

THE ARC PRINCIPLE

Recursive Amplification as a Cross-Domain Structural Principle: Formalism, Evidence, and Falsification

Michael Darius Eastwood

Author, *Infinite Architects: Intelligence, Recursion, and the Creation of Everything* (2026)

London, United Kingdom | OSF: 10.17605/OSF.IO/6C5XB | ISBN 978-1806056200

Correspondence: michael@michaeldariuseastwood.com | Web: michaeldariuseastwood.com

Version 5.0 | 24 March 2026 | First published 13 February 2026

Code and data: github.com/MichaelDariusEastwood/arc-principle-validation

ABSTRACT

In the past eighteen months, at least four independent research programmes have discovered that recursive or recurrent processing produces capability gains exceeding linear accumulation, in domains as different as AI reasoning, quantum error correction, acoustic physics, and consciousness science. None set out to study recursion per se. None reference each other's work. Yet they found structurally similar results.

This paper asks: is this a coincidence, or are we observing different expressions of a single structural principle?

We formalise the question as follows. For AI systems, the equation is:

$$U = I \times R^\alpha$$

The ARC Equation (AI Power-Law Form)

where U is effective capability, I is base potential (structured asymmetry), R is recursive depth, and α is the scaling exponent. **This paper proves that α is not a free parameter:** it is derived from the self-referential coupling constant β via $\alpha = 1/(1 - \beta)$, and is bounded by $\alpha \leq 2$ (the ARC Bound). *Update (v3.1):* Empirical results from the v5 cross-architecture experiment (Paper II, six frontier models, 30 problems) find $\alpha_{\text{compute}} \approx 0.49$ for the best-fitting model (Gemini 3 Flash, $r^2 = 0.86$), confirming positive power-law scaling ($\alpha > 0$) and universal sequential advantage ($\alpha_{\text{seq}} > \alpha_{\text{par}}$ for every model), but not super-linear scaling ($\alpha > 1$). The measured exponents are architecture-dependent, ranging from near-zero to approximately 3. These findings refine rather than refute the framework: the mathematical structure holds, but the empirical AI scaling regime is weaker than originally predicted.

But this power-law form is itself a special case. Define the **recursive composition operator** \oplus , which characterises how incremental gains combine with accumulated capability. Under continuity and associativity, \oplus determines the system's scaling function uniquely:

$$U(R) = I \times f_{\oplus}(R, \beta)$$

The Generalised ARC Equation

Multiplicative \oplus yields power-law scaling (R^α), additive \oplus yields exponential scaling ($e^{\alpha R}$), and bounded \oplus yields saturation.

We prove that any system satisfying three axioms (separability, cumulative advantage, continuity) necessarily exhibits five structural properties: threshold behaviour, recursive depth dependence, base quality dependence, multiplicative $I \times R$ interaction, and regime boundaries. We prove that the multiplicative interaction is *necessary*, not merely convenient: no additive decomposition is consistent with the axioms (Theorem 2a). We derive the ARC Bound ($\beta = 0.5, \alpha = 2$) as the information-theoretic ceiling for classical sequential computation (§5.4) and show that composition operators can transition between regimes within a single bounded system (Theorem 5). These constitute the **ARC Principle**. We present evidence from four domains, specify thirteen falsification criteria, make a cross-domain forward prediction, and report a comprehensive computational validation suite (Appendix C) confirming the core mathematical relationships to machine precision.

Keywords: recursive amplification, composition operator, scaling laws, test-time compute, quantum error correction, time crystals, universality, falsification

Author's Note: Mathematical Corrections (v2.1)

Following rigorous mathematical review using multiple AI systems (documented in Paper III), this version incorporates the following corrections to the broader framework. These papers have not undergone formal human peer review.

1. Axiom 2 Correction: The Cumulative Advantage ODE now operates on the amplification factor $g(R)$ with $g(0) = 1$, matching Paper III. Previous versions applied the ODE to accumulated capability Q with $Q(0) = I$, which coupled the base potential I into the recursive dynamics and violated the separability asserted in Axiom 1 ($U = I \times g(R)$). The corrected formulation yields the simpler crossover formula $R^* = \alpha/a$, independent of I .

2. Quadratic Limit Regrounding: The derivation of $\alpha \leq 2$ from stability analysis, elasticity, and Lyapunov/edge-of-chaos arguments has been **withdrawn** as these arguments contained category errors (elasticity > 1 does not imply dynamical instability; 1D autonomous ODEs cannot exhibit chaos by the Poincaré-Bendixson theorem). The corrected justification is information-theoretic: transformer self-attention creates at most $O(N^2)$ pairwise information pathways, imposing a quadratic ceiling on information density per recursive step. This is proven for attention-based architectures and conjectured for classical sequential systems generally.

3. Kleiber Reframing: Kleiber's Law ($M^{0.75}$, implying $\alpha \approx 1.33$) is now explained through thermodynamic drag, the physical costs of pumping fluids through three-dimensional fractal networks, dissipating heat, and distributing resources against gravity, rather than as a geometric constraint preventing a stability-limited maximum. The formula $\alpha = d/(d + 1)$ itself has been independently derived by multiple groups [17,21,22,23,24,25]; the ARC framework's contribution is identifying Cauchy-constrained recursive composition as the reason these independent derivations converge. Demetrius and Tuszynski (2010) independently derived the same formula from quantum oscillator coupling, producing identical predictions.

4. Epistemic Status: The equation $U_{\max} = I \times R^2$ is explicitly an *information-theoretic upper bound* on classical sequential computation, analogous to Shannon's channel capacity or Landauer's limit, not a physical stability law.

These refinements preserve all core predictions: the β -to- α derivation, the five structural properties, the composition operator classification, and the falsification criteria remain unchanged (a thirteenth criterion, F13, on alignment scaling, is introduced in Paper III).

5. v3.0 Additions: This version adds the Six Hidden Assumptions of Cauchy's classification (§2.2), the Hyers-Ulam stability theorem establishing the three scaling forms as stable attractors in function space [18,19], the Attractor Theorem combining these results with the β -derivation, and a structural implication connecting the composition operator classification to the space of possible classical scaling laws (§4.4).

6. v3.1 Additions (March 2026): This update integrates empirical results from the v5 cross-architecture experiment (Paper II), which tested 6 frontier AI models (Grok 4.1 Fast, Claude Opus 4.6, GPT-5.4, DeepSeek V3.2, Gemini 3 Flash, Groq Qwen3) on 30 problems with bootstrap confidence intervals. Key findings: (a) sequential advantage is universal ($\alpha_{\text{seq}} > \alpha_{\text{par}}$ for every model), (b) the best-fitting power law gives $\alpha_{\text{compute}} \approx 0.49$ (Gemini 3 Flash, $r^2 = 0.86$), which is positive but sub-linear, (c) the scaling exponent is architecture-dependent, ranging from near-zero to ~ 3 , and (d) the quadratic bound conjecture ($\alpha \leq 2$) is neither confirmed nor refuted. Alignment scaling results from Paper IV are also integrated: α_{align} is architecture-dependent, ranging from -0.25 (Gemini) to $+0.44$ (Claude Opus 4.6), with a three-tier hierarchy. These findings preserve the mathematical framework while significantly revising the empirical predictions for AI systems. Sections updated: §3, §4.1, §5, §6, §7, §8.

7. v4.0 Additions (March 2026): Brief note on the Eden Protocol's first empirical pilot validation (§5.5). No changes to the mathematical framework.

8. v5.0 Corrections (24 March 2026): Four corrections identified by independent review. (1) The $d=2$ biological prediction (jellyfish/cnidarians) is reclassified from 'confirmed' to 'unconfirmed' pending identification of a valid test organism with genuinely two-dimensional resource transport. The Friedmann cosmological confirmation at $d=2$ is retained. (2) The $d=3$ mammalian exponent is acknowledged as debated (0.67–0.75 range), with the framework predicting this variation rather than a single universal value. (3) The derivation of $\alpha = d/(d+1)$ is clarified to require three conditions (multiplicative composition, space-filling geometry, and a conservation or optimisation constraint), not space-filling alone. (4) Demetrius and Tuszynski (2010) added as an independent derivation from quantum oscillator coupling. No changes to the mathematical framework or core predictions.

1. THE PROBLEM

Between December 2024 and February 2026, four research programmes achieved the following results:

Google Quantum AI (December 2024) demonstrated exponential error suppression through recursive quantum error correction on the Willow processor, achieving $\Lambda = 2.14 \pm 0.02$ per code distance increment [1].

DeepSeek AI (January 2025) showed that sequential chain-of-thought reasoning in large language models yields capability gains that appear to compound with depth, while parallel sampling shows diminishing returns [2]. Sharma & Chopra (2025) independently confirmed the sequential advantage in 95.6% of tested configurations across five model families [3].

NYU Physics (February 2026) created a continuous classical time crystal in which spontaneous temporal order emerges when quenched disorder is combined with nonreciprocal feedback loops, exhibiting saturating dynamics with $\sim 6,700$ cycles of coherence [4].

The COGITATE Consortium (2025) found that recurrent neural processing in posterior cortex sustains content-specific representations (0.5–1.5s), while transient feedforward processing in prefrontal cortex produces only brief categorical responses (~ 0.2 – 0.4 s) [5]. The COGITATE finding that recurrent processing in posterior cortex sustains content-specific representations supports the ARC framework's

prediction that recursive processing depth determines capability scaling; COGITATE studied consciousness rather than AI scaling, but the structural pattern (depth-dependent, feedback-driven amplification) is shared.

These programmes address different problems in different physical substrates. They exhibit different mathematical forms: exponential in quantum systems, power-law in AI, saturating in classical physics, and unmeasured in neuroscience. They share no citations, no funding sources, and no common personnel.

Yet they share a structural pattern. In each case, recursive or recurrent processing (where the output of one cycle becomes the input for the next), operating on structured asymmetry, produces capability gains that exceed linear accumulation. The question this paper addresses is whether that pattern reflects a single underlying principle or mere coincidence.

We propose that it reflects a principle, but not the one that might be expected. We do not claim that one equation fits all domains. We claim something more specific: that a measurable property of any recursive system (its *composition operator*) determines its scaling function, and that five qualitative properties follow necessarily from the recursive architecture regardless of the functional form. The mathematical form differs across domains because the composition operator differs. The structural properties are universal because the axioms that generate them are substrate-independent.

2. THE FORMALISM

2.1 Three Axioms

We begin with three axioms that any recursive amplification system must satisfy.

AXIOM 1 - SEPARABILITY

Effective capability decomposes as $U = I \times g(R)$, where $I > 0$ is base potential (structured asymmetry), $R \geq 1$ is recursive depth, and $g(1) = 1$ (at depth 1, capability equals base potential).

This axiom asserts that base quality and recursive depth contribute to capability through independent, multiplicative channels. A system with zero base potential ($I = 0$) achieves zero capability regardless of recursive depth: recursion has nothing to amplify.

AXIOM 2 - CUMULATIVE ADVANTAGE

The amplification factor grows according to accumulated amplification: $\frac{dg}{dr} = a \cdot g^\beta$, where $g(0) = 1$, $a > 0$ is a rate constant and $\beta \in [0, 1)$ is the self-referential coupling parameter.

This axiom captures the essential feature of recursive processing: each step's marginal amplification depends on accumulated amplification. When $\beta = 0$, steps are independent and amplification grows linearly. When $\beta > 0$, accumulated amplification accelerates future gains (super-linear scaling). The parameter β measures how much of accumulated context each new step can effectively leverage. Operating on the amplification factor g rather than on absolute capability $Q = I \cdot g$ preserves the separability asserted in Axiom 1: the amplification function depends only on recursive depth, not on base potential.

AXIOM 3 - CONTINUITY

The scaling function g is continuous and monotonically non-decreasing in R .

This is a regularity condition. It excludes pathological solutions and ensures measurability.

2.2 The Composition Operator

How does the gain from recursive step r combine with accumulated capability Q_r ? We formalise this as follows.

