ARC PRINCIPLE - PAPER VII

Cauchy Unification

ARC/Cauchy Scaling Classification: A Structured Prediction Comparison Across 50 Domains in Five Evidence Tiers

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 3.0 | 24 March 2026 | First published 16 March 2026
v3.0: corrected space-filling overstatement (added conservation/optimisation constraint as third necessary condition); added 3/4 exponent debate acknowledgement (White & Seymour 2003; Glazier 2005, 2022); corrected cnidarian/jellyfish d=2 status to 'nearest-comparator match, not dimensional confirmation'; added new references
v2.0: expanded to 50 domains across five evidence tiers; replaced permissive R² criterion with proper AICc; honest miss analysis; tiered reporting
Supersedes: v2.0 (17 March 2026)
Extends: On the Origin of Scaling Laws | Foundational Paper

Validation code: GitHub (experiment) | Manifest
Research hub: michaeldariuseastwood.com/research

ABSTRACT

Cauchy's four functional equations - additive, multiplicative, exponential, and logarithmic - are the only continuous solutions to their respective composition constraints. We argue that this 200-year-old mathematical result has a physically testable consequence: under the stated axioms, it constrains scaling laws to one of three functional families (power law, exponential, or saturation curve). The family is determined by the composition operator of the underlying recursive amplification process.

The novel contribution is not Cauchy's mathematics itself, but the claim that Cauchy-type functional equations have a physically testable consequence for scaling-law classification across domains.

We present a structured prediction comparison across 50 real-world systems in five evidence tiers. For each empirical domain, the composition operator was classified from known physics before fitting, the scaling family was predicted from the operator class, published or standard reference data was loaded, and the prediction was tested by independent model fitting using AICc-based model selection.

Primary result: 19 out of 25 empirical curve-fit domains confirm the Cauchy-predicted family under strict AICc model selection ($p = 1.56 \times 10^{-5}$, binomial test against $p = 1/3$ chance). This is roughly 1 in 64,000 - about 300 times stronger than the conventional $p < 0.05$ threshold and equivalent to approximately 4.2σ. **Secondary results:** a baseline-20 rerun yields 15/20 ($p = 1.67 \times 10^{-4}$); published metabolic exponents match the nearest ARC/Cauchy comparator in 13/13 direct cases; analytic identities confirm 6/6. These tiers are reported separately and are never combined into a single blended number.

This result is strong enough to publish as exploratory evidence but not strong enough to claim as proven. No locked pre-registration artefact exists. Independent replication is needed.

Keywords: Cauchy functional equations, scaling laws, power laws, universality, cross-domain, recursive composition, AICc model selection, metabolic scaling, neural scaling, ARC Principle

1. The Question

Why do scaling laws exist at all? Why does a mouse's heart beat faster than an elephant's by the same mathematical relationship that governs how cities grow, how solar panels get cheaper, and how neural networks improve with size?

The standard answer is domain-specific. Biologists invoke fractal vascular networks. Physicists invoke renormalisation group flow. Economists invoke learning-by-doing. Each explanation works within its domain but offers no reason why the same mathematical forms recur across domains that share no physical mechanism.

This paper proposes and tests a domain-independent answer: the form of every scaling law is determined by the composition operator of the underlying recursive process, and Cauchy's functional equations constrain which forms are possible.

2. The Theory

2.1 Cauchy's four equations

Augustin-Louis Cauchy established in the early nineteenth century that four functional equations have unique continuous solutions:

Additive: $f(x+y) = f(x) + f(y)$ has solution $f(x) = cx$

Multiplicative: $f(xy) = f(x) \cdot f(y)$ has solution $f(x) = x^c$

Exponential: $f(x+y) = f(x) \cdot f(y)$ has solution $f(x) = a^x$

Logarithmic: $f(xy) = f(x) + f(y)$ has solution $f(x) = c \log x$

These are not four independent results. They are four faces of one constraint: the requirement that a function be compatible with a binary operation. The additive and multiplicative equations govern power laws. The exponential equation governs exponential growth and decay. The logarithmic equation governs saturation and diminishing returns.

