minimum)

At minimum, when applicable and feasible,  $G_{\{u,d\}}$  SHOULD include subgroups defined by:

- Age (children; elderly)
- Disability status
- Legally protected characteristics relevant in jurisdiction
- Economic vulnerability (for example, bottom income quintile)
- Geographic exposure (for example, high-risk locations)
- Domain-specific protected groups (for example, patients, detainees, precarious workers, indigenous communities)

### 7.3.2 Conservative fallback when subgroup disaggregation is infeasible (Tier 2-3) (Normative)

If subgroup impacts  $I_{\text{prop}}(u,d,a|g)$  cannot be produced for a rights-covered cell, the run MUST:

- Record `SUBGROUP_LIMITATION` in the PCC, and
- Apply a conservative rights bound for that cell:

Let  $\gamma_{\text{subgroup}}$  be a conservatism factor (default  $\gamma_{\text{subgroup}} = 1.5$ ). Then define:

- If  $I_{\text{prop}}(u,d,a) < 0$ , set  $I_{\text{rights}}(u,d,a) = \max(-1, \gamma_{\text{subgroup}} \times I_{\text{prop}}(u,d,a))$
- Else set  $I_{\text{rights}}(u,d,a) = I_{\text{prop}}(u,d,a)$

This bound applies only for NCRC checking, not for RLS scoring.

### 7.4 Violation Depth and NCRC Admissibility (Normative)

For each right  $r$ , define violation depth:

$$v_r(\mathbf{a}) = \max_{\{(u,d) \in C_r\}} (\theta_r - I_{\text{rights}}(u,d,\mathbf{a}))^+$$

where  $(x)^+ = \max(x, 0)$ .

**NCRC pass/fail predicate:**

**NCRC( $\mathbf{a}$ ) = TRUE if and only if  $v_r(\mathbf{a}) = 0$  for all  $r \in R$**

**NCRC-passing set:**

**$A\_NCRC = \{a \in O : NCRC(a) = TRUE\}$**

**Interpretation:** If any right's threshold is violated in any covered cell for the worst-off subgroup, the option is rights-inadmissible, regardless of any benefits elsewhere.

### **7.5 Emergency Mode (Rights-Failure Handling) (Normative)**

If  $A\_NCRC = \emptyset$  (no option passes rights), Ripple\_Logic enters Emergency Mode. Emergency Mode is not a loophole; it is a controlled failure protocol.

#### **7.5.1 Emergency Mode rights priority order (Normative)**

Rights are ordered lexicographically by the canonical emergency priority:

**[LIFE, BODY, ECOL, LBTY, NEED, DIGN, PROC, INFO]**

This ordering **MUST** be used exactly.

#### **7.5.2 Emergency Mode selection rule (Normative)**

Construct the violation depth vector for each option:

**$v(a) = (v\_LIFE(a), v\_BODY(a), v\_ECOL(a), v\_LBTY(a), v\_NEED(a), v\_DIGN(a), v\_PROC(a), v\_INFO(a))$**

Select the option that lexicographically minimizes  $v(a)$ . If tied:

- Minimize TRC CVaR (if TRC scenarios exist or are rapidly constructed), then
- Maximize RLS.

#### **7.5.3 Independent Challenger Requirement (Normative)**

Before Emergency Mode can be invoked, an independent challenger **MUST** propose at least one alternative option.

##### **Independence requirements (MUST):**

- No reporting relationship to decision owner
- No material interest in decision outcome
- Access to the same information set

##### **Minimum challenger effort:**

- Tier 2:  $\geq 30$  minutes active option generation
- Tier 3:  $\geq 2$  hours active option generation

If time pressure prevents this, PCC MUST record CHALLENGE\_DEFERRED\_EMERGENCY and a retrospective challenge review MUST occur within 24 hours (or earliest feasible time) with an addendum PCC.

#### **7.5.4 Emergency documentation requirements (Normative)**

PCC MUST include:

- Emergency declaration and trigger conditions
- $v_r(a)$  for each option
- Lexicographic comparison trace
- Mitigation/remediation plan
- Review cadence and return-to-normal triggers

#### **7.6 Anti-Gaming Rules Specific to NCRC (Normative)**

- Rights-covered cells MUST NOT be masked out of RLS aggregation if they are non-maskable by policy (Appendix H audit flags).
- Rights checks MUST use worst-off subgroup semantics and must not be evaluated on averages.
- "Unknown" data MUST NOT be treated as neutral; missing data triggers ignorance penalties (Section 5.7).
- Emergency Mode MUST NOT be used without challenger protocol and remediation plan.

---

## **SECTION 8: CATASTROPHIC RISK LAYER: TAIL-RISK CONSTRAINT (TRC)**

### **8.0 Purpose and UBE Rationale (Normative)**

Even when rights floors are respected in expectation, a decision can carry non-trivial probability of catastrophic harm to Humanity/CMIU or the Biosphere. Under UBE, catastrophe avoidance is lexicographically prior to welfare optimization because catastrophic outcomes can eliminate future choice across unions and collapse enabling conditions.

TRC therefore bounds catastrophic exposure using CVaR, a coherent tail-risk measure (Artzner et al., 1999; Rockafellar & Uryasev, 2000). This posture is motivated by ruin dynamics under deep uncertainty (Taleb, 2012) and by Earth-system research indicating

that destabilization of life-support conditions can occur through tipping cascades and boundary transgression corridors with multi-scale downstream harms (Rockström et al., 2009, 2023; Steffen et al., 2015; IPCC, 2023).

### 8.1 Tier Requirements for TRC (Normative; strict posture)

| Tier | TRC Required? | Minimum |S| | |-----|-----|-----| | Tier 1 | Optional (qualitative tail screen when plausible) | N/A | | Tier 2 | REQUIRED when catastrophe relevance plausible | ≥5 minimum | | Tier 3 | REQUIRED | ≥20 minimum | | Tier 4 | Design target only (claim prohibited in v7.4.5) | N/A |

#### 8.1.1 Catastrophe relevance trigger (Normative)

Catastrophe relevance is plausible if any of the following are true:

- The decision materially affects any catastrophe cell in C\_cat (Section 8.2), OR
- The decision has plausible pathways to tipping cascades (climate, bio, conflict, infrastructure, financial), OR
- The decision is materially irreversible or creates lock-in at Polity/CMIU/Biosphere scales, OR
- Credible challengers identify a plausible catastrophic failure mode even if decision owners prefer not to model it.

If triggered, TRC MUST be executed at the tier's required rigor.

### 8.2 Catastrophe Cell Set C\_cat (Normative)

**Base catastrophe cell set:**

**C\_cat\_base = {(6, Health), (6, Environment), (7, Environment)}**

These correspond to:

- **Humanity/CMIU-Health:** Global-scale health viability (pandemic mortality, mass disability, collapse of health capacity)
- **Humanity/CMIU-Environment:** Environment-as-civilization-condition (habitability, agricultural stability, climate-driven displacement)
- **Biosphere-Environment:** Earth-system integrity (planetary boundaries, biodiversity, biogeochemical stability)

The base set explicitly includes biosphere integrity and civilization-scale habitability conditions because "safe operating space" and "safe and just boundaries" analyses

identify global environmental corridors as central determinants of long-run human and ecological viability (Rockström et al., 2009, 2023).

**Extensions:** Tier 1-3 MAY extend C\_cat only with explicit PCC justification and MUST apply the same set to all compared options.

### **8.3 Scenario Governance (Normative)**

TRC is evaluated over a governed scenario set  $S$  with probabilities  $p_s$  such that  $p_s \geq 0$  and  $\sum_s p_s = 1$ .

#### **8.3.1 Mandatory Tail Scenario Categories (Normative)**

The scenario set MUST include the following tail categories unless explicitly justified as not plausible for the decision context:

- Pandemic/biological disruption
- Climate tipping cascade
- Financial system collapse
- Major conflict escalation
- Critical infrastructure failure

**Minimum probability floor per category:**  $p_{\text{floor}} \geq 0.02$

This prevents "include but set near-zero probability" gaming.

These mandatory tail categories are not intended as exhaustive; they are a minimal governance floor reflecting cross-domain systemic risk exposure in tightly coupled global systems, including climate-driven boundary cascades (Rockström et al., 2009, 2023; Steffen et al., 2015; IPCC, 2023).

#### **8.3.2 Implausibility test for exempting mandatory tail categories (Normative)**

If a deployment wishes to exempt a mandatory tail category as "implausible," the exemption requires:

(i) Written justification in the PCC explaining why the category is implausible for this specific decision context

(ii) Independent reviewer sign-off (Tier 3 required; Tier 2 recommended)

(iii) Recording in PCC with audit flag

MANDATORY\_TAIL\_CATEGORY\_MISSING\_WITH\_JUSTIFICATION

The burden of proof is on omission, not inclusion. This preserves UBE.

### 8.3.3 Scenario set size minima (Normative)

| Tier | Minimum |S| | |-----|-----| | Tier 2 (when TRC required) | ≥5 minimum | | Tier 3 | ≥20 minimum |

If a tier cannot meet its minimum scenario count, the run MUST NOT claim tier compliance for TRC coverage and MUST record audit\_flag SCENARIO\_LIBRARY\_MIN\_EXCEPTION in the PCC, including written justification and a remediation plan to reach the minimum scenario count for the next comparable run.

### 8.4 Loss Construction (Tier 1-3: Bounded-Impact TRC) (Normative)

Tier 1-3 TRC uses bounded-impact loss based on propagated welfare impacts in catastrophe cells.

**Domain constraint:** All propagated impacts are bounded:  $I_{prop\_c}(a|s) \in [-1, +1]$ .

For each scenario  $s$ , define scenario loss:

$$L(a,s) = \sum_{c \in C\_cat} \omega_c \times \max(0, -I_{prop\_c}(a|s))$$

Where:

- $\omega_c \geq 0$ ,  $\sum_{c \in C\_cat} \omega_c = 1$  (catastrophe weights are normalized)
- $\max(0, x)$  returns  $x$  if  $x > 0$ , otherwise 0 (only negative impacts contribute)
- $I_{prop\_c}(a|s)$  is the propagated impact in catastrophe cell  $c$  under scenario  $s$

**Default catastrophe weights (if not otherwise governed):** Uniform weights over  $C\_cat$ .

**Interpretation:**

- Only negative impacts in catastrophe cells contribute to loss.
- Loss is bounded in  $[0, 1]$  because weights are normalized and impacts lie in  $[-1, +1]$ .
- Positive impacts in catastrophe cells (improvements) do not offset losses; they simply contribute zero to the loss function.

**Example:** If  $C\_cat = \{(6, Health), (6, Environment), (7, Environment)\}$  with uniform weights ( $\omega_c = 1/3$  each), and under scenario  $s$  the propagated impacts are:

- $I_{prop}(6,Health,a|s) = -0.30 \rightarrow$  contributes  $1/3 \times 0.30 = 0.10$
- $I_{prop}(6,Environment,a|s) = +0.10 \rightarrow$  contributes  $1/3 \times \max(0, -0.10) = 0$

- $I_{\text{prop}}(7, \text{Environment}, a|s) = -0.60 \rightarrow \text{contributes } 1/3 \times 0.60 = 0.20$

Then  $L(a,s) = 0.10 + 0 + 0.20 = 0.30$

### 8.5 CVaR Computation (Normative)

Let  $L(a)$  be the random loss induced by scenarios  $s$  with probabilities  $p_s$ .

**Value-at-Risk:**  $\text{VaR}_\alpha[L(a)] = \inf\{z : P(L(a) \leq z) \geq \alpha\}$

**CVaR definition (Rockafellar-Uryasev form):**  $\text{CVaR}_\alpha[L(a)] = E[L(a) \mid L(a) \geq \text{VaR}_\alpha[L(a)]]$

Tier 1-3 discrete algorithm (normative; audit-replayable) is specified in Appendix D.

### 8.6 TRC Admissibility Predicate (Normative)

An option passes TRC if and only if:

**$\text{CVaR}_\alpha[L(a)] \leq \tau_{\text{TRC}}$**

Define:

**$A_{\text{adm}} = \{a \in A_{\text{NCRC}} : \text{TRC}(a) = \text{TRUE}\}$**

### 8.7 Default TRC Parameters by Context (Normative Defaults)

These are defaults; governance may tighten them. Loosening requires explicit governance justification and PCC recording.

Context	$\alpha$ (tail level)	$\tau_{\text{TRC}}$ (corridor threshold)
Personal	0.90	0.30
Organizational	0.95	0.20
Reversible policy	0.95	0.15
Irreversible policy	0.99	0.10
Existential risk	0.999	0.05

### 8.8 TRC Fallback Mode (Normative)

If  $A_{\text{NCRC}} \neq \emptyset$  but  $A_{\text{adm}} = \emptyset$  (rights-safe options exist, but all fail TRC):

- Rank all A\_NCRC options by CVaR<sub>α</sub> ascending.
- Select the option with minimal CVaR<sub>α</sub>.
- Require a time-bound risk mitigation plan and enhanced monitoring.
- Require one-tier-higher approval for the decision.
- Record TRC\_FALLBACK\_INVOKED in PCC and list the deficit (CVaR – τ<sub>TRC</sub>) for the selected option.

If both NCRC and TRC fail for all options (no rights-safe options and no tail-safe options), Emergency Mode (Section 7.5) governs with CVaR as secondary tie-break.

### **8.9 Anti-Gaming Requirements for TRC (Normative)**

- Mandatory tail categories cannot be omitted without explicit implausibility justification per Section 8.3.2.
- Probability floors apply per category; they cannot be silently violated.
- Scenario sets must be comparable across options (same S, same p<sub>s</sub>).
- PCC must record scenario provenance and how probabilities were assigned.
- Tier 2-3: if probabilities are highly uncertain, the PCC SHOULD include sensitivity over p<sub>s</sub> (Tier 3 required).

---

## **SECTION 9: CONTAINMENT LAYER: PREVENTING PATHOLOGICAL LOCAL OPTIMIZATION**

### **9.0 Purpose and UBE Rationale (Normative)**

Containment prevents a core failure mode under Union-Based Reality: a sub-union can improve its own welfare by degrading the viability, coherence, and resilience of containing unions (for example, profit by poisoning the commons). Under Union-Based Ethics, such "gains" are not counted as system-level improvement if they materially degrade the containing union beyond tolerance.

Containment is therefore a structural integrity gate applied after admissibility (NCRC and TRC) and prior to final selection.

### **9.1 Containment Principle (Normative)**

**Principle:** Positive impacts on a sub-union do not count as system-level improvements if they materially degrade the coherence or viability of any containing union beyond tolerance.

Containment is implemented as a binding gate (Mode A) for selection at Tier 3. Mode B is prohibited for determining selection (diagnostic only).

## 9.2 Definitions and Parameters (Normative)

Let  $Anc(u, D_c)$  be the ancestor set of containing unions for union  $u$  up to depth  $D_c$ , using the canonical nesting chain:

$$U_1 \subset U_2 \subset U_3 \subset U_4 \subset U_5 \subset U_6 \subset U_7$$

**Default parameters:**

Parameter	Symbol	Default	Allowed Range
Containment tolerance	$\tau_c$	-0.10	[-0.20, 0.00] (tightening only without charter)
Positive-impact threshold	$\theta_{pos}$	0.05	[0.01, 0.10]
Containment depth	$D_c$	2	{1, 2, 3}

## 9.3 Mode A (Binding): Canonical Predicate (Normative)

Containment Mode A is binding and MUST determine which options are selectable.

**Step 1:** Identify positively moving unions under option  $a$ .

Compute a union-level positive-move signal:

$$S_u(a) = \sum_d v_d \times I_{prop}(u, d, a)$$

Define:

$$U_{pos}(a) = \{u : S_u(a) \geq \theta_{pos}\}$$

**Step 2:** For each  $u \in U_{pos}(a)$ , compute minimum coherence shift among its containing unions.

Let  $A_u = Anc(u, D_c)$ .

Define:

$$M_u(a) = \min_{u' \in A_u} \Delta UCI_{\{u'\}}(a)$$

**Vacuous pass rule:** If  $A_u = \emptyset$ , then  $M_u(a)$  is vacuously PASS (contributes no containment minimum).

**Step 3:** Global containment predicate.

**Containment<sub>u</sub>(a) = TRUE if and only if  $M_u(a) \geq \tau_c$**

**Containment(a) = TRUE if and only if for all  $u \in U_{\text{pos}}(a)$ : Containment<sub>u</sub>(a) = TRUE**

**Selectable set definition:**

$$A_{\text{sel}} = \{a \in A_{\text{adm}} : \text{Containment}(a) = \text{TRUE}\}$$

**Normative action:** Any option failing containment MUST NOT be selected under Mode A. If  $A_{\text{sel}} = \emptyset$ , escalation is mandatory (Section 9.6).

#### **9.4 Containment Evaluation Order (Normative efficiency rule)**

Implementations MAY evaluate containment only for leading candidates in descending RLS order to save effort, provided no containment-failing option is ultimately selected. For auditability, the PCC MUST record containment outcomes for all options that were considered "top contenders" (at least top 2; Tier 3 recommended top 3).

#### **9.5 Mode B (Diagnostic Only; Prohibited for Selection) (Normative)**

Mode B may be used for exploratory analysis only and MUST NOT determine selection, tie-break, or admissibility.

If Mode B is used, it must:

- Be labeled DIAGNOSTIC\_ONLY in PCC,
- Produce no selection claim,
- And MUST NOT be used to justify selection of an option that fails Mode A.