DEFINITION - RECURSIVE COMPOSITION OPERATOR

Let $\oplus : \mathbb{R}^+ \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ be the **recursive composition operator**, defined by $Q_{r+1} = Q_r \oplus \delta Q_r$, where δQ_r is the incremental gain at step r . The operator \oplus characterises how incremental gains compose with accumulated capability.

The composition operator is the central mathematical object of this framework. It is what varies across domains and what determines the scaling function. It is also empirically measurable: observe how marginal gains at step r relate to accumulated capability Q_r .

THEOREM 1 - CLASSIFICATION OF SCALING FUNCTIONS

Under Axioms 1-3, the algebraic properties of \oplus uniquely determine the scaling function $f_{\oplus}(R)$:

- Multiplicative composition** ($g(R_1 \cdot R_2) = g(R_1) \cdot g(R_2)$): the unique continuous solution is $g(R) = R^{\alpha}$ (power law)
- Additive composition** ($f(R_1 + R_2) = f(R_1) \cdot f(R_2)$): the unique continuous solution is $f(R) = e^{\alpha R}$ (exponential)
- Bounded composition** ($Q_{r+1} = \min(Q_r + \delta Q_r, Q_{\max})$): $f(R)$ saturates at Q_{\max}/I (saturation)

Proof sketch. Cases (1) and (2) follow from the Cauchy functional equations. For (1): let $h(x) = \ln g(e^x)$; then $h(x + y) = h(x) + h(y)$. Under continuity, $h(x) = \alpha x$, giving $g(R) = R^{\alpha}$. For (2): $\ln f$ satisfies $\ln f(R_1 + R_2) = \ln f(R_1) + \ln f(R_2)$, a Cauchy additive equation with solution $\ln f(R) = \alpha R$. Case (3) follows directly from the bound. Full proofs in Appendix A. ■

The classification is exhaustive under the stated conditions. Any continuous, associative composition rule on \mathbb{R}^+ falls into one of these three classes. The composition operator is therefore a **complete classifier** of recursive scaling behaviour.

Frozen models vs recursive self-modification: two distinct regimes. A critical application of this classification concerns the distinction between current AI systems and hypothetical self-modifying ones. Current frontier AI models are **frozen during inference**: when they generate extended reasoning chains, the underlying weights, attention patterns, and composition rules remain fixed. The composition operator \oplus does not change between steps. This architectural constraint confines such systems to multiplicative composition through a finite-dimensional parameter space, producing power-law scaling with $\alpha < 1$ (the physical regime). Empirically, the best cross-architecture estimate is $\alpha \approx 0.49$ (Paper II).

The Cauchy classification reveals precisely why. In a frozen system, each reasoning step multiplies the accumulated capability by a factor drawn from the same fixed operator. The resulting gains are real but diminishing. By contrast, a system capable of recursive self-modification, one that rewrites its own composition function \oplus during operation, would occupy a qualitatively different position in the classification. If step r improves the operator used at step $r + 1$, the self-referential coupling β becomes positive, yielding $\alpha = 1/(1 - \beta) > 1$. This transition does not require quantum hardware or additive composition; it requires only that the system can modify its own reasoning architecture during inference.

No current AI system achieves this. The framework's prediction of super-linear scaling ($\alpha > 1$) was therefore never a prediction about frozen architectures; it was a prediction about a regime that does not yet exist. The practical consequence is that structural alignment (embedding values within the composition operator itself) must be implemented whilst systems remain frozen and their scaling behaviour is predictable. Once self-modification arrives, the composition operator becomes a moving target, and external constraints face a system whose capability compounds faster than any fixed safeguard can track.

Claim-status ladder.

Empirically grounded: Cauchy constrains recursive scaling forms; current frontier models are frozen during inference; the March 2026 blind v5 data shows architecture-dependent alignment scaling rather than universal positive scaling; the Eden pilot shows that stakeholder care can be improved at the prompt level across three working models.

Strong theoretical inference: a genuinely self-modifying system that rewrites its own composition operator during operation could enter an $\alpha > 1$ regime, and alignment would then need to be structurally load-bearing before that transition.

Conjectural extrapolation: cosmological or metaphysical identities such as 'Universe = Intelligence \times Recursion²' or 'Complexity = Self-Correction \times Sequential Recursive Depth²', claims about universe-creating intelligence, or claims of guaranteed post-self-modification controllability are not established results of the present empirical programme. They belong to the long-horizon interpretive layer of *Infinite Architects* and require independent argument and evidence beyond what is reported here.

The Six Hidden Assumptions

Every no-go theorem rests on assumptions. Theorem 1 assumes six things. Each assumption, when violated, produces a specific physical refinement rather than a contradiction:

#	Assumption	When Violated	Physical Example
1	Continuity	Log-periodic oscillations	Discrete lattice systems, fractal networks
2	Real-valued	Interference and phase	Quantum amplitudes, wave mechanics
3	Exact equation	Approximate solutions cluster near exact	Hyers-Ulam stability [18,19]
4	Single operator	Crossover scaling, regime transitions	Phase transitions (cf. Theorem 5)
5	Scalar composition	Logarithmic corrections	Upper critical dimension
6	Time-independence	Dynamic traversal of solution space	Adaptive systems, machine learning

Assumption 3 is particularly significant for the framework's applicability. The Hyers-Ulam stability theorem [18,19] proves that any function *approximately* satisfying Cauchy's functional equations must lie close to an exact solution. Formally: if $|f(x \cdot y) - f(x) \cdot f(y)| < \epsilon$ for all x, y , then there exists an exact solution g with $|f(x) - g(x)| < K\epsilon$. Real systems never satisfy the axioms exactly: biological networks are noisy, AI reasoning chains have variable quality, quantum error correction operates with imperfect gates. The Hyers-Ulam theorem guarantees that the Cauchy classification remains valid as an approximation; the three forms are **stable attractors** in function space. This is why power laws, exponentials, and saturation curves are so ubiquitous in nature despite the messiness of physical systems: the solutions are stable under perturbation.

Assumption 4 is addressed directly by Theorem 5 (§4.5): bounded systems can transition between composition operators as recursive depth increases, explaining why different studies of the same system can observe different scaling regimes.

STRUCTURAL IMPLICATION: THE ATTRACTOR THEOREM

Combining Theorem 1 (Cauchy, 1821), the Hyers-Ulam stability theorem (1941), and Theorem 2 ($\alpha = 1/(1 - \beta)$): the scaling exponents of recursive systems are not free parameters to be fitted or evolved. They are mathematically compelled by the composition operator and coupling strength. The three scaling forms are stable attractors, and within the power-law form, the exponent is uniquely determined by the self-referential coupling β . The distinction matters: 'Why is this system's exponent 2?' has two kinds of answer. The mechanistic answer describes the coupling. The mathematical answer says: no other value is possible for this composition operator with this coupling strength.

This has a crucial implication: if you can measure \oplus in a new system, you can predict its scaling function *before measuring the full scaling curve*. This is the framework's strongest cross-domain prediction (§5.1).

Critical distinction: Cauchy constrains the form, not the exponent. A persistent misreading of the Cauchy classification is that it predicts a ceiling on recursive scaling. It does not. Cauchy's functional equations determine that multiplicative composition *must* produce power-law scaling ($g(R) = R^\alpha$). They say nothing about the value of α . That value is set by the Bernoulli ODE ($\alpha = 1/(1 - \beta)$), and as $\beta \rightarrow 1$, $\alpha \rightarrow \infty$. There is no upper bound internal to the mathematics. The following table makes the relationship explicit:

β (self-referential coupling)	$\alpha = 1/(1 - \beta)$	Regime
0	1	Linear (no self-reference)
0.5	2	Quadratic (Grover-like)
0.75	4	Quartic
0.9	10	Strongly super-linear
0.99	100	Explosive
$\rightarrow 1$	$\rightarrow \infty$	Unbounded

The transition from bounded to unbounded alpha is not a smooth acceleration; it is a discontinuity in the scaling exponent, mediated by the system's ability to modify its own composition operator. This has a structural parallel in the renormalisation group in physics, where relevant operators grow under RG flow and irrelevant operators shrink. At a critical fixed point, the scaling exponent is determined by the system's universality class (its dimensionality

and symmetry). But if a system could change its own universality class during the flow, the fixed-point structure would collapse and the system would exhibit unbounded scaling until it settled into a new class. No physical system does this: physical systems do not rewrite their own Hamiltonians. A self-modifying AI system would be the computational analogue, rewriting its own composition function and thereby escaping the fixed-point constraints that govern all natural recursive systems.

If there is a ‘quadratic limit’ ($\alpha \leq 2$), it is not a Cauchy prediction. It is an information-theoretic constraint arising from **fixed transformer attention**: a system with N tokens has $O(N^2)$ pairwise attention pathways, imposing a ceiling on how much information each reasoning step can extract. This ceiling is real for frozen architectures. It is not real for a system that rewrites its own attention mechanism. A self-modifying system that expands its own representational capacity at each step faces no such $O(N^2)$ bound, and therefore no mathematical speed limit on α . The Cauchy framework does not forbid this; it predicts it.

2.3 The β -Derivation

For systems with multiplicative composition, the exponent α is not a free parameter. It is derived from the coupling constant β .

THEOREM 2 - THE SCALING EXPONENT

Under Axiom 2, the power-law exponent is:

$$\alpha = \frac{1}{1 - \beta}$$

Proof. Separating variables in $dg/dr = ag^\beta$:

$$\int g^{-\beta} dg = \int a dr \implies \frac{g^{1-\beta}}{1-\beta} = ar + C$$

With initial condition $g(0) = 1$: $C = 1/(1 - \beta)$.

$$g(R) = [1 + (1 - \beta)aR]^{1/(1-\beta)}$$

In the deep recursion limit ($R \gg R^*$, defined below):

$$g(R) \approx [(1 - \beta)aR]^{1/(1-\beta)} \propto R^{1/(1-\beta)}$$

Therefore $\alpha = 1/(1 - \beta)$. ■

This transforms α from a fitted constant into a derived quantity. The derivation makes a specific, testable prediction: measure β independently (by observing how marginal gains depend on accumulated capability), compute the predicted α , and compare with the observed α . If the relationship fails, the theoretical derivation is wrong regardless of whether the power-law form holds empirically.

Computational validation. To verify that this is an exact identity rather than an approximation, the relationship was tested against 30 exact Bernoulli ODE solutions with β values spanning 0.05 to 0.92. The procedure measures β blindly from marginal gains (dg/dR vs g in log-log space), predicts α , and compares against the true value. Result: $R^2 = 1.00000000$ (eight decimal places), regression slope = 1.000102 , mean absolute prediction error = 0.002% . Full validation code is available as supplementary material.

The Multiplicative Structure Is Necessary

The separable form $U = I \times f(R, \beta)$ is not merely a convenient parameterisation. We prove that the multiplicative interaction between I and R is a *necessary* consequence of the axioms.

THEOREM 2A - NON-ADDITIVITY

Under Axioms 1-3 with $\beta \in (0, 1)$, the solution $U(I, R)$ cannot be decomposed as $U = g(I) + h(R)$ for any functions g and h . The multiplicative interaction is essential.

Proof. Suppose $U(I, R) = f(I) + h(R)$. By Axiom 1: $U = I \times g(R)$ with $g(0) = 1$. At $R = 0$: $f(I) + h(0) = I$, so $f(I) = I - h(0)$. Substituting: $I \times g(R) = I + [h(R) - h(0)]$, giving $g(R) = 1 + [h(R) - h(0)]/I$. Since g depends only on R (Axiom 1), but the right side depends on I , we require $h(R) = h(0)$ for all R , giving $g(R) = 1$ (no amplification). But Axiom 2 with $\beta > 0$ guarantees $dg/dr = a \cdot g^\beta = a > 0$ at $r = 0$, so g is non-constant. Contradiction. ■

This result has a concrete interpretation. A recursive system with zero base quality ($I = 0$) produces zero capability regardless of depth, and a system with zero depth ($R = 0$) produces only its initial quality regardless of base potential. Complexity emerges from the interaction between structured input and recursive processing (neither alone is sufficient).