2.2 The prediction

Consider any system where a quantity U grows with recursive depth R. The system's *composition operator* determines how consecutive applications combine:

- **Multiplicative composition** ($U(R_1) \cdot U(R_2) \sim U(R_1 + R_2)$): Cauchy's multiplicative equation forces $U \propto R^c$ - a power law.
- **Additive composition** ($U(R_1) + U(R_2) \sim U(R_1 + R_2)$): Cauchy's exponential equation forces $U \propto a^R$ - an exponential.
- **Bounded composition** (output constrained to $[0, K]$): the compatible forms are saturation curves (logistic, Hill, Michaelis-Menten, hyperbolic). A formal proof of exhaustiveness for the bounded case is not provided here and remains an open problem.

> **THE PREDICTION**
>
> If you know the composition operator of a system, you can predict the functional form of its scaling law before seeing any data. The prediction follows from Cauchy's equations with no free parameters and no curve fitting.

3. Prediction Protocol

For each domain:

1. **Classify the composition operator from known physics** (before fitting). Determine whether the system's recursive amplification process combines multiplicatively, additively, or with a bound.
2. **Predict the scaling family** from the operator class using Cauchy's equations.
3. **Load published or standard reference data.**
4. **Independently fit candidate models** (power law, exponential, logistic, Hill, Michaelis-Menten, saturation-exponential) to the data using nonlinear least-squares regression.

5. **Select the best model by AICc** (corrected Akaike information criterion) and compare the prediction against the selected model's family.

> **STRUCTURED PREDICTION COMPARISON - NOT A BLIND TEST**
>
> The operator classification is made before the fitting procedure runs. However, the predictions, data, and fitting logic all reside in a single author-written script. There is no locked pre-registration artefact. This test should therefore be understood as a structured prediction comparison, not a formally pre-registered blind trial. A pre-registered replication with independent operator classification is planned.

3.1 Statistical method

The v2 test uses proper AICc (corrected Akaike information criterion) computed from residual sum of squares, sample size, and parameter count - not the R^2-based approximation used in the legacy v1 test. The v1 approximation ($\mathrm{AIC} \approx n \ln(1 - R^2) + 2k$) was adequate as a heuristic but should not have been described as strict AIC in the standard statistical sense. The v2 rewrite corrects this.

The primary significance test is a one-sided binomial test: under the null hypothesis that each domain's family is chosen uniformly at random from three families, the probability of 19 or more matches in 25 trials is $p = 1.56 \times 10^{-5}$.

4. Evidence Tiers

The 50-domain suite comprises five evidence tiers of decreasing strength. These tiers are never combined into a single blended number.

TIER	DESCRIPTION	COUNT	RESULT	SIGNIFICANCE
PRIMARY	Empirical curve-fit (AICc)	25	19/25	$p = 1.56 \times 10^{-5}$
SECONDARY	Baseline-20 rerun (AICc)	20	15/20	$p = 1.67 \times 10^{-4}$
SECONDARY	Published exponents (direct)	13	13/13	Nearest comparator match
NOTED	Published exponents (provisional)	6	3/6	Contested organisms
NOTED	Analytic identities	6	6/6	Tautological confirmation