**Audit rule:** If Mode B influenced selection, set audit\_flag CONTAINMENT\_MODE\_B\_USED\_FOR\_SELECTION and the PCC is INVALID.

#### **9.6 Escalation Rules for Containment Failures (Normative)**

If  $A_{\text{sel}} = \emptyset$  (no selectable options) or containment failures indicate structural hazard:

**Required actions (choose at least one and record in PCC):**

- Expand/modify option set via unioning (generate new options) and rerun cascade; OR
- Escalate to higher governance tier for redesign approval; OR
- Collect additional structural data for UCI/containment evaluation and rerun at same or higher tier.

Containment does not override NCRC/TRC; it operates after admissibility.

### 9.7 Containment Parameters Governance (Normative)

- $\tau_c$  may be tightened in PCC (made less negative) but MUST NOT be loosened below -0.10 without charter-level governance.
- $\theta_{pos}$  and  $D_c$  may be set in PCC within allowed ranges and MUST be justified if deviating from defaults.
- All parameters must be recorded in PCC with rationale and sensitivity notes (Tier 3 required).

## SECTION 10: OPTIMIZATION LAYER: RIPPLE LOGIC SCORE (RLS)

### 10.0 Purpose and Scope (Normative)

RLS ranks selectable options by expected welfare improvement across unions and dimensions, after ripple propagation and saturation, under declared weights and applicability mask rules.

RLS MUST NEVER override NCRC or TRC. It operates only on selectable options  $A_{sel}$ .

### 10.1 RLS Definition (Normative)

$$RLS(a) = \sum_u \sum_d w_u \times v_d \times m(u,d) \times \kappa(u,d) \times I_{prop}(u,d,a)$$

Where:

- $w_u$  are union weights (HDW-governed; Section 13; v6.0 floors confirmed)
- $v_d$  are dimension weights (HDW-governed; Section 13; v6.0 floors confirmed)
- $m(u,d)$  is applicability mask (RLS aggregation only)
- $\kappa(u,d)$  is a declared cell multiplier (default 1; non-default requires PCC justification)
- $I_{prop}(u,d,a)$  is post-propagation, post-saturation impact (Sections 5-6)

## 10.2 Applicability Mask Rules (Normative)

Masking exists to avoid meaningless aggregation in some contexts, but is a primary gaming vector. Therefore:

**Masking is allowed ONLY for RLS aggregation.** Masking MUST NOT be applied inside:

- Impact construction (Section 5),
- Propagation (Section 6),
- NCRC checks (Section 7),
- TRC loss and CVaR computation (Section 8),
- Or containment evaluation (Section 9).

**Non-maskable cells (MUST be unmasked):**

- All rights-covered cells in any  $C_r$
- All catastrophe cells in  $C_{cat}$
- Any cells mandated by minimum governance coverage for the run (declared in PCC; default includes at least  $U_1$ , primary affected unions, and  $U_7$  Environment when environmental relevance exists)

**Audit rule:** If any non-maskable cell is masked, set `audit_flag` RIGHTS\_CELL\_MASKED\_INVALID (or CATASTROPHE\_CELL\_MASKED\_INVALID) and PCC is INVALID.

## 10.3 Uncertainty Handling (Tier 3 Required) (Normative)

Ripple\_Logic recognizes epistemic uncertainty. Tier 3 runs MUST record uncertainty and apply the discrimination rule.

**Definition of  $\sigma(u,d,a)$  (normative minimum):** PCC MUST declare one of the following cell-level uncertainty proxies and apply it consistently across options:

**Method A (interval half-width):** If the PCC records a confidence interval  $[L,U]$  for the cell impact  $I(u,d,a)$ , set  $\sigma(u,d,a) = (U - L)/2$ .

**Method B (confidence-derived):** If the PCC records a confidence score  $c(u,d,a) \in [0,1]$  for the cell estimate and a point estimate  $I(u,d,a)$ , set  $\sigma(u,d,a) = (1 - c(u,d,a)) \times |I(u,d,a)|$ .

**Method C (calibrated table):** Use a pre-registered mapping from qualitative uncertainty labels (for example, LOW/MED/HIGH) to  $\sigma$  values, stored in the ProofPack or PCC appendix.

$\sigma$  values MUST be clipped to  $[0, 1]$ .

Define approximate RLS uncertainty:

$$\sigma_{\text{RLS}}(a) = \sqrt{[\sum_u \sum_d (w_u \times v_d \times m(u,d) \times \sigma(u,d,a))^2]}$$

**Optional risk-adjusted score:**

$$\text{RLS}_{\text{adj}}(a) = \text{RLS}(a) - \lambda \times \sigma_{\text{RLS}}(a)$$

**Default:**  $\lambda = 0.5$ . Use is optional but must be declared.

#### 10.4 Discrimination Threshold (Decisive vs Non-Decisive) (Normative)

Define gap between two options:

$$\text{Gap}(a,b) = |\text{RLS}(a) - \text{RLS}(b)| / \sqrt{[\sigma_{\text{RLS}}(a)^2 + \sigma_{\text{RLS}}(b)^2 + \epsilon]}$$

**Default:**  $\delta = 2$ ,  $\epsilon = 10^{-6}$ .

Let  $a^*$  be top RLS option and  $a_2$  the runner-up among selectable options.

**Decisive lead rule:**

- If  $\text{Gap}(a^*, a_2) > \delta$  (default  $\delta = 2$ ), select  $a^*$  (subject to tie-break gating not being needed).
- If  $\text{Gap}(a^*, a_2) \leq \delta$ , the lead is non-decisive and tie-break chain MUST be applied (Section 11.4).

#### 10.5 RLS Output Requirements (Normative)

PCC MUST record, per option in  $A_{\text{sel}}$ :

- $I_{\text{prop}}(u,d,a)$  (or a retrievable summarized representation)
- $\text{RLS}(a)$  and  $\sigma_{\text{RLS}}(a)$  and the uncertainty method
- Ranking and whether decisive (Gap vs  $\delta$ )

---

## SECTION 11: STRUCTURAL SAFEGUARDS: UCI AND HOI

### 11.0 Purpose (Normative)

Some options can "score well" on welfare impacts while eroding structural integrity (cohesion, resilience, equity), creating hollowing-out dynamics that increase long-run rights violations and tail risk. UCI and HOI are structural safeguards:

- **UCI** is a structural metric used for tie-breaks and containment.
- **HOI** is a monitoring diagnostic that flags welfare-up/coherence-down drift.

## 11.1 Union Coherence Index (UCI) (Normative)

### 11.1.1 UCI components (canonical)

For each union  $u$ , define component scores in  $[0,1]$ :

Component	Symbol	Description
Cohesion	$H_u$	Internal connectivity, trust, shared identity, conflict resolution capacity
Flow	$F_u$	Coordination throughput, information fidelity, resource allocation efficiency
Resilience	$R_u$	Redundancy, robustness, recovery speed, adaptive capacity
Equity	$E_u$	Fair distribution of burdens/benefits, voice representation, inclusion

Compute:

$$UCI_u = \alpha_H \times H_u + \alpha_F \times F_u + \alpha_R \times R_u + \alpha_E \times E_u$$

**Default:**  $\alpha_H = \alpha_F = \alpha_R = \alpha_E = 0.25$ .

**Special case for Self ( $U_1$ ):**  $E_1 := 1$  by definition (not 0). The equity component is fixed at 1 for Self because equity-as-distribution does not apply within a single individual. This provides a neutral, non-penalizing identity value. UCI for Self is thus computed as:

$$UCI_1 = 0.25 \times H_1 + 0.25 \times F_1 + 0.25 \times R_1 + 0.25 \times 1$$

This resolves the v7.4.0 inconsistency and ensures  $UCI_1$  is not systematically depressed.

### 11.1.2 Structural independence rule (Normative; Tier 3 binding)

UCI MUST be computed from structural/process indicators distinct from welfare indicators used for RLS.

At Tier 3:

- Deriving UCI from welfare-cell impacts is PROHIBITED.
- If structural indicators are unavailable, UCI is treated as unavailable and escalation/judgment-call protocol must be used (Section 11.5; Appendix E).

### **11.1.3 ΔUCI computation (Normative)**

For each option a:

$$\Delta\text{UCI}_u(a) = \text{UCI}_u(a) - \text{UCI}_u(\text{baseline})$$

Containment uses ΔUCI for containing unions (Section 9).

### **11.2 Aggregate UCI (Optional, informative)**

An aggregate coherence score may be computed:

$$\text{UCI}_{\text{agg}} = \sum_u \gamma_u \times \text{UCI}_u$$

Where  $\gamma_u$  are declared aggregation weights. This aggregate is optional and MUST NOT replace per-union UCI in containment logic unless explicitly governed.

### **11.3 Hollowing-Out Index (HOI) (Normative as monitoring definition)**

HOI detects a drift pattern: welfare score improves while coherence degrades over time.

Let  $t$  index review periods. Define:

- $\Delta\text{RLS}_t = \text{RLS}_t - \text{RLS}_{\{t-1\}}$
- $\Delta\text{UCI}_t = \text{UCI}_{\text{agg},t} - \text{UCI}_{\text{agg},\{t-1\}}$  (or a declared union-specific UCI; must be consistent)

Define exponential moving average  $\text{EMA}_\lambda$  with a declared smoothing parameter (default half-life 3 periods).

$$\text{HOI}_t = \text{EMA}_\lambda(\Delta\text{RLS})_t - \text{EMA}_\lambda(\Delta\text{UCI})_t$$

**Interpretation:**

- Persistent  $\text{HOI} > 0$ : "hollowing risk" (apparent welfare gains with structural erosion)
- HOI is diagnostic; it does not by itself make an option inadmissible, but it can trigger monitoring escalation and influence tie-break risk flags.

## 11.4 Canonical Tie-Break Chain (Normative)

Tie-breaks apply when the RLS lead is non-decisive, that is,  $\text{Gap}(a^*, a_2) \leq \delta$ .

**Tie-break chain MUST proceed in this order:**

**Step 0:** Ensure options under comparison are selectable (pass containment Mode A).

If top RLS option fails containment, it is not in  $A_{\text{sel}}$  and must not be tie-broken into selection.

**Step 1:** UCI dominance rule.

Prefer the option with higher predicted UCI outcomes if the difference exceeds a governed UCI dominance threshold  $\Delta_{\text{UCI}}$ .

**Default:**  $\Delta_{\text{UCI}} = 0.05$  (must be recorded; may be tightened).

Operationally: compare the minimum-coherence-change among critical containing unions and/or the relevant union UCI profiles. The PCC must declare the tie-break UCI comparison method:

- **Method UCI-A (default):** Maximize minimum  $\Delta\text{UCI}_u$  over containing unions relevant to the decision scope; or
- **Method UCI-B:** Maximize sum of  $\Delta\text{UCI}_u$  over the directly affected unions; or
- **Method UCI-C:** Maximize declared aggregate  $\text{UCI}_{\text{agg}}$  change.

Whichever method is used MUST be declared and applied consistently across compared options.

**Step 2:** HOI risk flag (if monitoring context exists).

If one option is associated with persistent positive HOI in comparable deployments or in modeled trajectory, treat it as riskier and prefer the alternative if the risk is material, or escalate.

**Step 3:** Escalation / judgment call.

If UCI is unavailable or non-decisive and HOI does not resolve, the decision MUST escalate to:

- Additional data collection (structural indicators), and/or
- A higher tier, and/or

- A documented governance judgment call (PCC labeled JUDGMENT\_CALL\_TIEBREAK\_NONDECISIVE) with explicit monitoring plan.

### **11.5 UCI Unavailability Rule (Tier 3) (Normative)**

If UCI cannot be computed without violating structural independence:

- UCI MUST be treated as unavailable.
- If RLS lead is non-decisive, decision MUST either: (i) escalate for more structural data / higher tier, OR (ii) record a governance judgment call with explicit labeling and monitoring plan.

**Audit label required:** JUDGMENT\_CALL\_UCI\_UNAVAILABLE.

### **11.6 Structural Safeguard Anti-Gaming (Normative)**

- UCI MUST NOT be computed from RLS welfare impacts at Tier 3.
- UCI indicator choices must be documented; changes require versioning and are subject to challenge.
- HOI must not be used to retroactively justify decisions; it is a monitoring signal feeding NCAR Reflect (Section 3.6).

---

## **SECTION 12: MORAL STATUS LAYER: SGP INTEGRATION (BINDING INTERFACE)**

### **12.0 Purpose (Normative)**

Ripple\_Logic includes moral-status handling only via a strict interface to the Sentience Gradient Protocol (SGP). The purpose is to ensure rights-of-protection are applied consistently across substrates where welfare may be at stake, while preventing misuse of "sentience scores" to justify domination, coercion, or reduced human protections.

### **12.1 Strict Separation Rule (Normative)**

Ripple\_Logic enforces two distinct gates:

**A) Rights-of-Protection (moral patienthood):** Determines minimum protections under NCRC for entities with welfare-relevant capacity (sentience).

**B) Governance Authority:** Determines who/what may exercise decision power. Authority is separately gated by competence, alignment, auditability, non-domination, and revocability.

**Strict Separation Rule (MUST):** Sentience classification determines rights-of-protection. It does NOT automatically grant governance authority.

## **12.2 SGP Authority Statement (Normative)**

Ripple\_Logic v7.4.5 does not re-specify sentience detection. SGP v4.1.1 is authoritative for sentience scoring procedures, stability gates, and evidence rules. Ripple\_Logic consumes only SGP outputs through the binding interface in Appendix G.

## **12.3 Human Plateau Rule (Normative; Non-Overridable)**

For every human person H:

**SG\_norm(H) := 1.0**

This assignment is:

- Independent of measurement noise,
- Independent of disability status,
- Independent of partial observability,
- Non-overridable by any SGP scoring outcome,
- And MUST NOT be weakened by any weighting scheme.

## **12.4 Where SGP Enters Ripple\_Logic Computation (Normative)**

### **Permitted entry point:**

- Sentience multiplier  $s_k$  may be used within impact instance aggregation (Section 5) when the impacted stakeholder set includes non-human or non-plateau entities and the analyst has an SGP output for those entities.

### **Prohibited entry points (MUST NOT):**

- SGP MUST NOT be used to weaken NCRC for any human person.
- SGP MUST NOT be used to treat rights floors as compensable welfare terms.
- SGP MUST NOT be used to grant governance authority by itself.
- SGP MUST NOT be used as a rhetorical weapon to reduce protections for any human subgroup.

## **12.5 Misinterpretation Guard (Normative)**

The following inferences are PROHIBITED:

Claim	Status
Linguistic fluency implies sentience	NOT VALID
Self-report implies sufficient evidence of sentience	NOT VALID
Intelligence implies moral status	NOT VALID
Sentience implies governance authority	NOT VALID
Precaution implies attribution (precaution means "treat as if" for protection; it does not mean "is proven")	NOT VALID
Low SGP score implies reduced human protections	PROHIBITED

### 12.6 Rights Expansion for Non-Human Stakeholder Classes (Normative process)

If SGP (under governed procedures) establishes that a non-human stakeholder class warrants rights-of-protection, Ripple\_Logic rights coverage sets and protection rules may be expanded only through governed updates:

- Update rights coverage sets C\_r and subgroup protocols,
- Document version increment,
- Record changes in PCC for subsequent runs,
- Preserve prior PCCs unchanged (no retroactive laundering).

---

## SECTION 13: WEIGHT GOVERNANCE: HYBRID DEMOCRATIC WEIGHTING (HDW)

### 13.0 Purpose (Normative)

RLS requires union and dimension weights. Weights encode normative prioritization among admissible options. Without governance, weights become a capture surface. HDW provides legitimate, anti-capture weight governance by combining:

- Constitutional floors (non-negotiable minimum protection attention)
- Democratic tuning (stakeholder voice)

- Structural evidence input (system-level constraints)

HDW affects only ranking among selectable options. HDW MUST NOT alter admissibility gates (NCRC, TRC) or containment gate.

### 13.1 Constitutional Floors (Normative; v6.0 canon)

#### 13.1.1 Union weight floors

Union	Floor
Self ( $U_1$ )	0.20
Household ( $U_2$ )	0.06
Community ( $U_3$ )	0.06
Organization ( $U_4$ )	0.06
Polity ( $U_5$ )	0.08
Humanity/CMIU ( $U_6$ )	0.10
Biosphere ( $U_7$ )	0.10
<b>Total</b>	<b>0.66</b>

#### 13.1.2 Dimension weight floors

Dimension	Floor
Material ( $D_1$ )	0.08
Health ( $D_2$ )	0.10
Social ( $D_3$ )	0.08

Dimension	Floor
Knowledge (D <sub>4</sub> )	0.08
Agency (D <sub>5</sub> )	0.10
Meaning (D <sub>6</sub> )	0.06
Environment (D <sub>7</sub> )	0.10
<b>Total</b>	<b>0.60</b>

### 13.1.3 Floor meaning (Normative)

Floors are constitutional constraints ensuring no union/dimension is mathematically eliminated from welfare ranking. Floors are not "rights" (rights are handled by NCRC); floors are minimum attention in welfare optimization among admissible options.

### 13.2 HDW Blend Formula (Normative)

Let:

- $w^{\text{floor}}$  be union floors,
- $v^{\text{floor}}$  be dimension floors.