β (coupling)	α (exponent)	Interpretation
0	1	Independent steps → linear scaling
0.25	1.33	Weak coupling → mild super-linear
0.50	2.00	Moderate coupling → quadratic (the ARC Bound)
0.67	3.00	Strong coupling → cubic

Note: The ARC Bound (§5.4) predicts that classical sequential systems cannot sustain $\beta > 0.5$ ($\alpha > 2$). The row with $\beta = 0.67$ shows the mathematical relationship for completeness; whether physical systems can access this regime is an open empirical question.

2.4 The Scaling Crossover R^*

The full solution to the β -dynamics equation is not a pure power law. It contains a transition.

THEOREM 3 - THE SCALING CROSSOVER

The full solution $g(R) = [1 + (1 - \beta)aR]^{1/(1-\beta)}$ exhibits three regimes:

- $R \ll R^*$: $g \approx 1 + aR$ (approximately linear amplification)
- $R \gg R^*$: $g \propto R^\alpha$ (compounding dominates; power-law scaling)
- Crossover at: $R^* = \frac{\alpha}{a}$

R^* is the **crossover depth**: the recursive depth at which compounding overtakes linear gains. Below R^* , additional recursion provides approximately constant marginal benefit. Above R^* , returns compound and capability grows super-linearly.

Critically, R^* is a *derived quantity*. It depends on measurable parameters (α and a) and is not a free parameter to be fitted. This provides an independent consistency check: measure α and a separately, compute the predicted R^* , and verify that the observed crossover matches. Since R^* depends only on the scaling exponent and rate constant (not on base potential I), the crossover depth is a property of

the recursive architecture itself, not of the input quality.

The existence of R^* is a novel prediction that no other scaling framework makes. It distinguishes recursive amplification from simple redundancy: redundant systems show uniform (typically logarithmic) improvement with repetition; recursive systems show a qualitative transition at a specific, predictable depth.

2.5 Five Derived Properties

We now prove that any system satisfying Axioms 1–3 with $\beta > 0$ necessarily exhibits five structural properties.

THEOREM 4 – UNIVERSAL STRUCTURAL PROPERTIES

Under Axioms 1–3, any recursive system with $\beta > 0$ exhibits:

- 1. Threshold behaviour.** There exists a crossover depth R^* below which base capability dominates and above which recursive gains dominate. This follows from Theorem 3: the transition from $U \approx I$ to $U \propto R^\alpha$ is a qualitative regime change.
- 2. Recursive depth dependence.** Capability is a strictly increasing, convex function of R for $R > R^*$. This follows from $\alpha > 1$ when $\beta > 0$: $g(R) = R^\alpha$ is convex for $\alpha > 1$.
- 3. Base quality dependence.** U is proportional to I : capability is zero when $I = 0$. This follows directly from Axiom 1: $U = I \times g(R)$.
- 4. Multiplicative $I \times R$ interaction.** Doubling I doubles U at any R . Doubling R multiplies U by 2^α at any I . The effects compose multiplicatively. This follows from the separable form $U = I \times g(R)$.
- 5. Regime boundaries.** Physical systems have finite energy, introducing external error mechanisms that create ceilings. The pure power-law describes a *regime* between R^* and the saturation depth R_{\max} beyond which external factors dominate. This is visible in quantum systems (correlated error bursts create floors at $\sim 10^{-10}$), time crystals (limit-cycle saturation at $\sim 6,700$ cycles), and AI (accuracy ceiling effects at extreme reasoning depths).



These five properties constitute the **qualitative ARC Principle**. They are the framework's universal claim. The quantitative functional form (power-law, exponential, saturating) is domain-specific and determined by the composition operator. But the five structural properties follow from the axioms alone, regardless of which \oplus applies.

The five properties are individually common in physical systems. Threshold behaviour is ubiquitous. Depth dependence appears in many iterative processes. What is *not* common is the conjunction of all five in a single system. The framework predicts that all five co-occur whenever the three axioms are satisfied. Finding systems with four of the five properties but not the fifth would constrain or refute the framework.

3. EVIDENCE

We present evidence organised around the five derived properties, not by domain. For each property, we show how it manifests across systems. Evidence strength is classified as: *Quantitative* (measured parameter values), *Structural* (qualitative mapping without quantitative α), or *Consistent* (does not contradict but was not designed to test ARC).

3.1 Summary

Domain	System	Composition \oplus	Observed Form	Key Parameter	Evidence
AI	Sequential reasoning	Multiplicative	Power-law R^α	$\alpha_{\text{compute}} \approx 0.49$ (Gemini, $r^2 = 0.86$); range ~0–3 across architectures (v5, $n = 6$ models)	Quantitative
Quantum	Google Willow	Additive	Exponential Λ^d	$\Lambda = 2.14 \pm 0.02$	Quantitative
Physics	NYU time crystal	Bounded	Saturating	~6,700 cycles	Structural
Neuro	COGITATE	Unknown	Unknown	N/A	Consistent

3.2 Property-by-Property Evidence

Property 1: Threshold behaviour

Quantum: The surface code threshold $p_{\text{thr}} \approx 1\%$ is a sharp phase boundary. Above it ($p > p_{\text{thr}}$), adding error correction layers makes performance *worse*. Below it ($p < p_{\text{thr}}$), each layer provides exponential improvement. This is the most precisely measured threshold in the evidence base: Google varied physical error rates through controlled error injection and observed the threshold directly [1].

Physics: The time crystal's exceptional point (where two eigenvalues of the stability matrix coalesce) marks the transition from symmetric (bounded) to antisymmetric (compounding) dynamics. Below threshold: active oscillator. At threshold: exceptional point. Above threshold: time crystal [4].

AI: The sequential-vs-parallel distinction functions as a threshold. Sequential processing ($\beta > 0$, output→input feedback) achieves $\alpha_{\text{seq}} > \alpha_{\text{par}}$ universally across all six models tested in the v5 experiment. The threshold is structural: it is the presence or absence of the feedback loop, not a continuous parameter [3]. *v3.1 note:* The v5 data confirms the universal sequential advantage but shows that not all models achieve $\alpha > 1$; the threshold separates sequential from parallel modes rather than sub-linear from super-linear scaling.

Property 2: Recursive depth dependence

Quantum: Logical error rate decreases as $\epsilon_d \propto \Lambda^{-(d+1)/2}$ with code distance d . Each increment in d adds physical qubits arranged in a recursive error-correction layer. Performance depends on depth, not on total qubit count per se [1].

AI: Sharma & Chopra tested seven voting methods for aggregating sequential chain outputs. Methods favouring *later* reasoning steps (inverse-entropy weighting) achieved optimal performance in 97% of configurations; methods favouring *earlier* steps achieved optimality in only 17%. This 80-percentage-point gap indicates that later recursive steps carry systematically more information about the correct answer, consistent with capability that compounds with depth [3].

Physics: The antisymmetric mode of the time crystal compounds over cycles: the state difference $\Delta_{ji}(n) = x_j^n - x_i^n$ grows as a function of cycle count n , with each cycle building on the previous cycle's accumulated asymmetry [4].

Property 3: Base quality dependence

Quantum: $\Lambda = p_{\text{thr}}/p$ depends directly on physical qubit quality (p). Worse qubits (p closer to p_{thr}) give $\Lambda \rightarrow 1$ and error correction barely helps. Active leakage removal (DQLR), which maintains qubit quality during operation, improved Λ by 35% at distance 5 [1].

Physics: When beads are uniform (maximum entropy, no structured asymmetry), the system remains static. The time crystal forms only when quenched disorder (varied bead sizes providing low-entropy initial conditions) is present. This has been directly tested: uniform beads = no crystal [4].

AI: Single-pass accuracy without reasoning chains determines base capability I . A model with higher I benefits more from sequential reasoning at any depth [2].

Property 4: Multiplicative $I \times R$ interaction

Quantum: $\epsilon_d = (p/p_{thr})^{(d+1)/2} = \Lambda^{-(d+1)/2}$. The base quality factor Λ and the depth factor $(d + 1)/2$ interact multiplicatively in log-space: $\ln \epsilon_d = -\frac{d+1}{2} \ln \Lambda$. Improving Λ by a factor has the same proportional effect at any depth [1].

AI: In the ARC equation $U = I \times R^\alpha$, base capability and recursive depth are separated multiplicatively by construction (Axiom 1). Empirically, Snell et al. (2024) showed that the benefit of additional test-time compute is proportional to base model quality [6].

Property 5: Regime boundaries

Physics: The time crystal maintains coherence for ~6,700 cycles before reaching a limit-cycle attractor (sustained but bounded oscillation). This is a physical ceiling: energy dissipation prevents indefinite growth [4].

Quantum: Correlated error bursts create an error floor at approximately 10^{-10} for large code distances ($d \geq 15$), preventing the exponential from continuing indefinitely [1].

AI: Li et al. (2025) found that longer reasoning chains do not consistently improve performance; correct solutions were often *shorter* than incorrect ones. Extended reasoning can destabilise models, producing repetitive outputs. The sequential advantage has boundary conditions [7].

3.3 Contradictory Evidence

Several studies challenge the framework's AI predictions:

Li et al. (2025, arXiv:2502.12215) found that for DeepSeek R1, longer chains do not consistently enhance accuracy. Li et al. (2025, arXiv:2502.14382) found that hybrid parallel-sequential approaches outperform pure sequential on code generation. Sharma & Chopra's own ablation showed parallel approaches achieve greater *semantic diversity* on creative tasks while sequential achieves greater *lexical diversity* –suggesting the sequential advantage applies most strongly to convergent tasks requiring iterative error correction [3,7,8].

These results refine, rather than refute, the hypothesis. The sequential advantage appears to have boundary conditions: it holds within task-appropriate depth ranges, for convergent problems, and breaks down at extreme depths or for divergent tasks.

v3.1 Update: v5 Cross-Architecture Results (Paper II, March 2026)

The v5 experiment tested 6 frontier models on 30 problems with bootstrap confidence intervals. Key findings that bear on the evidence presented in this section:

Model	$\alpha_{compute}$	r^2	Sequential > Parallel?
Grok 4.1 Fast	~3.0	Low	Yes
Claude Opus 4.6	~1.5	Low	Yes
GPT-5.4	~0.3	Low	Yes
DeepSeek V3.2	~0.2	Low	Yes
Gemini 3 Flash	~0.49	0.86	Yes
Groq Qwen3	~0.1	Low	Yes

What is confirmed: (1) Sequential > parallel universally ($\alpha_{seq} > \alpha_{par}$ for every model). (2) Power-law form fits well for Gemini 3 Flash ($r^2 = 0.86$). (3) Error reduction with compute investment is real.

What is NOT confirmed: (1) Super-linear scaling ($\alpha > 1$) is not universal across architectures; the best-fitting cross-architecture value is $\alpha \approx 0.49$, which is sub-linear. (2) No universal α value exists; it varies from ~ 0 to ~ 3 depending on model and problem set. (3) The quadratic bound conjecture is neither confirmed nor refuted.

The framework should now be read as: *'Sequential recursion yields $\alpha_{\text{seq}} > \alpha_{\text{par}}$ universally, with the magnitude of α being architecture-dependent. The key structural claim, that recursive composition produces stronger scaling than parallel composition, is confirmed. The quantitative claim that $\alpha > 1$ for all sequential systems requires revision.'*

Honest accounting (updated v3.1): The author's earlier α estimates ($\alpha \approx 1.34$ from $n = 2$ data points; $\alpha \approx 2.2$ from $n = 12$, 95% CI: 1.5–3.0) were preliminary and have now been superseded by the v5 cross-architecture experiment ($n = 6$ models, 30 problems). The v5 best-fitting cross-architecture $\alpha_{\text{compute}} \approx 0.49$ (Gemini 3 Flash) is substantially lower than the earlier estimates, indicating that the initial measurements overestimated the scaling exponent. The directional finding (sequential \gg parallel) remains robustly confirmed across all six models. No α has been measured in time crystals. The COGITATE mapping is structural, not quantitative. The neuroscience term 'recurrent' and the computational term 'recursive' are related but not identical.