5. Primary Result: 25 Empirical Curve-Fit Domains

#	DOMAIN	OPERATOR	PREDICTED	AICC BEST	MATCH
1	Kleiber's Law (Metabolic Scaling)	Multiplicative	Power law	Saturation-exp	No
2	Urban Scaling (GDP vs Population)	Multiplicative	Power law	Power law	Yes
3	Species-Area (Galapagos)	Multiplicative	Power law	Hill	No
4	Wright's Law (Solar PV)	Multiplicative	Power law	Power law	Yes
5	Heap's Law (Vocabulary)	Multiplicative	Power law	Power law	Yes
6	Zipf's Law (Word Frequency)	Multiplicative	Power law	Michaelis-Menten	No
7	Learning Curve (Cigar Rolling)	Multiplicative	Power law	Power law	Yes
8	Moore's Law (Transistor Count)	Additive	Exponential	Exponential	Yes
9	Radioactive Decay (P-32)	Additive	Exponential	Exponential	Yes
10	Gutenberg-Richter (Earthquakes)	Additive	Exponential	Exponential	Yes
11	Bacterial Growth (E. coli)	Bounded	Bounded	Logistic	Yes
12	O_2-Hemoglobin Curve	Bounded	Bounded	Hill	Yes
13	Epidemic SIR (Ebola 2014)	Bounded	Bounded	Hill	Yes
14	Amdahl's Law (CPU Scaling)	Bounded	Bounded	Michaelis-Menten	Yes
15	Muscle Force-Velocity (Hill 1938)	Bounded	Bounded	Exponential	No
16	Facebook MAU Growth	Bounded	Bounded	Logistic	Yes
17	Brownian Diffusion (MSD)	Multiplicative	Power law	Power law	Yes
18	Horton's Law (Stream Numbers)	Additive	Exponential	Exponential	Yes
19	Neural Scaling Laws (LLM Loss)	Multiplicative	Power law	Power law	Yes
20	Time Crystal Order (Rydberg Gas)	Bounded	Bounded	Power law	No
21	Stellar Mass-Luminosity	Multiplicative	Power law	Logistic	No
22	Heart Rate vs Body Mass	Multiplicative	Power law	Power law	Yes
23	Rent's Rule (VLSI Pin Count)	Multiplicative	Power law	Power law	Yes
24	Taylor's Power Law	Multiplicative	Power law	Power law	Yes
25	Hack's Law (Stream Length)	Multiplicative	Power law	Power law	Yes

PRIMARY RESULT: 19/25 EMPIRICAL DOMAINS CONFIRMED ($p = 1.56 \times 10^{-5}$)

Under strict AICc-based model selection, the Cauchy-predicted family is confirmed in 19 of 25 empirical curve-fit domains. The binomial probability of this result under random assignment to three families is $p = 1.56 \times 10^{-5}$ - roughly 1 in 64,000. This is approximately 300 times stronger than $p < 0.05$ and equivalent to about 4.2σ.

This is strong enough to publish as exploratory evidence. It is not strong enough to claim as proven. Independent replication with pre-registered operator classification is the necessary next step.

5.1 Baseline-20 rerun

As context, the original 20 baseline domains were rerun with the new proper AICc fitter. The result is 15/20 ($p = 1.67 \times 10^{-4}$). This is not directly comparable to the v1 result of 14/20, because the fitter changed (R^2-based AIC approximation to proper AICc) and some model selections shifted accordingly. The baseline-20 rerun is reported for transparency, not as evidence that the new fitter is more or less conservative than the old one.

6. The Six Empirical Misses

Six of 25 empirical domains produced a best-fit model outside the predicted family. Each miss has a specific, identifiable explanation. Honest analysis follows.

#	DOMAIN	PREDICTED	AICC BEST	EXPLANATION
1	Kleiber's Law	Power law	Saturation-exp	Classic contested case. Kleiber's original 1932 dataset has only 13 mammals. A bounded model can capture slight upper curvature at extreme body masses. The power-law exponent debate (0.67 vs 0.75) has persisted for decades.
3	Species-Area (Galapagos)	Power law	Hill	Small dataset (30 Galapagos islands) with high scatter. Several islands have very small areas with disproportionate species counts. The Hill function can absorb this scatter with its extra parameter.
6	Zipf's Law	Power law	Michaelis-Menten	Finite-corpus truncation. The Brown Corpus has a hard upper bound on word frequency (69,971 for 'the'). At low ranks, frequencies saturate against this ceiling, pulling AICc towards a bounded model. Zipf's law is a power law in the asymptotic limit; finite corpora violate this limit.
15	Muscle Force-Velocity	Bounded	Exponential	The Hill 1938 force-velocity data shows a monotonic decrease from maximum isometric force to zero at maximum velocity. The data does not reach or demonstrate an asymptotic plateau - it terminates where force reaches zero. An exponential decay fits this truncated range better than a saturation curve. The bounded prediction requires data spanning the full approach to saturation.
20	Time Crystal Order	Bounded	Power law	Only 4 data points. With so few observations, AICc penalises the extra parameters of saturation models. A power law, with fewer free parameters, wins on parsimony. The bounded prediction may be correct but is untestable with this dataset.
21	Stellar Mass-Luminosity	Power law	Logistic	The mass-luminosity relation shows significant scatter at high stellar masses where radiation pressure, convective instability, and different fusion pathways create deviations from a simple power law. A logistic model absorbs this scatter by fitting an apparent upper rollover.

