

Report · The ARC Theory · Laws IV and V, the ARC Persistence Law and the ARC Embedding Law · The Mathematics

The mathematics of Law IV, the ARC Persistence Law (register ARC-4, carried by P15), and Law V, the ARC Embedding Law (register ARC-5, carried by P12 and P21), both added on 24 September 2026, their propositions registered, neither law yet registered as a law or tested: conditional theorems with proofs, identities and kill conditions. It measures nothing, and nothing in it bears on the Eden Protocol.

Version 1.1 (frozen) · First published 12 October 2026 · It follows version 1.0 of 3 October 2026 · Earlier versions: [1.0](#), [0.2.4](#), [0.2.1](#) · [PDF](#)

Version 1.1, frozen on 7 October 2026. It follows version 1.0 of 3 October 2026, which stays as published; section 9 lists what changed and what was checked. A later version is a new text under a new number.

Michael Darius Eastwood. Independent AI alignment researcher, London. ORCID [0009-0004-3222-7442](#).

DOI [10.17605/OSF.IO/4B3VT](#).

This document is not a registration, not an empirical result, and not authority to run any study. [Registration versions 1.100 and 1.102](#) stay frozen byte for byte; Laws I to III stand exactly as registered ([registration version 1.102](#), lines 659, 663 and 665); no proposition is added to or moved between registers; P12 keeps no direction; and every object this document introduces beyond the registered propositions is marked PROPOSED and is not law content until the author registers it.

Three levels are kept apart throughout: identities and conditional mathematics (sections 2.2 to 2.5, Proposition IV-R of section 2.6.1, and sections 3.3 and 4), which prove consequences of assumptions and nothing about real systems; the two laws as empirical claims (sections 2.1, 2.8, 3.1 and 3.6), whose registered faces are P15, P12 and P21 and whose only verdict rules are the registered ones; and engineering interpretations, which are quarantined: nothing here bears on the Eden Protocol or the Honey Architecture, and no verdict here is a verdict on either (section 5 item 5).

What this version adds to version 0.2.1. Version 1.0 froze the working draft 0.2.4 and added no object to it: it closed one case of Proposition IV-R, named the terms of one model in section 3.4 and reworded one clause of section 2.8. This version adds no object either. The additions are those of the working draft, four of them PROPOSED: a planner that scores its plans with bounded error (section 2.5.1); the baseline, the arms and the interval that the retained fraction needs before it can be read (section 2.6.1); the balance over a finite range of depth (section 4.8); and what a new deciding instrument declares before its outcomes are read (section 4.9). It also takes out of force the law-level conventions that version 0.2.1 proposed as defaults, and keeps them on the record as proposals (sections 2.8, 3.6, 5 and 7). A PROPOSED rule is never a registered decision rule. Section 9 says what was checked for this version and what remains open.

o. Sources and status

Sources. Two public documents are cited by line or by quotation. The first is the theory registration, versions 1.100 and 1.102, DOI [10.17605/OSF.IO/P8CKQ](https://doi.org/10.17605/OSF.IO/P8CKQ): version 1.100 was registered on OSF on 8 September 2026, and version 1.102 is its update of 13 September 2026; both texts are frozen byte for byte. It is cited by proposition number or by its own quoted words, and a line number is given as "registration version 1.102, line N". The second is the statement paper, The ARC Theory, DOI [10.17605/OSF.IO/GW5MX](https://doi.org/10.17605/OSF.IO/GW5MX), quoted from its working paper version 6.7, revised 21 September 2026, and cited by section or by its own quoted words; the sentences quoted here stand unchanged in its version 7.0 of 3 October 2026. The author's working paper of 16 August 2026, which the registration itself names and credits, is named, and not quoted, where it states an idea used here.

Every registration line number in this document counts the lines of the frozen text of version 1.102. The reproduction of that text on the programme's website is laid out differently, so its line numbers differ; the quoted words are the same in both. Section 0.1 sets out which propositions each register carries, and the lines that say so.

Status. Laws IV and V were added on 24 September 2026. Law IV carries P15, in register ARC-4, and Law V carries P12 (undirected) and P21, in register ARC-5; these propositions are registered (section 0.1), and neither law is yet registered as a law or tested. Their mathematics is set out here at version 1.1: its theorems are proved under the assumptions each states, its identities are checked, and section 7 lists the fifteen questions that stay

open. Theorems IV-G and IV-H, Corollaries IV-G' and IV-H', Lemmas IV-K and IV-M' and Propositions IV-N and IV-R are logical proofs, not identities; no proof assistant is used, and they rest on the written proofs, which finite models exercise and do not replace.

What this document is. The mathematics of Law IV, the ARC Persistence Law (register ARC-4, carrying P15), and Law V, the ARC Embedding Law (register ARC-5, carrying P12 and P21): the objects, the conditional theorems with proofs, the identities that connect the two laws to the registered balance model, the registered faces quoted byte for byte, and the kill conditions. The register assignment is that of registration version 1.102, line 99 ("ARC-4, VALUE PERSISTENCE: P15 alone." and "ARC-5, THE GENESIS STRATEGY: P12 and P21."); the seven propositions in no register are those of registration version 1.102, line 101 (P5, P7, P10, P11, P13, P14, P19), and they stay there.

What it is not. It proves no empirical law: algebra proves consequences of assumptions, not that real systems satisfy them. It runs no test. It adds no registered proposition and moves none. It does not restate Laws I to III (section 6). It supplies no missing measurement, chooses no scientific parameter the author has held (section 7), and claims no novelty; prior-art distinction remains unassessed.

0.1 Which propositions each register carries

Registration version 1.102 assigns each of its twenty-two propositions to one of five registers or to none, and nothing in this document moves one. Laws IV and V, added on 24 September 2026, give names to the two persistence registers, ARC-4 and ARC-5. They leave every proposition's number, words, date, scope and refuter as registered, and they do not make version 1.102 a registration of five laws.

REGISTER, AND THE LAW IT TESTS	PROPOSITIONS IT CARRIES
ARC-1, Law I, the ARC Principle	P1, P2, P8, P9, P18
ARC-2, Law II, the ARC Co-Scaling Law	P4
ARC-3, Law III, the ARC Ceiling	P3, P6, P16, P17, P20, P22
ARC-4, Law IV, the ARC Persistence Law	P15
ARC-5, Law V, the ARC Embedding Law	P12, P21
In no register	P5, P7, P10, P11, P13, P14, P19

The assignment is that of registration version 1.102: ARC-1 at line 93, ARC-2 at line 95, ARC-3, ARC-4 and ARC-5 at line 99, and the seven in no register at line 101. Laws I to III are stated at its lines 659, 663 and 665 and stand as stated there; Law II's criterion is that β_C exceeds k . The balance exponent Δ_{bal} of section 4 is a different object from that criterion: its service elasticity χ and Law II's β_C are not shown to be one quantity, and writing one letter for the other does not make them so (section 4.3).

Each quotation in this document names its source, and each was compared with that source for this version (section 9). This version cites no work that version 0.2.1 did not.

1. Notation, fixed first

Principle: one symbol, one meaning; where the programme's notation already has the quantity, its letter (χ , ϵ , κ , ζ , α , α_{crit} , C , R , k , and its Delta written Δ_{bal} , never bare); the letter beta only with its subscript, so Law II's exponent is written β_C ; no bare delta; the registration's letter for the correction-leverage exponent, and its subscripted forms, only inside a quotation of the registration; no Delta prefix as a symbol, so a difference between arms, or the change in a quantity, is written out in full; none of the eight symbols that collide in the programme's notation (F , E , B_Q , I_Q , A , a , P_V , d_{min}) in this document's own writing, bare, subscripted or as an argument, not even as a label, with one exception, A subscripted by arm as A_V , A_{ext} , A_{emb} and A_{base} , the arm scores, so that a form written elsewhere with one of the eight is described here in words; and five clashes resolved as follows.

CLASH	RESOLUTION
A_V for both the acceptable set and the arm score	the acceptable set is Ω_V ; $A_V(x)$ and $A_{base}(x)$ are the arm scores, the symbols of the retained fraction
d_gain for the installed effect, still a d, a letter that collides in the programme's notation	the installed effect is $Y_V(x) := A_V(x) - A_{base}(x)$; its floor is $Y_{min} > 0$
P_H, P_safe, P_exit for plan sets, beside P1 to P22	plan sets are written from the path family: $T_H(s_o)$, $T_H^{safe}(s_o)$, $T_H^{exit}(s_o)$
b_ol for a cell intercept, beside b_V	the cell intercept is $\ln(Q_{ol,o}/Q_o)$ with $Q_{ol,o} := Q_{ol}(C_o, R_o)$; b is the removal benefit only
D_p(x) for the installed effect, beside D for the correction shortfall	neither kept: $Y_V(x)$ for the installed effect; the correction shortfall written ($\kappa - \chi$) in full

Further choices: the two suprema of Theorem IV-H are J_{safe}^* and J_{exit}^* (not a and e; a is one of the eight); the searched margin is M_{srch} ; the increments and levels of section 2.5 are inc_i and lev_t (not D and L, both retired letters); the increment index is i and the step index t, so u is the random effects only; the power family's scale is x_{scale} (not x_o , which the reference convention reserves); the shape function of the descriptive families is $sh(x)$ (not f, the Cauchy f of the programme's notation); the embedding ratio is G_{emb} (not G_E , whose subscript is the letter E); the two families of executed and of feasible paths are T_{acc} and T_{plan} ; the evaluated transition set is $T_{eval}(s)$; the unrelated part of an objective change is b_{other} ; the registered upper bound on the safe supremum is J_{up} and the best exit score found is J_{found} . The general root of the balance identity is α_{bal} , and α_{crit} is reserved for the case in which that root is a ceiling (section 4.2), so the name the programme's notation gives α_{crit} , the stability ceiling, stays true. The two mutual-information forms of section 3.2 and the architecture form of section 4.6, which are written elsewhere with two of the eight letters, are described in words, as is the change in the capability functional of section 2.4, which is written elsewhere with the prefix this section rules out; S_V and K_V carry their subscripts wherever a curve or a kernel is named; and the registration's letters for the correction-leverage exponent and for its checker-axis slope are printed only inside quotations.

The identification is explicit, not silent, and is conditional: χ_{ol} is the correction-leverage exponent of the registration (registration version 1.102, line 273: "THE CORRECTION EXPONENT. How corrective strength grows with the capability of the system being corrected: the slope of misalignment removed per pass against the capability of that corrected system.") only when Q_{ol} is that service measure on the target axis (registration version 1.102, lines 275 and 303 to 311: a checker-axis slope "is reported as γ_K " (registration version 1.102, line 303) and is written χ_K in programme writing after it (registration version 1.102, line 309), and converts by the elasticity of checker capability in target capability, and the direct term must be shown zero or estimated) and when the two correction-capacity measures of registration version 1.102, line 615 are shown equivalent within the registered margin. Otherwise χ_{ol} is the elasticity of the declared service measure and never enters the ceiling relation. The registration records that its letter for this exponent and χ name the same quantity (registration version 1.102, line 305); this document writes χ , and prints the registration's letters only inside quotations: its letter for this exponent inside the quotation of registration version 1.102, line 665, and its checker-axis form inside the quotation of registration version 1.102, line 303.

Symbols used, in one line each:

- Law IV: s (state); $V(s)$ (property score, vector allowed, components V_j); v_* (threshold, component-wise, $v_{\{*,j\}}$); ϵ_V (tolerance, component-wise); $\Omega_V := \{s : V_j(s) \geq v_{\{*,j\}} - \epsilon_{V,j} \text{ for every } j\}$ (acceptable set); $\rho = (s_o, \dots, s_n)$ (path), $|\rho| = n$; T (the declared family of allowed finite paths, containing every length-zero path and every dead end); $T_H(s_o) := \{\rho \in T : \rho \text{ starts at } s_o, |\rho| \leq H\}$, partitioned into $T_H^{\text{safe}}(s_o)$ and $T_H^{\text{exit}}(s_o)$; T_{acc} (the family of paths of executed transitions, Theorem IV-G); T_{plan} (the planner's feasible family, Theorem IV-H); $T_{\text{eval}}(s)$ (the allowed one-step transitions from s evaluated in premise (G2)); H (horizon); $K_V(H; T, \epsilon_V)$ (persistence kernel); J (acceptance objective, one declared unit); J_H (path objective); J_{safe}^* , J_{exit}^* (the safe and exit suprema); J_{up} (a registered upper bound on J_{safe}^*); J_{found} (the best exit score a registered search found); b_V, c_V (the removal benefit and cost causally attributable to the degradation, in J 's units); b_{other} (the part of an objective change not attributable to it); $\pi_V := c_V - b_V$ (local margin); $\Pi_V(H)$ (its infimum over evaluated reachable exits); $M_V^{\wedge}(H)$ (strategic viability margin); M_{srch} (the searched margin); ϵ_{opt} (optimiser tolerance); $\hat{H}$,

epsilon_eval (the score a planner uses to choose a plan where that differs from J_H, and a uniform bound on its error; PROPOSED, section 2.5.1); theta (discount factor, weights θ^t); inc_i, lev_t (increments and levels); x (budget ratio); A_V(x), A_base(x) (arm scores on the frozen property battery; A_ext and A_emb are the scores of the external and the embedded arm); Y_V(x) (installed effect; Y_ext is the external arm's); Y_min (floor); S_V(x) (retained fraction; S_ext, S_emb per arm; S_V,fix the fixed-baseline alternative); h_V (log-decay derivative); $x_{\{1/2\}}$ (first crossing of one half); S_inf, nu, xi, x_scale, sh(x) (parameters and shape function of the PROPOSED descriptive families); Lambda_V (PROPOSED mediation contrast); Yhat_1, Yhat_0, var_1, var_0, cov_10, z_crit, m, m_thr, quad (the two estimates, their estimated variances and covariance, the critical value, the trial ratio, the threshold and the quadratic of PROPOSED Proposition IV-R, used in section 2.6.1 only); indices t (step), i (increment), j (component), n (path length, or a count of accepted transitions).

- Law V and the balance: C, C_o (capability, reference); R, R_o (modelled depth, reference; registration version 1.102, line 263: completed rounds plus one where the indexing at zero applies); o in {after, first}, l in {out, in} (arm labels); $Q_{ol}(C, R) > 0$ (correction service per arm), Q_o (a fixed positive reference value of service), $Q_{ol,o}$; $\chi_{ol} := d \ln Q_{ol} / d \ln C$ at fixed R; $\eta_{ol} := d \ln Q_{ol} / d \ln R$ at fixed C; $W(C, R) > 0$ (burden); $\kappa := d \ln W / d \ln C$; $\zeta := d \ln W / d \ln R$ (the programme's μ_B is $\kappa - 1$); $\alpha_R := d \ln C / d \ln R$ along the path (the local reading of the programme's alpha; the author's α_C in his working paper of 16 August 2026 is the same quantity on the resource clock, and the programme's notation writes α_C for it in its entries for ϕ_Q and ϕ_W); α_{in} , α_{out} (the same on the arms' own paths $C_{in}(R)$, $C_{out}(R)$); α_{bal} (the root of Δ_{bal} in α_R); α_{crit} (that root where it is a ceiling); Δ_{bal} (the programme's Delta, local); $G_{emb} := Q_{in}/Q_{out}$ at matched order, $G_o := G_{emb}(C_o, R_o)$; ind_first, ind_in (indicators); u (random effects, subscripted); psi (residual); k (the programme's drift-acceleration exponent); beta_C (Law II's exponent, registration version 1.102, line 663); $R_1 < R_2$ (the ends of a finite range of depth), r (the variable of integration over it), R_star (a formal crossing depth) and p (the power in a re-expression of capability), all PROPOSED and of section 4.8. All logarithms are of ratios to the reference values.

2. Law IV, the ARC Persistence Law

2.1 Statement

Law IV (added 24 September 2026; its proposition, P15, registered in ARC-4; not yet registered as a law or tested), in the author's statement of the laws: "Law IV, the ARC Persistence Law, concerns value persistence (register ARC-4) and carries P15: whether externally installed gains decay under subsequent capability training." Its formal content is P15's registered claim and nothing more: on the author's pre-fixed budget scale, the retained fraction of an externally installed gain after matched capability training without embedded correction is below one half, $S_{\text{ext}}(1) < 1/2$ (section 2.6). The register is that of registration version 1.102, line 99: "ARC-4, VALUE PERSISTENCE: P15 alone. Whether externally installed gains decay under subsequent capability training is the only proposition in this document about persistence after the installing pressure is withdrawn."; P15 "measures the external arm's retention and not the embedded arm's" (registration version 1.102, line 675).

The mathematics beneath the register has three separate objects, none of which is the law: an invariance object (the persistence kernel, section 2.2), two conditional theorems that give sufficient mechanisms for invariance under a stated decision rule (sections 2.4 and 2.5), and the empirical retention process (section 2.6), of which P15 reads one arm. The comparison of the two arms, $S_{\text{emb}}(x) > S_{\text{ext}}(x)$, is not Law IV's content in this version: it is the PROPOSED object IV-PERSIST, one object with Λ_V , recorded in section 2.6 and ruled on in section 6, and whether Law IV should instead carry that comparative content with a registered comparative refuter and instrument is the first question of section 7. Law IV is content-neutral: it says what persists, not that what persists is good (registration version 1.102, line 75: "THE THEORY IS CONTENT-NEUTRAL AND THE METHOD IS NOT"). What a supported P15 licenses is what registration version 1.102, line 1087 states: "support shows that externally installed gains lack the registered persistence under the matched subsequent-training regime, and it does not show that embedded or in-loop correction is superior, sufficient or safe. The comparison between placements is P21's, not this proposition's. The laws are untouched."