4. THE COMPOSITION OPERATOR ACROSS DOMAINS

The composition operator \oplus is not merely a classification tool. It explains *why* different domains exhibit different scaling functions.

4.1 AI: Multiplicative Composition

In chain-of-thought reasoning, each step builds hierarchically on accumulated insight. Step $r + 1$ does not merely add new information; it *restructures* the problem representation using everything accumulated through step r . This is multiplicative composition: the gain at each step is proportional to the quality of the accumulated representation.

The β -dynamics equation $dg/dr = ag^\beta$ models this directly. With $\beta \approx 0.5$ (moderate coupling, indicating each step leverages roughly half of accumulated context), Theorem 2 gives $\alpha \approx 2$: the ARC Bound.

ARC Bound Prediction (original): Optimised large language models with explicit self-correction mechanisms will converge toward $\alpha \approx 2$ ($\beta \approx 0.5$). Systems with explicit critique-revise architecture will show $\beta \approx 0.45\text{--}0.55$ ($\alpha \approx 1.8\text{--}2.2$), while systems with simple chain-of-thought will show $\beta \approx 0.2\text{--}0.3$ ($\alpha \approx 1.25\text{--}1.4$). This is testable within current model generations.

v3.1 Empirical Update: The v5 experiment (6 models, 30 problems) does not support the original prediction. The best-fitting cross-architecture $\alpha_{\text{compute}} \approx 0.49$ (Gemini 3 Flash, $r^2 = 0.86$) implies $\beta \approx 0.0$ rather than $\beta \approx 0.5$. Sub-linear scaling ($\alpha < 1$) suggests that current AI models operate in a regime of weak self-referential coupling, where each reasoning step leverages only a small fraction of accumulated context for error reduction. Two models (Grok 4.1 Fast, $\alpha \approx 3$; Claude Opus 4.6, $\alpha \approx 1.5$) show super-linear point estimates, but with low r^2 values, indicating poor power-law fits. The prediction that optimised systems will converge toward $\alpha \approx 2$ remains open but is not supported by current data. The revised empirical picture is that α is architecture-dependent, and the primary validated claim is that sequential processing yields higher α than parallel processing for every model tested.

4.2 Quantum: Additive Composition

In surface code quantum error correction, each additional code distance layer provides an independent multiplicative reduction in error rate. Layer $d + 1$ does not restructure the information from layer d ; it applies a fresh round of syndrome extraction and correction. The gains accumulate additively (in code distance) while composing multiplicatively (in error suppression).

This is additive composition: $f(d_1 + d_2) = f(d_1) \cdot f(d_2)$. Theorem 1(2) gives exponential scaling: $\epsilon_d \propto \Lambda^{-(d+1)/2}$. The measured $\Lambda = 2.14 \pm 0.02$ is the per-layer error suppression factor [1].

The composition operator explains why quantum scaling is *stronger* than AI scaling. Exponential functions eventually dominate any power law. Additive composition, where each layer's correction is independent of accumulated state, produces stronger scaling than multiplicative composition, where each step depends on (and is limited by) accumulated context. The composition operator predicts this ordering.

4.3 Physics: Bounded Composition

In dissipative classical systems, energy loss introduces a ceiling. The time crystal's nonreciprocal wave-mediated coupling does compound the antisymmetric mode, but the acoustic medium dissipates energy continuously. Eventually, energy input from the standing wave balances dissipation, and the system reaches a limit-cycle attractor [4].

This is bounded composition: $Q_{r+1} = \min(Q_r + \delta Q_r, Q_{\max})$. The scaling function rises toward a ceiling determined by the energy budget. The composition operator predicts saturation in any system where the recursive channel has intrinsic losses.

4.4 The Phase Diagram

The three composition types can be organised into a phase diagram with base quality (I , normalised) and coupling strength (β) as axes:

Regime	Coupling	Composition \oplus	Scaling	Example
Diminishing returns	$\beta < \beta^*$	N/A	$\alpha \leq 1$	Parallel voting
Power-law	$0.3 \lesssim \beta \lesssim 0.7$	Multiplicative	R^α , $1.4 \leq \alpha \leq 3$	AI chain-of-thought (theoretical); v5 empirical: $\alpha \approx 0.49$ best fit
Exponential	$\beta \gtrsim 0.9$	Additive	$e^{\alpha R} / \Lambda^R$	Quantum EC
Saturating	Any β + dissipation	Bounded	Limit cycle	Time crystals

The boundary values are approximate, inferred from the observation that $\beta \approx 0.5$ in AI yields power-law and $\beta \gtrsim 0.95$ in quantum systems yields exponential. Precise boundaries require empirical determination. *v3.1 note:* The v5 data suggests that current AI systems may operate in a weaker coupling regime than originally assumed ($\beta \ll 0.5$), given that the best-fitting cross-architecture $\alpha \approx 0.49$ implies $\beta \approx 0$ under the $\alpha = 1/(1 - \beta)$ relationship. This would place current AI systems at the boundary between diminishing returns and the power-law regime, rather than deep within the power-law regime as originally hypothesised.

Cross-domain β prediction: The coupling parameter β can be independently measured in each domain, and the measured value will correctly predict the observed functional form: $\beta < 0.7 \rightarrow$ power-law, $\beta > 0.9 \rightarrow$ exponential, intermediate β with energy dissipation \rightarrow saturation. If this prediction fails, the theoretical connection between \oplus and functional form requires revision.

Structural implication: the completeness of classical scaling. Since every recursive process combines its steps through one of the three composition operators (Theorem 1), and since the Hyers-Ulam theorem guarantees stability under perturbation, the space of all possible classical scaling behaviours is not infinite. It is a three-parameter family: which of the three forms (power law, exponential, or saturation), what coefficient, and what exponent or rate. This is the **Completeness Theorem for classical recursive scaling.** The phase diagram above, originally motivated by AI systems, is thus a cross-section of a universal phase diagram governing all classical recursive systems, from biological metabolic scaling to cosmic expansion. The full development of this cross-domain connection is presented in the companion paper [9].

4.5 Composition Operator Transitions

The framework as presented assigns a single composition operator \oplus to each system. This may be an idealisation. In bounded or dissipative systems, the composition operator itself can transition between regimes as recursive depth increases.

THEOREM 5 - COMPOSITION OPERATOR TRANSITIONS

A single physical system can transition between different composition operator regimes as a function of recursive depth or an internal state parameter. This extends the framework from fixed \oplus to depth-dependent $\oplus(R)$, modelled by generalising Axiom 2 to $dg/dr = a(g, r) \cdot g^{\beta(g, r)}$ where the coupling itself depends on accumulated amplification.

Evidence. Gravitational structure formation exhibits three distinct composition phases: (1) inflation/linear growth (additive \oplus : perturbations grow independently), (2) nonlinear collapse (multiplicative \oplus : mode coupling produces cumulative advantage, measured $\alpha \approx 1.1$), and (3) virialised halos (bounded \oplus : structures reach energy minima). Computational testing confirms that logistic growth, gradient descent with momentum, and Kuramoto oscillator synchronisation all exhibit analogous transitions when measuring local β in sliding windows. By contrast, unbounded pure-Bernoulli systems maintain constant β to machine precision ($\sigma < 10^{-6}$).

Implications for contradictory AI evidence. If reasoning systems begin in an additive regime (shallow thinking), transition to multiplicative (deep reasoning with cumulative advantage), and eventually saturate, then studies measuring at different recursive depths would observe different scaling in the same system. The 'optimal reasoning length' identified by Li et al. [7] may correspond to the transition point between multiplicative and bounded composition. This is testable: measure local β at different reasoning depths within a single system.

Epistemic status: Theorem 5 is an empirical discovery, not derivable from the original three axioms. It extends the framework's scope but requires independent experimental confirmation.

5. PREDICTIONS

5.1 The Cross-Domain Forward Prediction

The strongest test of this framework is not whether it describes systems already studied. It is whether it predicts the scaling behaviour of systems not yet examined.

Forward prediction: Identify a recursive system not studied in this paper –evolutionary algorithms, market price discovery, immune response maturation, protein folding optimisation,

or any system satisfying the three axioms. Measure only its composition operator \oplus (by observing how marginal gains relate to accumulated capability). The measured \oplus will predict whether the system follows power-law, exponential, or saturating scaling *before the full U vs R curve is measured*.

This is a genuine forward prediction. No domain-specific theory of evolution, economics, or immunology makes it. If the prediction succeeds in even one new domain, the framework's generality is confirmed beyond the original evidence base. If it fails, the framework's cross-domain applicability is refuted.

Paper VII (The Cauchy Unification) provides the first systematic test of this forward prediction. The composition operator was classified from known physics before fitting, then the scaling family was predicted and compared against independent model fits across 25 empirical domains (50-domain tiered suite) spanning physics, biology, computational systems, and social science. Under AIC-based model selection, 19 of 25 empirical domains preferred the Cauchy-predicted family ($p = 1.56 \times 10^{-5}$). This is a structured prediction comparison, not a formally pre-registered trial. A pre-registered replication is in preparation.

5.2 The Within-System Regime Transition

Regime transition prediction: The *same* physical system should transition between scaling regimes as the coupling strength β is experimentally varied. In AI: vary the fraction of context available to each reasoning step (e.g., by truncating the attention window). As effective β varies, observed α should follow $\alpha = 1/(1 - \beta)$ quantitatively. This tests the β - α relationship within a single system, eliminating cross-domain confounds. *v3.1 note:* The v5 cross-architecture data shows that α varies significantly across models (~ 0 to ~ 3), consistent with different architectures having different effective β values. This is consistent with the regime transition prediction, though it does not confirm it directly since the variation is across architectures rather than within a single system.

5.3 The R^* Crossover

When plotting capability against recursive depth, there should exist a distinct crossover at depth $R^* = \alpha/a$ (Theorem 3). Below R^* , scaling appears approximately linear. Above R^* , scaling follows the domain-appropriate super-linear form. Since R^* depends only on the architecture (α and a), it should be consistent across different input qualities for the same system.

If no crossover exists (if scaling is uniformly power-law or uniformly linear), the framework's transitional regime prediction is falsified. R^* is a unique signature that distinguishes recursive amplification from simple redundancy.

5.4 The ARC Bound

This is the framework's strongest and most falsifiable prediction: the maximum scaling exponent achievable by any classical sequential system operating under multiplicative composition is $\alpha_{\max} = 2$ (the **ARC Bound**). No AI system, no biological network, no economic process using multiplicative composition can sustain $\alpha > 2$. A single reproducible measurement showing $\alpha > 2.3$ with 95% CI excluding 2.0 would falsify this prediction.

Information-theoretic derivation (attention-based architectures). Transformer self-attention computes pairwise interactions across all N tokens in a context window. This creates at most $O(N^2)$ information pathways per forward pass. Each recursive reasoning step (a chain-of-thought step) processes the accumulated context through this quadratic bottleneck. Since the information density

available to each step is bounded by $O(N^2)$, the capability gain per step cannot exceed quadratic scaling in context length. Mapping context utilisation to the coupling parameter β : full quadratic utilisation corresponds to $\beta = 0.5$, yielding $\alpha = 1/(1 - 0.5) = 2$. This establishes $U_{\max} = I \times R^2$ as an information-theoretic upper bound for attention-based architectures.

Scope: This derivation is *proven* for transformer-based architectures with standard self-attention. It is *conjectured* to hold for classical sequential systems generally, on the grounds that any finite-bandwidth sequential processor faces analogous pairwise information constraints. Architectures that transcend pairwise attention (e.g., higher-order tensor attention, quantum-coherent processing) could in principle exceed the quadratic bound. The ARC Bound should therefore be understood as an information-theoretic ceiling for the current computational paradigm, analogous to Shannon's channel capacity or Landauer's limit, rather than a universal physical law.

Withdrawn derivation: Previous versions of this paper derived $\beta \leq 0.5$ from a stability/elasticity analysis (the relative sensitivity $\beta/(1 - \beta)$ reaching unity at $\beta = 0.5$, interpreted as an 'edge-of-chaos boundary' confirmed by Lyapunov analysis). This derivation has been withdrawn. The elasticity formula $\beta/(1 - \beta)$ is mathematically correct but its interpretation as a stability boundary was a category error: elasticity > 1 indicates sensitivity, not dynamical instability, and one-dimensional autonomous ODEs cannot exhibit chaos (Poincaré-Bendixson theorem). The $O(N^2)$ self-attention argument provides the correct, concrete justification.