In summary: three misses involve small or truncated datasets (species-area, time crystal, Zipf); one involves data that does not span the full saturation range (muscle force-velocity); one involves a classically contested relationship (Kleiber); and one involves astrophysical complexity at extreme values (stellar mass-luminosity). None represents a case where the predicted family is fundamentally wrong in principle - but neither can any miss be dismissed. These are genuine failures that an independent replication must address.

7. Secondary Tier: Published Exponents

7.1 Direct published exponents (13/13)

Thirteen taxa with well-established metabolic scaling exponents were tested against the ARC/Cauchy dimensional prediction. The prediction requires three conditions: (a) multiplicative composition (Cauchy's equation constrains the form to a power law), (b) d-dimensional space-filling transport geometry, and (c) a conservation or optimisation constraint on resource flow (energy minimisation in

West et al.; supply-demand balance in Banavar et al.; steady-state energy balance in Demetrius). Under all three conditions, the exponent is constrained to $d/(d+1)$. Neither Cauchy alone, nor space-filling alone, is sufficient. Organisms with three-dimensional vascular transport and conserved flow should scale as $M^{3/4}$, those with two-dimensional transport as $M^{2/3}$, and those with one-dimensional filamentous transport as $M^{1/2}$.

NOTE ON THE 3/4 EXPONENT

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 & Seymour 2003; Glazier 2005, 2022). The largest modern dataset (619 species) gives a maximum-likelihood estimate of 0.687 (95% CI 0.674-0.701), which excludes 0.750. The $d/(d+1)$ prediction of 0.750 for $d=3$ matches the upper end of the empirical 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.

#	TAXON	PREDICTED	OBSERVED	NEAREST	CI INCLUDES?
26	Mammals	0.750	0.737	$d=3$	Yes
27	Birds	0.750	0.720	$d=3$	Yes
28	Fish	0.750	0.800	$d=3$	Yes
29	Reptiles	0.750	0.760	$d=3$	Yes
30	Insects	0.750	0.750	$d=3$	Yes
31	Amphibians	0.750	0.740	$d=3$	Yes
32	Crustaceans	0.750	0.730	$d=3$	Yes
33	Jellyfish	0.667	0.680	$d=2$	Yes
34	Cnidarians	0.667	0.700	$d=2$	Yes
35	Ctenophores	0.667	0.660	$d=2$	Yes
36	Ectomycorrhizal fungi	0.500	0.580	$d=1$	Yes
37	Marine fungi	0.500	0.530	$d=1$	Yes
38	Saprotrophic fungi (20°C)	0.500	0.530	$d=1$	Yes

All 13 published exponents fall nearest to the predicted dimensional comparator ($d/(d+1)$), and all 13 confidence intervals include the predicted value. This is strong corroborating evidence, but it tests a different claim (specific exponent) than the primary result (functional family), so it is reported separately.

CAVEAT ON D=2 BIOLOGICAL ENTRIES

The jellyfish, cnidarian, and ctenophore entries (rows 33-35) are classified as nearest-comparator matches to $d = 2$, meaning their observed exponents fall closer to 2/3 than to either 1/2 or 3/4. However, the $d = 2$ biological prediction remains untested in the strong sense because no known organism possesses a genuinely two-dimensional hierarchical space-filling transport network with conserved flow. These organisms lack the vascular architecture that the $d/(d+1)$ derivation assumes. The $d = 2$ confirmation exists in cosmology (Friedmann matter-era solution, exact) and physics (percolation, fragmentation), not in biology. These rows should be understood as empirical proximity to the $d = 2$ comparator, not as dimensional confirmation.

7.2 Provisional published exponents (3/6)

Six additional taxa with contested or proxy-based dimensional assignments were tested. Only 3 of 6 match the nearest predicted comparator. The misses are flatworms (disputed 2D vs 3D transport), bryozoans (colonial organisms with anomalous scaling near 1.0), and glass eels (elongated but retaining 3D vascular transport). These are noted for completeness but carry low evidential weight due to the contested dimensional assignments.