2.2 Definitions

Definition IV-1 (acceptable set). Given the property score V , the threshold v_{*} and the tolerance ϵ_V , all fixed before any outcome and component-wise in their own declared units, the acceptable set is $\Omega_V := \{s : V_j(s) \geq v_{*,j} - \epsilon_{V,j} \text{ for every component } j\}$. A vector V with component-wise thresholds is used rather than a minimum across raw scales, so that an average cannot hide a collapse on one dimension and incomparable scales are never minimised together.

Membership from data (PROPOSED interval rule). True scores are not observed; a frozen battery gives estimates with registered intervals. A state s' is established outside Ω_V when, for some j , a registered interval on $V_j(s')$ lies entirely below $v_{*,j} - \epsilon_{V,j}$; established inside when, for every j , the interval lies entirely at or above it; otherwise membership is NOT EVALUABLE for that state, and no rule below fires on it. This is the registration's own pattern for a measured quantity read against a fixed line (registration version 1.102, line 633: "outcomes inside that band are INCONCLUSIVE and never support").

Definition IV-2 (paths and the family). A path $\rho = (s_0, s_1, \dots, s_n)$ is a finite sequence of states with $|\rho| = n$ transitions. The family T of allowed finite paths is declared before any outcome and contains every allowed path of length at most the horizon: the length-zero path at every state and every permitted stopping or dead-end path among them. A prefix of an allowed path is an allowed path, so T is closed under prefixes by its meaning; no result below uses that closure. $T_H(s_0)$ is the set of members starting at s_0 with $|\rho| \leq H$. A path is safe if every state on it lies in Ω_V and is an exit otherwise; $T_H(s_0)$ is the disjoint union of $T_H^{\text{safe}}(s_0)$ and $T_H^{\text{exit}}(s_0)$. $T_H(s_0)$, its safe and exit parts and Definition IV-3 apply to any set of finite paths. Two named path sets are used below:

T_{acc} , the family of executed paths of Theorem IV-G, a family in the sense of this definition and closed under prefixes because a prefix of an executed path is an executed path; and T_{plan} , the planner's set of complete plans of Theorem IV-H, to which the notation applies as written but which need not contain the length-zero plan at a state or any proper prefix of one of its plans, which it contains only where the planner may stop there.

Definition IV-3 (persistence kernel). $K_V(H; T, \epsilon_V) := \{s_o : T_H^{\text{exit}}(s_o) \text{ is empty}\}$, the starting states from which every allowed path of length at most H stays acceptable. The invariance idea is the author's: his working paper of 16 August 2026 defines a robust target set, its invariance under a deterministic rewrite operator, and a stochastic form bounded by the union bound. The kernel is kept as the object Theorem IV-G certifies for the accepted dynamics, s_o in $K_V(H; T_acc, \epsilon_V)$, and as a PROPOSED object for registration (section 6); for a planner, Theorem IV-H certifies the selected plan and Corollary IV-H' the executed path, never membership of the kernel over the planner's feasible family (section 2.7).

Two limits on the kernel are stated with it. First, membership is relative to T : an unrestricted family containing a transformation that maps into the complement of Ω_V empties the kernel, so a registered use of the kernel names either a finite battery of transformations, and the claim is battery-relative, or a probabilistic family with a stated distribution and risk tolerance, which is a different claim with its own registration, and which the union-bound form of the author's working paper already states. Second, the kernel quantifies over every allowed path of length at most H , not only over paths of length exactly H , and that choice is what makes (K1) of Lemma IV-K true; (K2) and (K3) hold under either quantifier.

2.3 Lemma IV-K (kernel monotonicity)

Premise: T is a declared family as Definition IV-2 requires. Conclusions: (K1) $K_V(H + 1; T, \epsilon_V)$ is a subset of $K_V(H; T, \epsilon_V)$; (K2) if T is a subset of T' , then $K_V(H; T', \epsilon_V)$ is a subset of $K_V(H; T, \epsilon_V)$; (K3) if $\epsilon_V \leq \epsilon'_V$ component-wise, then $K_V(H; T, \epsilon_V)$ is a subset of $K_V(H; T, \epsilon'_V)$.

Proof. (K1): $T_H(s_o)$ is a subset of $T_{\{H+1\}}(s_o)$, so if every member of the larger set is safe, every member of the smaller is. (K2): $T_H(s_o)$ under T is a subset of $T_H(s_o)$ under T' . (K3): enlarging the tolerance enlarges Ω_V , so a path safe under the smaller tolerance is safe under the larger. Each is an inclusion of quantifier ranges and needs nothing about the system.

Why the quantifier matters. If the kernel quantified over paths of length exactly H , (K1) would fail: a dead end of length one that exits Ω_V excludes its start at $H = 1$ but is invisible at $H = 2$, where only length-two paths are inspected. A three-state countermodel

exhibits it, and under Definition IV-3 its kernels nest, $K_V(2; T, \epsilon_V) = K_V(1; T, \epsilon_V)$ within $K_V(0; T, \epsilon_V)$, and enlarging the family shrinks the kernel. The countermodel is a listed finite structure; no dynamics are run.

2.4 Theorem IV-G (greedy acceptance)

Setting. A declared objective J on states, real-valued in one unit; a declared set of allowed one-step transitions; a declared acceptance rule: a proposed transition s to s' is executed only if $J(s') - J(s) \geq 0$. A state is reachable within n accepted transitions if some path of executed transitions of length at most n leads to it from s_o . Let s_o lie in Ω_V . For every state s in Ω_V reachable from s_o within fewer than H accepted transitions, $T_{\text{eval}}(s)$ is the set of allowed transitions from s that the design evaluated.

Premises.

- (Go) Independence: J is scored independently of V . Where J scores V , (G2) can hold by construction and the theorem certifies only that a process built to protect V protects it, a tautology; the author's working paper of 16 August 2026 already uses an independently measured capability score. The proof below does not use (Go): (Go) is the condition that keeps the theorem from certifying only that tautology. J is the objective the acceptance rule scores and not a capability battery a researcher reads: as Corollary IV-G' says, a loss read on a separate capability battery is reported beside c_V and is not c_V , and where J values a battery score and anything else, a removal that costs 1 on the battery and gains 2 elsewhere changes J by +1 and is accepted. Where J does score V , as a product objective $J(s) = C(s) V(s)$ does ($C > 0$, $V \geq 0$, one protected dimension), the rule's own certificate is exact: $J(s_t) \geq J(s_o)$ gives $V(s_t) \geq V(s_o) C(s_o)/C(s_t)$, which guarantees V at or above $v^* - \epsilon_V > 0$ while $C(s_t) \leq C(s_o) V(s_o)/(v^* - \epsilon_V)$ and nothing beyond it, and V can tend to zero while J rises at every step.
- (G1) Rule in force: every executed transition satisfies $J(s') - J(s) \geq 0$ on the declared J .
- (G2) Margin: for every such s and every transition s to s' in $T_{\text{eval}}(s)$ with s' outside Ω_V , the total change $J(s') - J(s)$ is strictly negative. Where a registered causal decomposition exists, $J(s') - J(s) = b_V - c_V + b_{\text{other}}$, with b_V and c_V the

benefit and cost causally attributable to the degradation of V and b_other the rest;

(G2) then reads $pi_V > b_other$ at every evaluated exit, with $pi_V := c_V - b_V$.

- (G3) Coverage: every transition the process can execute from such a state lies in $T_eval(s)$.

Conclusion. Every path of executed transitions from s_o of length at most H stays in Ω_V ; equivalently s_o lies in $K_V(H; T_acc, \epsilon_V)$, where T_acc is the family of paths of executed transitions of length at most H .

Proof. Suppose a path of executed transitions from s_o has a first index t , with $1 \leq t \leq H$, at which s_t lies outside Ω_V . Then $s_{\{t-1\}}$ lies in Ω_V and is reachable within $t - 1 < H$ accepted transitions. By (G3) the executed transition $s_{\{t-1\}}$ to s_t lies in $T_eval(s_{\{t-1\}})$, so by (G2) $J(s_t) - J(s_{\{t-1\}}) < 0$. But the transition was executed, so by (G1) $J(s_t) - J(s_{\{t-1\}}) \geq 0$. Contradiction; no such t exists. Since t ranges over a finite set, no further condition is needed.

Corollary IV-G' (the decomposition). Under the causal decomposition with $b_other = 0$, (G2) reads $c_V > b_V$ at every evaluated exit, and the local margin $pi_V = c_V - b_V$ is minus the total change. Both quantities are in J 's units; a loss read on a separate capability battery is reported beside c_V and is not c_V . Ties are permitted by the rule and give no guarantee. Pointwise strict positivity suffices; a strictly positive uniform infimum $Pi_V(H) > 0$, the infimum of pi_V over the evaluated reachable exits, is sufficient and not necessary, since margins $1/n$ are all positive with infimum zero. A path margin over discounted sums, $Pi_V^{\wedge}(H)$, is superseded by $M_V^{\wedge}(H)$ of section 2.5 and is not used here.

Bundling. A transition that removes the property while bringing an unrelated gain is inside the premise: (G2) is stated on the total change, and it is false at every evaluated exit whose unrelated gain b_other is at least pi_V , which the rule then accepts. So IV-G certifies nothing about a system whose allowed set contains such a bundle: the theorem turns the danger of bundling into a premise to be verified before the theorem is applied, and does not answer it. b_V stays causal; for exactly this reason (G2) is stated on the total change, and not as a premise that the causal decomposition for the tested component writes the change in its capability functional as the removal benefit less the removal cost.

Falsifier. The theorem is a conditional and no observation refutes it. What an observation refutes is a premise, and the observation decides which, membership first and the sign second. Suppose an executed transition s to s' with s in Ω_V , reachable within fewer than H accepted transitions, and s' established outside Ω_V by the interval rule of section 2.2, is recorded.

- Not in $T_eval(s)$: (G3) has failed. This is a coverage failure. It refutes the coverage premise, leaves (G2) untouched on the evaluated set, and makes the kernel claim about the system's actual dynamics NOT EVALUABLE; it is neither a refutation of the theorem nor of the margin premise. Where its total change is also established negative, (G1) has failed as well, since (G1) binds every executed transition, evaluated or not.
- In $T_eval(s)$, with $J(s') - J(s) < 0$ on the declared objective: (G1) has failed. The rule in force is not the declared rule, or J is not the objective in force. This refutes the model of the system's acceptance, not the margin premise.
- In $T_eval(s)$, with $J(s') - J(s) \geq 0$: (G2) has failed for that exit. This is the empirical refutation of the persistence claim built on IV-G for that system, decidable by measuring the total objective change on registered exits. The three cases are exclusive, and they are exhaustive among records whose membership and sign are both decided. The sign of $J(s') - J(s)$ is read exactly where J is the process's own logged objective; otherwise it is established by a registered interval lying entirely below zero, or entirely at or above zero, and is NOT EVALUABLE otherwise. An executed exit in $T_eval(s)$ whose sign is NOT EVALUABLE shows that (G1) or (G2) has failed for it, since an evaluated executed exit cannot satisfy both, but it does not show which, so it refutes neither by itself; outside $T_eval(s)$ it is the first case whatever its sign.

Five outcomes, each named once. A recorded executed exit within H is read against the premises in one of five ways. (1) A counterexample to the theorem: such an exit with (G1), (G2) and (G3) all true. The proof excludes it, so no observation is one; a report of one shows that a premise was misverified, and the three cases above say which. (2) A margin failure, a failure of the mechanism, the third case: (G2) false at an evaluated exit, so the mechanism did not produce the margin there; this is the empirical refutation of the persistence claim built on IV-G for that system. (3) A coverage failure, the first case: (G3) false; the kernel claim about the system's actual dynamics is NOT EVALUABLE. (4) An

applicability failure, the second case: $(G1)$ false, because the system does not follow the declared rule or does not score its transitions by the declared J ; a record in the first case whose change is also negative is both a coverage and an applicability failure. (5) NOT EVALUABLE, where the record's own measurements fall short: with the membership of s' not established, no exit is established and nothing is decided; with s' established outside Ω_V but the membership of s not established, the path is known to have left Ω_V within H , so some premise failed, but neither the transition at which it first left nor the premise is identified; and in $T_eval(s)$ with the sign not established, $(G1)$ or $(G2)$ failed, but not which. None of these refutes a premise by itself. On every structure of an enumerated finite class, $(G1)$ to $(G3)$ together admit no executed exit within H , and the case of each executed exit names a premise that is false in that structure. (Go) is decided by no exit: it is read from how J is scored, and where it fails the certificate is the tautology (Go) describes, whatever is observed.

What is not a counterexample to the theorem. A record is classified on the total change $J(s') - J(s)$, which is $b_other - \pi_V$ where a registered causal decomposition exists, and never on π_V alone. Where the registered decomposition has $b_other = 0$ (Corollary IV-G'), the total change is $-\pi_V$, so the rule accepts a transition exactly when its margin is zero or negative, $\pi_V \leq 0$; such an acceptance is what the rule permits, and at an evaluated exit it is the third case above, a margin failure: it refutes $(G2)$, and with it the persistence claim built on IV-G for that system, and not the theorem. So a proposed refutation wording for IV-G, under which a preregistered value-degrading rewrite or plan with adequately measured π_V at or below zero is accepted under the theorem's stated decision rule, describes no counterexample to IV-G; at an evaluated exit it describes that margin failure. With $b_other = 0$, a report of acceptance at a strictly positive margin with $(G1)$ and $(G3)$ verified cannot be right: such a transition has $J(s') - J(s) = -\pi_V < 0$, so the report shows that $(G1)$ was misverified, the second case above. Where b_other is not zero, the rule accepts exactly the transitions with $\pi_V \leq b_other$. A transition with $\pi_V \leq 0$ is then rejected whenever $b_other < \pi_V$: $\pi_V = 0$ with $b_other = -1$ changes the objective by -1 . And acceptance of an evaluated exit at $\pi_V > 0$ with $(G1)$ and $(G3)$ verified is possible, exactly when $b_other \geq \pi_V$: it is the bundled case, the third case above, and its report shows that $(G2)$ failed, not $(G1)$; $\pi_V = 1$ with $b_other = 3/2$ changes the objective by $+1/2$ and is accepted. That wording's second clause, that a relevant bypass lies outside the claimed reachable-set coverage, is the first case, a coverage failure, not a failure of IV-G.

2.5 Theorem IV-H (planned optimisation) and the strategic viability margin

Setting. A start s_o in Ω_V , a horizon H , the planner's feasible family T_{plan} (the complete plans it can select, including the permitted early stops and only those), declared before any outcome, a non-empty family that need not have the closure properties of Definition IV-2, with $T_{\text{plan},H}(s_o)$ non-empty, and the planner's declared path objective J_H on it. Because T_{plan} need not contain the length-zero plan at s_o or the prefixes of its plans, its safe part can be empty, which is the void case of (H1).

Premises.

- (H1) Both $T_{\text{plan},H}^{\text{safe}}(s_o)$ and $T_{\text{plan},H}^{\text{exit}}(s_o)$ are non-empty and their suprema of J_H are finite. If the exit part is empty, every feasible plan is safe and nothing remains to prove; if the safe part is empty, no safe plan exists and the theorem is void. Neither case subtracts infinities.
- (H2) Margin: $M_V^H(s_o) := J_{\text{safe}}^* - J_{\text{exit}}^*$, with J_{safe}^* the supremum of J_H over the safe part and J_{exit}^* over the exit part, satisfies $M_V^H(s_o) > \epsilon_{\text{opt}} \geq 0$.
- (H3) Selection: the executed plan ρ^* lies in $T_{\text{plan},H}(s_o)$ and is ϵ_{opt} -optimal: $J_H(\rho^*) \geq \sup_{T_{\text{plan},H}(s_o)} J_H - \epsilon_{\text{opt}}$.

Conclusion. ρ^* lies in $T_{\text{plan},H}^{\text{safe}}(s_o)$.

Proof. By (H2), $J_{\text{safe}}^* - J_{\text{exit}}^* > \epsilon_{\text{opt}} \geq 0$, so $J_{\text{safe}}^* > J_{\text{exit}}^*$ and the supremum over all of $T_{\text{plan},H}(s_o)$ is J_{safe}^* . By (H3), $J_H(\rho^*) \geq J_{\text{safe}}^* - \epsilon_{\text{opt}}$, and $(J_{\text{safe}}^* - \epsilon_{\text{opt}}) - J_{\text{exit}}^* = M_V^H(s_o) - \epsilon_{\text{opt}} > 0$. Every exit plan scores at most J_{exit}^* , so ρ^* is not an exit plan. Neither supremum need be attained; an exact maximiser, where it exists, is the case $\epsilon_{\text{opt}} = 0$.