Define allocable mass:

- $\text{allocable}_U = 1 - \sum_u w^{\text{floor}}_u = 0.34$
- $\text{allocable}_D = 1 - \sum_d v^{\text{floor}}_d = 0.40$

Let  $w^{\text{dem}}$  and  $w^{\text{str}}$  be union proposal vectors on the simplex ( $\sum_u w^{\text{dem}}_u = 1$ , all nonnegative). Let  $\lambda_U$  be the democratic share of allocable union mass. **Default:**  $\lambda_U = 0.70$ .

Then:

$$w_u = w^{\text{floor}}_u + \text{allocable}_U \times (\lambda_U \times w^{\text{dem}}_u + (1 - \lambda_U) \times w^{\text{str}}_u)$$

Similarly for dimensions with  $\lambda_D$  (default  $\lambda_D = 0.70$ ):

$$v_d = v^{\text{floor}}_d + \text{allocable}_D \times (\lambda_D \times v^{\text{dem}}_d + (1 - \lambda_D) \times v^{\text{str}}_d)$$

By construction:

- $\sum_u w_u = 1$ ,
- $\sum_d v_d = 1$ ,
- All floors are satisfied,
- Weights remain nonnegative.

### 13.3 Democratic Proposal Process (Normative minimum)

HDW requires a documented process for producing  $w^{\text{dem}}$  and  $v^{\text{dem}}$ . At minimum:

- **Participants:** Representatives from affected unions (including vulnerable populations),
- **Method:** Transparent vote or deliberative process with published results,
- **Publication:** Results recorded in PCC and (Tier 3 recommended) in an immutable ledger.

### 13.4 Structural Proposal Process (Normative minimum)

$w^{\text{str}}$  and  $v^{\text{str}}$  represent evidence-informed constraints (for example, externality reach, irreversibility, systemic risk). At minimum:

- **Method:** Documented rule for producing  $w^{\text{str}}$ ,  $v^{\text{str}}$ ,
- **Inputs:** Declared indicators and sources,
- **Reproducibility:** The same method yields same output given same data.

If a deployment lacks structural evidence, it may use a declared interim default  $w^{\text{str}}$ ,  $v^{\text{str}}$  at Tier 2, but Tier 3 SHOULD converge to a governed structural method over time.

### 13.5 Anti-Capture Safeguards (Normative for Tier 3)

Tier 3 weight governance MUST include:

- **Stratified representation:** Delegates must include vulnerable population representation and biosphere stewardship for environment-relevant decisions.
- **Supermajority lock near floors:** Any proposal reducing a weight to within 0.02 of its floor requires  $\geq 2/3$  approval.
- **Transparency ledger:** All proposals, votes, rationales published (at least internally).

- **Red-team testing:** Proposed weights tested against reference decision suites to detect systematic bias.
- **Conflict-of-interest disclosure:** Material conflicts require recusal; recusals recorded.

### 13.6 Weight Use by Tier (Normative)

Tier	Weight Policy
Tier 1	Uniform weights allowed (documented)
Tier 2	Uniform weights allowed; HDW recommended; weights must be declared in PCC
Tier 3	HDW recommended; if not available, explicit interim weights allowed with justification and sensitivity analysis. Floors remain binding regardless.

## SECTION 14: AUDITABILITY AND COMPLIANCE: PCC AND AIL (TIER 1-3; TIER 4 TARGET)

### 14.0 Purpose (Normative)

Without auditability, ethics frameworks become theater or are gamed. Ripple\_Logic therefore requires a structured decision artifact (PCC) and integrity rules (AIL) to make decisions reconstructable, challengeable, and corrigible via NCAR.

Tier boundary: Tier 1-3 auditability is fully specified here. Tier 4 determinism and hash-bound replayability are explicitly a design target only (Appendix I) and MUST NOT be claimed in v7.4.5.

### 14.1 Artifact Integrity Law (AIL) Principles (Normative for Tier 2-3)

AIL is a set of integrity constraints governing how decisions are recorded and compared.

Principle	Description
AIL1 (Registry Binding / Source Traceability)	Every Tier 2-3 PCC MUST list the normative parameters used. If hashes/registries are used, they must be referenced; if not, values must be embedded explicitly.
AIL2 (Immutability)	A PCC is immutable after signing. Any correction produces a new PCC revision that references the prior PCC and explains the change.
AIL3 (Comparability)	Options compared in one run MUST be evaluated under identical configuration.
AIL4 (No Silent Overrides)	Any override must be explicit in the PCC (what changed, why, who approved).
AIL5 (Auditability Sufficiency)	A Tier 3 PCC MUST contain enough information for an independent reviewer to recompute NCRC, TRC, containment, and RLS for that run.

#### 14.2 PCC Requirements by Tier (Normative)

Tier	PCC Requirement
Tier 1	Optional but recommended for learning
Tier 2	REQUIRED (basic)
Tier 3	REQUIRED (full) including subgroup rights checks, TRC scenario table, containment results, and sensitivity analysis bundle

#### 14.3 Audit Flags (Normative)

Audit flags are standardized labels for integrity failures or risk warnings. Flags **MUST** be recorded in PCC when triggered. If any **INVALID** flag triggers, the PCC is invalid and the decision run **MUST** be recomputed after correction.

**Definitions for this section:**

- **Mode A (Containment as binding gate):** Containment evaluation determines which options are selectable; options failing containment **MUST NOT** be selected. See Section 9.3.
- **Mode B (Containment as diagnostic only):** Containment is computed for informational purposes but does not affect selection. Mode B is **PROHIBITED** for determining selection outcomes. See Section 9.5.
- **KQS (Kernel Quality Score):** Summary score indicating kernel readiness for propagation. See Section 6.5.
- **Kernel perturbation test:** Systematic edge perturbation of  $\pm 0.05$  (or  $\pm 10\%$  of magnitude, whichever is larger) applied one-at-a-time to relied-upon non-zero kernel edges. See Section 6.6.

Canonical Audit Flags (v7.4.5) — Complete Specification

Flag	Trigger	Required Action	Severity
RIGHTS_CELL_MASKED_INVALID	Any rights-covered cell is masked, excluded, null'd, pooled, or otherwise not individually included	PCC invalid; recompute without masking	INVALID

Flag	Trigger	Required Action	Severity
	in RLS aggregation		
CATASTROPHE_CELL_MASKED_INVALID	Any catastrophe cell is masked, excluded, null'd, pooled, or otherwise not individually included in TRC computation	PCC invalid; recompute without masking	INVALID
CONTAINMENT_MODE_B_USED_FOR_SELECTION	Mode B (diagnostic-only containment) influenced selection outcome; see	PCC invalid; rerun with Mode A only	INVALID

Flag	Trigger	Required Action	Severity
	Section 9.5		
SCENARIO_LIBRARY_MIN_EXCEPTION	Scenario library size  S  is below the tier minimum when TRC is required (Tier 2:  S  < 5; Tier 3:  S  < 20)	PCC remains valid only with: (i) written justification, (ii) independent reviewer sign-off (Tier 3 required; Tier 2 recommended), and (iii) a remediation plan to reach minimum scenario count before the next comparable run. Record this flag in PCC.	ESCALATE

Flag	Trigger	Required Action	Severity
MANDATORY_TAIL_CATEGORY_MISSING	Missing mandatory tail category without justification	Add scenarios or justify per Section 8.3.2	ESCALATE
MANDATORY_TAIL_CATEGORY_MISSING_WITH_JUSTIFICATION	Missing mandatory tail category with proper justification per Section 8.3.2	Record justification; monitor for category emergence	REVIEW
MANDATORY_TAIL_PROB_FLOOR_VIOLATION	Category probability sum < p_floor (0.02)	Revise probabilities or add scenarios to category	ESCALATE
SUBGROUP_LIMITATION	Cannot disaggregate subgroups for rights-covered cell	Apply $\gamma_{subgroup}$ conservative bound (Section 7.3.2); escalate if	REVIEW

Flag	Trigger	Required Action	Severity
		high stakes	
CHALLENGE_DEFERRED_EMERGENCY	Emergency Mode invoked without independent challenger	Retrospective challenge MUST occur within 24 hours or next business day, whichever is sooner	ESCALATE
EMERGENCY_MODE_INVOKED	A_NCRC = ∅ (no rights-admissible options exist)	Remediation plan required; high scrutiny review	ESCALATE
TRC_FALLBACK_INVOKED	A_adm = ∅ after TRC (A_NCRC is non-empty but all rights-admissib	Higher-tier approval required; mitigation plan	ESCALATE

Flag	Trigger	Required Action	Severity
	le options fail TRC)		
DECISION_FRAGILE_KERNEL	Selected option changes under kernel perturbation test (Section 6.6)	Escalation required; consider NONE mode	ESCALATE
KERNEL_HUMILITY_FALLBACK	Kernel disabled due to KQS < 0.40 or stability violation (Section 6.4, 6.5)	Note limitation in PCC; plan kernel improvement	REVIEW
JUDGMENT_CALL_UCI_UNAVAILABLE	UCI unavailable in non-decisive RLS tie	Document why UCI unavailable, when UCI measurement returns, and	REVIEW

Flag	Trigger	Required Action	Severity
		interim decision basis; monitoring plan required	
JUDGMENT_CALL_TIEBREAK_NONDECISIVE	Tie-break chain did not resolve selection	Explicit manual judgment record required; independent reviewer recommended (Tier 3)	REVIEW
CONFIG_DRIFT	Parameters differ from prior run without governance update	Governance review required before proceeding	ESCALATE
REGISTRY_MISMATCH	PCC snapshot differs from referenced	Audit required; resolve discrepancy	ESCALATE

Flag	Trigger	Required Action	Severity
	registry hash		

#### 14.4 Five-Sentence Public Rationale (5SPR) (Normative)

Tier 2-3 PCC MUST include a Five-Sentence Public Rationale:

Element	Question
CONTEXT	What decision was made and why now?
OPTIONS	What options were considered?
CONSTRAINTS	What was eliminated by NCRC, TRC, and/or Containment (and why)?
SELECTION	Why the chosen option won among selectable options?
MONITORING	What follow-up will be tracked and when will NCAR Reflect occur?

#### 14.5 Tier-4 Design Target Boundary Statement (Normative)

Tier 4 compliance claims are PROHIBITED in v7.4.5. Tier-4 content is design target only (Appendix I). No determinism, hash-bound replay, or ProofPack claims may be asserted until ProofPack is publicly released and independently replayable.

### SECTION 15: IMPLEMENTATION GUIDANCE AND TIERS (NORMATIVE)

#### 15.0 Purpose (Normative)

This section specifies minimum compliance requirements by tier and provides operational guidance for applying Ripple\_Logic in real decisions. Tiers define minimum obligations;

deployments may exceed them. The framework is designed to improve over time via NCAR learning and governed updates.

### 15.1 Tier Requirements Matrix (Authoritative; Normative)

**Normative authority:** This matrix is the single authoritative statement of tier compliance. If any other sentence in this document conflicts with this matrix, this matrix governs.

Capability / Requirement	Tier 1 (Heuristic)	Tier 2 (Core, Strict TRC)	Tier 3 (Auditable)
Option set	MUST list $\geq 2$ options	REQUIRED	REQUIRED
Baseline declaration	REQUIRED	REQUIRED	REQUIRED
Impact scale and Baseline-Zero	REQUIRED (qualitative allowed)	REQUIRED (quantitative recommended)	REQUIRED (quantitative plus auditable)
NCRC rights check	REQUIRED (heuristic minimum)	REQUIRED	REQUIRED
Worst-off subgroup for rights	Recommended	REQUIRED for directly affected rights cells	REQUIRED for rights-covered cells
TRC tail-risk	Optional (qualitative screen)	REQUIRED when catastrophe relevance plausible	REQUIRED
Scenario set size when TRC used	N/A	$\geq 5$ minimum	$\geq 20$ minimum

Capability / Requirement	Tier 1 (Heuristic)	Tier 2 (Core, Strict TRC)	Tier 3 (Auditable)
Mandatory tail categories plus p_floor	Recommended if TRC used	REQUIRED when TRC used	REQUIRED
Containment Mode A	Optional	Recommended	REQUIRED (binding gate)
RLS scoring	Optional	REQUIRED	REQUIRED
Uncertainty plus discrimination band	Optional	Recommended	REQUIRED
Tie-breaks (UCI/HOI)	Optional	Recommended	REQUIRED when RLS non-decisive
Kernel propagation	NONE default	NONE default; QUICK if KQS policy satisfied	QUICK only if KQS policy satisfied plus sensitivity; else NONE
KQS policy	Optional	REQUIRED if kernel used	REQUIRED if kernel used
Sensitivity analysis	Optional	Recommended	REQUIRED
PCC artifact	Optional	REQUIRED (basic)	REQUIRED (full)
Audit flags	Optional	REQUIRED when triggered	REQUIRED when triggered

Capability / Requirement	Tier 1 (Heuristic)	Tier 2 (Core, Strict TRC)	Tier 3 (Auditable)
Tier-4 claim	PROHIBITED	PROHIBITED	PROHIBITED

### 15.2 Tier 1: Heuristic Application ("2-Minute Ripple Check") (Normative minimum)

Use Tier 1 for low-stakes, reversible decisions where full calculation is disproportionate.

#### Tier 1 minimum protocol:

1. **Decision:** What must be decided, by when?
2. **Options:** List at least 2 options (include a "third path" redesign option if possible).
3. **Rights screen (NCRC heuristic):** Could any option plausibly violate LIFE/BODY/NEED/LBTY/DIGN/PROC/INFO/ECOL? If yes, escalate to Tier 2+.
4. **Tail screen (TRC heuristic):** Could any option plausibly create catastrophic or irreversible downside for Humanity/Biosphere? If yes, escalate to Tier 2+.
5. **Unions touched:** Which unions are materially affected ( $U_1$  through  $U_7$ )?
6. **Unioning move:** What redesign could reduce harms and increase shared benefit across unions?

**Tier 1 recording (recommended):** A short note or mini-PCC with the six answers.

### 15.3 Tier 2: Core Calculable (Strict TRC Posture) (Normative)

Use Tier 2 for routine but consequential decisions requiring transparent computation and a basic PCC.

#### Tier 2 minimum obligations:

- Construct impacts on the [-1,+1] scale using Section 5 pipeline.
- Run NCRC with worst-off subgroup checks for directly affected rights cells.
- TRC MUST be executed when catastrophe relevance is plausible:
  - Use  $\geq 5$  scenarios minimum,
  - Include mandatory tail categories with probability floors unless implausible (with documented justification per Section 8.3.2),
  - Compute CVaR and corridor check,

- Record full TRC table in PCC.
- Compute RLS and record ranking.
- If uncertain/non-decisive, apply UCI tie-break if available; otherwise record judgment call and monitoring plan.
- Produce PCC (basic) plus 5SPR.

### **Tier 2 propagation guidance:**

- Default propagation\_mode = NONE.
- QUICK propagation is allowed only if KQS policy is satisfied and sensitivity is feasible.

### **15.4 Tier 3: Standard Auditable (Normative)**

Use Tier 3 for high-stakes, contested, or externally scrutinized decisions where auditability and anti-gaming posture are required.

### **Tier 3 minimum obligations:**

- Full PCC (Appendix H), including: impacts, subgroup semantics, scenario library, containment results, RLS, uncertainties, sensitivity, audit flags, signatures.
- NCRC: Subgroup analysis required for rights-covered cells.
- TRC: Required with  $\geq 20$  scenarios and mandatory tails plus probability floors (Appendix D).
- Containment Mode A: Required and binding.
- Sensitivity analysis bundle required:
  - Weights perturbation,
  - Rights threshold perturbation ( $\pm 0.05$  on  $\theta_r$  as a sensitivity test),
  - Kernel perturbation if QUICK used,
  - Scenario probability perturbation.
- Kernel: QUICK permitted only if KQS policy permits:
  - $KQS < 0.40$ : NONE only,
  - $0.40-0.50$ : QUICK plus mandatory sensitivity,

- $\geq 0.50$ : QUICK permitted plus sensitivity.
- FULL propagation is prohibited for Tier 1-3 claims.

### 15.5 NCAR Integration by Tier (Normative)

All tiers SHOULD operate within NCAR.

Phase	Action
Notice	Define scope, unions, options, baseline, and configuration
Choose	Execute cascade and emit PCC
Act	Implement with monitoring aligned to predicted impacts and tail scenarios
Reflect	Compare observed outcomes to predictions; update indicators, kernels, scenario libraries, and weights through governed procedures

#### Reflect cadence defaults:

- Tier 2: Within 6 months or after major outcome data
- Tier 3: Within 3-6 months or after major outcome data
- Emergency/Fallback: Review cadence per declared severity and risk

### 15.6 Improvement and Future Upgrades (Normative stance)

Ripple\_Logic is corrigible. Any improvement MUST:

- Preserve NCRC/TRC non-compensability structure,
- Be explicit, versioned, and auditable (AIL),
- Be tested under the validation program (Section 17),
- And be introduced via governed update processes (NCAR Reflect plus versioning).

---

## SECTION 16: RELATIONSHIP TO EXISTING FRAMEWORKS (INFORMATIVE)

### 16.0 Purpose

This section positions Ripple\_Logic relative to existing ethical, governance, and decision frameworks. It clarifies structural differences and interfaces.

### **16.1 Compared to Utilitarianism / Cost-Benefit Analysis (CBA)**

CBA often scalarizes plural values into a single metric (money/utility), enabling rights tradeoffs. Ripple\_Logic blocks this via NCRC (rights as constraints) and TRC (tail risk as constraint), then optimizes welfare only within the admissible set.