7.3 Analytic identities (6/6)

Six domains defined by exact analytic formulae ($E = mc^2$, hydrogen energy levels, Arrhenius kinetics, Michaelis-Menten kinetics, Friedmann matter-era expansion, de Sitter dark-energy expansion) were included as sanity checks. All 6 confirm. These are tautological - fitting a curve to data generated by an exact formula must recover that formula - and carry no independent evidential weight. They confirm only that the fitter works correctly.

8. Significance Framing

The primary result ($p = 1.56 \times 10^{-5}$) deserves honest contextualisation.

- **In absolute terms:** roughly 1 in 64,000. If 64,000 researchers each picked scaling families at random for 25 domains, on average one would match 19 or more by chance.
- **Relative to convention:** about 300 times stronger than $p < 0.05$. Equivalent to approximately 4.2σ in Gaussian terms.
- **What it supports:** the result is strong enough to publish as exploratory evidence that Cauchy-type composition constraints predict scaling families at well-above-chance rates.
- **What it does not support:** it is not strong enough to claim the theory is proven. The particle-physics discovery threshold is 5σ ($p \approx 3 \times 10^{-7}$). The test lacks formal pre-registration. The operator classifications were made by the author, not an independent panel.
- **What is needed:** independent replication with pre-registered operator classification, ideally by researchers with no stake in the outcome.

9. Group Analysis

9.1 Multiplicative composition: power laws

Of 14 empirical domains classified as multiplicative, 11 produce power-law scaling as their AICc-best model. These span metabolic allometry (heart rate vs body mass), urban economics (Bettencourt 2007), technology learning curves (Wright's law), computational linguistics (Heap's law), skill acquisition (Crossman 1959), Brownian motion (Catipovic 2013), neural network scaling (Kaplan 2020), VLSI design (Rent's rule), population ecology (Taylor's power law), and fluvial geomorphology (Hack's law).

Three multiplicative domains miss: Kleiber's law (bounded model edges out power law on a 13-point dataset), species-area (high scatter on small island data), and stellar mass-luminosity (scatter at high mass). Zipf's law also misses due to finite-corpus truncation.

9.2 Additive composition: exponentials

All four domains with additive composition operators produce exponential scaling as their AICc-best model: Moore's law (transistor counts), radioactive decay (P-32), the Gutenberg-Richter earthquake frequency law, and Horton's law (stream numbers). This is a clean 4/4.

9.3 Bounded composition: saturation curves

Of seven domains with bounded composition operators, five produce bounded (saturation) models as AICc-best: bacterial growth (logistic), hemoglobin binding (Hill), Ebola epidemic (Hill), Amdahl's law (Michaelis-Menten), and Facebook growth (logistic). Two miss: muscle force-velocity (exponential decay fits the truncated range better) and time crystal order (only 4 data points).

10. Data Provenance

10.1 Provenance caveat

Three domains use theoretical or reference values rather than independent empirical measurements: radioactive decay (NIST reference half-life for P-32), Amdahl's law (theoretical formula with assumed 10% serial fraction), and Moore's law (Wikipedia-sourced transistor counts from manufacturer specifications). These confirm the framework trivially and should be weighted accordingly.

10.2 Citations for all 25 empirical domains

[1] Kleiber, M. (1932). Body size and metabolism. *Hilgardia* 6:315-353. Reproduced in Dodds, P. et al. (2001) *J. Theor. Biol.* 209:9-27.

[2] Bettencourt, L. et al. (2007). Growth, innovation, scaling, and the pace of life in cities. *PNAS* 104(17):7301-7306. Data: BEA + Census 2006, 25 largest US MSAs.

[3] Johnson, M. & Raven, P. (1973). Species number and endemism: the Galapagos archipelago revisited. *Science* 179:893-895. Data: R *faraway::gala* dataset, 30 islands.

[4] Nemet, G. (2009). Demand-pull, technology-push. *Energy Policy*; Farmer, J.D. & Lafond, F. (2016) *Research Policy*; IRENA; Our World in Data.

[5] Manning, C. et al. (2008). *Introduction to Information Retrieval*. Reuters RCV1 corpus.

[6] Kucera, H. & Francis, W. (1967). *Computational Analysis of Present-Day American English*. Brown Corpus.

[7] Crossman, E. (1959). A theory of the acquisition of speed-skill. *Ergonomics* 2(2):153-166. Reanalysis: Newell, A. & Rosenbloom, P. (1981).