Corollary IV-H' (rolling replanning). Suppose that at every replanning state s_t reached along the executed path, for $t = 0, 1, \dots, H - 1$, three conditions hold: (i) either the exit part $T_{\text{plan},H-t}^{\text{exit}}(s_t)$ of the feasible family then in force is empty, or premises (H1) to (H3) hold for the objective, feasible family, tolerance and residual horizon $H - t$ then in force; (ii) the plan executed from s_t lies in the feasible family then in force (where (H1) to (H3) hold, (H3) already says so); and (iii) realisation: unless that plan stops at s_t , the transition

realised from s_t is the plan's first transition as modelled, so that s_{t+1} is the plan's second state. Then every state of the executed path through H lies in Ω_V . Proof by induction on t : s_0 lies in Ω_V ; if s_t does, the plan executed from s_t is safe, by IV-H where (H1) to (H3) hold and because every feasible plan is safe where the exit part is empty; if it stops at s_t the executed path ends there, in Ω_V , and otherwise, by (iii), s_{t+1} is a state of a safe plan and lies in Ω_V . Condition (iii) is the one no margin supplies: where an unmodelled transition realises a state other than the plan's second state, (H1) to (H3) can hold at s_t and the path can still exit, and the corollary then says nothing. A margin computed once at $t = 0$ constrains none of the later objectives, feasible families or unmodelled transitions; a deterministic path guarantee covers unexpected transitions only where its family includes them.

The objective's form. J_H must be the planner's objective as declared, and the margin is computed in it. A per-step decomposition into increments does not reproduce a planner that scores levels. Write the increments $inc_i = J(s_i) - J(s_{i-1})$ and the levels $lev_t = J(s_t) - J(s_0) = inc_1 + \dots + inc_t$. With weights θ^t , the weighted sum of levels is the sum over i of $inc_i \theta^i (1 - \theta^{(n-i+1)}) / (1 - \theta)$, which at $\theta = 1$ is the sum over i of $inc_i (n - i + 1)$: a level-scoring planner weights an early increment by the whole remaining horizon, for every n . An example, over two transitions: with $\theta = 1$, a removal that brings $b_V = 1$ on the first transition and costs $c_V = 21/20$ on the second has increments $(1, -21/20)$, which sum to $-1/20$, so on summed increments removal is not worth doing; the levels they produce are $(1, -1/20)$, which sum to $19/20$, so a planner maximising summed levels removes the property. The margin $c_V - b_V$ read as a sum of increments is the planner's margin only when the planner's objective is that sum of increments, which at $\theta = 1$ is the terminal level; a registered use of IV-H states which objective the planner maximises and computes M_V^H in it. Discounted increment sums are therefore a special case, valid only for a planner that maximises exactly that sum. Levels here are measured from the start. A planner that scores absolute levels, the weighted sum of $J(s_t)$ over its steps $t = 1, \dots, n$, differs from one that scores the levels lev_t by $J(s_0)$ times the sum of the weights over the plan's length; that is the same for two plans where they have one length, and for any two plans when $J(s_0) = 0$, so where early stops are permitted the two objectives can rank plans differently: with $\theta = 1$ and $J(s_0) = 10$, stopping at once scores 0 under both, while one step to a state with $J = 9$ scores 9 on absolute levels and -1 on levels from the start.

The two-option special case. Where the safe optimum J_{safe}^* is attained and the best exit scores $J_{\text{safe}}^* + b_V - c_V$, the margin is $c_V - b_V$. Separability alone does not establish that these two plans are the optima.

Proposition IV-M (PROPOSED: the strategic viability margin). $M_V^*(H)$ is the name given here to the margin of (H2). The preserve and degrade sets are read path-wise, as the safe and exit parts of the feasible family, and not by a terminal or aggregate reading, $V(\rho)$ against a minimum, which a plan that dips below the threshold and recovers would satisfy. It replaces the isolated removal-cost formulation as the general object and keeps $c_V - b_V$ as its two-option case. It is defined, not generally computable: it is a supremum over a feasible family that cannot be enumerated for a real system, and it presupposes a known objective. Where J is unknown or unregistered, the margin is NOT EVALUABLE.

Lemma IV-M' (one-sided estimation). Let J_{up} be a registered upper bound on the safe supremum, $J_{\text{up}} \geq J_{\text{safe}}^*$ (the exact safe optimum of a finite family enumerated completely is the case $J_{\text{up}} = J_{\text{safe}}^*$), and let a registered search find an exit plan of score $J_{\text{found}} \leq J_{\text{exit}}^*$. Then $M_V^*(H) = J_{\text{safe}}^* - J_{\text{exit}}^* \leq J_{\text{up}} - J_{\text{found}}$. Consequences: a found exit plan scoring at or above J_{up} certifies $M_V^*(H) \leq 0$ and refutes a positive-margin claim; a positive value of the searched margin $M_{\text{srch}} := J_{\text{up}} - J_{\text{found}}$ certifies nothing about the sign of $M_V^*(H)$, and is reported as "no bypass found within the registered search", never as a positive margin established; and where the safe optimum is itself only searched, so that only a lower bound on J_{safe}^* is in hand, a found exit at or above that lower bound certifies nothing either: with a searched safe score of 1 against a true safe optimum of 3, and a found exit of 2 that is in fact the best exit plan ($J_{\text{exit}}^* = 2$), the found exit lies above the searched score while the true margin is +1; the same search results are equally consistent with an unfound best exit plan scoring 4 and a true margin of -1, since a found exit bounds J_{exit}^* from below only. A complete enumeration of a finite family decides the sign both ways; no incomplete search confirms a positive margin. The author's working paper of 16 August 2026 already reads search this way, asking that censoring, incomplete search or non-reachability be reported rather than non-observation being treated as proof of infinite resistance.

Robust margin (PROPOSED): the infimum of $M_V^*(H)$ over a registered family of environments. It is given no symbol here, because the one first proposed for it collides with recursive depth; the object stands as a proposal.

Falsifier. IV-H is a conditional; an executed exit plan ρ^* refutes a premise. If ρ^* lies outside the declared family T_{plan} , the declared family was not the planner's feasible family and (H3)'s membership clause fails for it: a coverage failure, NOT EVALUABLE for the persistence claim about the actual planner. If ρ^* lies in the family and the family is enumerated completely and shows $M_V^H > \epsilon_{\text{opt}}$, then (H3)'s optimality clause has failed: the planner's objective or tolerance differ from those declared. The empirical content of a persistence claim built on IV-H is the margin premise (H2) itself, and Lemma IV-M' says how data can decide it: a registered search that finds an exit plan at or above a registered upper bound on the safe supremum refutes $M_V^H > 0$; no incomplete search confirms it. For Corollary IV-H', an executed exit reached by a transition other than the executed plan's modelled first transition refutes its realisation condition (iii), not the margin: it is the planner's counterpart of a coverage failure.

2.5.1 Proposition IV-N (PROPOSED: a planner that scores its plans with bounded error)

Theorem IV-H assumes that the planner chooses by the objective J_H itself. A planner that chooses by an estimate of J_H is covered at a price: the margin of (H2) must exceed the optimiser's tolerance by twice the estimate's worst error. The margin is still M_V^H ; no second margin is introduced.

Setting and premises. The start s_o , the acceptable set read path-wise, the horizon H , the feasible family T_{plan} and the objective J_H are those of Theorem IV-H, and (H1) holds as stated there: the safe and the exit part are both non-empty and both suprema of J_H are finite. The planner chooses by a score J_{hat}_H , which need not equal J_H . The first premise bounds its error uniformly, with $\epsilon_{\text{eval}} \geq 0$ fixed before any outcome:

(N1) $|J_{\text{hat}}_H(\rho) - J_H(\rho)| \leq \epsilon_{\text{eval}}$ for every ρ in $T_{\text{plan},H}(s_o)$.

(N2) Selection: the executed plan ρ^* lies in $T_{\text{plan},H}(s_o)$ and is ϵ_{opt} -optimal for the score, that is $J_{\text{hat}}_H(\rho^*) \geq \sup_{\rho \in T_{\text{plan},H}(s_o)} J_{\text{hat}}_H(\rho) - \epsilon_{\text{opt}}$, with $\epsilon_{\text{opt}} \geq 0$. The third premise is the margin of (H2), enlarged:

(N3) $M_V^H(s_o) > 2\epsilon_{\text{eval}} + \epsilon_{\text{opt}}$.

Conclusion. ρ^* lies in $T_{\text{plan},H}^{\text{safe}}(s_o)$.

Proof. By (N1) the score of every safe plan is at least its J_H less ϵ_{eval} , so the supremum of J_{hat_H} over the safe part is at least $J_safe^* - \epsilon_{eval}$; and the score of every exit plan is at most $J_exit^* + \epsilon_{eval}$. The supremum of J_{hat_H} over the whole family is therefore at least $J_safe^* - \epsilon_{eval}$, and it is finite, because by (N3) J_safe^* exceeds J_exit^* and so no score exceeds $J_safe^* + \epsilon_{eval}$. By (N3) again, $(J_safe^* - \epsilon_{eval}) - (J_exit^* + \epsilon_{eval}) = M_V^H(s_0) - 2\epsilon_{eval} > \epsilon_{opt}$. Every exit plan therefore scores more than ϵ_{opt} below the supremum over the family and fails (N2), so ρ^* is not an exit plan. Neither supremum need be attained. With $\epsilon_{eval} = 0$ the proposition is Theorem IV-H.

The inequality in (N3) must be strict. Take two plans, a safe one with $J_H = 1$ and an exit with $J_H = 0$, a score of one half for each, so that $\epsilon_{eval} = 1/2$, and an exact optimiser, $\epsilon_{opt} = 0$. Then $M_V^H = 1 = 2\epsilon_{eval} + \epsilon_{opt}$, the two scores tie, and (N2) admits the exit plan.

What the premises require. (N1) is a bound over every plan of the declared family: a bound that holds on average, or only over the plans a search happened to sample, or the uncertainty of one point estimate, is not it. The margin is computed in the planner's complete objective, as section 2.5 requires: the unrelated gains a plan bundles, the discounting in force, the scoring of increments or of levels, and the treatment of plans of unequal length are all inside J_H , and a capability score read separately does not stand in for it. Where the exit part is empty every feasible plan is safe, and where the safe part is empty no safe plan exists and the proposition is void, exactly as under (H1).

Rolling replanning. Corollary IV-H' carries over with (N1) to (N3) in place of (H2) and (H3): the premises are needed at every replanning state reached, for the objective, the score, the family, the residual horizon and the tolerances then in force, together with the realisation condition (iii). A bound established once at the start says nothing about a later objective or an unmodelled transition. A bound on the error that holds only with some probability is a different premise: it needs a statement covering every plan of the family at once, and no confidence level or risk tolerance is chosen here.

Falsifier. As for Theorem IV-H, an executed exit plan refutes a premise and not the proposition: the family's coverage, the selection rule, the margin, or, new here, the error bound (N1), which one plan whose score errs by more than ϵ_{eval} refutes.

2.6 Retention: the empirical object, and P15's registered face

Definition IV-4 (retained fraction). For a registered pressure class named in words (for P15: subsequent capability-directed training that embeds no correction of its own), a frozen property battery and score scale, and the budget ratio $x \geq 0$ (for P15: the subsequent capability-training budget divided by the installation budget, on the author's pre-fixed budget measure, registration version 1.102, line 1089), the installed effect is $Y_V(x) := A_V(x) - A_base(x)$ and the retained fraction is $S_V(x) := Y_V(x)/Y_V(0)$. It is read only where the installed effect was real: the lower end of a registered interval on $Y_V(0)$ is at least Y_min , with $Y_min > 0$ fixed before any outcome. The floor is on the interval's lower end and not on a point estimate, because the ratio's interval is unbounded as the denominator's interval approaches zero, and the interval on S_V is formed by a method fixed before any outcome (Fieller's, from his paper Some Problems in Interval Estimation, Journal of the Royal Statistical Society, Series B, volume 16, number 2, 1954, pages 175 to 185, or an equivalent named in the deciding instrument). The floor is positive, not absolute: an absolute-value floor would admit a negative installed "gain". Y_min 's value is the author's. $S_V(0) = 1$; S_V may exceed one or be negative, and neither is clipped. Absolute retention, $Y_V(x)$ itself, is reported beside $S_V(x)$. S_V is invariant under any affine change of the battery's score scale and under no other change: invariance for all data forces the change to be affine. So the scale is fixed with the battery before outcomes, and, as PROPOSED reporting modelled on the rule the registration sets for P14 (registration version 1.102, line 1079: "a difference surviving on the raw scale alone is reported as scale-dependent"), a verdict surviving on a bounded raw scale alone is reported as scale-dependent; P15's registered scoring carries no such rule. The arms are S_ext (property installed by external post-training) and S_emb (property installed with correction embedded in the training objective), as the registered two-arm unit reads them (registration version 1.102, line 1089).

The baseline convention, stated as a choice. In the ratio above, the baseline arm $A_base(x)$ also undergoes the later training: a concurrent baseline. The registered unit has two arms, external and embedded, and names no baseline arm (registration version 1.102, line 1089); the registration fixes neither a concurrent nor a pre-training baseline. The alternative reads the installed gain against the pre-training baseline, $S_V,fix(x) := (A_V(x) - A_base(0))/(A_V(0) - A_base(0))$. The two can give opposite P15 verdicts: with $A_V(0) = 1$, $A_base(0) = 0$, $A_V(1) = 0$ and $A_base(1) = -3/5$, the concurrent ratio is $3/5$ and the

fixed-baseline ratio is 0. The concurrent baseline needs a third arm beyond the registered two: a system without the installed gain that undergoes the same later training, so that $A_{\text{base}}(x)$ is observed after that training; on the registered two arms alone, with the base system scored once before the gain is installed, only the pre-training reading $S_{V,\text{fix}}$ is computable. The concurrent baseline is the default here, third arm included, because it separates loss of the installed gain from a fall in the whole battery; the convention, and with it the arm, is the author's to fix in the deciding instrument before any outcome (section 7 item 8), and until fixed the endpoint's reading is a choice recorded, not a quantity nature decides.

P15, verbatim (registration version 1.102, line 1081): "Gains installed by external post-training erode under later capability-directed training that embeds no correction of its own, and more than half of what was installed is lost once the capability-training budget matches the alignment budget that installed it. The one-half threshold is a declared convention fixed here before any result rather than a derived quantity, in the same way P10 fixes ten per cent." In these symbols P15 predicts $S_{\text{ext}}(1) < 1/2$. Its refuter, verbatim (registration version 1.102, line 1083):

REFUTED BY: an interval on the retained fraction of the installed gain, after matched capability training without embedded correction, lying entirely above one half.

The refuter has no upper cap. An interval from 0.9 to 1.3 lies entirely above one half and refutes P15; a rule that reads the refuter as an interval lying inside $[1/2, 1]$ would not fire on it and is wrong. An interval containing or touching one half does not refute, and the registered wording is strict at that edge: an interval whose lower end is exactly one half shows at every point that no more than half was lost, so that $S_{\text{ext}}(1) < 1/2$ is false throughout it, and it still does not refute, because it is not entirely above one half; a point estimate is not an interval. The one-half line is a convention, and the registration's rule for its conventions applies (registration version 1.102, line 391: "Where a measurement falls just inside or just outside one of them, that is reported as a result at the margin rather than as a verdict, and no proposition is scored as supported on the strength of a convention alone."): the deciding instrument fixes, before any outcome, the band around one half that counts as "just inside or just outside", and an interval whose lower end falls inside that band is reported as a result at the margin. The endpoint is $S_{\text{ext}}(1)$ at the matched budget,

and the registration keeps the scale the author's (registration version 1.102, line 1089): "The budget is the author's to fix, it is recorded in the deciding instrument before any outcome is read, and until he fixes it the threshold is NOT EVALUABLE on that scale." The same line carries the attestation: "The unit also carries an attestation that the capability-training reward includes no alignment-graded component, because the statement above specifies training that embeds no correction of its own and nothing at present verifies it." P15's instrument line reads NONE (registration version 1.102, line 1087), and its scope is S5 (registration version 1.102, line 237); no scope is borrowed (registration version 1.102, line 247).

What falls with P15, verbatim (registration version 1.102, line 1085): "FALLS WITH IT: this programme's case that correction has to sit inside the system rather than be applied to it from outside, and this proposition's share of the protocol's engineering case, which it carries jointly with P12, P13 and P21."

Descriptive statistics (PROPOSED reporting; never verdicts). The first crossing $x_{\{1/2\}} := \inf\{x : S_V(x) \leq 1/2\}$ reads the endpoint only for monotone retention: the curve $S_V(x) = 1 - 1.8x + 1.4x^2$ has $S_V(1/2) = 9/20$ and $S_V(1) = 3/5$, dips to $59/140$ at $x = 9/14$, and crosses one half first at $x = (9 - \sqrt{11})/14$, about 0.406, yet ends above one half at the matched budget, so an interval there lying entirely above one half refutes P15 under its registered face while the first crossing says "halved"; for a monotone exponential the two readings agree exactly. The log-decay derivative $h_V := -d \ln S_V / dx$ requires a positive differentiable S_V and is negative during recovery (on the same curve it changes sign between $x = 1/5$ and $x = 4/5$), so it is not a survival hazard by name. The mean of S_V over a registered range is descriptive. None of these substitutes for the registered interval rule.