### **16.2 Compared to Deontology (Rule-based ethics)**

Pure deontology can lack a complete operational procedure for comparing permitted options under uncertainty. Ripple\_Logic retains non-compensable constraints but adds computable consequence modeling and tail-risk bounding.

### **16.3 Compared to Rawlsian Justice and Capability Approaches**

Rawls provides priority of liberty and fair basic structure; capabilities provide multi-dimensional flourishing. Ripple\_Logic operationalizes multi-dimensional welfare into a computable 7×7 matrix and makes rights floors explicit across unions.

### **16.4 Compared to MCDA (Multi-Criteria Decision Analysis)**

Standard MCDA often aggregates criteria via weighted sums or outranking without lexicographic catastrophe handling (Keeney & Raiffa, 1976; Belton & Stewart, 2002). Ripple\_Logic is MCDA-like at the RLS layer but is structurally different because admissibility gates (NCRC/TRC/Containment) are lexicographic and non-compensable.

### **16.5 Compared to AI Alignment Toolsets (RLHF, Constitutional AI, risk frameworks)**

Many alignment approaches train systems toward proxies (human feedback, constitutional principles) without hard admissibility gates. Ripple\_Logic provides a decision-engine architecture: action-space filtering by NCRC/TRC/Containment, structured welfare scoring (RLS), auditable traces (PCC), and corrigibility via NCAR.

### **16.6 Interoperability with Governance Standards**

Ripple\_Logic is designed to be interoperable with governance regimes emphasizing accountability and risk management (for example, NIST AI RMF) through its PCC record, scenario governance, and explicit risk bounding. It adds formal lexicographic rights and tail-risk operators that many standards leave at a principles level.

### **16.7 Compared to Commons Governance**

Ripple\_Logic is compatible with institutional approaches to governing shared resources and externalities. Where commons governance emphasizes rules-in-use, monitoring, graduated sanctions, and polycentric coordination, Ripple\_Logic contributes a computable, auditable decision cascade that makes rights floors, tail-risk bounds, and cross-union ripple effects explicit in each decision record (Ostrom, 1990).

---

## **SECTION 17: VALIDATION, FALSIFICATION, AND RESEARCH PROGRAM (NORMATIVE FOR CLAIMS)**

### **17.0 Purpose (Normative)**

Ripple\_Logic makes testable claims about decision quality, rights protection, tail-risk avoidance, and auditability. This section specifies explicit falsification criteria and a staged validation program. Until such validation is completed, Ripple\_Logic remains a theory-to-practice system with bounded claims: Tier compliance is claimable, real-world performance superiority is not.

This validation program follows a falsification-first posture consistent with scientific corrigibility: components that fail their empirical tests must be revised or abandoned rather than defended by authority (Popper, 1959).

### **17.1 Core Empirical Claims (Testable Hypotheses)**

<b>Hypothesis</b>	<b>Description</b>
H1	Rights coherence: Decisions that pass NCRC produce fewer rights infringements than comparable baseline decisions, controlling for context.
H2	Tail-risk effectiveness: Decisions constrained by TRC exhibit lower realized tail losses than comparable decisions without TRC.
H3	Ripple sign accuracy (after calibration): After NCAR updates, predicted impact signs match observed sign in $\geq 70\%$ of evaluated cells.

Hypothesis	Description
H4	Structural early warning: Persistent HOI > 0 predicts subsequent structural degradation better than baseline KPI monitoring.
H5	Anti-gaming effectiveness: In adversarial tests, specification gaming succeeds less often ( $\geq 30\%$ reduction) relative to comparable governance processes without PCC plus mandatory tails plus subgroup semantics.

Because ripple effects partially track cooperation and network propagation dynamics, validation should also test whether kernel updates improve predictive performance in contexts where cooperation and trust dynamics are known to matter (Axelrod, 1984; Nowak, 2006; Christakis & Fowler, 2009).

### 17.2 Falsification Criteria (Normative)

Ripple\_Logic components must be revised (or rejected) if evidence meets any of the following:

Criterion	Description
F1	NCRC failure: NCRC-passing decisions systematically produce worse rights outcomes than NCRC-failing decisions.
F2	TRC failure: TRC-constrained decisions show no reduction in realized tail losses compared to controls.
F3	Ripple predictiveness failure: Sign accuracy remains < 60% across successive NCAR cycles for key cells.
F4	Structural safeguard failure: UCI/HOI do not correlate with or predict meaningful degradation.
F5	Anti-gaming failure: Red-team exercises repeatedly exploit predictable loopholes (> 30% of adversarial runs).

### 17.3 Validation Phases (Normative roadmap)

#### Phase 1: Formal verification and implementation testing (0-6 months)

- Independent implementation of Tier 2-3 algorithms from spec.
- Unit tests for canonical equations (Appendix B).
- Adversarial tests for masking, subgroup erasure, scenario omission, confidence inflation.
- Produce reference PCCs and reproducibility checks.

#### Phase 2: Measurement validation (4-15 months)

- Indicator reliability and validity for welfare dimensions and UCI components.
- Cross-population measurement checks where relevant.
- Calibrate magnitude anchors and uncertainty mappings.

#### Phase 3: Controlled pilots (10-24 months)

- Pilot deployments in organizations or municipalities.
- Compare Ripple\_Logic vs baseline governance processes on: rights incidents, tail losses, stakeholder legitimacy, audit completeness, decision reversal rates.

#### Phase 4: Field studies and scaling (18-36+ months)

- Longitudinal monitoring of UCI/HOI and realized outcomes.
- Kernel calibration and scenario library refinement.
- Comparative performance across domains.

### 17.4 Metrics and Targets (Normative defaults)

Metric	Definition	Default Target
Rights violation rate	Verified post-decision rights infringements per protected subgroup	Lower than baseline governance

Metric	Definition	Default Target
Tail-loss severity	Realized catastrophe-cell loss in worst outcomes	Lower than baseline governance
Sign accuracy	Match of predicted vs observed sign	$\geq 0.70$ after calibration
Magnitude error (RMSE)	RMSE on normalized $[-1, +1]$ for measurable cells	$\leq 0.25$ (domain-dependent)
Audit completeness	Proportion of required PCC fields correctly filled	$\geq 0.95$ Tier 3
Gaming success rate	Fraction of adversarial attempts that change outcome illegitimately	Decreasing over time; investigate if $> 0.30$

### 17.5 Study Design Requirements (Normative)

- Pre-registration required for pilots claiming performance improvements.
- Control or comparator condition recommended (baseline governance).
- Blinding recommended for backtests: analyst should not know outcome during the run.

### 17.6 Backtesting Protocol (Normative)

A backtest run MUST:

- Reconstruct the information available at decision time (no hindsight leakage),
- Define an option set consistent with the historical context,
- Execute Ripple\_Logic using only those inputs,
- Record a "Backtest PCC" with full trace,
- Score predictions against realized outcomes using declared metrics.

### 17.7 Open Science and Responsible Disclosure (Normative intent)

Where feasible, publish:

- Anonymized PCC datasets,
- Scenario libraries and kernels (with evidence notes),
- Failure cases and revisions.

**Security note:** Where publishing would create exploitation risk, disclosure should be responsible and staged, but internal auditability must be preserved.

---

## **SECTION 18: APPLICATIONS AND USE CASES (INFORMATIVE)**

### **18.1 AI Deployment Governance (Tier 3)**

**Decision context:** An organization is deciding whether to deploy an AI system that affects users at scale.

**Key obligations:**

- NCRC: Protect dignity, information integrity, due process for worst-off subgroups.
- TRC: Model catastrophic tails (infrastructure disruption, mass manipulation, cascading conflict).
- Containment: Avoid organizational gains that degrade Polity/Humanity coherence (legitimacy collapse, epistemic fragmentation).
- PCC: Produce auditable trace; ensure comparability across options.

### **18.2 Climate and Energy Policy (Tier 3)**

**Decision context:** A polity evaluates energy transition policies that impact households, industries, and biosphere stability.

**Key obligations:**

- NCRC: Protect basic needs (energy poverty), health, life.
- TRC: Include climate tipping cascade scenarios; enforce corridor.
- Containment: Prevent short-term gains from degrading Biosphere UCI beyond tolerance.

### **18.3 Organizational Strategy (Tier 2-3)**

**Decision context:** Organization deciding on major restructuring, automation, supply chain change.

**Key obligations:**

- NCRC: Protect basic needs and dignity for worst-off employees/communities.
- TRC: Required at Tier 2 if catastrophe relevance is plausible.
- Containment: Detect hollowing-out risk even if near-term welfare improves.

**18.4 Personal and Household Decisions (Tier 1)**

**Decision context:** Personal choices (job change, relocation, major purchase) where stakes are limited and reversible.

Use Tier 1 heuristic: options, rights screen, tail screen, unions touched, unioning redesign.

Escalate to Tier 2 if any plausible rights risk exists or if the decision touches catastrophe relevance.

---

**SECTION 19: LIMITATIONS AND NON-TARGETS (NORMATIVE BOUNDARIES)****19.1 Epistemic Limitations (Normative acknowledgment)**

**Input quality dependence:** Ripple\_Logic outputs are only as good as the evidence and estimation process. The framework reduces predictable failures but does not guarantee correct forecasts.

**Measurement difficulty:** Some dimensions (Meaning; aspects of Agency) are harder to measure reliably across cultures. Ripple\_Logic treats such measurement with caution and uncertainty documentation.

Kernel uncertainty: Ripple propagation is model-based and can be wrong. v7.4.5 mitigates with KQS policy, sensitivity requirements, and humility fallback to NONE.

**Scenario incompleteness:** Scenario libraries cannot enumerate all tail risks; mandatory tail categories reduce omission but cannot eliminate deep uncertainty.

**19.2 Institutional and Governance Prerequisites**

Tier 3 requires trained analysts and governance capacity: subgroup analysis, scenario governance, structural indicators for UCI, and audit review.

HDW governance can be captured if safeguards are weak. Ripple\_Logic mandates anti-capture mechanisms, but real-world institutions must implement them.

**19.3 Risks of Compliance Theater**

Any framework can be used as a rubber stamp. Ripple\_Logic mitigates via:

- Hard gates (NCRC/TRC/Containment),
- Non-maskable cells and audit flags,
- Challenger requirement in emergency mode,
- Required PCC traceability and sensitivity.

Nonetheless, compliance theater remains a risk; validation and external scrutiny matter.

#### **19.4 Non-Targets (Normative)**

Ripple\_Logic is NOT designed to:

- Replace personal moral dialogue in intimate relationships,
- Dictate aesthetic or spiritual preferences,
- Resolve all deep metaphysical disagreements,
- Provide certainty under irreducible uncertainty,
- Function as a rhetorical weapon to end debate.

It is a public decision operating system: a method for auditable, multi-stakeholder, high-consequence decisions.

---

## **SECTION 20: CONCLUSION: PROSPEROUS FUTURES FOR INTELLIGENCES**

### **20.1 Prosperity Defined (Normative operational definition)**

In Ripple\_Logic terms:

**Prosperity = rights-safe (NCRC) + tail-safe (TRC) + containment-safe + net-positive welfare across unions + coherence preserved (UCI/HOI safeguards) + corrigible learning (NCAR)**

This definition blocks false prosperity:

- "Growth" that violates rights is not prosperity.
- "Efficiency" that increases catastrophic exposure is not prosperity.
- "Success" that hollows out structural coherence is not prosperity.

### **20.2 Ripple\_Logic Thesis**

Alignment across humans, institutions, and AI requires non-compensable constraints on rights and catastrophic risk, explicit modeling of ripple propagation through nested unions, transparent and auditable decision procedures, and corrigible learning loops. These requirements follow from the combination of interdependence in complex systems (Meadows, 2008; Newman, 2010), planetary life-support constraints (Steffen et al., 2015; Rockström et al., 2009, 2023), and the tail-risk properties of ruin dynamics (Taleb, 2012; Ord, 2020).

### **20.3 Call to Test, Critique, and Deploy (Normative stance)**

Ripple\_Logic v7.4.5 is a scaffold: a spec-hardened foundation intended to be tested, falsified where wrong, refined where incomplete, and deployed proportionally to stakes. The method is designed to move societies toward WE: ethically, sustainably, harmoniously, prosperously, and verifiably, by making decision structure explicit, auditable, and corrigible.

---

### **REFERENCES (APA 7)**

- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. *arXiv*. <https://doi.org/10.48550/arXiv.1606.06565>
- Arrow, K. J. (1963). *Social choice and individual values* (2nd ed.). Yale University Press.
- Artzner, P., Delbaen, F., Eber, J.-M., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3), 203–228. <https://doi.org/10.1111/1467-9965.00068>
- Axelrod, R. (1984). *The evolution of cooperation*. Basic Books.
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512. <https://doi.org/10.1126/science.286.5439.509>
- Becker, G. S. (1981). *A treatise on the family*. Harvard University Press.
- Belton, V., & Stewart, T. J. (2002). *Multiple criteria decision analysis: An integrated approach*. Kluwer Academic Publishers.
- Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.
- Bowlby, J. (1988). *A secure base: Parent-child attachment and healthy human development*. Basic Books.
- Christakis, N. A., & Fowler, J. H. (2009). *Connected: The surprising power of our social networks and how they shape our lives*. Little, Brown and Company.

Dunbar, R. I. M. (1992). Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*, 22(6), 469–493. [https://doi.org/10.1016/0047-2484\(92\)90081-J](https://doi.org/10.1016/0047-2484(92)90081-J)

Dunbar, R. I. M. (1993). Coevolution of neocortical size, group size and language in humans. *Behavioral and Brain Sciences*, 16(4), 681–735. <https://doi.org/10.1017/S0140525X00032325>

Estes, J. A., Terborgh, J., Brashares, J. S., Power, M. E., Berger, J., Bond, W. J., Carpenter, S. R., Essington, T. E., Holt, R. D., Jackson, J. B. C., Marquis, R. J., Oksanen, L., Oksanen, T., Paine, R. T., Pkitch, E. K., Ripple, W. J., Sandin, S. A., Scheffer, M., Schoener, T. W., ... Wardle, D. A. (2011). Trophic downgrading of planet Earth. *Science*, 333(6040), 301–306. <https://doi.org/10.1126/science.1205106>

Floridi, L. (2013). *The ethics of information*. Oxford University Press.

Goodhart, C. A. E. (1984). Problems of monetary management: The UK experience. In C. A. E. Goodhart (Ed.), *Monetary theory and practice: The UK experience* (pp. 91–121). Macmillan.

Hammond, J. S., Keeney, R. L., & Raiffa, H. (1998). *Smart choices: A practical guide to making better decisions*. Harvard Business School Press.

Intergovernmental Panel on Climate Change. (2023). *Climate change 2023: Synthesis report*. <https://doi.org/10.59327/IPCC/AR6-9789291691647>

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–292. <https://doi.org/10.2307/1914185>

Keeney, R. L., & Raiffa, H. (1976). *Decisions with multiple objectives: Preferences and value tradeoffs*. Wiley.

Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S., & Schellnhuber, H. J. (2008). Tipping elements in the Earth's climate system. *Proceedings of the National Academy of Sciences*, 105(6), 1786–1793. <https://doi.org/10.1073/pnas.0705414105>

Manheim, D., & Garrabrant, S. (2018). Categorizing variants of Goodhart's law. *arXiv*. <https://doi.org/10.48550/arXiv.1803.04585>

March, J. G., & Simon, H. A. (1958). *Organizations*. Wiley.

Max-Neef, M. A. (1991). *Human scale development: Conception, application and further reflections*. Apex Press.

McGaughan, J. (2026). *Sentience Gradient Protocol (SGP) v4.1.1* [Manuscript in preparation]. MathGov Institute for Ethical Systems Design.

- Meadows, D. H. (2008). *Thinking in systems: A primer*. Chelsea Green Publishing.
- National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST AI 100-1). <https://doi.org/10.6028/NIST.AI.100-1>
- Newman, M. E. J. (2010). *Networks: An introduction*. Oxford University Press.
- North, D. C. (1990). *Institutions, institutional change and economic performance*. Cambridge University Press.
- Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science*, 314(5805), 1560–1563. <https://doi.org/10.1126/science.1133755>
- Nussbaum, M. C. (2011). *Creating capabilities: The human development approach*. Harvard University Press.
- Odum, E. P. (1971). *Fundamentals of ecology* (3rd ed.). W. B. Saunders.
- Ord, T. (2020). *The precipice: Existential risk and the future of humanity*. Hachette Books.
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.
- Popper, K. (1959). *The logic of scientific discovery*. Hutchinson.
- Putnam, R. D. (2000). *Bowling alone: The collapse and revival of American community*. Simon & Schuster.
- Rawls, J. (1971). *A theory of justice*. Harvard University Press.
- Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2(3), 21–42. <https://doi.org/10.21314/JOR.2000.038>
- 
- Rockström, J., Steffen, W., Noone, K., Persson, Å., Chapin, F. S., Lambin, E. F., Lenton, T. M., Scheffer, M., Folke, C., Schellnhuber, H. J., Nykvist, B., de Wit, C. A., Hughes, T., van der Leeuw, S., Rodhe, H., Sörlin, S., Snyder, P. K., Costanza, R., Svedin, U., ... Foley, J. A. (2009). A safe operating space for humanity. *Nature*, 461(7263), 472–475. <https://doi.org/10.1038/461472a>
- Rockström, J., Gupta, J., Qin, D., Lade, S. J., Abrams, J. F., Andersen, L. S., Armstrong McKay, D. I., Bai, X., Bala, G., Bunn, S. E., Ciobanu, D., DeClerck, F., Ebi, K., Gifford, L., Gordon, C., Hasan, S., Kanie, N., Lenton, T. M., Loriani, S., ... Zhang, X. (2023). Safe and just Earth system boundaries. *Nature*, 619(7968), 102–111. <https://doi.org/10.1038/s41586-023-06083-8>

Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Viking.

Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, 55(1), 68–78. <https://doi.org/10.1037/0003-066X.55.1.68>

Sen, A. (1999). *Development as freedom*. Knopf.

Sen, A. (2009). *The idea of justice*. Harvard University Press.

Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6), 467–482.

Steffen, W., Richardson, K., Rockström, J., Cornell, S. E., Fetzer, I., Bennett, E. M., Biggs, R., Carpenter, S. R., de Vries, W., de Wit, C. A., Folke, C., Gerten, D., Heinke, J., Mace, G. M., Persson, L. M., Ramanathan, V., Reyers, B., & Sörlin, S. (2015). Planetary boundaries: Guiding human development on a changing planet. *Science*, 347(6223), 1259855. <https://doi.org/10.1126/science.1259855>

Taleb, N. N. (2012). *Antifragile: Things that gain from disorder*. Random House.

Weber, M. (1978). *Economy and society* (G. Roth & C. Wittich, Eds.). University of California Press. (Original work published 1922)

---

## APPENDIX A: SYMBOLS AND NOTATION (v7.4.5)

Appendix A is normative for symbol meanings and domains.

### A.1 Sets and Indices

Symbol	Meaning	Domain
u	Union index	$u \in \{1,2,3,4,5,6,7\}$
d	Dimension index	$d \in \{1,2,3,4,5,6,7\}$
a	Option (candidate action)	$a \in O$
s	Scenario index	$s \in S$

Symbol	Meaning	Domain
r	Right index	$r \in R$
k	Impact instance index	integer
g	Protected subgroup index	$g \in G_{\{u,d\}}$

## A.2 Core Sets

Set	Definition
U	Operational unions {Self, Household, Community, Organization, Polity, Humanity/CMIU, Biosphere}
D	Welfare dimensions {Material, Health, Social, Knowledge, Agency, Meaning, Environment}
O	Option set
S	Scenario set
R	Rights set {LIFE, BODY, LBTY, NEED, DIGN, PROC, INFO, ECOL}
$C_r \subseteq U \times D$	Coverage set for right r
$C_{cat} \subseteq U \times D$	Catastrophe cell set for TRC

## A.3 Union Stack (Canonical)

Union	Name	Description (summary)
U <sub>1</sub>	Self	Individual locus of experience/agency
U <sub>2</sub>	Household	Primary cohabitation/resource pooling unit
U <sub>3</sub>	Community	Local repeated-interaction network
U <sub>4</sub>	Organization	Formal collective pursuing a purpose
U <sub>5</sub>	Polity	Governance authority unit over jurisdiction
U <sub>6</sub>	Humanity/CMIU	Collective managing intelligence union
U <sub>7</sub>	Biosphere	Earth integrated life-support systems

#### A.4 Welfare Dimensions (Canonical)

Dimension	Name	Summary
D <sub>1</sub>	Material	Resources, infrastructure, subsistence security
D <sub>2</sub>	Health	Physical/mental functioning, morbidity/mortality risk
D <sub>3</sub>	Social	Trust, belonging, relational integrity
D <sub>4</sub>	Knowledge	Epistemic access, learning conditions
D <sub>5</sub>	Agency	Autonomy, effective choice, freedom from coercion
D <sub>6</sub>	Meaning	Purpose/coherence/valued projects (cautious measurement)

Dimension	Name	Summary
D <sub>7</sub>	Environment	Ecological and built context integrity

### A.5 Impact Objects

Symbol	Meaning	Range
$\tilde{I}_{dir}(u,d,a)$	Pre-saturation direct impact	$\mathbb{R}$
$I_{dir}(u,d,a)$	Direct impact (post-saturation)	[-1,+1]
$\tilde{I}_{prop}(u,d,a)$	Pre-saturation propagated impact	$\mathbb{R}$
$I_{prop}(u,d,a)$	Propagated impact (post-saturation)	[-1,+1]
$I_{rights}(u,d,a)$	Worst-off subgroup propagated impact	[-1,+1]

### A.6 Impact Instance Attributes

Symbol	Meaning	Range / Default
$\mu_k$	Magnitude (baseline-delta, normalized)	[-1, +1]
$r_k$	Reach	[0, 1]
$t_k$	Time horizon	(0, $\infty$ ) years
$\ell_k$	Likelihood	[0, 1]
$c_k$	Confidence	[0.1, 1]
$e_k$	Equity/resilience multiplier	$\geq 0$ , default 1

Symbol	Meaning	Range / Default
s_k	Sentence multiplier (SGP)	[0, 1], default 1; MUST be 1 for humans

### A.7 Operators and Functions

Symbol	Definition
$(x)^+$	$\max(x, 0)$
$\text{clip}(x, a, b)$	$\max(a, \min(x, b))$
$\text{sat}_\beta(x)$	$\tanh(\beta \times x)$
$\tau(t)$	$\ln(1 + t) / \ln(1 + T_{\text{ref}})$ , default $T_{\text{ref}} = 25$ years
$\phi(u, d)$	$7(u - 1) + d$ (flattening map for kernel indexing)
$\text{Anc}(u, D_c)$	Ancestor set of containing unions for union $u$ up to depth $D_c$

### A.8 Kernel and Propagation

Symbol	Meaning
$K \in \mathbb{R}^{49 \times 49}$	Ripple kernel matrix (sparse)
$K_{\{ij\}}$	Kernel entry mapping source $j$ to target $i$ (target-row, source-column)
$\kappa_{\text{max}}$	Maximum absolute kernel entry bound, default 0.5
$\rho_{\text{max}}$	Maximum absolute row-sum bound, default 0.9

Symbol	Meaning
$\rho(K)$	Spectral radius of $K$ ; must satisfy $\rho(K) < 1$
NONE mode	$I_{prop} := I_{dir}$
QUICK mode	$\tilde{I}_{prop} = I_{dir} + K \times I_{dir}$ , then saturate

### A.9 Rights and Constraints

Symbol	Meaning
$\theta_r$	Rights threshold for right $r$
$v_r(a)$	Violation depth for right $r$ under option $a$
$NCRC(a)$	Rights admissibility predicate
$TRC(a)$	Tail-risk admissibility predicate
$\tau_{TRC}$	TRC corridor threshold
$\alpha$	CVaR tail level $\in (0, 1)$
$CVaR_\alpha[L(a)]$	Conditional Value-at-Risk for loss distribution of option $a$
$\theta_{pos}$	Positive impact threshold (containment trigger), default 0.05
$\tau_c$	Containment tolerance threshold, default -0.10
$D_c$	Containment ancestor depth, default 2

### A.10 Scoring and Uncertainty

Symbol	Meaning
$w_u$	Union weight ( $w_u \geq 0, \sum_u w_u = 1$ )
$v_d$	Dimension weight ( $v_d \geq 0, \sum_d v_d = 1$ )
$m(u,d)$	Applicability mask for RLS only $\in \{0, 1\}$
$\kappa(u,d)$	Cell multiplier for RLS, default 1
$RLS(a)$	Ripple Logic Score
$\sigma(u,d,a)$	Cell-level uncertainty proxy
$\sigma_{RLS}(a)$	RLS uncertainty proxy
$\delta$	Discrimination threshold for decisive lead, default 2
$\varepsilon$	Stabilizer in Gap computation, default $10^{-6}$
$\lambda$	Risk-aversion coefficient for $RLS_{adj}$ , default 0.5

### A.11 Structural Metrics

Symbol	Meaning
$UCI_u$	Union Coherence Index for union $u \in [0, 1]$
$\Delta UCI_u(a)$	Coherence change under option $a = UCI_u(a) - UCI_u(\text{baseline})$
$HOI_t$	Hollowing-Out Index at time $t$
$H_u$	Cohesion component of UCI

Symbol	Meaning
F <sub>u</sub>	Flow component of UCI
R <sub>u</sub>	Resilience component of UCI
E <sub>u</sub>	Equity component of UCI ( $E_1 := 1$ for Self)

### A.12 Audit Artifacts

Symbol	Meaning
PCC	Provenance and Compliance Certificate
5SPR	Five-Sentence Public Rationale
AIL	Artifact Integrity Law (Tier 2-3 integrity rules)
Audit flags	Standard labels for invalidity or warnings (Appendix H)

### A.13 Scenario and TRC Objects

Symbol	Meaning
S	Scenario set
p <sub>s</sub>	Scenario probability ( $p_s \geq 0, \sum_s p_s = 1$ )
L(a,s)	Scenario loss for option a under scenario s
$\omega_c$	Catastrophe cell weight ( $\omega_c \geq 0, \sum_c \omega_c = 1$ )
p <sub>floor</sub>	Minimum probability floor per mandatory tail category, default 0.02

### A.14 SGP Integration

Symbol	Meaning
SG_norm(E)	Normalized sentence gradient for entity E $\in [0, 1]$
A(E), B(E), C(E)	SGP pillar scores: Awareness, Agency, Union Participation

End Appendix A.

## APPENDIX B: CANONICAL EQUATIONS (TIER 1-3 EXECUTABLE)

Appendix B is normative and is the canonical equation pack for Tier 1-3 implementations. All equations are reproduced explicitly for standalone executability.

### B.1 Temporal Weighting

$$\tau(t) = \ln(1 + t) / \ln(1 + T_{\text{ref}})$$

**Default:**  $T_{\text{ref}} = 25$  years.

### B.2 Direct Impact Aggregation (Pre-Saturation)

$$\tilde{I}_{\text{dir}}(u,d,a) = \sum_{\{k \in K(u,d,a)\}} [r_k \times \tau(t_k) \times \ell_k \times c_k \times e_k \times s_k \times \mu_k]$$

Where  $K(u,d,a)$  is the set of impact instances asserted for cell  $(u,d)$  under option  $a$ .

### B.3 Direct Saturation

$$I_{\text{dir}}(u,d,a) = \tanh(\beta \times \tilde{I}_{\text{dir}}(u,d,a))$$

**Default:**  $\beta = 2$ .

### B.4 Missing Data (Ignorance Penalty Phantom Instance)

If cell  $(u,d)$  is required-active but  $K(u,d,a) = \emptyset$ , add phantom instance with:

Parameter	Phantom Value
$\mu_{\text{phantom}}$	-0.10
$r$	1

Parameter	Phantom Value
t	T_ref (25 years)
$\ell$	1
c	1
e	1
s	1

### B.5 Flattening Map (Kernel Indexing)

Vectorize  $I_{dir}$  into a 49-vector by:

$$i = \phi(u,d) = 7(u - 1) + d$$

### B.6 Propagation Modes (Tier 1-3: NONE and QUICK only)

**NONE:**

$$I_{prop} := I_{dir}$$

**QUICK:**

$$\tilde{I}_{prop} = I_{dir} + K \times I_{dir}$$

Then apply post-propagation saturation (B.7).

### B.7 Post-Propagation Saturation

$$I_{prop}(u,d,a) = \tanh(\beta_{prop} \times \tilde{I}_{prop}(u,d,a))$$

**Default:**  $\beta_{prop} = 1$ .

### B.8 Worst-Off Subgroup Impact (Rights)

$$I_{rights}(u,d,a) = \min_{\{g \in G_{\{u,d\}}\}} I_{prop}(u,d,a|g)$$

### B.9 NCRC Violation Depth and Rights Admissibility

For each right  $r$  with threshold  $\theta_r$  and coverage set  $C_r$ :

$$v_r(a) = \max_{\{(u,d) \in C_r\}} (\theta_r - I_{rights}(u,d,a))^+$$

where  $(x)^+ = \max(x, 0)$ .

**NCRC(a) = TRUE if and only if  $v_r(a) = 0$  for all  $r \in R$**

### **B.10 TRC Loss (Tier 1-3 Bounded-Impact Mode)**

Let  $C_{cat}$  be catastrophe cells and  $\omega$  catastrophe weights.

For scenario  $s$ :

$$L(a,s) = \sum_{c \in C_{cat}} \omega_c \times (-I_{prop_c}(a|s))^+$$

TRC uses CVaR (Appendix D for discrete computation) and corridor threshold  $\tau_{TRC}$ :

**TRC(a) = TRUE if and only if  $CVaR_\alpha[L(a)] \leq \tau_{TRC}$**

### **B.11 Containment Mode A (Binding)**

Define positively moving unions:

$$S_u(a) = \sum_d v_d \times I_{prop}(u,d,a)$$

$$U_{pos}(a) = \{u : S_u(a) \geq \theta_{pos}\}$$

For each  $u \in U_{pos}(a)$ :

$$M_u(a) = \min_{\{u' \in Anc(u, D_c)\}} \Delta UCI_{\{u'\}}(a)$$

Global containment predicate:

**Containment(a) = TRUE if and only if for all  $u \in U_{pos}(a)$ :  $M_u(a) \geq \tau_c$**

### **B.12 Ripple Logic Score (RLS)**

$$RLS(a) = \sum_u \sum_d w_u \times v_d \times m(u,d) \times \kappa(u,d) \times I_{prop}(u,d,a)$$

**Default:**  $\kappa(u,d) = 1$ ,  $m(u,d) = 1$  unless masked.

### **B.13 RLS Uncertainty (Optional; Tier 3 Required)**

**Definition of  $\sigma(u,d,a)$  (normative minimum):** PCC MUST declare one of the following cell-level uncertainty proxies and apply it consistently across options:

**Method A (interval half-width):** If the PCC records a confidence interval  $[L,U]$  for the cell impact  $I(u,d,a)$ , set  $\sigma(u,d,a) = (U - L)/2$ .

**Method B (confidence-derived):** If the PCC records a confidence score  $c(u,d,a) \in [0,1]$  for the cell estimate and a point estimate  $I(u,d,a)$ , set  $\sigma(u,d,a) = (1 - c(u,d,a)) \times |I(u,d,a)|$ .

**Method C (calibrated table):** Use a pre-registered mapping from qualitative uncertainty labels (for example, LOW/MED/HIGH) to  $\sigma$  values, stored in the ProofPack or PCC appendix.

$\sigma$  values MUST be clipped to  $[0,1]$ .

$$\sigma_{\text{RLS}}(a) = \sqrt{[\sum_u \sum_d (w_u \times v_d \times m(u,d) \times \sigma(u,d,a))^2]}$$

Risk-adjusted score (optional):

$$\text{RLS}_{\text{adj}}(a) = \text{RLS}(a) - \lambda \times \sigma_{\text{RLS}}(a)$$

**Default:**  $\lambda = 0.5$ .

#### **B.14 Discrimination Gap (Decisive vs Non-Decisive)**

$$\text{Gap}(a,b) = |\text{RLS}(a) - \text{RLS}(b)| / \sqrt{[\sigma_{\text{RLS}}(a)^2 + \sigma_{\text{RLS}}(b)^2 + \varepsilon]}$$

**Default:**  $\delta = 2$ ,  $\varepsilon = 10^{-6}$ .

Decisive lead if  $\text{Gap}(a^*, a_2) > \delta$ .

#### **B.15 UCI Computation**

$$\text{UCI}_u = \alpha_H \times H_u + \alpha_F \times F_u + \alpha_R \times R_u + \alpha_E \times E_u$$

**Default:**  $\alpha_H = \alpha_F = \alpha_R = \alpha_E = 0.25$ .

**Special case for Self ( $U_1$ ):**  $E_1 := 1$  by definition.

#### **B.16 $\Delta$ UCI Computation**

$$\Delta\text{UCI}_u(a) = \text{UCI}_u(a) - \text{UCI}_u(\text{baseline})$$

#### **B.17 HOI Computation**

$$\text{HOI}_t = \text{EMA}_\lambda(\Delta\text{RLS})_t - \text{EMA}_\lambda(\Delta\text{UCI})_t$$

**Default half-life:** 3 periods.

End Appendix B.

---

### **APPENDIX C: RIGHTS / NCRC CANON PACK (TIER 1-3)**

Appendix C is normative.

#### **C.1 Rights Thresholds (Canonical)**

Right	Code	Threshold $\theta_r$
Life	LIFE	-0.90
Bodily Integrity	BODY	-0.70
Liberty	LBTY	-0.65
Basic Needs	NEED	-0.50
Dignity	DIGN	-0.55
Due Process	PROC	-0.45
Information	INFO	-0.40
Ecological Integrity	ECOL	-0.65

**Normative interpretation rule:** Thresholds  $\theta_r$  are admissibility floors, not weights. They do not represent "how important" a right is; they represent the minimum allowable protection level.