[8] Wikipedia: Transistor count. Public vendor specifications (Intel, AMD, TSMC).

[9] NIST Nuclear Data Center. P-32 half-life: 14.29 days. Exact theoretical values.

[10] USGS Earthquake Hazards Program; IRIS compilation. Global seismicity catalogue.

[11] Sezonov, G. et al. (2007). *E. coli* physiology in Luria-Bertani broth. *J. Bacteriol.* 189:8746-8749.

[12] Severinghaus, J. (1979). Simple, accurate equations for human blood O_2 dissociation computations. *J. Appl. Physiol.* 46(3):599-602.

[13] WHO/CDC Ebola Situation Reports (2014-2015); NEJM Ebola Response Team.

[14] Amdahl, G. (1967). Validity of the single processor approach to achieving large scale computing capabilities. *AFIPS*. See also Hill, M. & Marty, M. (2008) *IEEE Computer*.

[15] Hill, A. (1938). The heat of shortening and the dynamic constants of muscle. *Proc. Royal Soc. B* 126:136-195.

[16] Meta SEC filings; Zhang, C., Liu, J. & Xu, Y. (2015). *J. Comp. Sci. Tech.*

[17] Catipovic, M. et al. (2013). Improving the quantification of Brownian motion. *Am. J. Physics* 81:485.

[18] Singh, V. et al. (2021). Hydrological stream order analysis. *Applied Water Science* 11:151.

[19] Kaplan, J. et al. (2020). Scaling Laws for Neural Language Models. arXiv:2001.08361.

[20] Shen, Y. et al. (2025). Time crystal order in a Rydberg gas. *Nature Communications*. See also Kongkhambut, P. et al. (2024) *Science* 386(6720):670-674.

[21] Torres, G. et al. (2010). Accurate masses and radii of normal stars. *Astron. Astrophys. Rev.* 18:67-126. Eker, Z. et al. (2018) *MNRAS* 479:5491. IAU 2015 nominal solar values.

[22] Stahl, W. (1967). Scaling of respiratory variables in mammals. *J. Appl. Physiol.* 22:453. Calder, W. (1984) *Size, Function, and Life History*. Schmidt-Nielsen, K. (1984) *Scaling: Why Is Animal Size So Important?*

[23] Landman, B. & Russo, R. (1971). On a pin versus block relationship for partitions of logic graphs. *IEEE Trans. Comp.* C-20:1469. Christie, P. (2000); Stroobandt, D. (2001).

[24] Taylor, L. (1961). Aggregation, variance and the mean. *Nature* 189:732-735. Taylor, L. & Woiwod, I. (1980) *J. Anim. Ecol.* 49:879.

[25] Hack, J. (1957). Studies of longitudinal stream profiles in Virginia and Maryland. *USGS Professional Paper* 294-B. Mueller, J. (1973); Rigon, R. et al. (1996) *Water Resour. Res.* 32:3367.

[26] 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.

[27] Banavar, J. R., Maritan, A. & Rinaldo, A. (1999). Size and form in efficient transportation networks. *Nature* 399(6732):130-132.

[28] Banavar, J. R., Moses, M. E., Brown, J. H., Damuth, J., Rinaldo, A., Sibly, R. M. & Maritan, A. (2010). A general basis for quarter-power scaling in animals. *Proceedings of the National Academy of Sciences* 107(36):15816-15820.

[29] Demetrius, L. (2003). Quantum statistics and allometric scaling of organisms. *Physica A* 322:477-490.

[29b] Demetrius, L. (2006). The origin of allometric scaling laws in biology. *Journal of Theoretical Biology* 243(4):455-467.

[29c] 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.

[30] Zhao, J. (2022). Universal growth scaling law determined by dimensionality. *arXiv:2206.08094*.

[31] Bettencourt, L. M. A. (2013). The origins of scaling in cities. *Science* 340(6139):1438-1441.

[32] White, C. R. & Seymour, R. S. (2003). Mammalian basal metabolic rate is proportional to body mass$^{2/3}$. *Proceedings of the National Academy of Sciences* 100(7):4046-4049.