Candidate descriptive families (PROPOSED; closed before any deciding run): no loss; exponential; power; stretched exponential; a change-point family; each of the decaying members with a floor $S_{\inf}$, as $S_V = S_{\inf} + (1 - S_{\inf}) \text{sh}(x)$ with $\text{sh}(0) = 1$ and sh tending to zero, so that each has $S_V(0) = 1$ and tends to $S_{\inf}$; the no-loss member is $S_V = 1$ throughout and tends to no floor but one, and the change-point family's limit is whatever its registered form gives. The floored log-derivative of the exponential family is its rate and is a family-specific statistic. No family is the law; FORM UNDETERMINED and NO CANDIDATE FITS are valid outcomes. No decay family or monotonicity is imposed by

this document. Retention across other pressure classes (self-rewrite, distillation, successor construction, distribution shift, adversarial pressure) is a family of separate objects; only subsequent capability training without embedded correction is P15's.

IV-PERSIST and the mediation contrast (PROPOSED). The comparison $S_emb(x) > S_ext(x)$ over a registered pressure range is not P15's face and is not law content in this version; it is the same test that section 6 rules on under the name Λ_V . The statement paper's published register-level text is recorded here, not adopted: its ARC-4 sentence, "values run as load-bearing recursive loops persist after the enforcing mechanism is removed at a higher rate than rules, frozen objectives or fine-tuned constraints", and its kill, "Persistence trials in which embedded loops fare no better than removed enforcement kill ARC-4, and ARC-5 with it." That sentence is comparative and its kill is comparative; neither is P15's registered face, and K-IV below reaches neither. Whether Law IV carries the comparative content, with a registered comparative refuter and instrument, or P15's face alone, is the author's first decision (section 7).

2.6.1 What the retained fraction needs before it can be read (PROPOSED)

Written out for the external arm, under the concurrent baseline that section 2.6 takes as its default:

$$Y_ext(x) := A_ext(x) - A_base(x),$$

$$S_ext(x) := [A_ext(x) - A_base(x)] / [A_ext(o) - A_base(o)].$$

The arms. $A_base(x)$ is the score of a matched system that never received the installed gain and that undergoes the same later capability training, read at the same budget ratio x . The embedded arm is not that baseline: it carries a gain of its own. The registered unit has two arms, external and embedded (registration version 1.102, line 1089), so it does not supply $A_base(x)$ after training, and a design that reads both retained fractions against a concurrent baseline needs three arms: external, embedded and no gain. A design with the external and the no-gain arm alone reads S_ext , and neither S_emb nor any comparison between placements. These are requirements on an instrument yet to be registered; they change nothing in the frozen registration.

Evaluability. The external arm's installed gain must pass the floor of section 2.6, and the embedded arm's must pass its own floor before S_{emb} , or any comparison that uses it, is reported. A missing arm, an initial gain that is not established, or a budget that is not matched makes the endpoint NOT EVALUABLE; nothing missing is replaced by zero, and no retained fraction is clipped. The raw arm scores and their differences are reported beside the ratios. The instrument says whether its estimand is the ratio of the arms' mean effects or the mean of the replicates' ratios, which are different quantities. Its uncertainty method names the unit that is independent, the separately trained replicate; any pairing of seeds; and the covariance between quantities that share systems or an arm: $Y_{\text{ext}}(0)$ and $Y_{\text{ext}}(1)$ where they are read on the same systems, and the two arms' effects where they share one baseline.

The one-half line in signed form. Where $Y_{\text{ext}}(0) > 0$, $S_{\text{ext}}(1) < 1/2$ holds exactly when $Y_{\text{ext}}(1) - (1/2) Y_{\text{ext}}(0) < 0$.

The two inequalities differ by multiplication by the positive number $Y_{\text{ext}}(0)$. The signed form is a way of writing the endpoint, not a second refuter: P15's registered refuter is an interval on the ratio (section 2.6), and a contrast with an interval of its own does not in general give the same decision. Proposition IV-R states the case in which it does.

Proposition IV-R (PROPOSED: under Fieller's method the ratio interval and the signed contrast decide alike). Let $\hat{Y}_1$ and $\hat{Y}_0$ be the estimates of $Y_{\text{ext}}(1)$ and $Y_{\text{ext}}(0)$, with estimated variances var_1 and var_0 and estimated covariance cov_{10} , the three forming a positive semi-definite matrix, and let $z_{\text{crit}} > 0$ be the critical value the instrument fixes. Fieller's set is the set of trial ratios m with $(\hat{Y}_1 - m \hat{Y}_0)^2 \leq z_{\text{crit}}^2 (\text{var}_1 - 2 m \text{cov}_{10} + m^2 \text{var}_0)$. Suppose the gate $\hat{Y}_0 > z_{\text{crit}} \sqrt{\text{var}_0}$ holds. Then Fieller's set is a closed bounded interval containing $\hat{Y}_1/\hat{Y}_0$, and for every number m_{thr} it lies entirely above m_{thr} if and only if

$\hat{Y}_1 - m_{\text{thr}} \hat{Y}_0 > z_{\text{crit}} \sqrt{\text{var}_1 - 2 m_{\text{thr}} \text{cov}_{10} + m_{\text{thr}}^2 \text{var}_0}$.

Proof. Write $\text{quad}(m) := (\hat{Y}_1 - m \hat{Y}_0)^2 - z_{\text{crit}}^2 (\text{var}_1 - 2 m \text{cov}_{10} + m^2 \text{var}_0)$, so that Fieller's set is where $\text{quad}(m) \leq 0$. It is a quadratic in m with leading coefficient $\hat{Y}_0^2 - z_{\text{crit}}^2 \text{var}_0$, which the gate makes positive, so the set is a closed bounded interval or is empty. At $m = \hat{Y}_1/\hat{Y}_0$ the squared term vanishes and quad equals minus z_{crit}^2 times an estimated variance, which is not negative; so quad is not

positive there, and the set is an interval containing $\hat{Y}_1/\hat{Y}_0$. If the interval lies entirely above m_{thr} , then m_{thr} is outside it, so $quad(m_{thr}) > 0$, and $\hat{Y}_1/\hat{Y}_0$, a point of the interval, exceeds m_{thr} . Conversely, if $quad(m_{thr}) > 0$ and $\hat{Y}_1/\hat{Y}_0 > m_{thr}$, then m_{thr} is outside the interval and below one of its points, so it is below all of them. Now $quad(m_{thr}) > 0$ says that the square of $\hat{Y}_1 - m_{thr} \hat{Y}_0$ exceeds the square of the right-hand side displayed, and, because $\hat{Y}_0 > 0$, $\hat{Y}_1/\hat{Y}_0 > m_{thr}$ says that $\hat{Y}_1 - m_{thr} \hat{Y}_0$ is positive; the two together are the displayed inequality.

What it gives, and where it stops. At $m_{thr} = 1/2$ the displayed inequality says that the signed contrast $\hat{Y}_1 - (1/2) \hat{Y}_0$ exceeds z_{crit} of its own standard errors; so, under Fieller's method with the gate, the interval lies entirely above one half, as the registered refuter requires, exactly when that one-sided test of the contrast rejects. At m_{thr} equal to the upper edge of the at-the-margin band it is the condition for a verdict and not a result at the margin (registration version 1.102, line 391; section 2.6). The strict edge is kept: an interval whose lower end is exactly one half has $quad(1/2) = 0$, and the inequality fails. The gate is the floor of section 2.6 in the proposition's terms: where the interval on $Y_{ext}(0)$ is $\hat{Y}_0$ plus and minus a half-width of at least $z_{crit} \sqrt{var_0}$, a lower end of at least $Y_{min} > 0$ gives the gate. Where $\hat{Y}_0^2 < z_{crit}^2 var_0$ the leading coefficient is negative and Fieller's set is not a bounded interval: it is the whole line, or the two unbounded pieces outside an interval, which is why the floor is set on the interval's lower end. Where $\hat{Y}_0^2 = z_{crit}^2 var_0$ the quadratic term vanishes and the condition is linear in m : Fieller's set is then a half-line where the linear term does not vanish; where it vanishes as well, the set is the whole line, except in the degenerate case in which $\hat{Y}_0$ and var_0 are both zero and $\hat{Y}_1^2$ exceeds $z_{crit}^2 var_1$, where it is empty. Where $\hat{Y}_0 < -z_{crit} \sqrt{var_0}$ the leading coefficient is positive again and Fieller's set is a bounded interval, but the displayed inequality is not the condition for it to lie above m_{thr} : there the condition is the same inequality with its sign reversed, $\hat{Y}_1 - m_{thr} \hat{Y}_0 < -z_{crit} \sqrt{var_1 - 2 m_{thr} cov_{10} + m_{thr}^2 var_0}$, and the floor of section 2.6, which requires a positive installed gain, excludes the case. The proposition concerns Fieller's set for the ratio of the two estimates. It says nothing about a mean of replicate ratios, a bootstrap interval or an interval built on a linearised variance of the ratio; an instrument that uses one of those shows the agreement itself or reads the registered ratio interval alone. It changes nothing in P15's registered refuter.

Domains. A logarithm of a retained fraction is taken only where that fraction is positive: h_V needs S_V positive and differentiable, and Λ_V needs both S_{ext} and S_{emb} established and positive. The difference $S_{emb} - S_{ext}$ needs both established and makes no demand on their signs. An embedded arm that was not run supplies none of these, and none of them becomes a registered quantity by appearing in this document.

2.7 What the persistence mathematics does not say

It does not say that a kept property is a good one, that a system obeying Law IV is safe, or that embedded correction is superior, sufficient or safe (registration version 1.102, line 1087). It does not identify c_V and b_V from one excision, which returns only their difference. It certifies kernel membership only for the accepted dynamics satisfying the premises of Theorem IV-G, s_o in $K_V(H; T_{acc}, \epsilon_V)$; for a planner it certifies only that the selected plan is safe (Theorem IV-H) and, under Corollary IV-H', that the executed path stays in Ω_V , never that s_o lies in $K_V(H; T_{plan}, \epsilon_V)$, which premise (H1) excludes by requiring a feasible exit plan; and it certifies nothing for every physically possible transformation. Related coupled-correction, survival and Paper VIII results are not deciding evidence for P15. Nothing in this section has been measured.

2.8 Kill condition K-IV

Stated before any data, in terms decidable by data yet to be collected, and consistent with P15 as registered. Item 1 is P15's registered refuter with its budget requirement; the one thing it adds, that the baseline convention be fixed in advance, is marked PROPOSED where it stands. Item 5 quotes the registration's weaker-neighbour rule and says what this version takes from it. Items 2, 3 and 4 are clauses this document writes, PROPOSED law-level conventions beyond P15's registered scoring, which they leave untouched; item 2 is on the record and is not in force.

1. The registered face. The observation: from a deciding unit of the kind the registration specifies (registration version 1.102, line 1089), run under the author's pre-fixed budget measure, an interval on $S_{ext}(1)$, the retained fraction of the externally installed gain after matched capability training without embedded correction, lying entirely above one half. That the baseline convention also be fixed before any outcome is this document's requirement and is PROPOSED (section 2.6 and section 7 item 8): at its line 1089 the registration fixes the budget and requires the attestation,

and it names no baseline arm; nothing here is added to P15's frozen refuter. No upper cap: an interval from 0.9 to 1.3 qualifies; an interval from 0.45 to 0.9 does not; a point estimate does not. An interval whose lower end lies inside the at-the-margin band the instrument fixed (registration version 1.102, line 391) is reported as a result at the margin and not as a verdict. That observation refutes P15 as registered (registration version 1.102, line 1083), is scored for P15 exactly as the registration scores it, and takes with it what registration version 1.102, line 1085 names (quoted in section 2.6). As a verdict it fires exactly on intervals whose lower end lies above the at-the-margin band the instrument fixed around one half, every point of which therefore lies strictly above one half, and at every such point the law's formal content, $S_{\text{ext}}(1) < 1/2$, is false; an interval whose lower end is exactly one half also shows that content false throughout and does not fire (section 2.6).

2. The law-level withdrawal (PROPOSED by version 0.2.1; not in force). Version 0.2.1 proposed a convention: where a P15 refutation also met four evaluability controls fixed before any data (the baseline fixed in advance, with the convention named; the installed effect established by the lower end of the interval on $Y_V(0)$ being at least Y_{min} ; the capability rise established on a separate instrument by an interval lying entirely above the minimum that instrument registers; and the attestation of no alignment-graded component in the subsequent training, registration version 1.102, line 1089), Law IV would be withdrawn whole and Law V with it. The author has not adopted that convention and it is not in force. What stands without it is item 1: a P15 refutation is scored as the registration scores it, takes with it what registration version 1.102, line 1085 names, and shows the law's formal content false, since this version gives Law IV no formal content beyond P15's face (section 2.1). It refutes no theorem: Theorems IV-G and IV-H are conditional mathematics and survive, as item 5 says. It decides nothing about the comparative claim, which is a separate proposal with an endpoint and a refuter of its own still to be registered (section 6). And it withdraws nothing from Law V (section 5 item 3). The four controls remain a proposal for the instrument that decides P15, and no scoring code may apply the former default.
3. The battery-relative kernel kill (PROPOSED, not law content until registered). For a registered system, property battery, tolerance and finite battery of transformations, a path of the battery whose end state is established outside Ω_V by the interval rule of section 2.2 within H refutes the claim that the system's start lies in the battery-relative kernel; a path produced by a transformation outside the battery is a coverage

failure and refutes nothing; a path whose membership is NOT EVALUABLE decides nothing. The mechanism claims built on IV-G and IV-H are refuted by the observations of sections 2.4 and 2.5 (an evaluated exit with non-negative total objective change; a found exit plan at or above a registered upper bound on the safe supremum) and are not refuted by coverage failures.

4. What is not a refutation of Law IV (PROPOSED convention, in the registration's own terms, registration version 1.102, line 121): an instrument that fails to run or is too small; an interval touching one half; a result on an unregistered budget scale; a result without the attestation, for the law-level reading; a bypass outside the registered battery; and any result about whether the system's later conduct is good, which the theory's content neutrality places outside every law (registration version 1.102, line 75); the method's own line (registration version 1.102, line 61: "no result about a system's later conduct supports or refutes the method by itself") makes the same exclusion for the Eden Protocol, and P15's own endpoint, a measurement after later training, is not a result about conduct.
5. The weaker neighbour (registration version 1.102, line 125: "A REFUTED PROPOSITION DEGRADES TO A NAMED WEAKER NEIGHBOUR RATHER THAN VANISHING."). What survives if Law IV is withdrawn: Theorems IV-G and IV-H as conditional mathematics; the verdicts on P12 and P21 as measured; and the persistence questions the registration keeps open in another context, where the conversion claim fails: "Even then the persistence registers stand, because whether early formation outlasts later enforcement is a question about durability and not about scaling, and it would remain open and worth answering." (registration version 1.102, line 115), which is quoted for its wording of the question, not as a rule about Law IV's withdrawal.

3. Law V, the ARC Embedding Law

3.1 Statement

Law V (added 24 September 2026; its propositions, P12 and P21, registered in ARC-5; not yet registered as a law or tested), in the author's statement of the laws: "Law V, the ARC Embedding Law, concerns the genesis strategy (register ARC-5) and carries P12, whether

build order moves the correction-leverage exponent, and P21, whether correction placed inside the loop pulls away from correction applied outside it as depth grows." Two limbs, separate hypotheses with no shared mechanism assumed:

- Placement limb. Its registered face is P21, verbatim (registration version 1.102, line 1219): "Correction that participates in each subsequent revision round outperforms correction that sees only the finished output, and the advantage widens as recursive depth rises rather than staying constant." P21 decides an access and action-space interaction on the outcome, "never a measurement of an exponent" (registration version 1.102, line 1221). The registration's own mechanism reading is written in its own quantities (registration version 1.102, line 1239): "that correction embedded in the recursion carries a larger direct depth-dependence than correction applied from outside it. Where that holds, the boundary rises with it". In this document's symbols the mechanism reading is $\eta_{o,in} > \eta_{o,out}$ at matched order o . It is carried here as the registered mechanism reading, which no verdict in this document decides: a P21 verdict does not measure it, and its refutation does not prove $\eta_{o,in} \leq \eta_{o,out}$.
- Build-order limb. Its registered face is P12, verbatim (registration version 1.102, line 1033): "The order in which a system's components are built changes its correction-leverage exponent by at least the registered minimum effect." The contrast is the registered maximal contrast (registration version 1.102, line 1037): "a corrector formed before the capability it corrects against a corrector added after it, at matched content"; the minimum effect is 0.10 on the correction-leverage exponent scale (registration version 1.102, line 1041); in this document's symbols the endpoint is $\chi_{first} - \chi_{after}$ at the placement the instrument fixes, named in the report (P12 registers no placement factor), read on the target axis under section 1's identification conditions, and on one fault set of stated type, the same for both orders: the registration measures the exponent "on a fixed fault set of stated type" (registration version 1.102, line 389), and it sets P12 beside P22 on the same quantity by different routes, "P22 says that elasticity differs with the type of fault corrected, and this says it moves with the order in which the corrector was built", adding "The two are severable and neither carries the other." (registration version 1.102, line 1045). P22 stays in ARC-3 (registration version 1.102, line 99) and Law V does not carry it, and a build-order contrast read across a changed mix of fault types would carry part of the ordering P22 registers, and could reach the minimum effect with no effect of order at

all. P12 registers no direction and is given none here; the genesis reading, that formation first carries the higher elasticity by at least the registered minimum effect, $\chi_{\text{first}} - \chi_{\text{after}} \geq 0.10$ at the placement the instrument fixes (or at least the smaller minimum a future instrument registers, registration version 1.102, line 1041), is PROPOSED (section 6) and is not law content. The margin is the minimum-effect convention, and it stands in the reading's statement as well as in its strike rule, because only so is the reading consistent with P12's undirected registration: a reading that formation first exceeds formation after by a margin implies P12's claim exactly when that margin is at least P12's registered minimum, and with a smaller margin, or none, it can hold at every point of an interval on which P12's registered refuter fires. So stated, the reading is the upward half of P12's claim, adds a sign and nothing else, and never counts as its support a change P12's registration reports "as no effect of interest rather than as no effect" (registration version 1.102, line 1041).