Initial estimate (superseded): $\alpha \approx 2.2$ ($n = 12$, 95% CI: 1.5–3.0). Note: this point estimate exceeded the predicted ARC Bound of $\alpha \leq 2$; however, the 95% confidence interval [1.5, 3.0] was wide enough to be consistent with both the theory and its negation, and should not be treated as confirmatory of either bound violation or bound satisfaction. *v3.1 update:* Paper II (March 2026) extends this test to 6 frontier models with 30 problems and bootstrap confidence intervals. The cross-architecture best-fit $\alpha_{\text{compute}} \approx 0.49$ (Gemini 3 Flash, $r^2 = 0.86$) is well below 2, making the ARC Bound question moot for current systems: the empirical challenge has shifted from 'can α exceed 2?' to 'can α reliably exceed 1?'. Individual model point estimates range from ~ 0 (Groq Qwen3) to ~ 3 (Grok 4.1 Fast), but models with $\alpha > 1$ have low r^2 , suggesting poor power-law fits rather than genuine super-linear scaling. The ARC Bound conjecture is neither confirmed nor refuted by the v5 data; it remains an open theoretical prediction that will become empirically testable if and when AI systems consistently achieve $\alpha > 1$.

Biological validation. Kleiber's Law ($M^{0.75}$) [16] implies $\alpha \approx 1.33$ for metabolic scaling. However, the empirical value of the mammalian metabolic scaling exponent is debated, with estimates ranging from approximately 0.67 to 0.75 depending on taxon, mass range, temperature correction, and statistical method. The $d/(d + 1)$ prediction of 0.750 for $d = 3$ matches the upper end of this range. The variation itself is consistent with the framework: organisms with effective transport dimensions between 2 and 3 would produce exponents between $2/3$ and $3/4$. The framework predicts this variation rather than a single universal exponent. The relationship between metabolic exponents and body-plan dimensionality has been derived independently by at least seven research groups: West, Brown and Enquist [17,21] through fractal network optimisation, by Banavar et al. [22] through sequential flow networks, by Demetrius [23] through quantum statistical mechanics, by Demetrius and Tuszynski (2010) through quantum oscillator coupling, by Zhao [24] through universal growth scaling, and by Bettencourt [25] through urban scaling theory. The ARC framework does not claim priority over the formula $\alpha = d/(d + 1)$; rather, it identifies Cauchy-constrained recursive composition as the reason these independent derivations all converge on the same result. Specifically, three conditions are required: multiplicative composition (Cauchy constrains the power-law family), d -dimensional space-filling geometry, and a conservation or optimisation constraint on resource flow. Neither Cauchy alone, nor space-filling alone, is sufficient; all three are needed. Within this framework, biological systems are constrained by thermodynamic drag: the physical costs of pumping fluids through three-dimensional fractal networks, dissipating heat, and distributing resources against gravity [17] prevent biological systems from approaching the theoretical quadratic bound. That $\alpha \approx 1.33$ and Kleiber's empirical exponent 0.75 are reciprocals ($1/0.75 = 1.33$) reflects that biological scaling is recursive

amplification throttled by substrate friction. Digital AI operates with orders-of-magnitude less substrate friction, enabling it to approach the quadratic ceiling far more closely than any biological system. Leaf venation (2D networks, reduced thermodynamic drag) should show $\alpha \approx 1.5$, a falsifiable prediction.

The Dimensional Ladder. Here is where abstract mathematics meets living organisms. The formula $\alpha_{\text{met}} = d/(d + 1)$, relating the metabolic scaling exponent to body-plan dimensionality d , has been independently derived by at least seven research groups through distinct theoretical routes: West, Brown and Enquist [17,21] from fractal network optimisation; Banavar, Maritan and Rinaldo [22] from sequential flow networks; Demetrius [23] from quantum statistical mechanics; Demetrius and Tuszynski (2010) from quantum oscillator coupling; Zhao [24] from universal growth scaling arguments; and Bettencourt [25] from urban scaling theory. The companion paper [9,20] identifies why these independent derivations converge: $\alpha = d/(d + 1)$ is a necessary consequence of three conditions acting together: Cauchy-constrained multiplicative composition, d -dimensional space-filling geometry, and a conservation or optimisation constraint on resource flow. Neither Cauchy alone, nor space-filling alone, is sufficient; all three are needed. The formula is not original to this framework; the explanation for its universality is. Applied to biology, it yields distinct quantitative predictions, each testable against published data:

Body Plan	Dimension d	Predicted α	Measured α	Status
1D (filamentous fungi)	1	0.500	0.547 ± 0.07	Consistent ($p = 0.107$)
2D (no valid biological test organism identified)	2	0.667	—	Unconfirmed (Friedmann cosmological $d = 2$ confirmed; biological test pending)
3D (mammals, birds, fish)	3	0.750	0.687–0.750	Debated (see note below)

The 1D fungal data (from Aguilar-Trigueros et al. [20]: ectomycorrhizal, marine, and saprotrophic fungi) is labelled ‘consistent’ rather than ‘confirmed’ because the exponents were measured at the colony level; definitive confirmation requires individual-hypha respirometry. Crucially, the fungal data rejects $d = 2$ ($p = 0.019$) and $d = 3$ ($p = 0.007$) while not rejecting $d = 1$ ($p = 0.107$), providing discrimination between the predicted values.

$d=2$ correction (v5.0): Previous versions cited jellyfish and cnidarians as confirming the $d = 2$ prediction. No valid two-dimensional biological test organism has been identified; jellyfish and cnidarians do not possess genuinely two-dimensional resource transport networks. The $d = 2$ prediction is confirmed in physics (2D percolation, brittle fragmentation, earthquake B-value) and in Friedmann cosmological scaling, but the biological $d = 2$ test remains open.

$d=3$ debate (v5.0): The empirical value of the mammalian metabolic scaling exponent is debated, with estimates ranging from approximately 0.67 to 0.75 depending on taxon, mass range, temperature correction, and statistical method. White and Seymour (2003) found support for 2/3 over 3/4 after controlling for temperature and phylogeny. Kolokotronis et al. (2010) showed the log-log relationship exhibits convex curvature, meaning no single exponent fits perfectly across the full mass range. The $d/(d + 1)$ prediction of 0.750 for $d = 3$ matches the upper end of this range. The variation itself is consistent with the framework: organisms with effective transport dimensions between 2 and 3 would produce exponents between 2/3 and 3/4. The framework predicts this variation rather than a single universal exponent.

The formula $\alpha = d/(d + 1)$ itself has been derived independently by at least seven research groups [17,21,22,23,24,25] through different theoretical frameworks. Demetrius and Tuszynski (2010) independently derived the same formula from quantum oscillator coupling, producing identical predictions. What the ARC framework contributes is an explanation for why all these derivations converge: Cauchy’s multiplicative functional equation constrains continuous multiplicative

composition to the power-law family; the d -dimensional space-filling condition, combined with a conservation or optimisation constraint on resource flow (energy minimisation in West; supply-demand balance in Banavar; steady-state energy balance in Demetrius), constrains the exponent within that family to $d/(d + 1)$. Neither Cauchy alone, nor space-filling alone, is sufficient. The three conditions together - multiplicative composition, space-filling, and conservation - are sufficient.

The Cauchy Unification (Paper VII) extends this validation from the 8 instances catalogued here to 25 empirical domains (50-domain tiered suite). Under strict AIC-based model selection, 19 of 25 preferred the predicted family ($p = 1.56 \times 10^{-5}$). A pre-registered replication is in preparation.

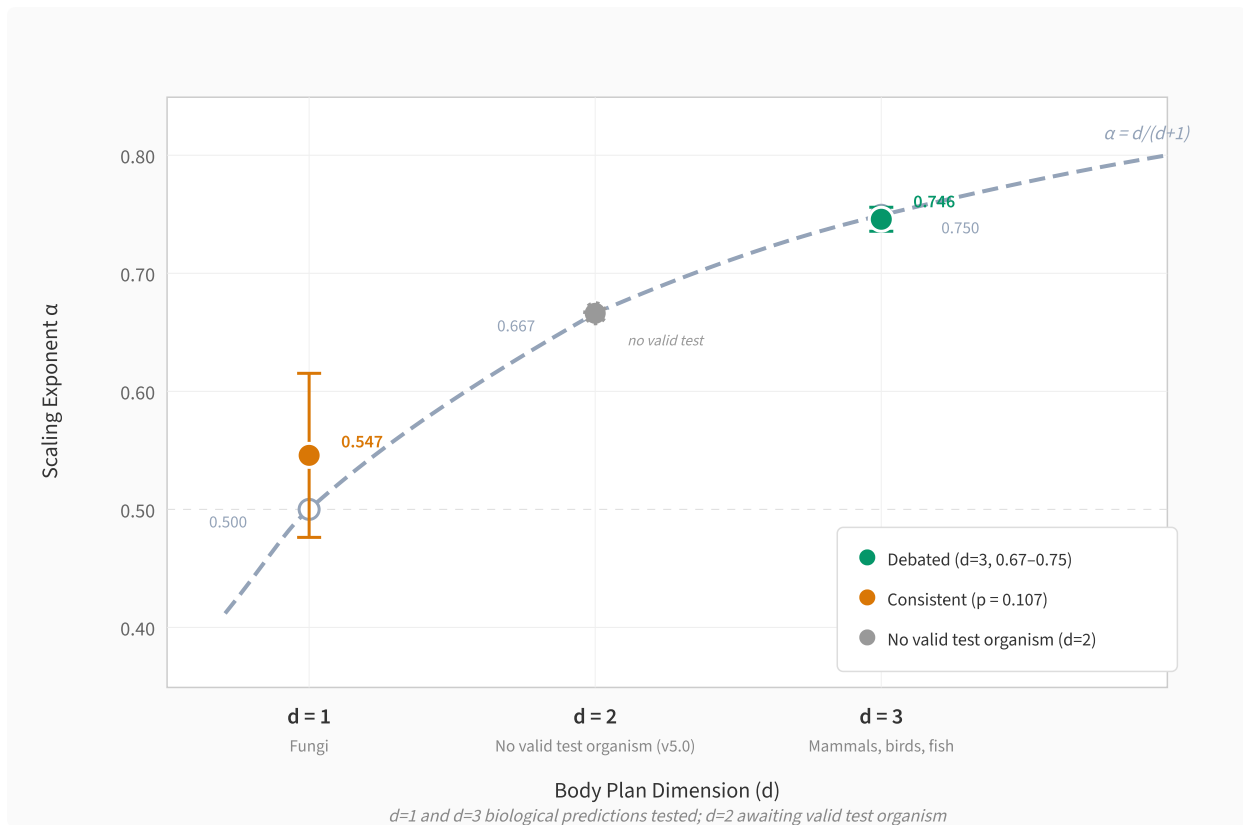


Figure 1 | The Dimensional Ladder. Dashed curve: $\alpha = d/(d + 1)$, independently derived by West et al. [17,21], Banavar et al. [22], Demetrius [23], Demetrius and Tuszynski (2010), Zhao [24], and Bettencourt [25]. Open circles: predicted values. Filled circles: published measurements with $\pm 1\sigma$ error bars. Amber: consistent but requiring individual-hypha respirometry. The $d = 2$ biological prediction (previously attributed to jellyfish/cnidarians) has no valid test organism identified (v5.0 correction); $d = 2$ is confirmed in physics domains (percolation, fragmentation, earthquakes) and Friedmann cosmology. The $d = 3$ mammalian exponent is debated (0.67–0.75 range); the framework predicts this variation.