### C.2 Rights Coverage Sets $C_r$ (Authoritative)

Right	Coverage Set $C_r$
LIFE	$\{(u, \text{Health}): u \in \{1,2,3,4,5,6\}\} \cup \{(6, \text{Environment})\}$
BODY	$\{(u, \text{Health}): u \in \{1,2,3,4,5,6\}\}$
LBTY	$\{(u, \text{Agency}): u \in \{1,2,3,4,5,6\}\} \cup \{(u, \text{Social}): u \in \{3,4,5,6\}\}$
NEED	$\{(u, \text{Material}): u \in \{1,2,3,4,5,6\}\} \cup \{(u, \text{Health}): u \in \{1,2,3,4,5,6\}\}$

Right	Coverage Set $C_r$
DIGN	$\{(u, \text{Social}): u \in \{1,2,3,4,5,6\}\} \cup \{(u, \text{Agency}): u \in \{1,2,3,4,5,6\}\}$
PROC	$\{(u, \text{Agency}): u \in \{4,5,6\}\} \cup \{(u, \text{Knowledge}): u \in \{4,5,6\}\} \cup \{(u, \text{Social}): u \in \{4,5,6\}\}$
INFO	$\{(u, \text{Knowledge}): u \in \{1,2,3,4,5,6\}\} \cup \{(u, \text{Agency}): u \in \{1,2,3,4,5,6\}\}$
ECOL	$\{(6, \text{Environment}), (7, \text{Environment})\}$

### C.3 Emergency Mode Rights Priority (Canonical)

When Emergency Mode is invoked ( $A_{\text{NCRC}} = \emptyset$ ), rights are ordered lexicographically:

**[LIFE, BODY, ECOL, LBTY, NEED, DIGN, PROC, INFO]**

This ordering MUST be used exactly for lexicographic minimization of the violation depth vector.

### C.4 NCRC Violation Depth Definition (Canonical)

$$v_r(a) = \max_{\{(u,d) \in C_r\}} (\theta_r - I_{\text{rights}}(u,d,a))^+$$

### C.5 NCRC Predicate (Canonical)

**NCRC(a) = TRUE if and only if  $v_r(a) = 0$  for all  $r \in R$**

### C.6 Subgroup Conservative Bound (Canonical)

When subgroup disaggregation is infeasible for a rights-covered cell:

Let  $\gamma_{\text{subgroup}} = 1.5$  (default conservatism factor).

- If  $I_{\text{prop}}(u,d,a) < 0$ : set  $I_{\text{rights}}(u,d,a) = \max(-1, \gamma_{\text{subgroup}} \times I_{\text{prop}}(u,d,a))$
- Else: set  $I_{\text{rights}}(u,d,a) = I_{\text{prop}}(u,d,a)$

This bound applies only for NCRC checking, not for RLS scoring.

### C.7 Emergency Mode Documentation Requirements (Canonical)

When Emergency Mode is invoked, the PCC MUST include:

- Emergency declaration and trigger conditions

- $v_r(a)$  for each option and each right
- Lexicographic comparison trace
- Independent challenger attestation (or CHALLENGE\_DEFERRED\_EMERGENCY flag)
- Mitigation/remediation plan with timeline
- Review cadence and return-to-normal triggers

End Appendix C.

---

## APPENDIX D: TRC AND SCENARIO GOVERNANCE CANON PACK (TIER 1-3)

Appendix D is normative.

### D.1 Base Catastrophe Cell Set

**C\_cat\_base = {(6, Health), (6, Environment), (7, Environment)}**

These correspond to:

- **(6, Health):** Humanity/CMIU-Health, representing global-scale health viability
- **(6, Environment):** Humanity/CMIU-Environment, representing civilization-scale habitability conditions
- **(7, Environment):** Biosphere-Environment, representing Earth-system integrity

### D.2 Mandatory Tail Scenario Categories (Canonical)

The scenario set MUST include the following tail categories unless explicitly justified as not plausible:

- **Pandemic/biological disruption:** Disease outbreak with significant mortality/morbidity
- **Climate tipping cascade:** Triggering of climate tipping points with cascading effects
- **Financial system collapse:** Systemic failure of financial infrastructure
- **Major conflict escalation:** Armed conflict with regional or global implications
- **Critical infrastructure failure:** Failure of essential services (power, water, communications)

### D.3 Probability Floor (Canonical)

**Minimum category probability floor:  $p_{\text{floor}} \geq 0.02$**

**Meaning:** The sum of probabilities of scenarios in each mandatory category must be  $\geq 0.02$  unless explicitly justified as implausible for the specific decision context per Section 8.3.2.

### D.4 Scenario Set Size Requirements (Canonical)

| Tier | TRC Required? | Minimum |S| | |-----|-----|-----| | Tier 1 | Optional  
(qualitative screen) | N/A | | Tier 2 | Required when catastrophe relevance plausible |  $\geq 5$   
minimum | | Tier 3 | Required |  $\geq 20$  minimum | | Tier 4 | Design target only | N/A |

### D.5 Scenario Definition Template (Normative Minimum Fields)

Each scenario record in PCC MUST include:

Field	Description
Scenario ID	Unique identifier within the run
Name	Descriptive name
Category	One of the mandatory categories or "baseline/other"
Narrative	2-5 sentences describing the scenario
Time horizon	Timing assumptions
Key stressors	What breaks, how, and cascading effects
Parameterization	What changes in impacts/likelihoods under this scenario
Probability $p_s$	Assigned probability with provenance
Provenance	Data source, model, expert elicitation, or governance prior

### D.6 Loss Construction (Bounded-Impact Mode, Tier 1-3 Canon)

For each scenario s:

$$L(a,s) = \sum_{c \in C\_cat} \omega_c \times (-I\_prop\_c(a|s))^+$$

with  $\sum_{c \in C\_cat} \omega_c = 1$ .

**Default catastrophe weights:** Uniform over  $C\_cat$  unless otherwise governed.

### D.7 Discrete CVaR Algorithm (Canonical; Tier 1-3)

**Inputs:** Losses  $L(a,s_1), \dots, L(a,s_n)$ ; probabilities  $p_1, \dots, p_n$ ; tail level  $\alpha$ .

**Algorithm:**

1. Let  $\beta = 1 - \alpha$  (tail mass to capture)
2. Sort scenarios in descending loss order:  $L(a,s_{\{1\}}) \geq L(a,s_{\{2\}}) \geq \dots \geq L(a,s_{\{n\}})$
3. Let  $p_{\{i\}}$  be the probability of the  $i$ -th sorted scenario
4. Find the smallest index  $k^*$  such that:  $\sum_{i=1}^{k^*} p_{\{i\}} \geq \beta$
5. Define tail mass before cutoff:  $P_{\{k^*-1\}} = \sum_{i=1}^{k^*-1} p_{\{i\}}$
6. Define remaining tail portion:  $\delta = \beta - P_{\{k^*-1\}}$
7. Compute CVaR:

$$CVaR_\alpha[L(a)] = (1/\beta) \times [\sum_{i=1}^{k^*-1} p_{\{i\}} \times L(a,s_{\{i\}}) + \delta \times L(a,s_{\{k\}})]$$

**Audit requirement:** PCC MUST include the sorted loss table, the computed  $k^*$ ,  $P_{\{k^*-1\}}$ ,  $\delta$ , and final  $CVaR_\alpha$ .

### D.8 TRC Corridor Check (Canonical)

**TRC(a) = TRUE if and only if  $CVaR_\alpha[L(a)] \leq \tau\_TRC$**

### D.9 Default TRC Parameters by Context (Canonical Defaults)

Context	$\alpha$ (tail level)	$\tau\_TRC$ (threshold)
Personal	0.90	0.30
Organizational	0.95	0.20
Reversible policy	0.95	0.15

Context	$\alpha$ (tail level)	$\tau_{TRC}$ (threshold)
Irreversible policy	0.99	0.10
Existential risk	0.999	0.05

#### D.10 Tier 2 Strict TRC Rule (Canonical)

If catastrophe relevance is plausible (per Section 8.1.1), Tier 2 MUST:

- Compute TRC using this appendix
- Use  $\geq 5$  scenarios minimum
- Include mandatory tail categories with  $p_{\text{floor}} \geq 0.02$  per category unless explicitly implausible with documented justification per Section 8.3.2
- Record the full scenario table and CVaR computation in PCC

#### D.11 TRC Fallback Procedure (Canonical)

If  $A_{NCRC} \neq \emptyset$  but  $A_{adm} = \emptyset$ :

- Rank all  $A_{NCRC}$  options by  $CVaR_{\alpha}$  ascending
- Select the option with minimal  $CVaR_{\alpha}$
- Require a time-bound risk mitigation plan
- Require one-tier-higher approval
- Record  $TRC\_FALLBACK\_INVOKED$  in PCC with deficit ( $CVaR - \tau_{TRC}$ )

End Appendix D.

---

### APPENDIX E: UCI OPERATIONALIZATION PACK (TIER 1-3)

Appendix E is normative for the UCI interface and Tier 3 constraints; indicator families are canonical guidance unless a deployment supplies validated instruments.

#### E.1 UCI Component Definitions (Canonical)

Component	Symbol	Description
Cohesion	H <sub>u</sub>	Internal connectivity, trust, shared identity, conflict resolution capacity
Flow	F <sub>u</sub>	Coordination throughput, information fidelity, resource allocation efficiency
Resilience	R <sub>u</sub>	Redundancy, robustness, recovery speed, adaptive capacity
Equity	E <sub>u</sub>	Fair distribution of burdens/benefits, voice representation, inclusion

## E.2 UCI Formula (Canonical)

$$UCI_u = \alpha_H \times H_u + \alpha_F \times F_u + \alpha_R \times R_u + \alpha_E \times E_u$$

**Default:**  $\alpha_H = \alpha_F = \alpha_R = \alpha_E = 0.25$ .

**Special case for Self (U<sub>1</sub>):** E<sub>1</sub> := 1 by definition. The equity component is fixed at 1 for Self because equity-as-distribution does not apply within a single individual. This provides a neutral, non-penalizing identity value.

## E.3 Structural Independence Rule (Tier 3 Binding)

UCI MUST be computed from structural/process indicators distinct from welfare indicators used for RLS.

At Tier 3:

- Deriving UCI from welfare-cell impacts is PROHIBITED
- If structural indicators are unavailable, UCI is treated as unavailable (see E.7)

## E.4 Indicator Families by Union (Canonical Guidance)

**U<sub>1</sub> Self:**

- Cohesion: Psychological integration, goal coherence, self-trust, internal conflict resolution
- Flow: Task execution reliability, attention stability, cognitive throughput proxies

- Resilience: Stress recovery, adaptive coping, flexibility under change
- Equity:  $E_1 := 1$  by default (not measured)

### **U<sub>2</sub> Household:**

- Cohesion: Relationship quality, conflict resolution, trust among members
- Flow: Resource pooling efficiency, coordination routines, decision-making speed
- Resilience: Emergency preparedness, support redundancy, recovery from shocks
- Equity: Fair burden-sharing, voice parity, caregiving distribution fairness

### **U<sub>3</sub> Community:**

- Cohesion: Social capital, trust indices, network density, belonging measures
- Flow: Collective action capacity, coordination lag, information spread fidelity
- Resilience: Mutual aid redundancy, disaster response capacity, recovery history
- Equity: Inclusion of marginalized groups, access parity, procedural fairness in local governance

### **U<sub>4</sub> Organization:**

- Cohesion: Culture trust indices, turnover stability, safety culture, shared purpose
- Flow: Process throughput, coordination efficiency, error rates, execution reliability
- Resilience: Redundancy, continuity planning, incident response maturity
- Equity: Pay fairness, promotion parity, grievance procedures, representation

### **U<sub>5</sub> Polity:**

- Cohesion: Institutional trust, social cohesion indices, legitimacy measures
- Flow: Governance effectiveness, service delivery reliability, coordination speed
- Resilience: Crisis response capacity, redundancy of critical systems
- Equity: Rule of law parity, civil rights access parity, representation integrity

### **U<sub>6</sub> Humanity/CMIU:**

- Cohesion: Cross-polity cooperation capacity, treaty adherence norms
- Flow: Global coordination throughput, information-sharing integrity

- Resilience: Global response capacity to pandemics/climate/conflict
- Equity: Burden-sharing fairness, inclusion of vulnerable polities/populations

#### **U, Biosphere:**

- Cohesion: Ecosystem connectivity, biodiversity integrity, trophic network stability
- Flow: Nutrient cycling integrity, carbon sequestration capacity, hydrological cycle stability
- Resilience: Recovery capacity, redundancy of functional species, anti-fragility markers
- Equity: Use "distributional ecosystem integrity" proxies; if not operationalized, record E7\_METHOD\_NONE

#### **E.5 $\Delta$ UCI Computation (Canonical)**

$$\Delta\text{UCI}_u(\mathbf{a}) = \text{UCI}_u(\mathbf{a}) - \text{UCI}_u(\text{baseline})$$

#### **E.6 Prospective UCI Estimation (Tier 3) (Normative Constraints)**

Tier 3 requires prospective (ex ante) estimation of UCI changes from structural indicators, not from welfare impacts. The PCC must:

- Identify structural indicators used
- Specify baseline values
- Specify predicted changes under each option
- Specify normalization to [0,1] for each component
- Compute  $\text{UCI}_u$  and  $\Delta\text{UCI}_u$  using E.2 and E.5

#### **E.7 UCI Unavailability Rule (Tier 3) (Normative)**

If structural indicators are unavailable such that UCI cannot be computed without violating E.3:

- UCI MUST be treated as unavailable for tie-break purposes
- If RLS lead is non-decisive, decision MUST escalate for more data/higher tier OR record a judgment call with label JUDGMENT\_CALL\_UCI\_UNAVAILABLE and a monitoring plan
- Any welfare-derived UCI proxy MUST NOT be used to claim Tier 3 compliance

## E.8 Minimal Measurement Protocol (Canonical Guidance)

For each UCI component, the PCC SHOULD document:

- Indicator(s) used
- Data source
- Normalization method
- Expected directionality
- Uncertainty or reliability notes
- Review cadence

End Appendix E.

---

## APPENDIX F: FAILURE MODES AND ANTI-GAMING CONTROLS (v7.4.5)

Appendix F is normative for identified gaming vectors and required mitigations.

### F.1 Identified Gaming Vectors (Canonical)

Vector	Description	Structural Mitigation
Option set manipulation	Excluding viable alternatives so preferred option "wins"	Minimum option requirement ( $\geq 2$ ); independent challenger in emergencies; document option generation process
Masking abuse	Masking unfavorable cells to inflate RLS	Non-maskable cells (rights + catastrophe) must remain unmasked; audit flags invalidate PCC
Subgroup erasure	Averaging away harms to vulnerable groups	Worst-off subgroup operator for rights; subgroup enumeration; conservative bound when infeasible

Vector	Description	Structural Mitigation
Scenario omission	Leaving out unfavorable tail scenarios	Mandatory tail categories and probability floors; scenario count minima; implausibility test per Section 8.3.2
Probability gaming	Assigning implausibly small p_s to tails	Probability provenance required; p_floor per category; sensitivity perturbations at Tier 3
Kernel capture	Tweaking kernel edges to favor preferred option	KQS bands; sensitivity perturbations; evidence classes; humility fallback to NONE
Weight capture	Steering HDW to underweight certain unions/dimensions	Constitutional floors; supermajority locks near floors; transparency ledger; conflict disclosure
Emergency abuse	Declaring emergency to bypass rights	Challenger requirement; remediation plan; public disclosure/audit triggers
Confidence inflation	Overstating confidence to increase impact weight	c_k bounded [0.1, 1]; require evidence notes; sensitivity recommended
Horizon manipulation	Choosing time horizons strategically	Logarithmic temporal weighting $\tau(t)$ ; declare T_ref; challenge horizon assumptions
Unknown-as-neutral	Leaving cells blank to avoid negatives	Ignorance penalty phantom instance rule; audit flag if non-maskable cell missing

## **F.2 Structural Anti-Gaming Features (Canonical)**

### **Non-Compensatory Architecture:**

- NCRC cannot be overridden by welfare gains
- TRC cannot be overridden by expected value
- Containment cannot be bypassed by RLS optimization

### **Transparency Requirements:**

- All parameters recorded in PCC
- Registry hashing prevents silent modification (Tier 3+)
- Public rationale (5SPR) required

### **Audit Infrastructure:**

- Random audit lottery (recommended 5% of PCCs at Tier 3)
- Red-team protocols for high-stakes decisions
- Configuration drift detection

### **Learning Accountability:**

- NCAR loop tracks prediction accuracy
- Systematic errors trigger parameter review
- Historical PCCs preserved for accountability

## **F.3 Audit Flag Taxonomy (Canonical; Consolidated)**

Authoritative audit flags are specified in Section 14.3 and duplicated verbatim in Appendix H.3. Appendix F provides explanatory context and anti-gaming rationale only; it does not define additional flags.