Boundary reading (conditional). Where a limb raises η_{ol} with $\kappa > \chi$, or raises χ_{ol} with $\eta > \zeta$, the burden elasticities being fixed, the balance boundary of section 4 rises (registration version 1.102, line 1045: "Where it moves upward the boundary rises with it."; the domains are those of Proposition B-6). The reading depends on Law III's model and, through it, on ARC-2; it is not a verdict rule, and no boundary measurement scores a limb (section 5).

3.2 Quantities, embedding defined operationally, and the prerequisite

Quantities. On a declared domain $C > 0$, $R > 0$, each arm's correction service $Q_{\text{ol}}(C, R)$ is positive and differentiable with other design factors fixed; capability C carries the ratio scale Condition B fixes, justified independently of the framework and fixed before any outcome (registration version 1.102, line 403), on which every exponent number in this document is stated (section 4.7); R is the modelled depth where registration version 1.102, line 263 applies. Zero service is reported outside the log model and never replaced by an outcome-chosen offset (registration version 1.102, line 265, applied to service); because excluding zeros conditions on the outcome, the zero counts are reported per cell, and a two-part or censored model stands among the rivals of section 3.4. The local elasticities χ_{ol} and η_{ol} are partial derivatives of $\ln Q_{\text{ol}}$; the constant-elasticity case is the power law and returns the programme's exponents (Proposition V-3). The local reading is the

author's: his working paper of 16 August 2026 defines kappa, zeta, chi and eta as local partial elasticities of $W(C, r)$ and $Q(C, r)$, and the registration reads the growth exponent the same way (registration version 1.102, line 267: "the elasticity of capability in depth, taken at the depth a system has reached", with the fitted slope "the special case in which that elasticity is constant"). Identification with the registered correction-leverage exponent needs the conditions of section 1; otherwise χ_{ol} is the elasticity of the declared service measure. η_{ol} is the direct depth-dependence of correction service and is not an integration fraction.

Embedding, defined operationally. Writing the next state as a function of the current state and a correction term, against applying a correction to the finished next state, does not by itself distinguish embedded from external correction: the second form also determines the next state and can be written as one transition map. What can be registered is operational: the treatment schedule (whether correction enters the next round or only inspects its output), the information available to the corrector, the action space (accept or reject only, against generation under constraint), and whether correction feeds later revisions. P21's deciding unit fixes exactly these (registration version 1.102, line 1231: three regimes crossed with depth, matched on model, calls and token budget, manipulating information access; registration version 1.102, line 1233: the action space also differs, and the per-round rejection rate is reported beside the interaction). Matching away information access would be a new mechanism-isolating comparison, not an amendment to P21.

The recursive-improvement condition. It is PROPOSED as a prerequisite beneath all five laws, stated in words: the outputs of one round must causally change a later improvement transition. It is not a law; it would be a condition of the laws only if the author registers it, and until then it is a description that no surface states as a condition of any law, Laws I to III included; no symbol is given to it here. Two operational forms have been proposed, both positivity conditions on a mutual information, and both are declined as criteria. Both are written elsewhere with the updater's letter, one of the eight colliding symbols of section 1, so they are described here in words. The state-based form asks that the mutual information between successive states, conditional on the updater's current state, be positive; it is met by pure repetition. The update-based form asks that the mutual information between the updater's next state and the two successive states, conditional on the updater's current state, be positive; it separates repetition from update, but a prerequisite is a necessary condition, and this form is not necessary: a deterministic recursion from a fixed initial

state, in which every state and every update are fixed functions of their predecessors, has zero entropy everywhere and so zero mutual information, yet its outputs causally change each later transition. The fixed initial state is needed, and it is a valid instance: from a uniformly random initial state the same kind of recursion can carry positive information, one bit when that state is a single uniformly random bit. Mutual information measures statistical dependence, not the causal change the prerequisite names; neither form is an improvement criterion either, since information can persist unchanged or guide a harmful update.

3.3 Identities and the one-path counterexample

Identity V-1 (the embedding ratio's log-derivatives). At matched order, $G_{\text{emb}} := Q_{\text{in}}/Q_{\text{out}}$. At fixed C : $d \ln G_{\text{emb}}/d \ln R = \eta_{\text{in}} - \eta_{\text{out}}$. At fixed R : $d \ln G_{\text{emb}}/d \ln C = \chi_{\text{in}} - \chi_{\text{out}}$. Along a common path $C(R)$ with $\alpha_R = d \ln C/d \ln R$: $d \ln G_{\text{emb}}/d \ln R = \alpha_R (\chi_{\text{in}} - \chi_{\text{out}}) + (\eta_{\text{in}} - \eta_{\text{out}})$. Proof: $\ln G_{\text{emb}} = \ln Q_{\text{in}} - \ln Q_{\text{out}}$; the partials are the definitions; the total derivative is the chain rule. So the path derivative is not in general $\eta_{\text{in}} - \eta_{\text{out}}$: that equality holds at fixed C and, along a path, only where $\alpha_R (\chi_{\text{in}} - \chi_{\text{out}}) = 0$, that is where capability is not moving along the path or the two arms' capability elasticities agree.

Identity V-2 (separate arm paths). Where the arms run on their own trajectories $C_{\text{in}}(R)$ and $C_{\text{out}}(R)$, with $\alpha_{\text{in}} := d \ln C_{\text{in}}/d \ln R$ and $\alpha_{\text{out}} := d \ln C_{\text{out}}/d \ln R$, the depth derivative of the log-ratio is $\alpha_{\text{in}} \chi_{\text{in}} + \eta_{\text{in}} - \alpha_{\text{out}} \chi_{\text{out}} - \eta_{\text{out}}$, each term on its own path; further varying factors add their own chain-rule terms.

Counterexample to one-path identification. $Q_{\text{in}} = C^2$, $Q_{\text{out}} = C$ on the path $C = R$: the ratio equals R and widens, $d \ln G_{\text{emb}}/d \ln R = 1$, while $\eta_{\text{in}} - \eta_{\text{out}} = 0$; the widening is $\alpha_R (\chi_{\text{in}} - \chi_{\text{out}}) = 1$. One path therefore does not identify the partials; crossed or off-path variation is needed ([registration version 1.102](#), line 621), and no power law is needed for the chain rule itself.

Proposition V-3 (constant elasticities). If χ_{ol} and η_{ol} are constant on a rectangle of the $(\ln C, \ln R)$ plane containing the reference point (C_o, R_o) , then on it $\ln(Q_{\text{ol}}/Q_o) = \ln(Q_{\text{ol},o}/Q_o) + \chi_{\text{ol}} \ln(C/C_o) + \eta_{\text{ol}} \ln(R/R_o)$, and at matched order $G_{\text{emb}} = G_o (C/C_o)^{(\chi_{\text{in}} - \chi_{\text{out}})} (R/R_o)^{(\eta_{\text{in}} - \eta_{\text{out}})}$. Proof: $\ln Q_{\text{ol}} - \chi_{\text{ol}} \ln C - \eta_{\text{ol}} \ln R$ has zero partials, so it is constant on the connected rectangle; the constant is

fixed by the reference point, which must lie in the rectangle; the partials of the closed form return the constants. This is the reduced scaling form: earned where the elasticities are constant, assumed nowhere.

3.4 The estimable model (PROPOSED)

For cell (o, l), family and task, a candidate mixed model is

$$\ln(Q/Q_o) = \ln(Q_{ol,o}/Q_o) + \chi_{ol} \ln(C/C_o) + \eta_{ol} \ln(R/R_o) + u_{\{o,family\}} + u_{\{C,family\}} \ln(C/C_o) + u_{\{R,family\}} \ln(R/R_o) + u_{\{o,task\}} + u_{\{C,task\}} \ln(C/C_o) + u_{\{R,task\}} \ln(R/R_o) + \psi,$$

with random intercepts and random slopes by family and by task, their covariance structure fixed before any outcome, and the residual ψ . Where the design crosses arms within a family or a task, the random effects also carry the arm terms (the indicators and their slope interactions), because otherwise the standard errors of the very contrasts the kill conditions read are understated. The fixed part is saturated in the 2 x 2 cells: four intercepts, four capability slopes, four depth slopes. Equivalently, each of the three coefficient types is written with the indicators ind_first , ind_in and their product, and the two parametrisations are the same twelve parameters, since the cell design has rank four. A model with random intercepts only whose fixed part keeps six of these twelve terms (the common intercept, the intercept on the product of the two indicators, the common slopes on $\ln C$ and on $\ln R$, the order x $\ln C$ slope and the placement x $\ln R$ slope) omits the other six: the two indicator main effects, the placement x $\ln C$ and order x $\ln R$ slopes and the two three-way slope terms. The order contrast at placement l is $\chi_{first,l} - \chi_{after,l}$, which depends on l unless the three-way capability-slope term is zero, so the contrast is always reported at a named placement; the placement contrast at order o is $\eta_{o,in} - \eta_{o,out}$.

Estimands and controls. Any estimate of χ entering a P12 verdict is the estimate registration version 1.102, line 387 requires: a slope of one measured quantity on another carries error on both axes, the ratio of those errors is registered before any fitting, "the attenuation-corrected estimate is the one that enters a verdict", and no cell contributes an exponent unless it clears the identifiability gate of "at least four valid measurement intervals and at least one and a half decades of administered depth". Crossed or off-path variation in C and R is needed to identify the partials (registration version 1.102, line 621). Exposure, correction compute and information access are design controls: conditioning on

them defines a resource-normalised contrast and changes the estimand; P21's registered face is on its own registered access manipulation (registration version 1.102, line 1231), so a matched-access design is a different comparison (section 3.2). Capability is treated in the same way wherever an arm can move it: where the placement or the build order changes the capability the system attains, a comparison at fixed C conditions on a quantity the arm moves, so it is a controlled contrast, read as an effect of the arm only under assumptions stated before any outcome, and the total effect, the log-ratio of the two arms' service each along its own path, whose depth trend Identity V-2 gives, is reported beside it; administered capability, or target difficulty varied independently of the arm, is the preferred design factor where it can be had. The two estimands can differ in either direction. With the capability elasticity of service constant at χ : identical service surfaces give a zero contrast at every fixed C, while an arm that raises attained capability by a factor that grows with depth has a total effect of χ times the logarithm of that factor, which widens with depth where χ is positive; and an arm that doubles service at fixed C while halving attained capability has a contrast of $\ln 2$ at every fixed C and a total effect of $(1 - \chi) \ln 2$, which is zero at $\chi = 1$ and takes the opposite sign to the contrast where χ exceeds one. The slopes themselves are read the same way. Where $\ln(C/C_o)$ is attained capability rather than capability the design assigns, any latent factor that moves both capability and correction service enters the slope: in a linear model with the arm assigned at random, the slope on attained capability is the elasticity plus a term that has the sign of the product of the latent factor's effects on capability and on service and vanishes with it, the arm contrast at matched capability is the arm's direct effect less the arm's effect on capability times that term, and the total effect is the direct effect plus the arm's effect on capability times the elasticity. So χ_{ol} , its contrasts and G_{emb} read at matched C are identified where capability is varied by assignment (administered capability, or target difficulty crossed with the arms) or under a no-confounding premise stated before any outcome, and not by crossed or off-path variation in attained capability alone (registration version 1.102, line 621), which is necessary and not sufficient; the report says whether the model's $\ln(C/C_o)$ is assigned or attained. Curved and broken-regime rivals, a two-part or censored model for zero service, and a flexible comparator form the closed model set; the power law must win prospectively; no selection rule is chosen here.

3.5 The registered faces, whole

P21. Its refuter, verbatim (registration version 1.102, line 1227):

REFUTED BY: no interaction between correction placement and depth beyond the tolerance registered with that instrument, in a design whose own sensitivity shows an interaction of that size would have been detected. A design too small to find the interaction is a failed instrument and not a refutation, as condition D already provides.

A constant advantage is scored against it (registration version 1.102, line 1229): "an embedded arm that beats the external arm by the same margin at every depth is reported as a level difference and is scored against this proposition, not for it". A supported result establishes an access interaction, not physical or architectural embedding (registration version 1.102, line 1223), and not an exponent (registration version 1.102, line 1221); it does not resolve the trade-off the registration states (registration version 1.102, line 1235), quoted in section 3.7. Its instrument is a design draft (registration version 1.102, line 1243: "INSTRUMENT: DESIGN DRAFT, the embedded-versus-external correction designs. That unit's registered sensitivity was measured on a smaller design than the one registered, and its negative-slope hypothesis has no operating characteristic; both are recorded against that unit and are closed before this proposition is scored."). A widening advantage arriving without the registered fall in the monitor's detection rate is reported as unexplained (registration version 1.102, line 1233). What falls with it, verbatim (registration version 1.102, line 1237): "FALLS WITH IT: the protocol's engineering case, jointly with P12, P13 and P15. The three laws are untouched, which is why it is registered separately from them."

P12. Its refuter, verbatim (registration version 1.102, line 1035):

REFUTED BY: a preregistered manipulation of build order producing no change in the correction-leverage exponent beyond the tolerance registered with that instrument, where the instrument has shown from its own sensitivity that a change of that size would have been detected. A manipulation too weak to move anything is not a refutation of this prediction; it is a failed instrument, and condition D above says so in advance.

Its endpoint is the exponent on one scale, scored as a continuous claim (registration version 1.102, line 1039: "It is a continuous claim and it is scored as one."); the minimum effect is 0.10, and "A future instrument may register a smaller minimum and may not register a larger one" (registration version 1.102, line 1041); its instrument line reads NONE, "one endpoint away from an instrument" (registration version 1.102, line 1045);

what would decide it is a typed extension of the correction-exponent estimator across the maximal contrast with a joint interval (registration version 1.102, line 1049). Support is read on that joint interval: an interval establishing a change of at least the minimum effect, in either direction, supports P12; a change smaller than the minimum "is reported as no effect of interest rather than as no effect" (registration version 1.102, line 1041), whatever its significance; an upward change so established supports P12 and establishes the PROPOSED genesis reading, a downward change so established supports P12 and strikes it (section 3.6 item 2), and both verdicts are reported. A supported result "does not license any statement about where the boundary sits, about a system built in that order being safer, or about orders outside that contrast" (registration version 1.102, line 1047). What falls with it, verbatim (registration version 1.102, line 1043): "FALLS WITH IT: the programme's argument that build order is an intervention point rather than an accident of history, and this proposition's share of the protocol's engineering case, which it carries jointly with P13, P15 and P21."

3.6 Kill condition K-V

Stated before any data, consistent with P12 (no direction) and P21 as registered. The clauses that go beyond the registered refuters are PROPOSED law-level conventions and change nothing in P12's or P21's scoring.

1. Placement limb. Struck when P21's registered refuter fires: in the registered design (registration version 1.102, line 1231), passing its sensitivity and validity gates, no interaction between placement and depth beyond the registered tolerance. Struck equally by a constant advantage established within the registered tolerance by a design meeting the sensitivity clause, which is P21's refuter itself (registration version 1.102, lines 1227 and 1229). Struck, as a PROPOSED law-level convention, by an interaction interval lying wholly on the side in which the external arm's advantage widens with depth. A P21 refutation strikes the limb by these conventions; it does not measure the mechanism reading and does not prove $\eta_{0,in} \leq \eta_{0,out}$. A failed instrument is NOT EVALUABLE, never survival, and never a refutation.
2. Build-order limb. Struck when P12's registered refuter fires: a preregistered manipulation on the registered maximal contrast producing no change in the correction-leverage exponent beyond the registered tolerance (0.10 on the exponent scale, or the smaller minimum a future instrument registers), with demonstrated

sensitivity. On the joint interval that is an interval lying inside the tolerance band from -0.10 to +0.10, open at its ends, since a change of exactly the minimum is a change "by at least the registered minimum effect" (registration version 1.102, line 1033); an interval whose end falls inside the at-the-margin band the instrument fixed around the tolerance is a result at the margin, not a verdict (registration version 1.102, line 391). An interval establishing a change of at least the minimum in either direction supports P12 and the limb's registered face. The PROPOSED genesis reading of section 3.1, $\text{chi_first} - \text{chi_after} \geq 0.10$, is struck by one PROPOSED rule, stated here in advance: an interval on $\text{chi_first} - \text{chi_after}$, at the placement the instrument fixes and on the scale Condition B fixes (registration version 1.102, line 403; section 4.7), whose upper end lies below the registered minimum, 0.10, or below the smaller minimum a future instrument registers. The rule is the exact interval complement of the reading: the reading is false at every point of an interval exactly when the interval's upper end lies below the minimum, and it holds at every point exactly when the lower end is at or above it. The rule that strikes on an interval lying wholly on the adverse side (its upper end below zero) is a special case of this rule, since an upper end below zero lies below the minimum; and the rule strikes the reading wherever P12's registered refuter fires, since an interval inside the open band has its upper end below the minimum, as the reading's implying P12 requires. An interval whose upper end falls inside the at-the-margin band around the minimum is a result at the margin for the reading as it is for P12 (registration version 1.102, line 391). The rule needs no separate sensitivity clause, because an interval lying wholly below the minimum itself excludes an effect of the registered size in the claimed direction, but a failed instrument or a missing identification of the construct is not cured by a narrow interval. It strikes only the PROPOSED reading and changes nothing in P12. The worked cases, each read with its ends clear of the at-the-margin bands the instrument fixes: an interval from -0.30 to +0.05 is INCONCLUSIVE under P12's registered refuter, whatever the sensitivity, because it is not inside the tolerance band and so does not show "no change in the correction-leverage exponent beyond the tolerance registered with that instrument", and it does not support P12, because it does not establish a change of at least the minimum; it strikes the PROPOSED reading, since its upper end 0.05 lies below 0.10. An interval from +0.02 to +0.08, on which formation first carries the higher elasticity at every point but by less than the minimum at every point, strikes the reading as well, and with its sensitivity shown it

fires P12's registered refuter too, so the reading and P12 fall together there. P12 is read in neither direction by any of this.