[33] Glazier, D. S. (2005). Beyond the '3/4-power law': variation in the intra- and interspecific scaling of metabolic rate in animals. *Biological Reviews* 80(4):611-662.

[34] Glazier, D. S. (2022). Variable metabolic scaling breaks the law: a large-scale data synthesis. *Biological Reviews* 97(3):1056-1067.

[35] Kolokotrones, T., Savage, V., Deeds, E. J. & Fontana, W. (2010). Curvature in metabolic scaling. *Nature* 464(7289):753-756.

11. Prior Art

An exhaustive prior-art investigation found no previous work that invokes all four Cauchy functional equations as a unifying cross-domain principle constraining the forms of scaling laws.

The closest existing work includes:

- **Bingham & Ostaszewski (2015):** connect Cauchy's equation to regular variation theory. Pure mathematics. No cross-domain physical science application.
- **West, Brown & Enquist (1997):** derive a 3/4 metabolic exponent from fractal vascular networks with energy-minimised flow. Domain-specific. Does not predict other domains. Note that the empirical exponent is debated (see Warning box in Section 7.1).
- **Banavar, Maritan & Rinaldo (1999, 2010):** derive the same $d/(d+1)$ form from geometric constraints on transportation networks combined with supply-demand conservation, without requiring fractal geometry. Independently confirmed the quarter-power exponent from different starting assumptions. Banavar et al. (2010) showed that the exponent ranges from 2/3 to 3/4 depending on how velocity scales, with 3/4 as the optimised maximum - demonstrating that space-filling geometry alone is insufficient without a conservation or optimisation constraint.
- **Demetrius (2003, 2006) and Demetrius & Tuszynski (2010):** derive equivalent scaling from quantum statistical mechanics, applying the Debye model of thermal properties to coupled energy-transducing oscillator networks in d-dimensional space. A completely independent route to the same exponent from different physics (quantum oscillator coupling rather than network geometry). The convergence of this derivation with West's network geometry derivation is explained by Cauchy's multiplicative functional equation applied to multiplicative composition in d-dimensional space with a conservation constraint (see companion paper, On the Origin of Scaling Laws v3.1).
- **Bettencourt (2013):** derives city scaling exponents from network dimensionality, extending the $d/(d+1)$ form beyond biology to urban systems.
- **Zhao (2022):** derives the formula as a universal growth scaling law from network optimisation, providing another independent derivation.
- **He & Chen (2003):** derive metabolic scaling exponents from optimal vascular network geometry with flow conservation constraints.
- **He & Zhang (2004):** extend the derivation to heterogeneous branching networks, recovering the $d/(d+1)$ form under energy-minimisation assumptions.
- **Maino et al. (2014):** derive metabolic scaling from thermodynamic first principles applied to ontogenetic growth, independently obtaining the same exponent family.
- **Geoffrey West, *Scale* (2017):** discusses scaling across biology and cities. No mention of Cauchy's functional equations.
- **Aczél (1966):** standard reference for all four Cauchy equations. Applications to information theory but not to cross-domain scaling.

The gap is genuine. The $d/(d+1)$ formula has been independently derived by at least seven research groups through different mathematical frameworks - fractal networks, geometric constraints with supply-demand balance, quantum metabolism, urban scaling, network optimisation, optimal vascular geometry, and thermodynamic first principles. Every known derivation requires three conditions: multiplicative composition (which Cauchy constrains to the power-law family), d-dimensional space-filling geometry, and a conservation or optimisation constraint on resource flow. The mathematical theorems are 200 years old. The biological data is 90 years old. The novel contribution is not the

$d/(d+1)$ formula itself (which is well established), nor Cauchy's mathematics, but the claim that Cauchy-type functional equations explain *why* these independent derivations converge on the same form, and that the Cauchy constraint has a physically testable consequence for scaling-law classification across domains.