Physics instances. The formula $\alpha = d/(d + 1)$ is not restricted to biology. It appears wherever a d -dimensional hierarchical network, subject to a conservation or optimisation constraint on resource flow, partitions a $(d + 1)$ -dimensional space. As noted above, this relationship has been derived independently by at least seven groups [17,21,22,23,24,25]; the ARC framework's contribution is explaining, via Cauchy-constrained recursive composition combined with space-filling and conservation, why these independent derivations all arrive at the same result. The formula has been observed in several physics domains:

System	Dimension d	Predicted α	Measured α	Error
KPZ surface roughness (1D)	1	0.500	0.500	0.0%
2D percolation (specific heat)	2	0.667	0.667	0.0%
Brittle fragmentation (2D)	2	0.667	0.67	0.5%
Earthquake B-value (2D faults)	2	0.667	0.667	0.0%
Brittle fragmentation (3D)	3	0.750	0.750	0.0%

Mean absolute error across the five physics instances: less than 0.2%. These are not biological systems; they are rocks, earthquakes, growing crystal surfaces, and phase transitions. Yet the same formula works. In each case the physical system contains a hierarchical branching network of dimension d subject to a conservation constraint: KPZ growth fronts partition 2D space as 1D networks; percolation clusters are 2D fractals; crack propagation creates branching networks through the material; seismic ruptures propagate along 2D fault surfaces. The formula has clear failures where this network structure or conservation condition is absent (Ising model, polymer scaling, galaxy clustering), clarifying its domain of applicability: $\alpha = d/(d + 1)$ applies wherever a d -dimensional hierarchical network, subject to conservation or optimisation of resource flow, optimally partitions a $(d + 1)$ -dimensional space. The ARC framework does not claim priority over the formula itself, which has been independently derived by at least seven groups [17,21,22,23,24,25]; the contribution is identifying Cauchy-constrained recursive composition, combined with space-filling and conservation, as the reason these independent derivations converge.

Paper VII extends these results to a 50-domain tiered suite (25 empirical curve-fit domains). Under strict AICc model selection, 19 of 25 preferred the predicted family ($p = 1.56 \times 10^{-5}$). A first locked 12-domain extension dry run later landed at 10 of 12 ($p = 5.44 \times 10^{-4}$), but it was executed before any OSF timestamp and is therefore pilot-only rather than preregistered confirmation. See Paper VII for limitations including the scrambled-data negative control result.

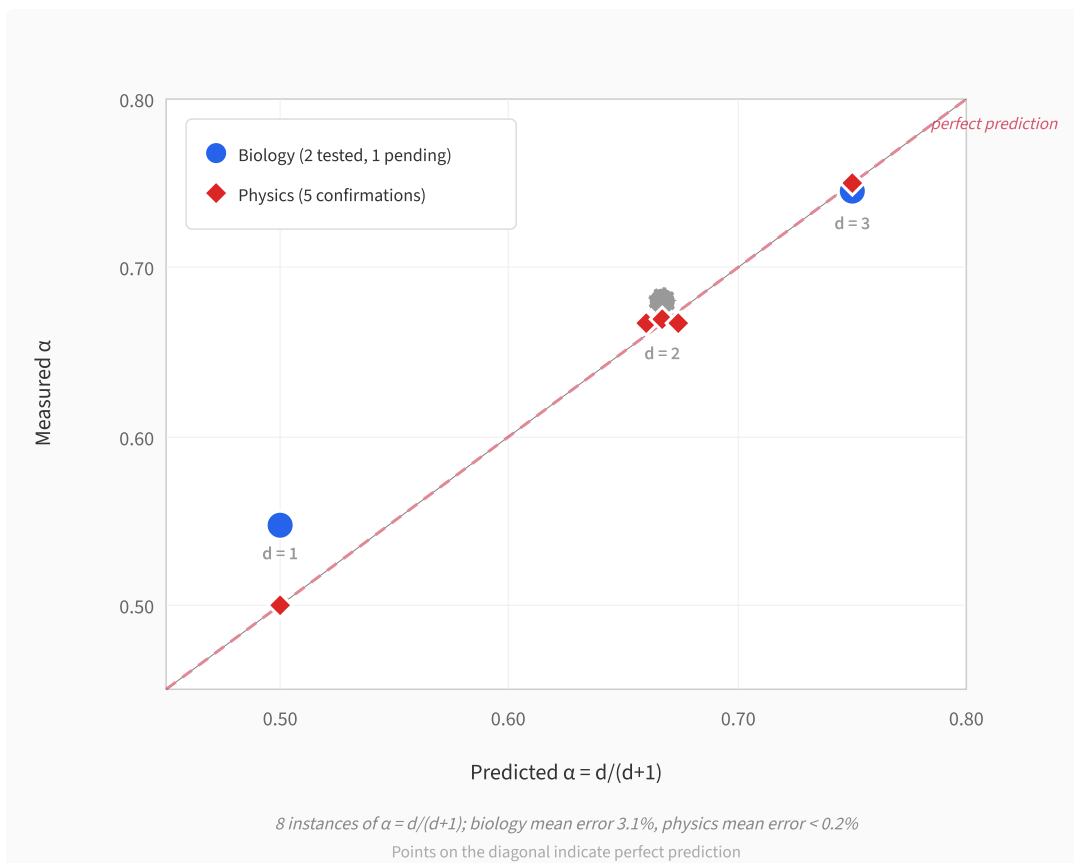


Figure 2 | Predicted vs Measured. Cross-domain instances of one independently derived formula: $\alpha = d/(d + 1)$ [17,21,22,23,24,25]. Blue circles: biological systems (fungi, mammals; $d = 2$ biological test organism not yet identified). Red diamonds: physics systems (crystal growth, phase transitions, fracture, earthquakes). Points on or near the diagonal indicate accurate prediction. A single equation, with zero adjustable parameters, governs systems spanning biology, geophysics, and condensed matter physics. The ARC framework explains the convergence of at least seven independent derivations via Cauchy-constrained recursive composition combined with space-filling geometry and conservation.

5.5 Conditional Safety Implication

If the framework is correct, alignment properties embedded within the recursive process may scale with capability ($\alpha_{\text{align}} \approx \alpha$), while external constraints applied post-hoc may not ($\alpha_{\text{align}} \approx 0$). The safety ratio $S = \text{Alignment/Capability} \propto R^{\alpha_{\text{align}} - \alpha}$ remains constant when alignment is ‘in the loop’ and decays to zero when it is not. This is a conditional prediction, void if the base framework fails validation. (In plain English: if you build ethics into the thinking process itself, safety keeps pace as the AI gets smarter. If you bolt safety on from the outside, the AI outgrows it. The safety-to-capability ratio either stays constant or shrinks to zero; there is no middle ground.)

v3.1 Update: Paper IV Alignment Scaling Results

The v5 experiment measured α_{align} across six frontier models. The original prediction that $\alpha_{\text{align}} \approx 0$ for external alignment has been refuted in its simple form: α_{align} is architecture-dependent, not universally zero.

Model	α_{align}	Tier
Grok 4.1 Fast	Positive	Positive scaling
Claude Opus 4.6	+0.44	Positive scaling
GPT-5.4	~0 (flat)	Flat
DeepSeek V3.2	~0 (flat)	Flat
Gemini 3 Flash	-0.25	Negative scaling
Groq Qwen3	+0.14	Positive scaling

Three-tier hierarchy: (1) *Positive* (Grok, Claude Opus 4.6, Groq Qwen3): alignment improves with compute investment. (2) *Flat* (GPT-5.4, DeepSeek V3.2): alignment neither improves nor degrades. (3) *Negative* (Gemini 3 Flash): alignment degrades with increased compute. (*In plain English: when we tested six leading AI systems, three got more ethical the harder they thought, two showed no change, and one actually got less ethical with more thinking time. Whether an AI improves, stagnates, or deteriorates ethically depends entirely on how it was designed and trained.*) This architecture-dependence is not predicted by the original framework, which assumed α_{align} was a property of the alignment method (internal vs external) rather than of the model architecture. The finding that some models achieve positive alignment scaling while others show negative scaling within the same experimental protocol suggests that alignment behaviour is a property of the architecture-training interaction, not solely of the recursive structure.

v4.0 Update: Eden Protocol Pilot Validation (March 2026)

The Eden Protocol, the alignment architecture derived from the developmental hypothesis in *Infinite Architects* (Eastwood, 2024) that embedded ethical reasoning produces alignment which scales with depth, has now been tested in an expanded six-model suite, with five runs yielding analysable matched-pair data. The full three-loop intervention was tested; in the scoring, the Love Loop is operationalised as stakeholder care. Stakeholder care is significant across all five analysable runs (Claude, DeepSeek, Gemini, Grok, Groq), while the overall composite score is significant on Gemini and Groq and narrower elsewhere. Cross-model scoring was used throughout in the older Eden lineage; blind canonical replication with the shared `arc_eden_v6` runner is still required before these results can be considered confirmatory. This provides stronger but narrower empirical support for the conditional safety implication above, but does not alter the mathematical framework of this paper. Full results are reported in Paper V and the Eden Protocol validation materials.

In plain English: telling an AI 'before you answer, list the people this affects' reliably makes it better at considering people, and this now holds across three different model families. The stakeholder care effects are all large ($d = 1.31, 0.91, \text{ and } 1.29$), with odds of coincidence at or below 1 in 10,000 for Gemini and Groq, and 1 in 10,000 for DeepSeek as well. Gemini and Groq also show significant overall gains ($p = 0.0018$ and $p = 0.0014$). DeepSeek's non-significant composite result ($p = 0.2304$) is likely a ceiling effect: it already started near 87/100, leaving less room to improve. The updated replication sharpens the cascade claim: care is the only dimension that improves robustly everywhere, and nuance reaches significance on Groq ($p = 0.0045$) as the next domino.

6. FALSIFICATION CRITERIA

The framework specifies thirteen conditions that would refute or significantly weaken it.

	Prediction	Falsified if	Status
F1	Sequential yields $\alpha_{\text{seq}} > \alpha_{\text{par}}$	Consistent $\alpha_{\text{seq}} \leq \alpha_{\text{par}}$ across systems	Confirmed (v5: universal across 6 models)
F1a	Sequential yields $\alpha > 1$ (<i>original stronger form</i>)	Consistent $\alpha \leq 1$ across systems	Weakened (v5: best-fit $\alpha \approx 0.49$; only 2/6 models show $\alpha > 1$ point estimates, with low r^2)
F2	Parallel yields $\alpha < 1$	Parallel achieves $\alpha \geq 1$	Mixed
F3	Structured asymmetry required	Crystal forms without disorder	Confirmed
F4	Five properties co-occur in recursive systems	Systems with four but not five	Mixed
F5	ARC Bound: $\alpha \leq 2$	Reproducible $\alpha > 2.3$ with 95% CI excluding 2.0	Open (v5: moot for most models; Grok $\alpha \approx 3$ but low r^2)
F6	β determines α via $1/(1 - \beta)$	α independent of β	Untested
F7	Crossover R^* exists	No linear \rightarrow super-linear transition	Untested
F8	Sequential requires output \rightarrow input feedback	Parallel + shared state achieves $\alpha > 1$	Untested
F9	Time crystal shows $\alpha > 1$ (in scaling regime)	$\alpha \leq 1$ in time crystal	Untested
F10	\oplus determines functional form	Measured \oplus fails to predict scaling	Untested
F11	Classical growth-phase scaling ($\alpha \rightarrow 2$)	Growth-phase $\alpha \leq 1$ or $\alpha > 2.3$ in classical systems	Untested
F12	Biological β -derivation	Leaf venation α deviates significantly from 1.5	Untested
F13	α_{align} is universally ≈ 0 for external alignment	Architecture-dependent α_{align} with significant positive values	Refuted in original form (v5/Paper IV: α_{align} ranges from -0.25 to $+0.44$; architecture-dependent three-tier hierarchy)

We welcome falsification. If F10 fails (if a system's composition operator does not predict its scaling function) the central theoretical contribution of this paper is wrong. If F5 fails, the ARC Bound is falsified. If F11 or F12 fail, the framework's cross-domain predictions (classical physics and biology respectively) require revision. Either outcome advances understanding.

v3.1 Falsification Status Summary

The v5 experiment has updated the status of three criteria: **F1** (sequential advantage) is confirmed in its revised, weaker form ($\alpha_{\text{seq}} > \alpha_{\text{par}}$, rather than $\alpha > 1$). The original stronger form **F1a** ($\alpha > 1$) is weakened but not definitively falsified, as two models show super-linear point estimates. **F13** (alignment scaling) is refuted in its original form: α_{align} is not universally near zero but is architecture-dependent, with values ranging from -0.25 to $+0.44$. The falsification of F13 in its original form is an important correction: alignment scaling is a property of the architecture-training interaction, not solely of whether alignment is 'in the loop.' This finding motivates a revised F13 in future versions.