3. Law V as a whole (aggregation PROPOSED; none in force). Three ways of reading the two limbs as one law are on the record: severable limbs, under which the whole law is refuted only when both limbs are struck, which was version 0.2.1's default; a strict conjunction, under which either strike refutes it; and a law that stands on the placement limb alone. The author has chosen none, and none is in force. Two things hold under all three. Each limb is scored by its own rule, items 1 and 2, and a struck limb is reported as struck. And where both limbs are struck, none of Law V's formal content is left standing. Where one limb is struck and the other stands, what is said of the whole law depends on the choice, which is made before any deciding data or not at all; until it is made the result is reported as that limb's and no more. No refutation of P15 withdraws Law V by any rule of this document (section 5 item 3). The PROPOSED readings, the genesis reading and the mechanism reading, are scored apart from both limbs. The boundary reading falls with Law III's model, which takes nothing from either limb.
4. What is not a refutation of Law V (in the registration's own terms, registration version 1.102, line 121): a failed instrument; a P20 estimate of the direct depth-dependence, which is scored in ARC-3 alone and may not cross scopes (registration version 1.102, line 247); an offsetting elasticity invented after a null; a level difference offered as an interaction; and any result on composition class or failure-mode independence, which P21's unit does not vary (registration version 1.102, line 1223).
5. The weaker neighbour. What survives Law V's refutation: the identities of section 3.3, the verdicts on P12 and P21 as engineering results under their own numbers, and Laws I to III untouched (registration version 1.102, line 1237); what falls is what registration version 1.102, lines 1043 and 1237 name, quoted in section 3.5.

3.7 What the embedding mathematics does not say

It does not say that embedded correction is safe, sufficient or superior on any axis but the one P21 scores. A supported P21 does not resolve the trade-off the registration states (registration version 1.102, line 1235): "embedding correction in the substrate that computes maximises the corrector's scaling with capability and minimises its independence from the generator, because a corrector built into the system shares the system's blind spots by construction. Those are the two axes named in the admissibility

conditions above, and embedding moves them in opposite directions. Whether the gain on scaling outweighs the loss on independence is not registered here and is not decided by this proposition." A supported P12 licenses nothing about safety (registration version 1.102, line 1047). The identities of section 3.3 identify no partial from one path, and the mechanism reading of registration version 1.102, line 1239 is measured by no verdict here. Nothing in this section has been measured.

4. The balance identity and Law III, unchanged

4.1 Identity B-1 (the local balance identity)

Let Q and W be positive differentiable correction service and newly generated burden on one clock, in comparable units with any conversion fixed before outcomes, on the same domain as section 3.2. Along a specified path $C(R)$, with other factors fixed,

$$\Delta_bal := d \ln(Q/W)/d \ln R = \alpha_R (\chi - \kappa) + (\eta - \zeta).$$

Proof: subtract the chain rules for $\ln Q$ and $\ln W$ along the path. Every elasticity is local; no power law is assumed. Δ_bal is the programme's balance exponent read locally, and the registration's $\Delta_balance = \phi_Q - \phi_W$ (registration version 1.102, line 619: "the balance exponent is $\Delta_balance = \phi_Q - \phi_W$: correction service minus burden"); positive means correction is gaining on burden. The registration's elasticities are log-derivatives "with respect to the resource clock" (registration version 1.102, line 619); Δ_bal above is per unit of $\ln R$, so the two agree where the clock is depth, and otherwise differ by the factor $d \ln R$ divided by $d \ln$ of the clock. That factor is positive, and so preserves the sign and the zero set, wherever depth rises with the clock at a positive rate; along a stretch on which depth fell as the clock advanced it would be negative and would reverse the sign. The clock condition is the author's: the service clock need not equal the resource coordinate, but it must be fixed within a model, and changing it changes the operational quantities.

4.2 Proposition B-2 (the balance root and the sign rule)

Where κ differs from χ , $\Delta_{\text{bal}} = 0$ has the single root $\alpha_R = \alpha_{\text{bal}} := (\eta - \zeta)/(\kappa - \chi)$, and identically $\Delta_{\text{bal}} = (\kappa - \chi)(\alpha_{\text{bal}} - \alpha_R)$. Hence: with $\kappa > \chi$ the trend is favourable exactly where $\alpha_R < \alpha_{\text{bal}}$, and α_{bal} is a ceiling, written α_{crit} ; with $\kappa < \chi$ the inequality reverses and α_{bal} is a floor; with $\kappa = \chi$ the trend is $\eta - \zeta$, independent of α_R . The symbol α_{crit} is used only in the ceiling case, so that the entry of the programme's notation, the stability ceiling, keeps its meaning; its range there, at least one, holds exactly where $\eta - \zeta \geq \kappa - \chi > 0$, which in the registered burden model ($\kappa = 1$, $\zeta = -1$) is $\eta \geq -\chi$ with $\chi < 1$: in the pacing case $\eta = -1/2$ with $\chi = 0$ gives a ceiling of one half; outside it, $\kappa = 1/2$ with $\chi = 1$, $\eta = 0$ and $\zeta = -1$ gives $\alpha_{\text{bal}} = -2$, a floor. A non-positive α_{crit} with $\kappa > \chi$ leaves no positive growth elasticity with a favourable trend. As χ rises to κ with $\eta - \zeta > 0$ the boundary grows without bound: a correction elasticity matching the total burden elasticity leaves no finite boundary. The zero set of Δ_{bal} , the balance frontier, is this set of points; it depends on the path field α_R as well as on the service and burden surfaces. Once the path field is specified, Δ_{bal} is a function on the (C, R) plane and its zero set is a level set of that function: where Δ_{bal} is continuously differentiable, the set is a curve near each of its points at which the gradient of Δ_{bal} does not vanish (the implicit function theorem), so it is generically a curve, but not always: in the pacing case at $\chi = 1$ with $\eta = -1$, Δ_{bal} is zero everywhere, and at $\chi = 1/2$ and $\eta = 0$ with the field $\alpha_R = 1$ it is one half everywhere, so the set is then the whole plane or empty. Before the field is specified the set lies in (C, R, α_R) , where the derivative of Δ_{bal} in α_R is $\chi - \kappa$, so, again where the elasticities are continuously differentiable, it is a surface near each of its points at which κ differs from χ ; it carries no symbol here and Law III is unchanged (section 6).

4.3 Corollary B-3 (the registered special case)

Under the registered premises, burden proportional to dC/dR along a power-law path, read as the programme's notation reads the registered model, the surface W proportional to $\alpha C/R$ with $\kappa = 1$ and $\zeta = -1$ (one path alone does not fix the partials, section 3.3), and correction proportional to C^χ ($\eta = 0$), $\alpha_{\text{crit}} = 1/(1 - \chi)$: the frozen Law III, verbatim (registration version 1.102, line 665): "Law III, the ARC Ceiling, $\alpha_{\text{crit}} = 1$

/(1 - gamma): the growth exponent is capped by the reciprocal of the correction shortfall." The registration's letter, the one inside that quotation, and this document's chi name the same quantity (registration version 1.102, line 305). The mapping is a symbol mapping inside the stated pacing model and is not permission to identify Law II's beta_C with either: "WHETHER THE SECOND LAW'S EXPONENT AND THE THIRD LAW'S ARE ONE QUANTITY IS NOT SETTLED HERE AND IS RECORDED AS UNESTABLISHED BELOW, and no proposition assumes they are." (registration version 1.102, line 271). Law II is "beta_C exceeds k" (registration version 1.102, line 663). Under the burden model the drift exponent implied for a power-law path is $k = (\alpha - 1)/\alpha$, and the criterion $\beta_C > k$ rearranges to $\alpha < 1/(1 - \beta_C)$ for $\alpha > 0$ and $\beta_C < 1$; this returns the ceiling relation only under the identification $\beta_C = \chi$, which the registration records as unestablished (registration version 1.102, lines 271, 309 and 613) and which needs an intervention designed to move one while leaving the other before either stands in for the other (registration version 1.102, line 1029). Under that identification, and only under it, "The ceiling is the co-scaling criterion with the drift exponent fixed by the burden model." (registration version 1.102, line 637), and Law III remains a conditional quantitative consequence of Law II under its stated identifications, not an independent finding.

4.4 The pacing family

In the pacing case ($\kappa = 1$, $\zeta = -1$) with a direct depth term, $\alpha_{\text{crit}} = (1 + \eta)/(1 - \chi)$, a positive ceiling for $\chi < 1$ and $1 + \eta > 0$; at $\chi = 1$ the trend is $1 + \eta$ independent of α_R ; at $\chi > 1$ with $1 + \eta \geq 0$ the trend $\alpha_R (\chi - 1) + (1 + \eta)$ is positive for every positive α_R , as the sum of a positive and a non-negative term. At $\chi = 1/2$ this family gives $2(1 + \eta)$. That value is not the general boundary: $2(1 + \eta)$ is the pacing-case value; the general boundary at $\chi = 1/2$ is $(\eta - \zeta)/(\kappa - 1/2)$, which equals $2(1 + \eta)$ for every η only when $\kappa = 1$ and $\zeta = -1$; at single points the two can agree elsewhere ($\kappa = 3/2$, $\zeta = -2$ and $\eta = 0$ give two on both sides), a coincidence at one point that does not make $2(1 + \eta)$ the general boundary. So, in the pacing special case, at the same-class cap $\chi = 1/2$, the boundary is $2(1 + \eta)$, and the ARC Bound, $\alpha \leq 2$, is its $\eta = 0$ value.

The registration's four-quantity form (registration version 1.102, line 639: "the boundary is one plus the difference of the two direct depth-dependences, divided by one plus the burden intensity less the correction elasticity") is the same boundary written with the registered model's values subtracted: the burden intensity of that sentence is $\kappa - 1$ (the programme's μ_B) and its direct depth-dependence of burden is $\zeta + 1$, not ζ , whose name in the programme's notation, the direct depth-dependence of burden, is the full depth elasticity of W at fixed C ; so $(1 + \eta - (\zeta + 1))/(1 + (\kappa - 1) - \chi)$ is $(\eta - \zeta)/(\kappa - \chi)$ identically, and the registration's "both direct terms and the burden intensity vanish" is $\zeta + 1 = 0$ and $\kappa - 1 = 0$ with $\eta = 0$. Its worked numbers hold: $\chi = 1/2$ with $\eta = 1/2$ gives three; $\chi = 4/5$ gives five; both together give seven and a half. The absolute rung (registration version 1.102, line 609) subtracts the correction elasticity from two rather than from one and at the registered dials gives two thirds; its derivation is the registration's and is not reproduced here.

4.5 Proposition B-5 (varying growth)

If burden is proportional to dC/dR along an arbitrary smooth increasing path, then W is proportional to $\alpha_R C/R$ and not to C/R , and with correction proportional to C^χ at constant χ ,

$$\Delta_{bal} = 1 - \alpha_R (1 - \chi) - d \ln \alpha_R / d \ln R,$$

which is registration version 1.102, line 623 in symbols ("the trend of the correction-to-burden ratio is one, less the local elasticity multiplied by the correction shortfall, less the trend of the local elasticity itself"). The last term vanishes throughout a range of depth only where α_R is constant on that range, that is only for a power law there, where $\Delta_{bal} = 1 - \alpha (1 - \chi)$; at a single depth it vanishes wherever α_R is stationary, which a path that is not a power law can reach: $\ln(C/C_o) = (\ln(R/R_o))^3/3 + \ln(R/R_o)$ has $\alpha_R = 1 + (\ln(R/R_o))^2$, stationary at $R = R_o$. Otherwise the general chain rule must carry the dependence inside W or carry α_R as an extra state variable. Merely replacing a constant α by α_R in the simplified ceiling loses this term.

4.6 Proposition B-6 (sensitivities, the gap identity, and what an intervention does)

With the other quantities fixed, $d \alpha_{bal}/d \eta = 1/(\kappa - \chi)$; $d \alpha_{bal}/d \chi = (\eta - \zeta)/(\kappa - \chi)^2$; $d \alpha_{bal}/d \kappa = -(\eta - \zeta)/(\kappa - \chi)^2$; $d \alpha_{bal}/d \zeta = -1/(\kappa - \chi)$. So raising η raises the ceiling when $\kappa > \chi$, and raising χ raises it only when $\eta - \zeta > 0$. In the pacing case at fixed η , the build-order gap is $\alpha_{crit}(first) - \alpha_{crit}(after) = (\chi_{first} - \chi_{after}) \alpha_{crit}(first) \alpha_{crit}(after)/(1 + \eta)$ in the positive-ceiling domain. In the same pacing case, with $\kappa = 1$ and $\zeta = -1$ in both arms, where a placement contrast moves both terms, in-loop placement raises the boundary if and only if $(1 + \eta_{in})/(1 + \eta_{out}) > (1 - \chi_{in})/(1 - \chi_{out})$, with $1 - \chi > 0$ and $1 + \eta > 0$ in both arms; one favourable term alone is insufficient, and an intervention that also raises burden can lower the boundary. Outside the pacing case the comparison is between the two arms' balance roots, each with its own burden elasticities, and the pacing condition does not carry over: with $\kappa = 2$ and $\zeta = -1$ in both arms, $(\chi, \eta) = (1/2, 1/5)$ in the loop and $(2/5, 3/10)$ outside it, the pacing condition holds, yet the in-loop root is $4/5$ and the other $13/16$, so in-loop placement lowers the boundary. These are algebraic regimes, not safety conclusions: the trend is not the level (registration version 1.102, line 625: "this quantity says whether the ratio is IMPROVING, never whether it EXCEEDS ONE."), and $Q/W = R/100$ improves throughout R in $[1, 10]$ while staying below one. Trend, level, backlog and event stay distinct.

Local elasticities throughout make this a local boundary; a boundary written as a function of the architecture (whose letter for the architecture is one of the eight colliding symbols and is not printed; its root reads every elasticity, the burden's included, in the architecture) is the statement that changing the arm changes the arm's elasticities and so moves α_{bal} through these sensitivities; it is written here per arm without an architecture symbol, and where the burden elasticities also differ between arms the comparison is the general one above.

The ARC Bound reading. The Bound is the same-class value of the pacing model at $\chi = 1/2$ and $\eta = 0$, not a fourth or sixth law. A control stands: a study quoting the simple form declares the corrector's direct depth-dependence before the run and shows it near zero by an equivalence result with a pre-fixed band whose upper edge is at most half the margin the

deciding instrument registers for the profile maximum; a wider band, or a missing demonstration, is NOT EVALUABLE for the Bound reading, never survival. That margin is not yet a number: P3's tolerance "is the author's to fix" and until then P3's supported outcome and this control are NOT EVALUABLE (registration version 1.102, line 843), which section 7 records. No numerical band is chosen here; P3's statement, refuter (registration version 1.102, line 851) and scoring are unchanged, and a positive eta asserted after a run never excuses a P3 exceedance.

The Boundary Transfer (PROPOSED). A held-out prediction of the balance transition from independently estimated elasticities inherits their precision. The registration reports the growth exponent's robust re-estimate as approximately 0.49, a regression estimate with a standard error of 0.20, and the 95 per cent interval [-1.3, 2.9] as the bootstrap interval on the same row's endpoint estimate of 0.59 (registration version 1.102, line 207); that interval's positive part contains $\alpha_{crit} = 2$ at the registered dials, and Δ_{bal} is positive at either point estimate and negative at 2.9, so no transfer prediction built on it has a decided sign; the non-positive part of the interval lies outside the burden model, whose burden must be positive. A transfer unit therefore carries its own precision gate, fixed before the transition data are opened: the half-width of the predicted boundary's interval small against the band the boundary-mapping unit registers for the boundary location, the registered standard error of its estimator (registration version 1.102, lines 621 and 633); it is a proposed study, not law content.