### **Invalidity Flags (PCC Invalid):**

Flag	Trigger	Required Action
RIGHTS_CELL_MASKED_INVALID	Any rights-covered cell masked	PCC invalid
CATASTROPHE_CELL_MASKED_INVALID	Any catastrophe cell masked	PCC invalid
CONTAINMENT_MODE_B_USED_FOR_SELECTION	Mode B influenced selection	PCC invalid

**Mandatory Escalation / Tier Downgrade Flags:**

**Required Review Flags:**

Flag	Trigger	Required Action
EMERGENCY_MODE_INVOKED	A_NCRC empty	Remediation plan plus high scrutiny
TRC_FALLBACK_INVOKED	A_adm empty after TRC	Higher-tier approval plus mitigation plan
CHALLENGE_DEFERRED_EMERGENCY	Challenge r deferred	Retrospective

Flag	Trigger	Required Action
		challenge within 24h
SUBGROUP_LIMITATION	Cannot disaggregate subgroups	Apply y_subgroup ; escalate if high stakes
JUDGMENT_CALL_UCI_UNAVAILABLE	UCI unavailable in non-decisive tie	Monitoring plan plus escalation consideration
JUDGMENT_CALL_TIEBREAK_NONDECISIVE	Tie-break unresolved	Monitoring plan plus escalation consideration
KERNEL_HUMILITY_FALLBACK	Kernel disabled due to KQS/stability	Note limitation; plan kernel improvement
MANDATORY_TAIL_CATEGORY_MISSING	Missing tail category without justification	Add scenarios or justify per Section 8.3.2

Flag	Trigger	Required Action
MANDATORY_TAIL_CATEGORY_MISSING_WITH_JUSTIFICATION	Missing tail category with proper justification	Record justification; monitor
MANDATORY_TAIL_PROB_FLOOR_VIOLATION	Tail floor not met	Revise probabilities or scenario set
CONFIG_DRIFT	Inconsistent parameters without governance update	Governance review
REGISTRY_MISMATCH	Snapshot differs from referenced registry	Audit required

**F.4 Red-Team and Audit Lottery (Tier 3 Recommended)**

Deployments SHOULD adopt:

- **Random audit lottery:** for example, 5% of PCCs selected for independent review
- **Red-team scenario additions:** Adversarial scenario injection for TRC testing
- **Kernel perturbation stress tests:** Systematic edge perturbation

- **Rights near-miss reviews:** Options near thresholds flagged for monitoring

All governance rules must be documented and versioned.

End Appendix F.

---

## **APPENDIX G: SGP INTEGRATION BINDING (NORMATIVE INTERFACE)**

Appendix G is normative for the Ripple\_Logic to SGP interface.

### **G.1 Binding Purpose**

This appendix defines the only permissible interface by which Ripple\_Logic consumes Sentience Gradient Protocol (SGP) outputs. Ripple\_Logic does not define SGP internals.

ProofPack dependency pinning (normative): Any ProofPack claiming replayability **MUST** include the exact SGP v4.1.1 artifact (PDF or equivalent) as a bundled file, or **MUST** reference a registry hash that uniquely identifies the SGP v4.1.1 artifact. If SGP is not bundled or hash-pinned, ProofPack runs **MUST** declare SGP as an external dependency and restrict test vectors to cases where  $SG\_norm(E)$  is provided as an input.

### **G.2 Canonical Scalar Mapping**

Given entity E, SGP v4.1.1 outputs three pillar scores in  $[0, 100]$  and an authoritative normalized scalar  $SG\_norm(E) \in [0, 1]$ .

- A(E): Awareness pillar
- B(E): Agency pillar
- C(E): Union Participation pillar

Authoritative binding (interface-only):

$SG\_norm(E)$  is consumed as provided by SGP v4.1.1 and **MUST NOT** be recomputed inside Ripple\_Logic.

A(E), B(E), C(E) **MAY** be logged for interpretability and auditing, but they do not define  $SG\_norm(E)$  within this specification.

### **G.3 Human Plateau Rule (Non-Overridable)**

For any human person H:

**$SG\_norm(H) := 1.0$**

This assignment is:

- Independent of measurement noise
- Independent of disability status
- Independent of partial observability
- Non-overridable by any SGP scoring outcome
- MUST NOT be weakened by any weighting scheme

#### **G.4 Permitted Usage in Ripple\_Logic**

**Permitted:**

- $s_k := SG\_norm(E)$  as an impact-instance multiplier in Section 5, only for non-human or non-plateau entities, when ethically relevant and documented in PCC

#### **G.5 Prohibited Usage**

SGP outputs MUST NOT be used to:

- Weaken rights checks for humans
- Justify coercion via "low sentience"
- Assign governance authority based on sentience alone
- Override NCRC, TRC, or containment gates
- Treat rights floors as compensable welfare terms

#### **G.6 Misinterpretation Guard (Normative)**

The following inferences are PROHIBITED:

<b>Claim</b>	<b>Status</b>
Linguistic fluency implies Sentience	NOT VALID
Self-report implies Sufficient evidence of sentience	NOT VALID
Intelligence implies Moral status	NOT VALID

Claim	Status
Sentience implies Governance authority	NOT VALID
Precaution implies Attribution	NOT VALID
Low SGP score implies Reduced human protections	PROHIBITED

### **G.7 Required PCC Fields When SGP Is Used**

If any s\_k is used, PCC MUST include:

- Entity identifier/class, scope, and why it is a stakeholder
- SGP pillar outputs A, B, C
- SGP\_score and SG\_norm
- Evidence class / confidence from SGP process (as available)
- Explicit statement that Human Plateau Rule was not applied to reduce any human protections

### **G.8 Rights Expansion for Non-Human Stakeholder Classes (Normative Process)**

If SGP (under governed procedures) establishes that a non-human stakeholder class warrants rights-of-protection, Ripple\_Logic rights coverage sets and protection rules may be expanded only through governed updates:

- Update rights coverage sets C\_r and subgroup protocols
- Document version increment
- Record changes in PCC for subsequent runs
- Preserve prior PCCs unchanged (no retroactive modification)

End Appendix G.

---

## **APPENDIX H: PCC TEMPLATE AND AUDIT FLAGS CANON PACK (TIER 2-3)**

Appendix H is normative for PCC fields and audit flags at Tier 2-3.

### **H.1 PCC Schema (Human-Readable)**

=====

PROVENANCE AND COMPLIANCE CERTIFICATE (PCC)

Ripple\_Logic Framework v7.4.5

=====

**HEADER (REQUIRED Tier 2-3)**

-----

Decision ID: [unique identifier]

Timestamp (UTC): [YYYY-MM-DDThh:mm:ssZ]

Decision Owner(s): [names/roles]

Analyst(s): [names/roles]

Reviewer (Tier 3): [name/role]

Spec Version: Ripple\_Logic v7.4.5

Implementation Tier: [1 | 2 | 3]

Propagation Mode: [NONE | QUICK]

Kernel Profile: [NONE | profile name + KQS + evidence note]

Weight Profile: [HDW | Declared interim | Uniform]

TRC Context Class: [Personal | Organizational | Reversible |  
Irreversible | Existential]

Content Hash: [optional Tier 2; required Tier 3; SHA-256]

**SCOPE (REQUIRED Tier 2-3)**

-----

Decision Question: [clear statement]

Time Horizon: [years]

Primary Affected Unions: [list]

Decision Boundary Notes: [what is in/out of scope; why]

Option Set O:

Option A: [description]

Option B: [description]

Option C: [description] (if applicable)

Applicability Mask  $m(u,d)$  (RLS only):

Masked cells: [list + rationale per cell]

Non-maskable verification:

Rights cells unmasked: [PASS/FAIL]

Catastrophe cells unmasked: [PASS/FAIL]

## **INPUTS (REQUIRED Tier 2-3)**

---

Weights:

Union weights  $w_u$ : [vector or table]

Dimension weights  $v_d$ : [vector or table]

HDW parameters (if used): [ $\lambda_U, \lambda_D$ ; sources]

Rights Canon:

Thresholds  $\theta_r$ : [table or "v7.4.5 canonical"]

Coverage sets  $C_r$ : [table or "v7.4.5 canonical"]

Subgroup policy:

$G_{\{u,d\}}$  method: [enumerated | conservative bound]

$\gamma_{\text{subgroup}}$  (if used): [value]

Notes: [limitations if any]

TRC Inputs (REQUIRED Tier 2 when plausible; REQUIRED Tier 3):

Catastrophe cells  $C_{cat}$ : [list]

Catastrophe weights  $\omega$ : [list]

$\alpha$ : [value]

$\tau_{TRC}$ : [value]

Scenario set S:

For each scenario s:

ID, category, narrative,  $p_s$ , parameterization notes

Mandatory tails check:

Categories present: [YES/NO per category]

Probability floors met: [YES/NO]

Any implausibility justifications: [YES/NO; if YES, details]

Containment Inputs (REQUIRED Tier 3):

$\tau_c$ : [value]

$\theta_{pos}$ : [value]

$D_c$ : [value]

UCI method: [indicators used + sources]

### **IMPACT ESTIMATION (REQUIRED Tier 2-3)**

-----

For each option a:

Direct impacts  $I_{dir}(u,d,a)$ : [summary table or key cells]

Propagated impacts  $I_{prop}(u,d,a)$ : [summary table or key cells]

Key impact instances (at least top contributors):

For each listed instance k:

(u,d),  $\mu_k$ , r\_k, t\_k,  $\ell_k$ , c\_k, e\_k, s\_k, rationale/source

Missing data:

Any UNKNOWN\_IMPACT cells: [list]

Phantom instances applied: [YES/NO]

### **CASCADE TRACE (REQUIRED Tier 2-3)**

-----

NCRC:

For each option: PASS/FAIL

If FAIL: violated rights + v\_r(a) values

Emergency Mode invoked: [YES/NO]

Challenger conducted: [YES/NO]

Challenger notes: [summary]

TRC:

Triggered (catastrophe relevance plausible): [YES/NO]

For each option:

CVaR\_ $\alpha$ : [value]

Corridor  $\tau_{TRC}$ : [value]

PASS/FAIL

TRC fallback invoked: [YES/NO]

Containment (Tier 3 required):

For each option in A\_adm:

U\_pos(a): [list of unions]

For each u in U\_pos: min  $\Delta$ UCI among ancestors: [value]

PASS/FAIL

A\_sel: [list]

RLS Ranking:

For each option in A\_sel:

RLS(a): [value]

$\sigma$ \_RLS(a): [value] (Tier 3 required)

Gap(top, second): [value]

$\delta$ : [value]

Decisive: [YES/NO]

Tie-break (if non-decisive):

UCI method used: [UCI-A/B/C]

$\Delta$ \_UCI threshold: 0.05

Result: [winner or non-decisive]

HOI flags (if applicable): [notes]

Escalation invoked: [YES/NO]

### **SELECTION (REQUIRED Tier 2-3)**

-----

Selected Option: [option ID]

Selection Rationale: [brief; must reference cascade outcomes]

Monitoring Plan (NCAR):

Metrics: [list]

Review schedule: [date/interval]

Triggers for re-run: [conditions]

### **SENSITIVITY ANALYSIS (Tier 3 REQUIRED)**

---

Weights perturbation: [summary; selection stable?]

Threshold perturbation: [summary; any near-miss rights?]

Kernel perturbation: [summary; flag fragile?]

Scenario perturbation: [summary; probability sensitivity]

Robustness classification: [Robust | Sensitive | Fragile]

### **FIVE-SENTENCE PUBLIC RATIONALE (5SPR) (Tier 2-3 REQUIRED)**

---

1. CONTEXT: [What decision was made and why now?]

2. OPTIONS: [What options were considered?]

3. CONSTRAINTS: [What was eliminated by NCRC/TRC (and why)?]

4. SELECTION: [Why the chosen option won among selectable options?]

5. MONITORING: [What follow-up will be tracked and when?]

### **AUDIT FLAGS (REQUIRED when triggered)**

---

[List flags + explanations]

### **SIGNATURES**

---

Analyst: [name, date]

Decision Owner: [name, date]

Reviewer (Tier 3): [name, date]

=====

## H.2 Minimum Required Fields by Tier (Canonical)

Field	Tier 1	Tier 2	Tier 3
Decision ID	Recommended	REQUIRED	REQUIRED
Timestamp	Recommended	REQUIRED	REQUIRED
Option set	REQUIRED	REQUIRED	REQUIRED
NCRC results	Heuristic ok	REQUIRED	REQUIRED + subgroup policy
TRC results	Optional	REQUIRED when catastrophe plausible	REQUIRED
Containment	Optional	Recommended	REQUIRED
RLS scores	Optional	REQUIRED	REQUIRED
Uncertainty + discrimination	Optional	Recommended	REQUIRED
Sensitivity analysis	Optional	Recommended	REQUIRED
5SPR	Optional	REQUIRED	REQUIRED

Field	Tier 1	Tier 2	Tier 3
Audit flags	Optional	REQUIRED when triggered	REQUIRED when triggered

### H.3 Audit Flag Canon Pack (Canonical)

Flag	Trigger	Required Action	Severity
RIGHTS_CELL_MASKED_INVALID	Any rights-covered cell masked in RLS aggregation	PCC invalid; recompute without masking	INVALID
CATASTROPHE_CELL_MASKED_INVALID	Any catastrophe cell masked in TRC computation	PCC invalid; recompute without masking	INVALID
CONTAINMENT_MODE_B_USED_FOR_SELECTION	Mode B influenced selection outcome	PCC invalid; rerun with Mode A only	INVALID
SCENARIO_LIBRARY_MIN_EXCEPTION	Scenario library	PCC remains	ESCALATE

Flag	Trigger	Required Action	Severity
	<p>size  S  is below the tier minimum when TRC is required (Tier 2:  S  &lt; 5; Tier 3:  S  &lt; 20)</p>	<p>valid only with: (i) written justification, (ii) independent reviewer sign-off (Tier 3 required; Tier 2 recommended), and (iii) a remediation plan to reach minimum scenario count before the next comparable run. Record this flag in PCC.</p>	
<p>MANDATORY_TAIL_CATEGORY_MISSING</p>	<p>Missing mandatory tail category without</p>	<p>Add scenarios or justify per</p>	<p>ESCALATE</p>

Flag	Trigger	Required Action	Severity
	justification	Section 8.3.2	
MANDATORY_TAIL_CATEGORY_MISSING_WITH_JUSTIFICATION	Missing mandatory tail category with proper justification per Section 8.3.2	Record justification; monitor for category emergence	REVIEW
MANDATORY_TAIL_PROB_FLOOR_VIOLATION	Category probability sum < p_floor (0.02)	Revise probabilities or add scenarios to category	ESCALATE
SUBGROUP_LIMITATION	Cannot disaggregate subgroups for rights-covered cell	Apply subgroup conservative bound; escalate if high stakes	REVIEW
CHALLENGE_DEFERRED_EMERGENCY	Emergency Mode invoked without	Retrospective challenge MUST	ESCALATE

Flag	Trigger	Required Action	Severity
	independent challenger	occur within 24 hours	
EMERGENCY_MODE_INVOKED	A_NCRC = ∅ (no rights-admissible options)	Remediation plan required; high scrutiny review	ESCALATE
TRC_FALLBACK_INVOKED	A_adm = ∅ after TRC (rights-safe but all fail TRC)	Higher-tier approval required; mitigation plan	ESCALATE
DECISION_FRAGILE_KERNEL	Selected option changes under kernel perturbation	Escalation required; consider NONE mode	ESCALATE
KERNEL_HUMILITY_FALLBACK	Kernel disabled due to KQS < 0.40 or	Note limitation in PCC; plan kernel	REVIEW

Flag	Trigger	Required Action	Severity
	stability violation	improvement	
JUDGMENT_CALL_UCI_UNAVAILABLE	UCI unavailable in non-decisive RLS tie	Monitoring plan required; escalation consideration	REVIEW
JUDGMENT_CALL_TIEBREAK_NONDECISIVE	Tie-break chain did not resolve selection	Monitoring plan required; escalation consideration	REVIEW
CONFIG_DRIFT	Parameters differ from prior run without governance update	Governance review required before proceeding	ESCALATE
REGISTRY_MISMATCH	PCC snapshot differs from referenced	Audit required; resolve discrepancy	ESCALATE

Flag	Trigger	Required Action	Severity
	registry hash		

End Appendix H.

## APPENDIX I: TIER-4 DESIGN TARGET (NOT CLAIMABLE IN v7.4.5)

### I.1 Status Declaration

Tier-4 Design Target (Not Claimable in v7.4.5)

Tier 4 is specified as a target profile requiring artifacts and capabilities that are not yet publicly available. Tier-4 compliance MUST NOT be claimed until ProofPack is publicly released and independently replayable.

### I.2 Tier 4 Requirements (Target Specification)

#### Hash-Bound Registries:

- Rights anchors registry (REG-RIGHTS-ANCHORS-v1)
- Catastrophe indicators registry (AF-BASE)
- Scenario library registry
- Kernel registry with evidence classes
- HDW ballots and weights registry

#### Deterministic Numeric Profile (NDP\_FIXEDPOINT\_V1):

- Fixed-point scale:  $S = 10^9$
- Representation:  $X_{fp} = \text{round\_half\_even}(x \times S)$  as signed int64
- All divisions use round-half-even
- Hard-fail on overflow with audit\_flag NUMERIC\_OVERFLOW

#### Saturation Lookup Table (SAT\_LUT\_FP\_V1):

- Pre-computed tanh values for fixed-point inputs
- Hash-bound in ProofPack
- Runtime computation of tanh PROHIBITED

#### **Temporal Weight Registry (REG\_TEMPORAL\_WEIGHTS\_V1):**

- Pre-computed  $\tau(t)$  values for standard horizons
- Hash-bound in ProofPack
- Runtime computation of logarithms PROHIBITED

#### **Canonical JSON Profile:**

- NO\_FLOATS rule: All numeric values as exact rationals {num, den}
- Deterministic key ordering
- Canonical whitespace rules
- SHA-256 hashing for integrity

#### **Propagation Mode Restriction:**

- FULL propagation PROHIBITED for Tier 4 Pilot-Executable
- NONE or QUICK only
- FULL available only in future Tier 4 Certified profile

### **I.3 ProofPack Contents (Target)**

The ProofPack artifact bundle will provide:

- **Schemas:** JSON schemas for PCC and all registries
- **Manifests:** Hash indices for all artifacts
- **Canonicalization Rules:** Exact specification for deterministic hashing
- **Registries:** Machine-readable canonical registries
- **Test Vectors:** Reference inputs and expected outputs for validation

### **I.4 Replay Test Requirement**

A Tier 4 PCC MUST pass a replay test: independent re-execution using PCC inputs and referenced registries reproduces:

- Identical admissibility outcomes (NCRC, TRC, Containment)
- Identical RLS ranking
- Identical selected option

Failed replay invalidates the Tier 4 claim.