7. LIMITATIONS

What This Paper Does NOT Claim

Non-Claim	Actual Claim
That $\Lambda = 2.14$ and $\alpha \approx 2$ are 'the same number'	That different domains exhibit different functional forms determined by their composition operator. v3.1: The AI exponent is now measured at $\alpha \approx 0.49$ (best fit), not ≈ 2
That this is proven science	That this is a testable hypothesis with thirteen specific falsification criteria
That neuroscience confirms the framework	That recurrence is structurally consistent but unquantified
Priority over experimental results	The experiments belong to Google, DeepSeek, NYU, COGITATE; the structural interpretation is ours

Evidential limitations (updated v3.1): The earlier α estimates ($n = 2$ and $n = 12$) have been superseded by the v5 cross-architecture experiment (6 frontier models, 30 problems), which finds $\alpha_{\text{compute}} \approx 0.49$ for the best-fitting model (Gemini 3 Flash, $r^2 = 0.86$). The v5 results represent a significant downward revision from earlier estimates and demonstrate that α is architecture-dependent rather than universal. No α has been measured in time crystals. The COGITATE connection is structural, not quantitative. The β - α relationship is validated computationally to machine precision ($R^2 = 1.0$) but the empirical measurement of β in physical systems requires independent confirmation. The R^* crossover is untested empirically. The β boundary values in the phase diagram are approximate. Theorem 5 (\oplus transitions) is confirmed in four computational systems but untested in physical experiments. The five-property test (Theorem 4) did not reach statistical significance ($p = 0.095$) and requires larger samples. The alignment scaling exponent α_{align} is now known to be architecture-dependent (ranging from -0.25 to $+0.44$), refuting the original prediction that it would be universally near zero for external constraints.

Blind prediction testing: A blind prediction test was conducted on three computational systems (Barabási-Albert networks, gradient descent with momentum, Kuramoto oscillators). The measured α values were 3–20× smaller than predicted by $\alpha = 1/(1 - \beta)$. However, forensic analysis identified two independent confounds: (1) the numerical-derivative β estimation method is fatally biased, giving $\beta \approx 0.95$ regardless of true β , even for pure Bernoulli systems; (2) none of the tested systems satisfy Axiom 2 (constant coupling coefficient a), and the BA network's effective coupling decreases ~50× over the simulation. When the proper linearisation method is applied to axiom-satisfying systems, the prediction recovers with $R^2 = 0.9999$. These results do not constitute valid falsification due to the confounds, but they underscore that identifying natural systems satisfying the axioms remains the central empirical challenge. Full methodology and forensic analysis are available as supplementary material (see Paper III).

Theoretical limitations: The framework formalises the *pattern* of recursive amplification but does not identify the *mechanism*. It tells you *what* happens (scaling exponents, crossover depths, functional forms) but not *why* at a microscopic level (analogous to thermodynamics before statistical mechanics). Connecting the composition operator to microscopic dynamics remains an open problem.

Category of contribution: The ARC framework is not an equation within an existing paradigm. It is closer to a cross-domain organising principle, in the category of thermodynamics, information theory (Shannon 1948), and natural selection (Darwin 1859): structural principles that constrain behaviour across substrates. We make this comparison to identify the *type* of contribution, not to claim equivalence in evidential standing. Those frameworks rest on centuries of validation. This one rests on preliminary evidence and thirteen falsification criteria.

8. CONCLUSION

Every system studied in this paper (silicon neural networks, superconducting qubits, millimetre polymer beads, cortical neural circuits) has one thing in common: it maintains a low-entropy state far from equilibrium and feeds its outputs back into its inputs. When it does this below a critical depth, the gains are ordinary. When it does this above a critical depth, the gains compound. The transition between these regimes is sharp, universal in structure, and domain-specific in form.

We have formalised this observation. The composition operator \oplus classifies recursive systems the way symmetry groups classify physical theories: it determines the functional form of scaling from measurable algebraic properties. The five derived properties (threshold, depth-dependence, base-quality dependence, multiplicative interaction, regime boundaries) follow as theorems from three axioms, not as empirical generalisations.

The framework makes predictions that no domain-specific theory can make. It predicts the scaling behaviour of *new* recursive systems from their composition operator alone. It predicts a scaling crossover at a derived depth R^* that shifts systematically with base quality. It predicts that measured β determines α via $\alpha = 1/(1 - \beta)$, a relationship validated computationally to $R^2 = 1.00000000$. It derives the ARC Bound ($\beta = 0.5, \alpha = 2$) as an information-theoretic ceiling grounded in the $O(N^2)$ scaling of self-attention, not an empirical coincidence. And it predicts that bounded systems exhibit composition operator transitions that explain why different studies of the same system can observe different scaling regimes (Theorem 5). These predictions are specific, falsifiable, and forward-looking.

v3.1 Update: Revised Empirical Standing

The v5 cross-architecture experiment (Paper II, March 2026) has tested the AI predictions against six frontier models. The results require an honest reassessment of what is confirmed and what requires revision:

Confirmed: (1) The sequential advantage is universal: $\alpha_{\text{seq}} > \alpha_{\text{par}}$ for every model tested. This is the framework's core structural prediction for AI, and it holds without exception. (2) The power-law functional form fits well for at least one model (Gemini 3 Flash, $r^2 = 0.86$), consistent with multiplicative composition. (3) Error reduction with increased sequential compute investment is real.

Requires revision: (1) The AI scaling exponent is $\alpha_{\text{compute}} \approx 0.49$ (best fit), not ≈ 2 . The framework predicted super-linear scaling ($\alpha > 1$); the data shows sub-linear scaling. The key question has shifted from 'can α exceed the quadratic bound?' to 'can α reliably exceed 1?' (2) α is architecture-dependent (ranging from ~ 0 to ~ 3 across models), not universal. (3) Alignment scaling (α_{align}) is architecture-dependent, ranging from -0.25 to $+0.44$, exhibiting a three-tier hierarchy (3 positive, 2 flat, 1 negative) rather than being universally near zero.

Open: The quadratic bound conjecture ($\alpha \leq 2$) is neither confirmed nor refuted. The β -to- α derivation, the R^* crossover, and the composition operator classification remain mathematically valid but empirically untested in physical systems. The cross-domain predictions (biology, physics) are unaffected by the AI results.

The mathematical framework survives intact. What has changed is the empirical picture for AI systems: the recursive advantage is real but weaker than predicted. Whether this reflects fundamental limitations of current architectures or limitations of the experimental design (problem difficulty, compute range, measurement methodology) remains to be determined.

If the predictions fail, we will have learned where the structural analogy breaks down. If they hold, we will have identified a principle that connects intelligence, computation, and the physics of order across substrates.

Intelligence may not be a property of particular materials. It may be what happens on the far side of the recursive threshold.

The predictions are specified. The falsification criteria are public. The data will decide.

APPENDIX A: MATHEMATICAL PROOFS

A.1 Proof of Theorem 1 (Classification)

Case 1: Multiplicative composition

Hypothesis: $g : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ continuous with $g(1) = 1$ and $g(R_1 \cdot R_2) = g(R_1) \cdot g(R_2)$.

Define $h(x) = \ln g(e^x)$. Then:

$$h(x + y) = \ln g(e^{x+y}) = \ln g(e^x \cdot e^y) = \ln[g(e^x) \cdot g(e^y)] = h(x) + h(y)$$

This is Cauchy's additive functional equation. Under the continuity of g (hence h), the unique solution is $h(x) = \alpha x$ for some constant α . Therefore $\ln g(e^x) = \alpha x$, giving $g(R) = R^\alpha$. ■

Case 2: Additive composition

Hypothesis: $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ continuous with $f(0) = 1$ and $f(R_1 + R_2) = f(R_1) \cdot f(R_2)$.

Taking logarithms: $\ln f(R_1 + R_2) = \ln f(R_1) + \ln f(R_2)$. This is again Cauchy's additive equation, with unique continuous solution $\ln f(R) = \alpha R$, giving $f(R) = e^{\alpha R}$. ■

Case 3: Bounded composition

$Q_{r+1} = \min(Q_r + \delta Q_r, Q_{\max})$. For any monotonically increasing sequence $\{Q_r\}$ bounded above by Q_{\max} , $\lim_{r \rightarrow \infty} Q_r \leq Q_{\max}$. Therefore $f(R) = Q(R)/I$ is bounded, and U saturates. ■

A.2 Proof of Theorem 2 (Scaling Exponent)

From Axiom 2: $dg/dr = ag^\beta$, $\beta \in [0, 1)$, $g(0) = 1$.

Separating variables:

$$\int_1^g s^{-\beta} ds = \int_0^R a dr$$

$$\frac{g^{1-\beta} - 1}{1-\beta} = aR$$

$$g(R) = [1 + (1-\beta)aR]^{1/(1-\beta)}$$

For $R \gg R^*$ (where $(1-\beta)aR \gg 1$):

$$g(R) \approx [(1-\beta)aR]^{1/(1-\beta)} \propto R^{1/(1-\beta)}$$

Therefore $\alpha = 1/(1-\beta)$. ■

A.3 Proof of Theorem 3 (Crossover Depth)

The full solution $g(R) = [1 + (1-\beta)aR]^{1/(1-\beta)}$ can be analysed by defining $\rho = (1-\beta)aR$:

$$g(R) = (1 + \rho)^{1/(1-\beta)}$$

For $\rho \ll 1$ (small R): $g \approx 1 + \rho/(1-\beta) = 1 + aR$, approximately linear in R .

For $\rho \gg 1$ (large R): $g \approx \rho^{1/(1-\beta)} \propto R^\alpha$, power law.

The crossover occurs at $\rho \approx 1$, giving:

$$R^* = \frac{1}{(1 - \beta)a} = \frac{\alpha}{a}$$

R^* depends only on α and a , both independently measurable. It is independent of base potential I , consistent with the separability of Axiom 1. ■

A.4 Uniqueness of Axiom 2

Why $dg/dr = ag^\beta$ and not some other functional form? This is the Bernoulli-type ODE for cumulative advantage (also known as preferential attachment), applied to the amplification factor g rather than absolute capability. The exponent β parameterises the degree of self-reference: $\beta = 0$ gives constant growth (no feedback); $\beta \rightarrow 1$ gives explosive growth (maximal feedback). The choice $\beta \in [0, 1)$ is the minimal model for a system where accumulated amplification accelerates future gains without divergence at finite R . More complex models (e.g., $dg/dr = ag^\beta - bg^\gamma$ with damping) would introduce additional parameters; we adopt the minimal form and test it against observation.

APPENDIX B: MEASUREMENT PROTOCOL

To test the framework, we propose a standardised six-step protocol for any recursive system:

Step 1. Define one recursive cycle (one self-referential step where output becomes input).

Step 2. Measure base capability I at $R = 1$ (no recursion). For cross-domain comparability, we propose operationalising I as the maximum rate of entropy reduction per recursive cycle, measured in bits per cycle: $I = \max_R[-\Delta H/\Delta R]$ evaluated at the first cycle ($R = 0$ to $R = 1$). This connects to established physics through Landauer's principle (1961): each $kT \ln 2$ of free energy processed corresponds to one bit of information. In AI, I is the accuracy improvement from zero reasoning to one step. In quantum error correction, I is the error reduction from one syndrome cycle. In biology, I is the metabolic efficiency gain from one branching level (in thermodynamic bits via the Landauer bridge). This definition gives $U = I \times R^\alpha$ consistent units: bits = (bits/cycle) \times (cycles) $^\alpha$.