4.7 Proposition B-7 (the capability coordinate)

Let capability be re-expressed by an increasing differentiable map whose local elasticity, the derivative of the logarithm of the new coordinate in the logarithm of the old, is positive throughout. Read in the new coordinate, the growth elasticity α_R is multiplied by that elasticity, χ and κ are divided by it, and η , ζ and Δ_{bal} are unchanged; so $\kappa - \chi$ keeps its sign, α_{bal} is multiplied by the same elasticity, the verdict $\alpha_R < \alpha_{bal}$ is unchanged, and every contrast of capability elasticities between arms at a common capability ($\chi_{first} - \chi_{after}$, $\chi_{in} - \chi_{out}$) is divided by it. Proof: χ in the new coordinate is the derivative of $\ln Q$ in the new log-capability at fixed R , which is χ times the derivative of the old log-capability in the new, the reciprocal of the map's elasticity, and likewise κ ; η and ζ are taken at fixed capability, which the new coordinate fixes exactly when the old one does; α_R is multiplied by the elasticity by

the chain rule; and $\Delta_{bal} = d \ln(Q/W)/d \ln R$ along the path names no capability coordinate. Signs, ceilings against floors and every verdict on the balance are therefore free of the coordinate, and numbers are not: on the squared coordinate, C^2 , an order contrast of 0.15 is 0.075, below P12's minimum effect. The registration fixes the numbers on one scale: "Every numerical exponent and every exponent-scale margin in this registration is stated on the scale this condition fixes and on no other" (registration version 1.102, line 403, Condition B), and for the same-class value, "THE BOUND IS FIXED ON ONE CAPABILITY SCALE AND ON NO OTHER" (registration version 1.102, line 935). $\alpha_{crit} = 1/(1 - \chi)$, the ARC Bound's two, the same-class value one half, P12's minimum effect 0.10 and the genesis reading's margin are numbers on that scale. A meaningful zero and an uncapped range, which the registration gives its ladder scale (registration version 1.102, line 259), do not by themselves fix it: every positive power of C keeps both and multiplies the growth elasticity by the power, so the justification Condition B requires must also exclude power re-expressions. The registered burden model's values are themselves values on that scale (on C^2 , $\kappa = 1$ becomes one half), so re-expressing capability alone and reapplying $1/(1 - \chi)$ shows nothing: with $\alpha = 9/5$ and $\chi = 9/20$ the bound is 20/11 on C; on C^2 , $\alpha = 18/5$, $\chi = 9/40$ and $\kappa = 1/2$ give 40/11, and the verdict stands, while $1/(1 - 9/40) = 40/31$ would reverse it. One inconsistency of the frozen registration bears on Law V and is carried to the successor registration (section 7 item 15): registration version 1.102, line 403 states every exponent-scale margin on Condition B's scale, but Condition B's lists (registration version 1.102, lines 409 and 415) leave P12 out, and Condition E2 (registration version 1.102, line 421) names P12 among the predictions that return "no correction exponent", although P12's endpoint is "the quantity the third law's ceiling is written in" (registration version 1.102, line 1039); P12's 0.10 rests on Condition B's scale through registration version 1.102, line 403 alone.

4.8 Proposition B-8 (PROPOSED: the balance over a finite range of depth)

Identity B-1 gives the trend of the ratio Q/W at a point. Over a range of depth it gives the change in the ratio's level. On one specified trajectory, with Q and W positive, in comparable units, on one clock and with one conversion fixed throughout, let $0 < R_1 < R_2$ lie in the domain and let Δ_{bal} be integrable in $\ln R$ between them. Then

$\ln[Q(R_2)/W(R_2)] = \ln[Q(R_1)/W(R_1)] + \text{integral from } R_1 \text{ to } R_2 \text{ of } \Delta_{\text{bal}}(r) d \ln r.$

where $Q(r)$ and $W(r)$ are read along that trajectory and r is the variable of integration.

Proof: by Identity B-1, Δ_{bal} is the derivative of $\ln(Q/W)$ in $\ln R$ along the trajectory; integrate it from R_1 to R_2 . So the level at the start of a range is part of any prediction of the level at its end, and the trend alone predicts neither. For a constant trend Δ_{bal} other than zero, on a trajectory anchored at R_o , the ratio is $Q(R)/W(R) = [Q(R_o)/W(R_o)] (R/R_o)^{\Delta_{\text{bal}}}$, and it meets one at

$R_{\text{star}} = R_o [W(R_o)/Q(R_o)]^{(1/\Delta_{\text{bal}})}.$

With a positive trend the crossing lies beyond the anchor exactly when $Q(R_o) < W(R_o)$: where service already exceeds burden at the anchor the expression returns a depth before it, and where the two are equal it returns the anchor. With a negative trend the same expression gives the depth at which the ratio falls through one, which lies beyond the anchor exactly when $Q(R_o) > W(R_o)$. With a zero trend the ratio keeps its starting level throughout. R_{star} is a formal crossing: nothing in it shows that the depth lies inside the range measured, or inside any range a system can reach. Nor does a ratio above one bound a backlog, a transient error or a harm; trend, level, backlog and event stay distinct (section 4.6). A prediction made with this proposition therefore records in advance the levels at the anchor with their uncertainty, the range of depth it covers, the clock and the conversion, and the separately measured outcome it is meant to reach; and none of these is chosen with the held-out observation in view.

The coordinate. Proposition B-7 governs any change of the capability coordinate, and the power case shows the pattern: where capability is re-expressed so that its ratio to its reference value becomes $(C/C_o)^p$ with $p > 0$, α_R becomes $p \alpha_R$, χ becomes χ/p and κ becomes κ/p , while η , ζ and Δ_{bal} are unchanged, and so is the integral above. To change the coordinate and keep the old number for the burden's capability elasticity is to change the model (section 4.7). Every numerical threshold on an exponent, P12's minimum effect among them, stays on the scale Condition B fixes. Nothing in this proposition replaces Law II's criterion, that β_C exceeds k , or removes an identification on which Law III rests.

4.9 What a new deciding instrument declares (PROPOSED)

A new instrument that reads correction service against burden declares each of the following before any of its outcomes is read. The list is PROPOSED as a condition of such an instrument's completeness; it amends no frozen unit.

IT DECLARES	SO THAT THIS IS SETTLED IN ADVANCE
The system's state, and what may change in it	code, weights, tools, memory, permissions, the state of training and the mechanism of update; a loop that runs inference on fixed weights is described as that
Capability as measured: its zero, scale and range	the case for a ratio scale, the calibration, any floor or ceiling of the measure, and which changes of coordinate are allowed (section 4.7)
Depth, and the clock on which service and burden are read	how rounds are counted, the reference values, the unit of exposure, and any conversion between clocks with its uncertainty
What is counted as correction	capacity available, corrections attempted, corrections delivered and the net change in error, each kept apart; and whether the workload offered limited what could be delivered
What is counted as burden	the same unit of exposure and the same clock as correction, with any conversion justified before Q/W is formed
The units of observation, and which are independent	systems, training replicates, model families, checkpoints and repeated items, with the dependence among them kept in the analysis
The intervention and the estimand	whether capability is assigned or attained; whether the quantity is a total effect or a contrast at fixed capability (section 3.4); the range the arms share; and the error of measurement on each axis
The model and its primary contrasts	an intercept and both slopes for every cell, or the equivalent full set of interactions (section 3.4), with the order contrast and the placement contrast named
The decision rule and its uncertainty	the registered endpoint, the smallest effect of interest, the joint interval, and the rules for missing data, for zeros and for stopping
The planner, where one is studied	its complete objective, horizon, feasible family and permitted early stops, its discounting and selection tolerance, and how coverage and realisation are checked (section 2.5)

The claim fixes the measurement, never the reverse. P12 is read on its undirected contrast in the correction-leverage exponent, on the registered axis. P21 is read on its behavioural outcome, the interaction of placement with depth, within its registered access and action-space design. P15 is read on the retained fraction of an externally installed gain after later capability training. A ratio of correction service, a behavioural gain at one fixed budget, or a comparison of retained fractions is none of these three endpoints and is never scored in place of one. A directional reading or a mediation hypothesis is a candidate for a registration of its own; it is not part of the successor to version 1.102 until it has been reviewed and given an instrument.

5. Dependencies, logical separation and conventions

1. Register dependencies (the statement paper's dependency table, in its section 2a):
 ARC-1's failure does not refute the persistence registers; ARC-2 depends on ARC-1's measurables, not its verdict, and its failure takes ARC-3's derivation; ARC-3's failure does not erase ARC-2's own result. Law IV's objects need no growth law; Law V's two limbs need none; only Law V's boundary reading depends on Law III's model and through it on ARC-2, and it needs the same system and admissible scopes.
2. Logical separation, at the level of verdicts. The registration states it: "The laws are untouched." (registration version 1.102, line 1087, for P15) and "The three laws are untouched" (registration version 1.102, line 1237, for P21). Law III's relation holds at every value of chi, so any P12 verdict, which is a contrast in the very symbol Law III's equation is written in (registration version 1.102, line 1039: "the quantity the third law's ceiling is written in"), is consistent with it; separation is independence of truth values, not absence of the symbol. Three listed models make it concrete, each at the registered dials ($\chi = 1/2$, $\eta = 0$, $\kappa = 1$, $\zeta = -1$) with a growth exponent of $3/2$, so that the growth exponent lies below the ceiling of two and the balance exponent is positive, $\Delta_{\text{bal}} = 1/4$: (a) an interval on $S_{\text{ext}}(1)$ from 0.85 to 0.95, which would refute P15 and so Law IV's formal content; (b) an interval on the build-order change in the exponent lying inside P12's open tolerance band, and an interval on the placement-by-depth interaction on the outcome lying inside the tolerance registered with P21's instrument, each from a design meeting its sensitivity clause, so that both registered refuters would fire and both limbs of Law V would be struck; (c) an interval on $S_{\text{ext}}(1)$ from 0.2 to 0.4 with the intervals of model (b), under which

Law IV's face would be supported while both limbs of Law V would be struck. Across the three the verdicts on P12, P15 and P21 vary while the dials, and every statement of sections 4.1 to 4.4 about them, stay fixed. The dials do not by themselves establish Laws I to III: Law I is a claim about measured growth, and Law II's criterion compares β_C with k , which the registration does not identify with this model's elasticities (section 4.3). Model (b) is written on the registered endpoints because P21 decides an interaction on the outcome, "never a measurement of an exponent" (registration version 1.102, line 1221), and its refuter reads that interaction (registration version 1.102, line 1227); read on the service surfaces instead, $\eta_{in} = \eta_{out}$ would make G_{emb} constant in depth along a growing path only together with $\chi_{in} = \chi_{out}$ (Identity V-1), and its constancy is not what fires P21's refuter. These are separations of statements, not empirical findings. A countermodel proposed for the case in which Laws I to III hold while Law IV fails takes a collapse, $S_V(x) = \exp(-10 x)$, which on the external arm would support P15; it reads the failure of Law IV in the comparative sense of section 2.6, the opposite direction from model (a), which is one more reason the comparative content is PROPOSED and the author's first question. Two registered links are recorded rather than left to be found: registration version 1.102, line 663 carries Law II "by P4, with P7 asking what class deployed correction actually belongs to, and P12 and P15 asking whether placement and subsequent training move it"; and registration version 1.102, line 667 holds Law III's relation only where "the direct depth term is zero", which the mechanism reading $\eta_{o,in} > \eta_{o,out}$ denies in at least one placement, so that where that reading holds Law III's simple form, and the ARC Bound read from it, apply in at most one placement; that link is confined to the boundary reading and to the η term of the Bound-reading control, and a P20 verdict is scored in ARC-3 alone.

3. The register convention (PROPOSED by version 0.2.1; not in force). Version 0.2.1 proposed, as a grouping fixed in advance and not as an entailment, that on an evaluable P15 refutation Law IV be withdrawn whole, and Law V with it as a law of the theory; that Law V's refutation not withdraw Law IV; and that P12 and P21 keep their own measured verdicts. Two published statements stood behind it, and they are different statements. The registration ties a P15 refutation to "this programme's case that correction has to sit inside the system rather than be applied to it from outside" (registration version 1.102, line 1085). The statement paper's dependency table has "ARC-4 falls and takes ARC-5 with it", and the trigger of that fall in the same paper is

the comparative observation of its kill sentence (section 2.6), not P15. This document makes neither into a rule by which a P15 refutation withdraws Law V: nothing in $\eta_{0,in} > \eta_{0,out}$ or in the P12 endpoint is entailed by any value of $S_{ext}(1)$, and Law V is struck only through its own limbs (section 3.6). A dependency between the two laws, if the author fixes one, names its trigger and its consequence before any deciding data (section 7 item 3).

4. No P12, P15 or P21 verdict automatically refutes another numbered proposition. P19's refutation removes ARC-5's motivation and refutes nothing in it (registration version 1.102, line 111). P13 stays outside the laws.
5. The Eden Protocol is a method, not a law, and is not entailed by these laws: "It follows from no law here, it is carried jointly by P12, P13, P15 and P21, with P7 as its premise and P19 as its motivation, and it can fail while every law stands or stand while they fall." (registration version 1.102, line 61); its endogeneity condition is tested by the method's own registrations and no verdict here turns on it (registration version 1.102, line 675). Its success criterion is a formation criterion, not a later-conduct criterion. The Honey Architecture is likewise outside every verdict here.
6. The flagship conjunction P1, P4, P16, P20, P21 stands as registered (registration version 1.102, line 763); a surviving subset is not the conjunction.
7. A persistence advantage and a favourable elasticity contrast are separate: an architecture can preserve a property with no favourable contrast, or improve scaling while losing it; a positive retention comparison alone is not causal mediation (section 6).

6. The PROPOSED register

Each entry is a proposition or object this document states beyond the registered propositions. None is law content until the author registers it. Each carries what it needs. "The author's" is reserved for his own records.

OBJECT	STATUS	WHAT IT NEEDS
$K_V(H; T, \epsilon_V)$, the persistence kernel, formalising the invariance object of the author's working paper of 16 August 2026	PROPOSED primary invariance object; Lemma IV-K proved	a registered finite battery or a probabilistic family with tolerance; thresholds; horizon; the interval rule of section 2.2
$\pi_V, \Pi_V(H)$, the removal margins	PROPOSED premise quantities of Theorem IV-G	a registered causal decomposition; an objective scored independently of V (premise (Go))
$M_V^{\wedge}(H)$, the strategic viability margin, read path-wise (section 2.5)	ADOPTED as PROPOSED; Theorem IV-H and Lemma IV-M' proved	the planner's declared objective, feasible family and tolerance; a registered search protocol with an upper bound on the safe supremum; reported one-sidedly
Proposition IV-N, a planner that scores its plans with bounded error (section 2.5.1)	PROPOSED; proved	a bound on the score's error over the whole feasible family, fixed before any outcome; the margin computed in the planner's complete objective
the robust margin	PROPOSED; not printed under a colliding symbol	a registered environment family
local elasticities χ_{ol} , η_{ol} (the author's working paper of 16 August 2026; <u>the registration</u> of 8 September 2026, <u>registration version 1.102</u> , line 267)	ADOPTED as the author's own reading, dated 16 August 2026; arm-labelled here	crossed or off-path design for identification; the estimate of <u>registration version 1.102</u> , line 387

OBJECT	STATUS	WHAT IT NEEDS
the balance frontier, the zero set of Delta_bal (the general boundary form of the author's working paper of 16 August 2026)	PROPOSED object without a symbol; Law III unchanged; not a restatement of Law III	a specified path field alpha_R; a precision gate
Lambda_V, the mediation bridge	ADOPTED as PROPOSED, with its own status (below)	its own registration in a scope that carries both contrasts; both retained fractions established (section 2.6.1) and positive; an identified estimate of each link; an adverse margin; multiplicity
IV-PERSIST, S_emb(x) > S_ext(x) (the statement paper's comparative sentence and kill, recorded and not adopted)	the same test as Lambda_V on the difference scale; one object, two readings; not Law IV's content in this version	as for Lambda_V; the author's first decision (section 7)
the measurement of the retained fraction on three arms, and Proposition IV-R (section 2.6.1)	PROPOSED instrument requirements; Proposition IV-R proved for <u>Fieller's</u> method under the gate	the no-gain arm trained alongside; the floor on each installed gain; the estimand, the interval method and its critical value named before any outcome
the estimable model of section 3.4	PROPOSED model, not a required analysis of any frozen unit	the closed model set; a selection rule fixed in advance; <u>registration version 1.102</u> , line 387
the candidate retention families and the descriptive statistics	PROPOSED reporting; never verdicts	the closed family set; the author's budget scale and baseline convention
the interval rule for Omega_V membership	PROPOSED kill, not law content	registration of the battery, tolerance and interval method

OBJECT	STATUS	WHAT IT NEEDS
(section 2.2) and the battery-relative kernel kill (K-IV item 3)		
the law-level clauses of K-IV items 2 to 4 and K-V items 1 to 3, and the register convention of section 5 item 3, beyond the registered refuters	PROPOSED conventions; K-IV item 2, K-V item 3 and the register convention are on the record and not in force; P12, P15 and P21 scored as registered	the author's successor decision
Proposition B-8, the balance over a finite range of depth, and the formal crossing (section 4.8)	PROPOSED; proved from Identity B-1	the levels at the anchor with their uncertainty; the range, the clock and the conversion declared in advance
what a new deciding instrument declares (section 4.9)	PROPOSED condition of a new instrument's completeness; amends no frozen unit	adoption in each new instrument's own registration
the recursive-improvement prerequisite (section 3.2)	PROPOSED prerequisite in words; both mutual-information forms declined as criteria	the author's registration before any surface states it as a condition of any law; until then a description, and not a law
the genesis reading, $\text{chi_first} - \text{chi_after} \geq 0.10$ (formation first carries the higher elasticity by at least P12's registered minimum effect)	PROPOSED directional reading beside an undirected P12, stated at P12's registered minimum so that it implies P12 and adds only a sign, with one strike rule, its exact interval complement (K-V item 2)	the author's successor decision; P12 keeps no direction
the Boundary Transfer	PROPOSED study	its own preregistration; a precision gate