12. Limitations

- **No pre-registration artefact.** The predictions, data arrays, and fitting logic all reside in a single author-written script. There is no locked pre-registration artefact. This test should be understood as a structured prediction comparison, not a formally pre-registered blind trial.
- **Operator classification by author.** The composition operator classification is done by the author, not by an independent assessor. A pre-registered replication with independent operator classification would be more rigorous.
- **Bounded composition exhaustiveness is open.** The proof that only saturation-family curves are compatible with bounded composition is not provided and remains an open problem. The bounded category currently groups logistic, Hill, Michaelis-Menten, and hyperbolic forms by empirical convention rather than mathematical proof.
- **Legacy AIC was an approximation.** The v1 test used an R^2-based AIC approximation ($\text{AIC} \approx n \ln(1-R^2) + 2k$), not a full likelihood-based nonlinear AIC. The v2 test corrects this with proper AICc computed from residual sum of squares.
- **Three datasets use reference values.** Radioactive decay (NIST half-life), Amdahl's law (theoretical formula), and Moore's law (manufacturer specs) use theoretical or reference values rather than independently collected empirical measurements.
- **Family prediction only.** The framework predicts the functional family but not the specific exponent. Predicting exponents requires additional structure (dimensionality, constraint type) developed in the companion Foundational paper.
- **Six empirical misses.** The misses are not dismissible. They identify real limits of the framework: small datasets, truncated ranges, high-scatter regimes, and domains where competing physical effects (radiation pressure, corpus boundaries) perturb the ideal prediction.
- **Operator classification is descriptive, not explanatory.** The framework does not explain *why* a particular domain has a multiplicative rather than additive composition operator. It takes the operator as given and predicts the consequence.
- **Baseline-20 rerun is not cleanly comparable to v1.** The fitter changed from R^2-based AIC approximation to proper AICc. The model set expanded. A direct comparison of 14/20 (v1) to 15/20 (v2 baseline rerun) is not meaningful because the selection criterion changed.
- **Publication bias risk.** The 5 new expansion domains were selected after the baseline-20 test. While they were chosen to span new physical regimes (stellar astrophysics, allometric heart rate, VLSI, ecology, geomorphology), the selection was not randomised.
- **Finite model set.** The fitter compares a finite set of candidate models (power law, exponential, logistic, Hill, Michaelis-Menten, saturation-exponential). The true best model for any domain may not be in this set.
- **Sample size.** Twenty-five empirical domains is a reasonable exploratory sample but not definitive. Expansion to 100+ domains with pre-registered selection criteria would strengthen or weaken the result.

13. Implications

If the framework holds under independent pre-registered replication, it would suggest:

1. The forms of scaling laws are not coincidences but mathematical consequences of composition structure, under the stated axioms.
2. The same principle that governs how a mouse's heart beats governs how neural networks improve with scale.
3. The renormalisation group's power-law predictions and Cauchy's multiplicative equation may be related theorems - a connection that, to the author's knowledge, has not been formally articulated. A formal proof of this equivalence remains open.
4. Any new system's scaling family can be predicted from first principles by classifying its composition operator, without any data fitting.
5. Two hundred years of domain-specific explanations for scaling laws - fractal networks, learning curves, critical phenomena - may be special cases of a single mathematical constraint that Cauchy identified in 1821.

These implications are conditional. The current evidence is exploratory, not definitive.

14. Conclusion

Plant an acorn and, given centuries, you get an oak. But plant that oak's acorn, and its acorn, recursively across millennia, and you get a forest that shapes the climate of continents. What we plant in these systems will compound across scales we cannot imagine. The seed determines the forest.

Cauchy proved 200 years ago that the seed also determines the *form* of the forest. Multiplicative seeds produce power laws. Additive seeds produce exponentials. Bounded seeds produce saturation. Twenty-five empirical domains tested against this prediction yield 19 confirmations at $p = 1.56 \times 10^{-5}$. Thirteen published metabolic exponents fall where the dimensional theory predicts.

The mathematics was always there. We just had not read it as a prediction about the physical world. Whether that reading holds under independent scrutiny is now a question for the scientific community to answer.

Raise AI with care.

15. Reproducibility

The complete validation suite - including all 50 domain definitions, the canonical manifest, and the AICc-based fitting code - is available at:

- Experiment code: github.com/.../experiments/cauchy-unification__Paper-VII
- 50-domain manifest: canonical_50_domain_manifest.json

The script requires Python 3.10+, numpy, and scipy. All data is embedded in the manifest with full provenance citations. Output is deterministic.

Research hub: michaeldariuseastwood.com/research | OSF: 10.17605/OSF.IO/6C5XB |