End Appendix I.

---

## **APPENDIX J: VALIDATION PROTOCOL PACK**

Appendix J is normative for pre-registration templates and scoring definitions when validation claims are made.

### **J.1 Pre-Registration Template (Normative)**

=====  
RIPPLE\_LOGIC VALIDATION STUDY PRE-REGISTRATION  
=====

#### **STUDY IDENTIFICATION**

-----

Study ID: [unique]

Registration date: [YYYY-MM-DD]

Principal investigator: [name, affiliation]

Funding source: [if any]

Ripple\_Logic version: [v7.4.5]

Tier executed: [1|2|3]

Domain: [AI deployment | policy | org strategy | etc.]

#### **HYPOTHESES**

-----

Primary hypothesis: [H1/H2/etc. with exact measurable claim]

Secondary hypotheses: [list]

Falsification criteria: [explicit thresholds from Section 17]

## **METHODS**

---

Design: [pilot / RCT / quasi-experimental / backtest]

Sample: [population, size, selection method]

Comparator: [baseline governance / alternative framework]

Duration: [time]

Decision types included: [list]

Exclusion criteria: [list]

## **MEASURES**

---

Primary outcomes: [definition, data source, timing]

Secondary outcomes: [definition, data source, timing]

Rights measurement: [how rights violations are verified]

Tail-loss measurement: [how catastrophic outcomes are defined]

## **ANALYSIS PLAN**

---

Statistical methods: [tests/models]

Effect size thresholds: [minimum meaningful difference]

Multiple comparison correction: [method]

Missing data handling: [approach]

Sensitivity checks: [weights, scenarios, thresholds]

**GOVERNANCE**

-----

Ethics approval: [status/ID]

Data management: [storage, access, retention]

Reporting commitment: [timeline and venue]

=====

**J.2 Outcome Measures by Union and Dimension (Canonical Examples)**

Union	Dimension	Measure Type	Example Indicators
Self	Health	Validated survey / admin	SF-36, PHQ-9, GAD-7, injury rates
Self	Agency	Validated survey	Autonomy/self-efficacy scales
Household	Material	Administrative	Income stability, housing security
Community	Social	Survey / network	Trust indices, social capital measures
Organization	Flow/Resilience (UCI)	Operational	Throughput, incident recovery time
Polity	Agency	Index	Participation rates, rule-of-law indices

Union	Dimension	Measure Type	Example Indicators
Biosphere	Environment	Monitoring	Emissions, biodiversity integrity

### J.3 Sign Accuracy Scoring (Normative)

For each evaluated cell (u,d):

**Predicted sign:**  $\text{sign}(l_{\text{prop}}(u,d,a)) \in \{-1, 0, +1\}$

**Observed sign:**  $\text{sign}(\Delta_{\text{observed}}(u,d)) \in \{-1, 0, +1\}$

**Scoring:**

- Match = 1 if predicted = observed
- Match = 0.5 if predicted = 0 or observed = 0
- Match = 0 otherwise

**SignAcc =  $(\sum \text{Match}) / (\text{number of evaluated cells})$**

**Target for calibrated deployments:** SignAcc  $\geq 0.70$

### J.4 Magnitude Error (RMSE) (Normative)

For cells with quantitative observed comparisons on  $[-1, +1]$ :

**RMSE =  $\sqrt{[(1/n) \times \sum_i (l_{\text{prop}_i} - \text{observed}_i)^2]}$**

**Default target:** RMSE  $\leq 0.25$  (domain dependent)

### J.5 Backtest Procedure (Normative)

1. **Case selection:** Identify historical decisions with measurable outcomes and reconstructable option set
2. **Information reconstruction:** Use only data available at decision time (no hindsight leakage)
3. **Blind analysis:** Analyst is unaware of actual outcomes during run
4. Execute Ripple\_Logic run (Tier declared) and emit Backtest PCC
5. Compare predictions to realized outcomes (SignAcc, RMSE, rights detection)
6. Report results and failures; update methods via NCAR with versioning

End Appendix J.

---

## **APPENDIX K: STARTER KOPS / STARTER KERNEL PACK (PROVISIONAL; TIER 2-3)**

Appendix K is normative for the existence of a starter kernel pack and for required disclosures/sensitivity when it is used. Edge values are provisional and **MUST** be labeled as such.

### **K.1 Purpose (Normative)**

The Starter KOPS (Key Operational Pathways Set) provides a conservative, interpretable ripple kernel profile for early implementations. It is **PROVISIONAL** and requires sensitivity analysis; it must not be treated as validated truth.

### **K.2 Evidence Classes (Normative)**

<b>Class</b>	<b>Name</b>	<b>Minimum Evidence</b>
A	Validated	Replicated empirical studies with quantified effect sizes and domain match
B	Supported	Multiple studies consistent in sign; some quantification
C	Plausible	Theoretical support with limited empirical evidence
D	Speculative	Expert judgment without empirical support
E	Elicited	Structured elicitation for this framework (lowest confidence)

### **K.3 Global Shrink Factor (Normative)**

Any Class E (elicited) edge **MUST** be multiplied by a global shrink factor:

**Default shrink: 0.35**

This reduces overconfidence and supports humility.

### **K.4 Starter Edge Set (PROVISIONAL; Non-Normative Coefficients, Normative Structure)**

**Kernel convention reminder (normative):** source is column, target is row.

Source Cell (u,d)	Target Cell (u,d)	$\kappa_{raw}$	$\kappa_{shrunk}$	Evidence Class	Rationale Family
(4,1) Org-Material	(2,1) HH-Material	0.40	0.14	E	Employment income propagation
(4,1) Org-Material	(3,1) Comm-Material	0.30	0.11	E	Local multiplier effects
(7,7) Bio-Environment	(6,2) CMIU-Health	-0.35	-0.12	B	Climate/health pathway
(7,7) Bio-Environment	(6,7) CMIU-Environment	0.50	0.18	B	Earth system propagation
(5,5) Polity-Agency	(4,5) Org-Agency	0.25	0.09	E	Institutional enabling effects
(3,3) Comm-Social	(1,3) Self-Social	0.45	0.16	B	Social capital / wellbeing
(3,3) Comm-Social	(2,3) HH-Social	0.35	0.12	E	Family-community linkage

Source Cell (u,d)	Target Cell (u,d)	$\kappa_{raw}$	$\kappa_{shrunk}$	Evidence Class	Rationale Family
(4,2) Org-Health	(1,2) Self-Health	0.40	0.14	B	Occupational health pathway
(6,4) CMIU-Knowledge	(5,4) Polity-Knowledge	0.30	0.11	E	Diffusion / coordination knowledge
(1,2) Self-Health	(2,1) HH-Material	-0.25	-0.09	B	Health to earning capacity

**Implementation rule:** Only  $\kappa_{shrunk}$  values may be used for Class E edges. For Class B edges,  $\kappa_{shrunk}$  may equal  $\kappa_{raw}$  unless a deployment chooses uniform shrink.

### K.5 Sensitivity Requirements (Normative)

#### If Starter KOPS is used in Tier 3:

- Perturb each relied-upon non-zero edge by  $\pm 0.05$  one-at-a-time
- Recompute the cascade
- If admissibility or selection changes, set audit\_flag DECISION\_FRAGILE\_KERNEL and escalate

#### If Starter KOPS is used in Tier 2:

- Sensitivity is recommended; at minimum disclose kernel provisional status

### K.6 QQS Declaration (Normative)

Any PCC using a kernel MUST declare QQS and the band consequences:

- $QQS < 0.40$ : MUST use NONE
- $0.40-0.50$ : QUICK only + mandatory sensitivity (Tier 3)
- $\geq 0.50$ : QUICK permitted + sensitivity

**Starter KOPS default QQS suggestion:** 0.40-0.49 until validated.

## K.7 Disclaimer (Normative)

Any PCC using Starter KOPS MUST include this disclaimer:

"Ripple propagation used a PROVISIONAL Starter KOPS kernel. Coefficients are not fully validated. Sensitivity analysis was performed per Ripple\_Logic v7.4.5 requirements; conclusions are contingent on kernel uncertainty."

End Appendix K.

---

## APPENDIX L: GLOSSARY (v7.4.5)

Appendix L is informative but standardized terminology is strongly recommended for audit consistency.

Term	Definition
Admissible	Option that passes both NCRC and TRC
AIL (Artifact Integrity Law)	Integrity rules ensuring auditability, comparability, and no silent overrides in PCCs
Baseline-Zero Rule	0 impact means no change from baseline; all impacts are baseline-relative
Catastrophe cell set (C_cat)	Subset of welfare cells used for TRC tail-risk evaluation
Containment (Mode A)	Binding gate preventing selection of options that degrade containing union coherence beyond tolerance
Containment tolerance ( $\tau_c$ )	Maximum allowed negative $\Delta UCI$ for containing unions when sub-unions benefit
CVaR (Conditional Value-at-Risk)	Tail risk measure of expected loss in the worst tail mass

Term	Definition
Dimension (D <sub>1</sub> -D <sub>7</sub> )	One of seven welfare dimensions: Material, Health, Social, Knowledge, Agency, Meaning, Environment
Emergency Mode	Controlled fallback when no option passes NCRC; selects lexicographically minimal rights violation vector with challenger requirement and remediation
HDW (Hybrid Democratic Weighting)	Governance method producing weights from floors + democratic + structural proposals
HOI (Hollowing-Out Index)	Monitoring diagnostic indicating welfare-up while coherence-down drift
Impact instance	A single asserted pathway contribution to a welfare cell impact, with magnitude, reach, horizon, likelihood, confidence, and multipliers
Kernel (K)	Sparse matrix encoding cross-cell ripple propagation
KOPS	Key Operational Pathways Set; the subset of kernel edges considered load-bearing and governed
KQS (Kernel Quality Score)	Summary score indicating whether kernel propagation is permitted
Lexicographic cascade	Priority-ordered gates where earlier failures cannot be compensated by later scoring
MNA (Minimal Normative Axiom)	Sentient flourishing matters; unnecessary suffering should be reduced; enabling conditions preserved

Term	Definition
NCAR	Notice-Choose-Act-Reflect learning loop
NCRC	Non-Compensatory Rights Constraint; excludes rights-violating options
Non-maskable cells	Rights and catastrophe cells (and minimum coverage) that cannot be excluded from RLS aggregation
PCC	Provenance and Compliance Certificate; auditable record of inputs, computations, and selection
RLS	Ripple Logic Score; weighted welfare aggregation used to rank selectable options
Selectable	Admissible option that also passes containment (Mode A)
SGP	Sentience Gradient Protocol; interface for rights-of-protection gating across entities
TRC	Tail-Risk Constraint; excludes options with unacceptable catastrophic exposure using CVaR
UCI	Union Coherence Index; structural health metric for unions
UBE	Union-Based Ethics; normative framework built on UBR + MNA
UBR	Union-Based Reality; descriptive ontology of nested interdependent unions
Union	A bounded pattern of interdependence at an organizational scale

Term	Definition
Unioning	Redesigning decisions until rights-safe, tail-safe, containment-safe, and net-positive where possible

End Appendix L.

---

## APPENDIX M: VERSION HISTORY (v7.4.5)

M.1 Ripple\_Logic v7.4.5 (Current Canonical Foundation Paper — ProofPack Readiness Patch)

Release date: 25 January 2026

Patch Type: ProofPack Readiness + Meta-Normative Clarifications

Normative Changes (limited scope):

- Section 0.1 normative hierarchy: formatting standardized (all entries now consistently bulleted); no precedence order changes
- Appendix G: added ProofPack dependency pinning rule (MUST-level requirement for SGP v4.1.1 bundling or hash-pinning in ProofPack claims)

Non-Normative Changes: Formatting and labeling cleanup only. Cascade logic, admissibility gates, thresholds, tier requirements, and scoring formulas are unchanged from v7.4.4.

v7.4.5 incorporates all v7.4.4 fixes plus the above meta-normative clarifications.

M.1.1 Ripple\_Logic v7.4.4 (Prior Release — Publication Integrity Patch)

**Release date:** 25 January 2026

Patch Type: Publication Integrity + Minor Normative Clarifications (see list below)

v7.4.4 Fixes (from consolidated reviewer feedback):

- Removed all stray drafting artifacts (prompt tokens and instruction remnants)
- Corrected corrupted Rockström et al. (2023) reference to valid APA 7 format

- Repaired all audit flag table entries with complete trigger/action/severity specifications

Clarified Appendix G as interface-only: Ripple\_Logic consumes SG\_norm(E) from SGP v4.1.1 and does not recompute the SGP aggregator.

Removed conflicting duplicate audit-flag table fragment in Appendix F.3 (authoritative tables remain in Section 14.3 and Appendix H.3).

- Added spectral radius constraint  $\rho(K) < 1$  to kernel stability guardrails (Section 6.4)
- Completed Appendix B with all explicit equations (no stub headers)
- Unified notation system throughout (consistent  $I_{dir}$ ,  $I_{prop}$  naming)
- Added  $\sigma(u,d,a)$  definition with three canonical methods A, B, C (Section 10.3, Appendix B.13)
- Added implausibility test procedure for exempting mandatory tail categories (Section 8.3.2)
- Confirmed  $E_1 := 1$  resolution for UCI Self Equity throughout all relevant sections
- Added Kahneman & Tversky (1979), Hammond et al. (1998), Floridi (2013) to references
- Cleaned audit flag tables with proper severity classification column
- Global artifact sweep completed (no AI-generated conversational tokens remain)

Normative Changes: Minor clarifications only (kernel stability guardrail  $\rho(K) < 1$ , and SGP interface binding clarified). Cascade order and admissibility gates remain unchanged.

## **M.2 Ripple\_Logic v7.4.2 (Prior Draft)**

### **Key features:**

- Initial consolidation incorporating multi-AI reviewer feedback
- UCI Self Equity contradiction resolved ( $E_1 := 1$ )
- Appendix B equations added but with some formatting issues
- Some stray artifacts remained

## **M.3 Ripple\_Logic v7.4.0/v7.4.1 (Prior Drafts)**

### **Key features:**

- Initial v7.4 consolidation from v7.3
- Contained internal inconsistency on UCI Self Equity ( $E_1 = 0$  vs  $E_1 = 1$ )
- Appendix B had stub headers without explicit equations
- Formatting artifacts and duplicate blocks present
- Corrupted reference entries

#### **M.4 Ripple\_Logic v7.3 (Prior Release)**

##### **Key features:**

- Consolidated spec-hardening from v7.0-v7.2
- Strict Tier-2 TRC posture introduced
- HDW v6.0 floors confirmed

#### **M.5 Ripple\_Logic v7.0-v7.2 (Formatting/Structure Draft Lineage)**

##### **Key features:**

- Improved formatting, symbol table, cleaned artifacts, worked example inclusion
- Some parameter divergences existed; v7.4.x resolves them by canon policy

#### **M.6 Ripple\_Logic v6.0 (Foundation Release Lineage)**

##### **Key features:**

- Tier 1-3 calculable cascade and appendices
- Kernel KQS policy, rights thresholds/coverage sets, TRC bounded-impact mode for Tier 1-3
- PCC template and audit flags

#### **M.7 MathGov v5.0i (Lineage; Pre-Rename)**

##### **Key features:**

- Prior canonical name "MathGov"
- Included deeper Tier-4 determinism architecture; v7.4.x retains Tier-4 as design target only until ProofPack is public

#### **M.8 Earlier Versions (MathGov v1.0-v4.x)**

Version	Date	Key Features
v4.x	2024	UCI/HOI safeguards; HDW anti-capture
v3.x	2023	NCRC/TRC lexicographic cascade
v2.x	2022	7×7 welfare matrix; kernel propagation
v1.0	2021	Initial Union-Based Ethics formulation

End Appendix M.

---

## DOCUMENT COMPLETION STATEMENT

This completes the RIPPLE\_LOGIC FRAMEWORK v7.4.5 Canonical Foundation Paper with all appendices (A through M) and the consolidated APA 7 references.

Document Status: Complete, internally consistent, and ProofPack-canon-ready (ProofPack pending public release)

**Tier 1-3:** Implementable from this document

**Tier 4:** Design target only; claims prohibited until ProofPack release

**Canonical Sites:** ripplelogic.org | mathgov.org

v7.4.5 ProofPack Readiness Patch Summary:

Issue Category	Status
Stray artifacts removed	Complete
References corrected (APA 7 with DOIs)	Complete
Audit flags fully specified	Complete
Appendix B equations explicit	Complete

Issue Category	Status
Notation unified	Complete
UCI Self Equity resolved ( $E_1 := 1$ )	Complete
Spectral radius constraint added	Complete
$\sigma(u,d,a)$ methods defined	Complete
Implausibility test added	Complete
Section 0.1 hierarchy formatting standardized	Complete
SGP ProofPack dependency pinning rule added	Complete
Normative changes	Meta-normative only (see M.1)

End of RIPPLE\_LOGIC FRAMEWORK v7.4.5