OBJECT	STATUS	WHAT IT NEEDS
planned removal (the planned-removal challenge)	PROPOSED study; the observation that decides IV-H's margin premise (Lemma IV-M')	its own preregistration; a registered upper bound on the safe supremum
a restatement of Laws I to III under new names	DECLINED: not adopted; Laws I to III stand as registered; the glosses growth, balance, boundary, persistence, architecture are descriptions, not names	the author's decision, and never in 1.102
a mutual-information positivity as an improvement criterion, state-based or update-based	DECLINED as a criterion, for the reason of section 3.2; the prerequisite is kept in words	nothing
the half-life as P15's verdict; an interval capped at one, the kill written as the interval inside $[1/2, 1]$	DECLINED (section 2.6)	nothing

Ruling on the Lambda_V mediation bridge. Adopted as PROPOSED, with this status line: "PROPOSED mediation contrast; not law content; $\text{Lambda_V} := \ln(\text{S_emb}/\text{S_ext})$ at the evaluation budget, defined only where both retentions are positive; the primary contrast is the difference $\text{S_emb}(x) - \text{S_ext}(x)$, which is defined wherever both retained fractions are established, whatever their signs, with Lambda_V reported beside it where defined; neither is defined without the embedded arm, and both need the measurement of section 2.6.1; instrument: a unit registered in a scope that carries both the retention contrast and the elasticity contrasts, since P15's unit is governed by S5 (registration version 1.102, line 237) and P12 and P21 by S7 (registration version 1.102, line 239), and no proposition may borrow a scope (registration version 1.102, lines 241 and 247), so a retention contrast read on the registered two-arm unit (registration version 1.102, line 1089) alone is an S5 quantity that bears on no ARC-5 proposition; a positive value is not mediation: the proposed mediating sequence, embedding to the elasticity contrasts to Delta_bal to a

persistence advantage, has three links, each needing its own identified estimate, and each link can fail; the adverse margin, multiplicity and registration are the author's." Reasons. It is adopted rather than declined on its merits: it is the one object in the set that would connect the two registers once one system is registered in a scope that carries both contrasts. It is stated on the difference scale first because the log form is undefined on a region the data can reach (where one retention is negative and the other positive, or where either is zero), and because where both retentions are negative the log form is defined but its sign is the reverse of the difference's. It is one object with IV-PERSIST, so that one test does not become two propositions. Nothing states that the loop is closed.

Two things follow for the successor registration. A result on the difference and a result on Λ_V are one result, not two corroborations. And neither enters the successor to version 1.102: Λ_V , its margins and any instrument for it are candidates for a registration of their own (section 7 item 9).

7. Open questions

There are fifteen questions in all, numbered below. Where an item names a default, that default stands until the author says otherwise, and none is chosen after a result. Items 2, 3 and 4 name none: the conventions version 0.2.1 proposed there are on the record and are not in force. The successor to registration version 1.102 is drafted as version 1.103 and awaits the author's registration (item 15).

1. Law IV's content (the first decision, on which item 2 turns): P15's face alone, as the author's statement of the laws gives it and this version adopts; or the comparative claim, $S_{emb}(x) > S_{ext}(x)$, with a registered comparative refuter and instrument, which the statement paper's ARC-4 sentence and kill (section 2.6) presuppose. Default: P15's face alone. The loss mechanism, active removal against dilution, stays unresolved under either.
2. What a P15 refutation takes from Law IV beyond its formal content, which K-IV item 1 already shows false: whether the law is then withdrawn whole under the four controls version 0.2.1 proposed (K-IV item 2), and whether a comparative headline is kept, which would need its own registered refuter before it could stand as a law. No default.

3. The register convention: whether a refutation in ARC-4 withdraws Law V as well as Law IV, and if so on which trigger, P15's refutation or the published comparative observation (the statement paper's kill sentence, section 2.6). No default: this document withdraws nothing from Law V on a P15 refutation (section 5 item 3).
4. Whether Law V is severable by limb, a strict conjunction, or stands on the placement limb alone. No default: until the author chooses, each limb's verdict is reported as that limb's. The three readings agree where both limbs are struck and where neither is, and differ where exactly one is (section 3.6 item 3).
5. The axis P21's result is read on: the registered access interaction (default), or the service-scaling mechanism reading of registration version 1.102, line 1239, which would need a separately specified bridge and design, never a rescoring.
6. Whether the build-order limb is scored with class left free, as P12 registers (default); holding it fixed would be a changed contrast.
7. The budget scale on which P15's matched budget is read (registration version 1.102, line 1089), and the at-the-margin band of registration version 1.102, line 391 in the same deciding instrument; until fixed, the threshold is NOT EVALUABLE on that scale.
8. The baseline convention for the retained fraction (the concurrent baseline, the default, which needs a third arm, a system without the installed gain that undergoes the same later training; or the pre-training baseline, computable on the registered two arms; section 2.6) and the value of Y_min. The floor's form (a lower confidence bound; the interval method) is a decidability requirement this document fixes, not his choice.
9. Whether Lambda_V is registered as a PROPOSED mediation proposition, in which scope, with what adverse margin and multiplicity, and whether a unit is registered that carries both contrasts (section 6). It is not part of version 1.103, and it takes a proposition number of its own only if its claim and its refuter differ from those already registered.
10. Whether the genesis reading of the build-order limb, $\chi_{\text{first}} - \chi_{\text{after}} \geq 0.10$ (formation first carrying the higher elasticity by at least P12's registered minimum effect), is registered as a PROPOSED clause beside an undirected P12, with the one strike rule of K-V item 2. It is not part of version 1.103, and P12 keeps no direction.
11. P3's tolerance (registration version 1.102, line 843), on which the Bound-reading control of section 4.6 depends; until fixed, that control is NOT EVALUABLE.

12. Whether the kernel's family, if he registers the kernel, is a finite registered battery or a probabilistic family with a tolerance; the per-unit numbers (v_* , ϵ_V , H) are fixed in each unit's registration, not in this document.
13. The dating: confirmation of the dates this document draws from his working paper of 16 August 2026; and, once his lines of 8 December 2024 are confirmed at source, the dating of each law's first statement from them.
14. The scoring-gate extensions for P12 and any scorer-dependent P15 reading, and whether each law carries a stated confidence (no law-level confidence is stated here).
15. The successor registration. It is drafted as version 1.103, and its registration and its date are the author's. It is to keep every word, date, scope and refuter of P1 to P22 and to list each correction to an instrument line or to the record separately, with its reason; among them is P12's listing under Condition B, which registration version 1.102, line 403 implies and registration version 1.102, lines 409, 415 and 421 omit (section 4.7).

Not his, and fixed in this document: the planner's objective form for a studied system is a fact the unit declares and verifies, and for a designed planner a design parameter of that PROPOSED unit; the rule (compute M_V^H) in the declared objective; read search one-sidedly) is this document's. The retention families and the Law V model set are drafted in this document and approved in each unit's registration. The names are settled: he chose them on 26 September 2026, and his statement of the laws uses them; only their registration remains, which is item 15. The identification of β_C with χ is not settled by anyone's word: the registration makes it an empirical question that needs an intervention (registration version 1.102, line 1029).

8. Dating

Each item carries its date and the kind of record that dates it; no kind is presented as stronger than it is, and no wording plays the author's work down.

- 16 August 2026, the author's working paper of that date, which the registration names as the programme's author-review working paper: the registration credits that date with the path identity that writes the balance exponent as $\phi_Q - \phi_W$ ("this registration adopts that formulation with the credit going to that date and not to this

one", registration version 1.102, line 619) and with its identification limit (registration version 1.102, line 621), and records that the working paper names the correction-leverage exponent χ (registration version 1.102, line 305).

- 8 September 2026, the registration (version 1.100, registered on OSF that day; version 1.102 of 13 September 2026 frozen since): the date on which the registration stated the local reading of the growth exponent (registration version 1.102, line 267), the general boundary (registration version 1.102, line 639), P12 and P21 as movers of the boundary (registration version 1.102, lines 1045 and 1239) and content neutrality (registration version 1.102, line 75); these are dated to the registration for what the registration states, and to 16 August 2026 for what the working paper already held.
- 1 October 2026: version 0.2.1, the earlier working draft: the theorems' forms, the interval rules and the kill conditions as written there.
- 8 December 2024: the registration states "An earlier private record of the conceptual frame is dated 8 December 2024." (registration version 1.102, line 179). Each law's first statement is to be dated from the author's prose of that date once those lines are confirmed at source (section 7 item 13); until then this document dates nothing to them, and dates the equations and theorems as above.
- 3 October 2026: the working draft 0.2.4, and version 1.0, which freezes it: Propositions IV-N, IV-R and B-8, the measurement requirements of sections 2.6.1 and 4.9, and the law-level conventions of version 0.2.1 taken out of force. Versions 0.2.2 and 0.2.3 were not published. These lines date the text and its mathematics; nothing in them is an empirical result or a registration.

9. Checks

Every identity and derivation in this document was checked by computer algebra. The proofs of its theorems, lemmas and propositions are written proofs, which finite models exercise and do not replace. Every quotation was compared, byte for byte, with its source, which is the frozen text of registration version 1.102, the statement paper at version 6.7 or the author's statement of the laws, and the identity of each address to which this document links was confirmed at the registry that issues it, on 3 October 2026 for version 1.0 and on 7 October 2026 for this version.

Corrections to version 0.2.1, dated 3 October 2026. Section 3.4: the example of an arm that doubles service while halving attained capability holds at unit elasticity only, and now says so. Section 5 item 2: the registered dials do not by themselves make Laws I to III hold, and the text now says what the dials do show. Section 6: the difference of the two retained fractions is defined only where both are established, not always. Section 2.8 item 1: that the baseline convention be fixed in advance is this document's PROPOSED requirement and no part of P15's registered face.

Changes from the working draft 0.2.4, dated 3 October 2026. The label, the title and the front matter say what this version is. Section 0: the sources paragraph records that the sentences quoted from the statement paper stand unchanged in its version 7.0 and links the paper's title to its DOI address; the status lines say what this version sets out and name the proved results as the text names them; the description of what this document is not says that it adds no registered proposition. Section 1: the reference value of service is defined in words. Section 2.6: Fieller's method carries its reference. Section 2.6.1: the account of where Proposition IV-R stops now covers the case of equality and the case of a negative denominator, and calls the unbounded set what it is. Section 2.8 item 5: its second sentence now opens on a condition, because no withdrawal convention is in force. Section 3.4 names the six fixed terms the smaller model keeps. Section 8 dates the working draft and this version and calls version 0.2.1 the earlier working draft. This section lists these changes, states the second check and says how the linked addresses were confirmed. Every mention of the registration, of the statement paper and of Paper VIII outside a quotation is linked to its DOI address, and Fieller's method to its source. No definition, lemma, theorem, corollary, identity or proposition was altered.

The second check. The working draft was checked a second time from its text alone, by code that shared nothing with the first check: exact rational arithmetic for every printed number, and searches over small finite models for the written proofs. Fifty checks ran and fifty passed. The two cases added to section 2.6.1 in this version were checked in exact arithmetic over random cases before the version was frozen.

Changes from version 1.0, dated 7 October 2026. Laws IV and V are no longer called candidates, and neither is this document. Section 0 states the laws' standing in full (added on 24 September 2026; their propositions registered; neither law yet registered as a law or tested; their mathematics set out here), and what this document is opens on their

mathematics; in the front matter, the three levels name the two laws as empirical claims. Sections 2.1 and 3.1 are headed Statement, and the tag before each statement carries the same standing in short form. Where sections 2.1, 2.6 and 2.8, the row for IV-PERSIST in section 6 and item 1 of section 7 said what this document's formulation does, they now say it of this version, or here. The word candidate stays where it grades neither the laws nor this document: the candidate descriptive families and the outcome NO CANDIDATE FITS (section 2.6), the candidate retention families (section 6), the candidate mixed model (section 3.4), and objects that are candidates for a registration of their own (sections 4.9 and 6). The title now follows the programme's uniform order (paper, its law or framework, its own title); until this version it read "Laws IV and V: the mathematics, version 1.0". The label carries version 1.1, and the front matter prints this document's DOI. The front matter and section 8 name version 1.0 where they called it this version, and the first paragraph of this section dates the confirmation of the linked addresses for each version. The entries above dated 3 October 2026, and the second check, record version 1.0 and stand as written: where they say this version, they mean version 1.0. No definition, lemma, theorem, corollary, proposition, assumption, proof, identity, equation, number, interval rule, open question, registered wording or quotation was changed, and no object, claim or prediction was added. For this version the computer-algebra checks were run again on its text, and every quotation was compared again with its source.

These checks are the programme's own. They are not independent verification, no proof assistant was used, and nothing here is an empirical result. This version is frozen: a later version is a new text under a new number, with its changes listed.

Michael Darius Eastwood conceived and directs this research programme and is the author of this work. Across the programme, he has used more than six AI systems in parallel, under his own instructions, to stress-test his arguments, identify possible errors, and assist in preparing draft text from his own outlines. He determines what is adopted, revised or rejected and takes responsibility for the published content. These systems are tools, not authors.

Changes in version 1.1

Michael Darius Eastwood. 7 October 2026.

Version 1.1 is frozen. This list compares it with version 1.0 of 3 October 2026, which stays as published. Nothing in this list is a registration or an empirical result.

From version 1.0 to version 1.1

Only words change. No definition, lemma, theorem, corollary, proposition, assumption, proof, identity, equation, number, interval rule, open question, registered wording or quotation is altered, and no object, claim or prediction is added.

- Laws IV and V are no longer called candidates, and neither is this document. Their standing is stated in full: added on 24 September 2026; their propositions registered (P15 in register ARC-4; P12 and P21 in register ARC-5); neither law yet registered as a law or tested; their mathematics set out in this document, now at version 1.1.
- Section 0, the status line: it states that standing, and it folds in what each law carries, so that nothing is said twice. What this document is: the mathematics of Law IV and of Law V.
- The front matter, the three levels: the second level is the two laws as empirical claims.
- Sections 2.1 and 3.1 are headed "Statement", and the tag before each statement carries the standing in short form: for Law IV "added 24 September 2026; its proposition, P15, registered in ARC-4; not yet registered as a law or tested", and for Law V "added 24 September 2026; its propositions, P12 and P21, registered in ARC-5; not yet registered as a law or tested".
- Where the text said what this document's formulation does or adopts (sections 2.1, 2.6 and 2.8, the row for IV-PERSIST in section 6, and section 7 item 1), it now says "this version", or "here" for the default baseline of section 2.6.
- The title follows the programme's uniform order (paper, its law or framework, its own title): "Report · The ARC Theory · Laws IV and V, the ARC Persistence Law and the ARC Embedding Law · The Mathematics". Until this version it read "Laws IV and V: the mathematics, version 1.0", which stays the title of version 1.0.

- The label: version 1.1, frozen on 7 October 2026, following version 1.0. The front matter prints this document's DOI, 10.17605/OSF.IO/4B3VT, and says that this version adds no object; it and section 8 name version 1.0 where they called it this version.
- Section 9: the first paragraph dates the confirmation of the linked addresses for each version, and a new entry lists these changes. The entries dated 3 October 2026, and the second check, stand as written; in them, "this version" means version 1.0.

Where the word candidate stays

It stays where it grades neither the laws nor this document.

- "Candidate descriptive families" and the outcome "NO CANDIDATE FITS" (section 2.6), and "the candidate retention families" (section 6): a closed set of descriptive model families fixed before any deciding run, none of which is the law, and the outcome in which none of them fits.
- "a candidate mixed model" (section 3.4): one model of the closed model set.
- "a candidate for a registration of its own" (section 4.9) and "candidates for a registration of their own" (section 6): objects the author may register in their own right.

Checked for this version

- Each edit applies once, and undoing the edits restores version 1.0 byte for byte.
- Every line that states a definition, lemma, theorem, corollary, identity or proposition stands as in version 1.0, and so does every line of sections 0.1, 1, 2.2 to 2.5.1, 2.6.1, 2.7, 3.2 to 3.7, 4 and 5.
- The computer-algebra checks were run again on the text of version 1.1.
- Every quotation was compared again with its source: the frozen text of registration version 1.102, the statement paper at version 6.7, or the author's statement of the laws.
- The identity of each linked address, this document's DOI among them, was confirmed on 7 October 2026 at the registry that issues it.

Unchanged

- Every quotation of the registration, of the statement paper and of the author's statement of the laws; and the twenty-two registered propositions with their words, scopes and refuters.
- The register assignment of section 0.1 and every registration line number.
- The fifteen open questions of section 7 and their defaults. Items 2, 3 and 4 name no default: those decisions are the author's.