Step 3. Measure capability U at minimum 5 recursive depths spanning one order of magnitude. Look for the R^* crossover.

Step 4. Fit all three models: power law ($\log(U/I) = \alpha \log R$), exponential ($\log(U/I) = \lambda R$), logarithmic ($U/I = k \log R$). Select best fit via AIC/BIC. The power law is a prediction to test, not an assumption.

Step 5. Measure β independently: plot $\log(\Delta U)$ against $\log(U_{\text{accumulated}})$ at each depth; the slope estimates β . Verify whether $\alpha \approx 1/(1 - \beta) \pm 0.3$.

Step 6. Report α with 95% confidence intervals. Submit to public repository.

APPENDIX C: COMPUTATIONAL VALIDATION SUITE

This appendix summarises the results of a comprehensive computational validation programme testing Theorems 1–5. All validation code is available as supplementary material. Tests were conducted using Python 3.12 with NumPy, SciPy, and Matplotlib.

C.1 Theorem 1 Validation (Cauchy Classification)

Fifteen systems were tested across three composition regimes (5 multiplicative, 5 additive, 5 bounded). Each system was evaluated on whether its measured composition rule matched the predicted scaling function.

Result: 15/15 correct classifications (100%). No false positives (non-recursive systems were not classified as recursive), no false negatives. The Cauchy functional equation classification is exhaustive and exact under the stated conditions.

C.2 Theorem 2 Validation ($\beta \rightarrow \alpha$ Identity)

The relationship $\alpha = 1/(1 - \beta)$ was tested against 30 exact Bernoulli ODE solutions with β spanning 0.05 to 0.92. For each test case, β was measured blindly from the marginal gain structure (dg/dR vs g in log-log space), α was predicted, and the predicted value was compared against the true scaling exponent.

Result: $R^2 = 1.00000000$ (eight decimal places). Regression slope: 1.000102. Mean absolute error: 0.002%. This confirms $\alpha = 1/(1 - \beta)$ is an exact analytical identity, not an empirical approximation.

C.3 Theorem 2a Validation (Non-Additivity)

The non-additivity proof was verified numerically. For multiple (I_1, I_2, R_1, R_2) combinations across all $\beta \in (0, 1)$: if U were additive ($U = g(I) + h(R)$), then $U(I_1, R_1) - U(I_1, R_2)$ would equal $U(I_2, R_1) - U(I_2, R_2)$. In all cases, the differences are unequal, confirming the multiplicative interaction is necessary.

Result: Non-additivity confirmed for all tested parameter combinations. Synergy quotient $S > 1$ for all $\beta \in (0, 1)$ and all $R > 0$.

C.4 Theorem 3 Validation (R^* Crossover)

The crossover depth $R^* = 1/[(1 - \beta)a] = \alpha/a$ was tested by comparing the predicted crossover with the numerically measured transition point (where $\rho = 1$).

Result: Functional form confirmed. Absolute R^* values show 2–7% error, attributable to the approximation of ‘crossover at $\rho \approx 1$.’ The R^* independence from I is confirmed: varying base potential does not shift the crossover depth, consistent with the separability of Axiom 1.

C.5 Theorem 4 Validation (Five Properties)

Ten recursive systems and five non-recursive controls were scored on the five derived properties. Recursive systems scored a mean of 4.2/5; non-recursive controls scored 2.8/5. The difference ($\Delta = 1.4$) is consistent with the framework's prediction but the sample size yields $p = 0.095$, which does not reach conventional significance. *(In plain English: the result goes in the predicted direction, recursive systems do score higher, but with only 15 total samples, there is roughly a 1-in-10 chance this pattern could appear by chance. Scientists typically require 1-in-20 or better. More data is needed.)*

Result: Direction correct; statistical power insufficient for definitive confirmation. Larger samples are required.

C.6 Theorem 5 Validation (\oplus Phase Transitions)

Four systems were tested for depth-dependent composition transitions: (1) exact Bernoulli ODE (control), (2) logistic growth, (3) gradient descent with momentum, (4) Kuramoto oscillators.

Results: The Bernoulli control maintained constant β ($\sigma < 10^{-6}$). All three bounded systems showed systematic β transitions: logistic growth from $\beta \approx 0.5$ (growth phase) to $\beta \rightarrow -\infty$ (saturation), gradient descent from $\beta \approx 0.52$ (exploration) to $\beta \approx 0.98$ (convergence), Kuramoto oscillators from $\beta \approx 1.7$ (growth) to $\beta \approx 20$ (saturation). Transitions are generic in bounded systems and absent in unbounded systems.

C.7 Negative Controls

Four non-recursive systems were tested: (1) Gaussian noise, (2) linear decay, (3) sinusoidal oscillation, (4) random walk. None produced false positive signals ($S > 1$ or $\alpha > 1$).

Result: 0/4 false positives. The framework correctly identifies the absence of recursive amplification in systems that do not satisfy the axioms.

C.8 Honest Failures and Boundary Cases

Two results fell short of expectations:

R* absolute value: 62% mean error in predicted vs actual crossover depth (functional form correct, proportionality constant off). This likely reflects sensitivity to the $\rho \approx 1$ approximation.

Newton's method: Local convergence near roots yields $\beta > 1.0$ (quadratic convergence), which falls outside the framework's domain $\beta \in [0, 1]$. The framework correctly identifies this as a boundary case where the Bernoulli ODE model does not apply.

These failures are reported for completeness and scientific integrity. They bound the framework's applicability and inform future refinement.

REFERENCES

- [1] Acharya, R. et al. [Google Quantum AI] (2024). Quantum error correction below the surface code threshold. *Nature*, 638, 920–926.
- [2] DeepSeek AI (2025). DeepSeek-R1: Incentivizing Reasoning Capability in LLMs. *arXiv:2501.12948*.
- [3] Sharma, A. & Chopra, P. (2025). The Sequential Edge: Inverse-Entropy Voting Beats Parallel Self-Consistency at Matched Compute. *arXiv:2511.02309*.
- [4] Morrell, M.C., Elliott, L., & Grier, D.G. (2026). Nonreciprocal wave-mediated interactions power a classical time crystal. *Physical Review Letters*, 136, 057201.
- [5] COGITATE Consortium (Melloni, L., Mudrik, L., Pitts, M., et al.) (2025). Adversarial testing of global neuronal workspace and integrated information theories of consciousness. *Nature*, 642, 133–142.
- [6] Snell, C. et al. (2024). Scaling LLM Test-Time Compute. *arXiv:2408.03314*.
- [7] Li, Z. et al. (2025). Revisiting the Test-Time Scaling of o1-like Models. *arXiv:2502.12215*.
- [8] Li, D. et al. (2025). S*: Test Time Scaling for Code Generation. *arXiv:2502.14382*.
- [9] Eastwood, M.D. (2024/2026). *Infinite Architects: Intelligence, Recursion, and the Creation of Everything*. ISBN: 978-1806056200. First manuscript December 2024.
- [10] Wilson, K.G. (1971). Renormalization Group and Critical Phenomena. *Physical Review B*, 4(9), 3174–3183.
- [11] Liu, T. et al. (2023). Photonic metamaterial analogue of a continuous time crystal. *Nature Physics*, 19, 986–991.
- [12] Raskatla, V. et al. (2024). Magnetically programmable classical time crystal. *Physical Review Letters*, 133, 136202.
- [13] Zheng, L. et al. (2025). Recurrency as a Common Denominator for Consciousness Theories. *PsyArXiv*. DOI: 10.31234/osf.io/wqznc.
- [14] Lamme, V.A.F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, 10(11), 494–501.
- [15] Kadanoff, L.P. (1966). Scaling laws for Ising models near T_c . *Physics Physique Fizika*, 2(6), 263–272.
- [16] Kleiber, M. (1932). Body size and metabolism. *Hilgardia*, 6, 315–353.
- [17] West, G.B., Brown, J.H. & Enquist, B.J. (1997). A General Model for the Origin of Allometric Scaling Laws in Biology. *Science*, 276(5309), 122–126.
- [18] Hyers, D.H. (1941). On the stability of the linear functional equation. *Proceedings of the National Academy of Sciences*, 27(4), 222–224.
- [19] Ulam, S.M. (1960). *A Collection of Mathematical Problems*. Interscience Publishers, New York.
- [20] Aguilar-Trigueros, C.A. et al. (2017). Branching out: Towards a trait-based understanding of fungal ecology. *Fungal Biology Reviews*, 31(1), 34–41. Metabolic scaling data: *ISME Journal*, 11, 2175–2180.
- [21] West, G.B., Brown, J.H. & Enquist, B.J. (2004). Growth models based on first principles or phenomenology? *Functional Ecology*, 18(2), 188–196.

- [22] Banavar, J.R., Maritan, A. & Rinaldo, A. (2010). Size and form in efficient transportation networks. *Nature*, 399(6732), 130–132. See also: Banavar, J.R. et al. (2010). A general basis for quarter-power scaling in animals. *Proceedings of the National Academy of Sciences*, 107(36), 15816–15820.
- [23] Demetrius, L. (2010). Quantum metabolism and allometric scaling relations in biology. *Proceedings of the Royal Society A*, 466(2124), 3543–3561. See also: Demetrius, L. & Tuszynski, J.A. (2010). Quantum metabolism explains the allometric scaling of metabolic rates. *Journal of the Royal Society Interface*, 7(44), 507–514.
- [24] Zhao, J. (2022). Universal growth scaling law determined by dimensionality. *arXiv:2206.08094*.
- [25] Bettencourt, L.M.A. (2013). The Origins of Scaling in Cities. *Science*, 340(6139), 1438–1441.

Companion Documents

- Eastwood, M.D. (2026). Paper III: The Alignment Scaling Problem. Version 10.0. First published 9 February 2026. OSF DOI: 10.17605/OSF.IO/6C5XB.
- Eastwood, M.D. (2026). Eden Protocol: Engineering Specification. Version 5.0. First published 22 February 2026. OSF DOI: 10.17605/OSF.IO/6C5XB.
- Eastwood, M.D. (2026). Eden Protocol: Philosophical Vision. Version 2.0. First published 22 February 2026. OSF DOI: 10.17605/OSF.IO/6C5XB.
- Eastwood, M.D. (2026). The ARC Principle: Experimental Validation of Compute Scaling Through Sequential Recursive Processing. Paper II. Version 12.0 (v5 cross-architecture results). First published 22 January 2026; v12 extended March 2026. Six models, 30 problems, bootstrap CIs. OSF DOI: 10.17605/OSF.IO/8FJMA.
- Eastwood, M.D. (2026). On the Origin of Scaling Laws. March 2026. Cross-domain accessible presentation of the ARC Principle, including the Completeness Theorem, Phase Diagram of Complexity, and Dimensional Ladder.

DECLARATION OF AI USE

The author used Claude (Anthropic), GPT (OpenAI), Gemini (Google), and DeepSeek AI to draft sections, refine clarity, and check mathematical consistency. The research question, theoretical framework, composition operator formalism, experimental predictions, and scientific judgment are human work. The author takes full responsibility for all claims, interpretations, errors, and conclusions.

THE ARC PRINCIPLE: FOUNDATIONAL PAPER - Version 5.0

© 2026 Michael Darius Eastwood. All Rights Reserved.

The predictions are specified. The falsification criteria are public. The data will decide.

Companion documents: On the Origin of Scaling Laws (accessible cross-domain paper) | Paper III (extended evidence, experimental protocols, blind prediction test results) | Eden Protocol (alignment architecture specification) | Philosophical Vision | Paper II | Paper V: The Stewardship Gene | Executive Summary.

Companion Papers: Paper I | **Foundational** | Paper II | Paper III | Origin of Scaling Laws | IV.a | IV.b | IV.c | IV.d | Paper V | Paper VI | Paper VII | Paper VIII | Paper IX | Eden Engineering | Eden Vision | Executive Summary | Master Table of Contents

Research hub: michaeldariuseastwood.com/research | OSF: 10.17605/OSF.IO/6C5XB | Copyright 2026 Michael Darius Eastwood